The Julia Language

The Julia Project

August 17, 2017

Contents

Part I

Home

Part II

Julia Documentation

Chapter 1

Manual

- Introduction
- Getting Started
- Variables
- Integers and Floating-Point Numbers
- Mathematical Operations and Elementary Functions
- Complex and Rational Numbers
- Strings
- Functions
- Control Flow
- Scope of Variables
- Types
- Methods
- Constructors
- Conversion and Promotion
- Interfaces
- Modules
- Documentation
- Metaprogramming
- Multi-dimensional Arrays
- Linear Algebra
- Networking and Streams
- Parallel Computing
- Date and DateTime

- Running External Programs

- Calling C and Fortran Code

- Handling Operating System Variation

- Environment Variables

- Interacting With Julia

- Embedding Julia

- Packages

- Profiling

- Stack Traces

- Performance Tips

- Workflow Tips

- Style Guide

- Frequently Asked Questions

- Noteworthy Differences from other Languages

- Unicode Input

Chapter 2

Standard Library

- Essentials
- Collections and Data Structures
- Mathematics
- Numbers
- Strings
- Arrays
- Tasks and Parallel Computing
- Linear Algebra
- Constants
- Filesystem
- I/O and Network
- Punctuation
- Sorting and Related Functions
- Package Manager Functions
- Dates and Time
- Iteration utilities
- Unit Testing
- C Interface
- C Standard Library
- Dynamic Linker
- Profiling
- StackTraces
- SIMD Support

Chapter 3

Developer Documentation

Part III

Manual

Chapter 4

Introduction

Scientific computing has traditionally required the highest performance, yet domain experts have largely moved to slower dynamic languages for daily work. We believe there are many good reasons to prefer dynamic languages for these applications, and we do not expect their use to diminish. Fortunately, modern language design and compiler techniques make it possible to mostly eliminate the performance trade-off and provide a single environment productive enough for prototyping and efficient enough for deploying performance-intensive applications. The Julia programming language fills this role: it is a flexible dynamic language, appropriate for scientific and numerical computing, with performance comparable to traditional statically-typed languages.

Because Julia's compiler is different from the interpreters used for languages like Python or R, you may find that Julia's performance is unintuitive at first. If you find that something is slow, we highly recommend reading through the Performance Tips section before trying anything else. Once you understand how Julia works, it's easy to write code that's nearly as fast as C.

Julia features optional typing, multiple dispatch, and good performance, achieved using type inference and just-in-time (JIT) compilation, implemented using LLVM. It is multi-paradigm, combining features of imperative, functional, and object-oriented programming. Julia provides ease and expressiveness for high-level numerical computing, in the same way as languages such as R, MATLAB, and Python, but also supports general programming. To achieve this, Julia builds upon the lineage of mathematical programming languages, but also borrows much from popular dynamic languages, including Lisp, Perl, Python, Lua, and Ruby.

The most significant departures of Julia from typical dynamic languages are:

- The core language imposes very little; the standard library is written in Julia itself, including primitive operations like integer arithmetic

- A rich language of types for constructing and describing objects, that can also optionally be used to make type declarations

- The ability to define function behavior across many combinations of argument types via multiple dispatch

- Automatic generation of efficient, specialized code for different argument types

- Good performance, approaching that of statically-compiled languages like C

Although one sometimes speaks of dynamic languages as being "typeless", they are definitely not: every object, whether primitive or user-defined, has a type. The lack of type declarations in most dynamic languages, however, means that one cannot instruct the compiler about the types of values, and often cannot explicitly talk about types at all. In static languages, on the other hand, while one can – and usually must – annotate types for the compiler, types exist only at compile time and cannot be manipulated or expressed at run time. In Julia, types are themselves run-time objects, and can also be used to convey information to the compiler.

While the casual programmer need not explicitly use types or multiple dispatch, they are the core unifying features of Julia: functions are defined on different combinations of argument types, and applied by dispatching to the most specific matching definition. This model is a good fit for mathematical programming, where it is unnatural for the first argument to "own" an operation as in traditional object-oriented dispatch. Operators are just functions with special notation – to extend addition to new user-defined data types, you define new methods for the + function. Existing code then seamlessly applies to the new data types.

Partly because of run-time type inference (augmented by optional type annotations), and partly because of a strong focus on performance from the inception of the project, Julia's computational efficiency exceeds that of other dynamic languages, and even rivals that of statically-compiled languages. For large scale numerical problems, speed always has been, continues to be, and probably always will be crucial: the amount of data being processed has easily kept pace with Moore's Law over the past decades.

Julia aims to create an unprecedented combination of ease-of-use, power, and efficiency in a single language. In addition to the above, some advantages of Julia over comparable systems include:

- Free and open source (MIT licensed)

- User-defined types are as fast and compact as built-ins

- No need to vectorize code for performance; devectorized code is fast

- Designed for parallelism and distributed computation

- Lightweight "green" threading (coroutines)

- Unobtrusive yet powerful type system

- Elegant and extensible conversions and promotions for numeric and other types

- Efficient support for Unicode, including but not limited to UTF-8

- Call C functions directly (no wrappers or special APIs needed)

- Powerful shell-like capabilities for managing other processes

- Lisp-like macros and other metaprogramming facilities

Chapter 5

Getting Started

Julia installation is straightforward, whether using precompiled binaries or compiling from source. Download and install Julia by following the instructions at https://julialang.org/downloads/.

The easiest way to learn and experiment with Julia is by starting an interactive session (also known as a read-eval-print loop or "repl") by double-clicking the Julia executable or running `julia` from the command line:

```
$ julia
               _
   _       _ _(_)_     |  A fresh approach to technical computing
  (_)     | (_) (_)    |  Documentation: https://docs.julialang.org
   _ _   _| |_  __ _   |  Type "?help" for help.
  | | | | | | |/ _` |  |
  | | |_| | | | (_| |  |  Version 0.5.0-dev+2440 (2016-02-01 02:22 UTC)
 _/ |\__'_|_|_|\__'_|  |  Commit 2bb94d6 (11 days old master)
|__/                   |  x86_64-apple-darwin13.1.0

julia> 1 + 2
3

julia> ans
3
```

To exit the interactive session, type `^D` – the control key together with the d key or type `quit()`. When run in interactive mode, `julia` displays a banner and prompts the user for input. Once the user has entered a complete expression, such as `1 + 2`, and hits enter, the interactive session evaluates the expression and shows its value. If an expression is entered into an interactive session with a trailing semicolon, its value is not shown. The variable `ans` is bound to the value of the last evaluated expression whether it is shown or not. The `ans` variable is only bound in interactive sessions, not when Julia code is run in other ways.

To evaluate expressions written in a source file `file.jl`, write `include("file.jl")`.

To run code in a file non-interactively, you can give it as the first argument to the `julia` command:

```
$ julia script.jl arg1 arg2...
```

As the example implies, the following command-line arguments to `julia` are taken as command-line arguments to the program `script.jl`, passed in the global constant ARGS. The name of the script itself is passed in as the global PROGRAM_FILE. Note that ARGS is also set when script code is given using the `-e` option on the command line (see the `julia` help output below) but PROGRAM_FILE will be empty. For example, to just print the arguments given to a script, you could do this:

```
$ julia -e 'println(PROGRAM_FILE); for x in ARGS; println(x); end' foo bar
```

15

```
foo
bar
```

Or you could put that code into a script and run it:

```
$ echo 'println(PROGRAM_FILE); for x in ARGS; println(x); end' > script.jl
$ julia script.jl foo bar
script.jl
foo
bar
```

The `--` delimiter can be used to separate command-line args to the scriptfile from args to Julia:

```
$ julia --color=yes -O -- foo.jl arg1 arg2..
```

Julia can be started in parallel mode with either the `-p` or the `--machinefile` options. `-p n` will launch an additional n worker processes, while `--machinefile file` will launch a worker for each line in file `file`. The machines defined in `file` must be accessible via a passwordless `ssh` login, with Julia installed at the same location as the current host. Each machine definition takes the form `[count*][user@]host[:port] [bind_addr[:port]]`. user defaults to current user, `port` to the standard ssh port. `count` is the number of workers to spawn on the node, and defaults to 1. The optional `bind-to bind_addr[:port]` specifies the ip-address and port that other workers should use to connect to this worker.

If you have code that you want executed whenever Julia is run, you can put it in `~/.juliarc.jl`:

```
$ echo 'println("Greetings! ! ?")' > ~/.juliarc.jl
$ julia
Greetings! ! ?

...
```

There are various ways to run Julia code and provide options, similar to those available for the `perl` and `ruby` programs:

```
julia [switches] -- [programfile] [args...]
 -v, --version            Display version information
 -h, --help               Print this message

 -J, --sysimage <file>    Start up with the given system image file
 --precompiled={yes|no}   Use precompiled code from system image if available
 --compilecache={yes|no}  Enable/disable incremental precompilation of modules
 -H, --home <dir>         Set location of `julia` executable
 --startup-file={yes|no}  Load ~/.juliarc.jl
 --handle-signals={yes|no} Enable or disable Julia's default signal handlers

 -e, --eval <expr>        Evaluate <expr>
 -E, --print <expr>       Evaluate and show <expr>
 -L, --load <file>        Load <file> immediately on all processors

 -p, --procs {N|auto}     Integer value N launches N additional local worker processes
                          "auto" launches as many workers as the number of local cores
 --machinefile <file>     Run processes on hosts listed in <file>

 -i                       Interactive mode; REPL runs and isinteractive() is true
 -q, --quiet              Quiet startup (no banner)
 --color={yes|no}         Enable or disable color text
 --history-file={yes|no}  Load or save history
```

```
--compile={yes|no|all|min}Enable or disable JIT compiler, or request exhaustive compilation
-C, --cpu-target <target> Limit usage of cpu features up to <target>
-O, --optimize={0,1,2,3}  Set the optimization level (default is 2 if unspecified or 3 if
    specified as -O)
-g, -g <level>            Enable / Set the level of debug info generation (default is 1 if
    unspecified or 2 if specified as -g)
--inline={yes|no}         Control whether inlining is permitted (overrides functions declared as
     @inline)
--check-bounds={yes|no}   Emit bounds checks always or never (ignoring declarations)
--math-mode={ieee,fast}   Disallow or enable unsafe floating point optimizations (overrides
    @fastmath declaration)

--depwarn={yes|no|error}  Enable or disable syntax and method deprecation warnings ("error"
    turns warnings into errors)

--output-o name           Generate an object file (including system image data)
--output-ji name          Generate a system image data file (.ji)
--output-bc name          Generate LLVM bitcode (.bc)
--output-incremental=no   Generate an incremental output file (rather than complete)

--code-coverage={none|user|all}, --code-coverage
                          Count executions of source lines (omitting setting is equivalent to "
    user")
--track-allocation={none|user|all}, --track-allocation
                          Count bytes allocated by each source line
```

5.1   Resources

In addition to this manual, there are various other resources that may help new users get started with Julia:

- Julia and IJulia cheatsheet

- Learn Julia in a few minutes

- Learn Julia the Hard Way

- Julia by Example

- Hands-on Julia

- Tutorial for Homer Reid's numerical analysis class

- An introductory presentation

- Videos from the Julia tutorial at MIT

- YouTube videos from the JuliaCons

Chapter 6

Variables

A variable, in Julia, is a name associated (or bound) to a value. It's useful when you want to store a value (that you obtained after some math, for example) for later use. For example:

```julia
# Assign the value 10 to the variable x
julia> x = 10
10

# Doing math with x's value
julia> x + 1
11

# Reassign x's value
julia> x = 1 + 1
2

# You can assign values of other types, like strings of text
julia> x = "Hello World!"
"Hello World!"
```

Julia provides an extremely flexible system for naming variables. Variable names are case-sensitive, and have no semantic meaning (that is, the language will not treat variables differently based on their names).

```julia
julia> x = 1.0
1.0

julia> y = -3
-3

julia> Z = "My string"
"My string"

julia> customary_phrase = "Hello world!"
"Hello world!"

julia> UniversalDeclarationOfHumanRightsStart = ""
""
```

Unicode names (in UTF-8 encoding) are allowed:

```julia
julia> δ = 0.00001
1.0e-5

julia>  = "Hello"
"Hello"
```

In the Julia REPL and several other Julia editing environments, you can type many Unicode math symbols by typing the backslashed LaTeX symbol name followed by tab. For example, the variable name δ can be entered by typing \delta-*tab*, or even α by \alpha-*tab*-\hat- *tab*-_2-*tab*. (If you find a symbol somewhere, e.g. in someone else's code, that you don't know how to type, the REPL help will tell you: just type ? and then paste the symbol.)

Julia will even let you redefine built-in constants and functions if needed:

```julia
julia> pi
π = 3.1415926535897...

julia> pi = 3
WARNING: imported binding for pi overwritten in module Main
3

julia> pi
3

julia> sqrt(100)
10.0

julia> sqrt = 4
WARNING: imported binding for sqrt overwritten in module Main
4
```

However, this is obviously not recommended to avoid potential confusion.

6.1   Allowed Variable Names

Variable names must begin with a letter (A-Z or a-z), underscore, or a subset of Unicode code points greater than 00A0; in particular, Unicode character categories Lu/Ll/Lt/Lm/Lo/Nl (letters), Sc/So (currency and other symbols), and a few other letter-like characters (e.g. a subset of the Sm math symbols) are allowed. Subsequent characters may also include ! and digits (0-9 and other characters in categories Nd/No), as well as other Unicode code points: diacritics and other modifying marks (categories Mn/Mc/Me/Sk), some punctuation connectors (category Pc), primes, and a few other characters.

Operators like + are also valid identifiers, but are parsed specially. In some contexts, operators can be used just like variables; for example (+) refers to the addition function, and (+) = f will reassign it. Most of the Unicode infix operators (in category Sm), such as , are parsed as infix operators and are available for user-defined methods (e.g. you can use const = kron to define as an infix Kronecker product).

The only explicitly disallowed names for variables are the names of built-in statements:

```julia
julia> else = false
ERROR: syntax: unexpected "else"

julia> try = "No"
ERROR: syntax: unexpected "="
```

Some Unicode characters are considered to be equivalent in identifiers. Different ways of entering Unicode combining characters (e.g., accents) are treated as equivalent (specifically, Julia identifiers are NFC-normalized). The Unicode characters ɛ (U+025B: Latin small letter open e) and µ (U+00B5: micro sign) are treated as equivalent to the corresponding Greek letters, because the former are easily accessible via some input methods.

6.2 Stylistic Conventions

While Julia imposes few restrictions on valid names, it has become useful to adopt the following conventions:

- Names of variables are in lower case.

- Word separation can be indicated by underscores ('_'), but use of underscores is discouraged unless the name would be hard to read otherwise.

- Names of `Types` and `Modules` begin with a capital letter and word separation is shown with upper camel case instead of underscores.

- Names of `functions` and `macros` are in lower case, without underscores.

- Functions that write to their arguments have names that end in `!`. These are sometimes called "mutating" or "in-place" functions because they are intended to produce changes in their arguments after the function is called, not just return a value.

For more information about stylistic conventions, see the Style Guide.

Chapter 7

Integers and Floating-Point Numbers

Integers and floating-point values are the basic building blocks of arithmetic and computation. Built-in representations of such values are called numeric primitives, while representations of integers and floating-point numbers as immediate values in code are known as numeric literals. For example, 1 is an integer literal, while 1.0 is a floating-point literal; their binary in-memory representations as objects are numeric primitives.

Julia provides a broad range of primitive numeric types, and a full complement of arithmetic and bitwise operators as well as standard mathematical functions are defined over them. These map directly onto numeric types and operations that are natively supported on modern computers, thus allowing Julia to take full advantage of computational resources. Additionally, Julia provides software support for Arbitrary Precision Arithmetic, which can handle operations on numeric values that cannot be represented effectively in native hardware representations, but at the cost of relatively slower performance.

The following are Julia's primitive numeric types:

- **Integer types:**

Type	Signed?	Number of bits	Smallest value	Largest value
Int8		8	-2^7	2^7 - 1
UInt8		8	0	2^8 - 1
Int16		16	-2^15	2^15 - 1
UInt16		16	0	2^16 - 1
Int32		32	-2^31	2^31 - 1
UInt32		32	0	2^32 - 1
Int64		64	-2^63	2^63 - 1
UInt64		64	0	2^64 - 1
Int128		128	-2^127	2^127 - 1
UInt128		128	0	2^128 - 1
Bool	N/A	8	false (0)	true (1)

- **Floating-point types:**

Additionally, full support for Complex and Rational Numbers is built on top of these primitive numeric types. All numeric types interoperate naturally without explicit casting, thanks to a flexible, user-extensible type promotion system.

Type	Precision	Number of bits
Float16	half	16
Float32	single	32
Float64	double	64

7.1   Integers

Literal integers are represented in the standard manner:

```julia
julia> 1
1

julia> 1234
1234
```

The default type for an integer literal depends on whether the target system has a 32-bit architecture or a 64-bit architecture:

```julia
# 32-bit system:
julia> typeof(1)
Int32

# 64-bit system:
julia> typeof(1)
Int64
```

The Julia internal variable `Sys.WORD_SIZE` indicates whether the target system is 32-bit or 64-bit:

```julia
# 32-bit system:
julia> Sys.WORD_SIZE
32

# 64-bit system:
julia> Sys.WORD_SIZE
64
```

Julia also defines the types `Int` and `UInt`, which are aliases for the system's signed and unsigned native integer types respectively:

```julia
# 32-bit system:
julia> Int
Int32
julia> UInt
UInt32

# 64-bit system:
julia> Int
Int64
julia> UInt
UInt64
```

Larger integer literals that cannot be represented using only 32 bits but can be represented in 64 bits always create 64-bit integers, regardless of the system type:

```
# 32-bit or 64-bit system:
julia> typeof(3000000000)
Int64
```

Unsigned integers are input and output using the `0x` prefix and hexadecimal (base 16) digits `0-9a-f` (the capitalized digits `A-F` also work for input). The size of the unsigned value is determined by the number of hex digits used:

```
julia> 0x1
0x01

julia> typeof(ans)
UInt8

julia> 0x123
0x0123

julia> typeof(ans)
UInt16

julia> 0x1234567
0x01234567

julia> typeof(ans)
UInt32

julia> 0x123456789abcdef
0x0123456789abcdef

julia> typeof(ans)
UInt64
```

This behavior is based on the observation that when one uses unsigned hex literals for integer values, one typically is using them to represent a fixed numeric byte sequence, rather than just an integer value.

Recall that the variable ans is set to the value of the last expression evaluated in an interactive session. This does not occur when Julia code is run in other ways.

Binary and octal literals are also supported:

```
julia> 0b10
0x02

julia> typeof(ans)
UInt8

julia> 0o10
0x08

julia> typeof(ans)
UInt8
```

The minimum and maximum representable values of primitive numeric types such as integers are given by the `typemin()` and `typemax()` functions:

```julia
julia> (typemin(Int32), typemax(Int32))
(-2147483648, 2147483647)

julia> for T in [Int8,Int16,Int32,Int64,Int128,UInt8,UInt16,UInt32,UInt64,UInt128]
           println("$(lpad(T,7)): [$(typemin(T)),$(typemax(T))]")
       end
   Int8: [-128,127]
  Int16: [-32768,32767]
  Int32: [-2147483648,2147483647]
  Int64: [-9223372036854775808,9223372036854775807]
 Int128: [-170141183460469231731687303715884105728,170141183460469231731687303715884105727]
  UInt8: [0,255]
 UInt16: [0,65535]
 UInt32: [0,4294967295]
 UInt64: [0,18446744073709551615]
UInt128: [0,340282366920938463463374607431768211455]
```

The values returned by `typemin()` and `typemax()` are always of the given argument type. (The above expression uses several features we have yet to introduce, including for loops, Strings, and Interpolation, but should be easy enough to understand for users with some existing programming experience.)

Overflow behavior

In Julia, exceeding the maximum representable value of a given type results in a wraparound behavior:

```julia
julia> x = typemax(Int64)
9223372036854775807

julia> x + 1
-9223372036854775808

julia> x + 1 == typemin(Int64)
true
```

Thus, arithmetic with Julia integers is actually a form of modular arithmetic. This reflects the characteristics of the underlying arithmetic of integers as implemented on modern computers. In applications where overflow is possible, explicit checking for wraparound produced by overflow is essential; otherwise, the `BigInt` type in Arbitrary Precision Arithmetic is recommended instead.

Division errors

Integer division (the `div` function) has two exceptional cases: dividing by zero, and dividing the lowest negative number (`typemin()`) by -1. Both of these cases throw a `DivideError`. The remainder and modulus functions (`rem` and `mod`) throw a `DivideError` when their second argument is zero.

7.2   Floating-Point Numbers

Literal floating-point numbers are represented in the standard formats:

```
julia> 1.0
1.0

julia> 1.
1.0

julia> 0.5
0.5

julia> .5
0.5

julia> -1.23
-1.23

julia> 1e10
1.0e10

julia> 2.5e-4
0.00025
```

The above results are all Float64 values. Literal Float32 values can be entered by writing an f in place of e:

```
julia> 0.5f0
0.5f0

julia> typeof(ans)
Float32

julia> 2.5f-4
0.00025f0
```

Values can be converted to Float32 easily:

```
julia> Float32(-1.5)
-1.5f0

julia> typeof(ans)
Float32
```

Hexadecimal floating-point literals are also valid, but only as Float64 values:

```
julia> 0x1p0
1.0

julia> 0x1.8p3
12.0

julia> 0x.4p-1
0.125

julia> typeof(ans)
Float64
```

Half-precision floating-point numbers are also supported (`Float16`), but they are implemented in software and use `Float32` for calculations.

```
julia> sizeof(Float16(4.))
2

julia> 2*Float16(4.)
Float16(8.0)
```

The underscore _ can be used as digit separator:

```
julia> 10_000, 0.000_000_005, 0xdead_beef, 0b1011_0010
(10000, 5.0e-9, 0xdeadbeef, 0xb2)
```

Floating-point zero

Floating-point numbers have two zeros, positive zero and negative zero. They are equal to each other but have different binary representations, as can be seen using the `bits` function: :

```
julia> 0.0 == -0.0
true

julia> bits(0.0)
"0000000000000000000000000000000000000000000000000000000000000000"

julia> bits(-0.0)
"1000000000000000000000000000000000000000000000000000000000000000"
```

Special floating-point values

There are three specified standard floating-point values that do not correspond to any point on the real number line:

Float16	Float32	Float64	Name	Description
Inf16	Inf32	Inf	positive infinity	a value greater than all finite floating-point values
-Inf16	-Inf32	-Inf	negative infinity	a value less than all finite floating-point values
NaN16	NaN32	NaN	not a number	a value not == to any floating-point value (including itself)

For further discussion of how these non-finite floating-point values are ordered with respect to each other and other floats, see Numeric Comparisons. By the IEEE 754 standard, these floating-point values are the results of certain arithmetic operations:

```
julia> 1/Inf
0.0

julia> 1/0
Inf

julia> -5/0
-Inf
```

```
julia> 0.000001/0
Inf

julia> 0/0
NaN

julia> 500 + Inf
Inf

julia> 500 - Inf
-Inf

julia> Inf + Inf
Inf

julia> Inf - Inf
NaN

julia> Inf * Inf
Inf

julia> Inf / Inf
NaN

julia> 0 * Inf
NaN
```

The `typemin()` and `typemax()` functions also apply to floating-point types:

```
julia> (typemin(Float16),typemax(Float16))
(-Inf16, Inf16)

julia> (typemin(Float32),typemax(Float32))
(-Inf32, Inf32)

julia> (typemin(Float64),typemax(Float64))
(-Inf, Inf)
```

Machine epsilon

Most real numbers cannot be represented exactly with floating-point numbers, and so for many purposes it is important to know the distance between two adjacent representable floating-point numbers, which is often known as machine epsilon.

Julia provides `eps()`, which gives the distance between `1.0` and the next larger representable floating-point value:

```
julia> eps(Float32)
1.1920929f-7

julia> eps(Float64)
2.220446049250313e-16

julia> eps() # same as eps(Float64)
2.220446049250313e-16
```

These values are 2.0^-23 and 2.0^-52 as `Float32` and `Float64` values, respectively. The `eps()` function can also take a floating-point value as an argument, and gives the absolute difference between that value and the next representable floating point value. That is, eps(x) yields a value of the same type as x such that x + eps(x) is the next representable floating-point value larger than x:

```julia
julia> eps(1.0)
2.220446049250313e-16

julia> eps(1000.)
1.1368683772161603e-13

julia> eps(1e-27)
1.793662034335766e-43

julia> eps(0.0)
5.0e-324
```

The distance between two adjacent representable floating-point numbers is not constant, but is smaller for smaller values and larger for larger values. In other words, the representable floating-point numbers are densest in the real number line near zero, and grow sparser exponentially as one moves farther away from zero. By definition, eps(1.0) is the same as eps(Float64) since 1.0 is a 64-bit floating-point value.

Julia also provides the `nextfloat()` and `prevfloat()` functions which return the next largest or smallest representable floating-point number to the argument respectively:

```julia
julia> x = 1.25f0
1.25f0

julia> nextfloat(x)
1.2500001f0

julia> prevfloat(x)
1.2499999f0

julia> bits(prevfloat(x))
"00111111100111111111111111111111"

julia> bits(x)
"00111111101000000000000000000000"

julia> bits(nextfloat(x))
"00111111101000000000000000000001"
```

This example highlights the general principle that the adjacent representable floating-point numbers also have adjacent binary integer representations.

Rounding modes

If a number doesn't have an exact floating-point representation, it must be rounded to an appropriate representable value, however, if wanted, the manner in which this rounding is done can be changed according to the rounding modes presented in the IEEE 754 standard.

```julia
julia> x = 1.1; y = 0.1;
```

```julia
julia> x + y
1.2000000000000002

julia> setrounding(Float64,RoundDown) do
           x + y
       end
1.2
```

The default mode used is always `RoundNearest`, which rounds to the nearest representable value, with ties rounded towards the nearest value with an even least significant bit.

> **Warning**
>
> Rounding is generally only correct for basic arithmetic functions (`+()`, `-()`, `*()`, `/()` and `sqrt()`) and type conversion operations. Many other functions assume the default `RoundNearest` mode is set, and can give erroneous results when operating under other rounding modes.

Background and References

Floating-point arithmetic entails many subtleties which can be surprising to users who are unfamiliar with the low-level implementation details. However, these subtleties are described in detail in most books on scientific computation, and also in the following references:

- The definitive guide to floating point arithmetic is the IEEE 754-2008 Standard; however, it is not available for free online.

- For a brief but lucid presentation of how floating-point numbers are represented, see John D. Cook's article on the subject as well as his introduction to some of the issues arising from how this representation differs in behavior from the idealized abstraction of real numbers.

- Also recommended is Bruce Dawson's series of blog posts on floating-point numbers.

- For an excellent, in-depth discussion of floating-point numbers and issues of numerical accuracy encountered when computing with them, see David Goldberg's paper What Every Computer Scientist Should Know About Floating-Point Arithmetic.

- For even more extensive documentation of the history of, rationale for, and issues with floating-point numbers, as well as discussion of many other topics in numerical computing, see the collected writings of William Kahan, commonly known as the "Father of Floating-Point". Of particular interest may be An Interview with the Old Man of Floating-Point.

7.3 Arbitrary Precision Arithmetic

To allow computations with arbitrary-precision integers and floating point numbers, Julia wraps the GNU Multiple Precision Arithmetic Library (GMP) and the GNU MPFR Library, respectively. The `BigInt` and `BigFloat` types are available in Julia for arbitrary precision integer and floating point numbers respectively.

Constructors exist to create these types from primitive numerical types, and `parse()` can be used to construct them from `AbstractStrings`. Once created, they participate in arithmetic with all other numeric types thanks to Julia's type promotion and conversion mechanism:

```julia
julia> BigInt(typemax(Int64)) + 1
9223372036854775808
```

```julia
julia> parse(BigInt, "123456789012345678901234567890") + 1
123456789012345678901234567891

julia> parse(BigFloat, "1.23456789012345678901")
1.234567890123456789010000000000000000000000000000000000000000000000000000000004

julia> BigFloat(2.0^66) / 3
2.459565876494606882133333333333333333333333333333333333333333333333333333333344e+19

julia> factorial(BigInt(40))
815915283247897734345611269596115894272000000000
```

However, type promotion between the primitive types above and BigInt/BigFloat is not automatic and must be explicitly stated.

```julia
julia> x = typemin(Int64)
-9223372036854775808

julia> x = x - 1
9223372036854775807

julia> typeof(x)
Int64

julia> y = BigInt(typemin(Int64))
-9223372036854775808

julia> y = y - 1
-9223372036854775809

julia> typeof(y)
BigInt
```

The default precision (in number of bits of the significand) and rounding mode of BigFloat operations can be changed globally by calling setprecision() and setrounding(), and all further calculations will take these changes in account. Alternatively, the precision or the rounding can be changed only within the execution of a particular block of code by using the same functions with a do block:

```julia
julia> setrounding(BigFloat, RoundUp) do
           BigFloat(1) + parse(BigFloat, "0.1")
       end
1.100000000000000000000000000000000000000000000000000000000000000000000000000003

julia> setrounding(BigFloat, RoundDown) do
           BigFloat(1) + parse(BigFloat, "0.1")
       end
1.099999999999999999999999999999999999999999999999999999999999999999999999999986

julia> setprecision(40) do
           BigFloat(1) + parse(BigFloat, "0.1")
       end
1.1000000000004
```

7.4 Numeric Literal Coefficients

To make common numeric formulas and expressions clearer, Julia allows variables to be immediately preceded by a numeric literal, implying multiplication. This makes writing polynomial expressions much cleaner:

```
julia> x = 3
3

julia> 2x^2 - 3x + 1
10

julia> 1.5x^2 - .5x + 1
13.0
```

It also makes writing exponential functions more elegant:

```
julia> 2^2x
64
```

The precedence of numeric literal coefficients is the same as that of unary operators such as negation. So `2^3x` is parsed as `2^(3x)`, and `2x^3` is parsed as `2*(x^3)`.

Numeric literals also work as coefficients to parenthesized expressions:

```
julia> 2(x-1)^2 - 3(x-1) + 1
3
```

Additionally, parenthesized expressions can be used as coefficients to variables, implying multiplication of the expression by the variable:

```
julia> (x-1)x
6
```

Neither juxtaposition of two parenthesized expressions, nor placing a variable before a parenthesized expression, however, can be used to imply multiplication:

```
julia> (x-1)(x+1)
ERROR: MethodError: objects of type Int64 are not callable

julia> x(x+1)
ERROR: MethodError: objects of type Int64 are not callable
```

Both expressions are interpreted as function application: any expression that is not a numeric literal, when immediately followed by a parenthetical, is interpreted as a function applied to the values in parentheses (see Functions for more about functions). Thus, in both of these cases, an error occurs since the left-hand value is not a function.

The above syntactic enhancements significantly reduce the visual noise incurred when writing common mathematical formulae. Note that no whitespace may come between a numeric literal coefficient and the identifier or parenthesized expression which it multiplies.

Syntax Conflicts

Juxtaposed literal coefficient syntax may conflict with two numeric literal syntaxes: hexadecimal integer literals and engineering notation for floating-point literals. Here are some situations where syntactic conflicts arise:

- The hexadecimal integer literal expression `0xff` could be interpreted as the numeric literal `0` multiplied by the variable `xff`.

- The floating-point literal expression 1e10 could be interpreted as the numeric literal 1 multiplied by the variable e10, and similarly with the equivalent E form.

In both cases, we resolve the ambiguity in favor of interpretation as a numeric literals:

- Expressions starting with 0x are always hexadecimal literals.

- Expressions starting with a numeric literal followed by e or E are always floating-point literals.

7.5   Literal zero and one

Julia provides functions which return literal 0 and 1 corresponding to a specified type or the type of a given variable.

Function	Description
zero(x)	Literal zero of type x or type of variable x
one(x)	Literal one of type x or type of variable x

These functions are useful in Numeric Comparisons to avoid overhead from unnecessary type conversion.

Examples:

```julia
julia> zero(Float32)
0.0f0

julia> zero(1.0)
0.0

julia> one(Int32)
1

julia> one(BigFloat)
1.000000000000000000000000000000000000000000000000000000000000000000000000000000
```

Chapter 8

Mathematical Operations and Elementary Functions

Julia provides a complete collection of basic arithmetic and bitwise operators across all of its numeric primitive types, as well as providing portable, efficient implementations of a comprehensive collection of standard mathematical functions.

8.1 Arithmetic Operators

The following arithmetic operators are supported on all primitive numeric types:

Expression	Name	Description
`+x`	unary plus	the identity operation
`-x`	unary minus	maps values to their additive inverses
`x + y`	binary plus	performs addition
`x - y`	binary minus	performs subtraction
`x * y`	times	performs multiplication
`x / y`	divide	performs division
`x \ y`	inverse divide	equivalent to `y / x`
`x ^ y`	power	raises x to the yth power
`x % y`	remainder	equivalent to `rem(x,y)`

as well as the negation on Bool types:

Expression	Name	Description
`!x`	negation	changes `true` to `false` and vice versa

Julia's promotion system makes arithmetic operations on mixtures of argument types "just work" naturally and automatically. See Conversion and Promotion for details of the promotion system.

Here are some simple examples using arithmetic operators:

```
julia> 1 + 2 + 3
6

julia> 1 - 2
-1

julia> 3*2/12
0.5
```

(By convention, we tend to space operators more tightly if they get applied before other nearby operators. For instance, we would generally write -x + 2 to reflect that first x gets negated, and then 2 is added to that result.)

8.2   Bitwise Operators

The following bitwise operators are supported on all primitive integer types:

Expression	Name
~x	bitwise not
x & y	bitwise and
x | y	bitwise or
x y	bitwise xor (exclusive or)
x >>> y	logical shift right
x >> y	arithmetic shift right
x << y	logical/arithmetic shift left

Here are some examples with bitwise operators:

```julia
julia> ~123
-124

julia> 123 & 234
106

julia> 123 | 234
251

julia> 123  234
145

julia> xor(123, 234)
145

julia> ~UInt32(123)
0xffffff84

julia> ~UInt8(123)
0x84
```

8.3   Updating operators

Every binary arithmetic and bitwise operator also has an updating version that assigns the result of the operation back into its left operand. The updating version of the binary operator is formed by placing a = immediately after the operator. For example, writing x += 3 is equivalent to writing x = x + 3:

```julia
julia> x = 1
1

julia> x += 3
4

julia> x
4
```

The updating versions of all the binary arithmetic and bitwise operators are:

```
+=  -=  *=  /=  \=  ÷=  %=  ^=  &=  |=  =  >>>=  >>=  <<=
```

> **Note**
>
> An updating operator rebinds the variable on the left-hand side. As a result, the type of the variable may change.
>
> ```julia
> julia> x = 0x01; typeof(x)
> UInt8
>
> julia> x *= 2 # Same as x = x * 2
> 2
>
> julia> typeof(x)
> Int64
> ```

8.4 Vectorized "dot" operators

For *every* binary operation like `^`, there is a corresponding "dot" operation `.^` that is *automatically* defined to perform `^` element-by-element on arrays. For example, `[1,2,3] ^ 3` is not defined, since there is no standard mathematical meaning to "cubing" an array, but `[1,2,3] .^ 3` is defined as computing the elementwise (or "vectorized") result `[1^3, 2^3, 3^3]`. Similarly for unary operators like `!` or `√`, there is a corresponding `.√` that applies the operator elementwise.

```julia
julia> [1,2,3] .^ 3
3-element Array{Int64,1}:
  1
  8
 27
```

More specifically, a `.^ b` is parsed as the "dot" call `(^).(a,b)`, which performs a broadcast operation: it can combine arrays and scalars, arrays of the same size (performing the operation elementwise), and even arrays of different shapes (e.g. combining row and column vectors to produce a matrix). Moreover, like all vectorized "dot calls," these "dot operators" are *fusing*. For example, if you compute `2 .* A.^2 .+ sin.(A)` (or equivalently `@. 2A^2 + sin(A)`, using the `@.` macro) for an array A, it performs a *single* loop over A, computing `2a^2 + sin(a)` for each element of A. In particular, nested dot calls like `f.(g.(x))` are fused, and "adjacent" binary operators like `x .+ 3 .* x.^2` are equivalent to nested dot calls `(+).(x, (*).(3, (^).(x, 2)))`.

Furthermore, "dotted" updating operators like `a .+= b` (or `@. a += b`) are parsed as `a .= a .+ b`, where `.=` is a fused *in-place* assignment operation (see the dot syntax documentation).

Note the dot syntax is also applicable to user-defined operators. For example, if you define `⊗(A,B) = kron(A,B)` to give a convenient infix syntax `A ⊗ B` for Kronecker products (kron), then `[A,B] .⊗ [C,D]` will compute `[A⊗C, B⊗D]` with no additional coding.

8.5 Numeric Comparisons

Standard comparison operations are defined for all the primitive numeric types:

Here are some simple examples:

Operator	Name
==	equality
!=, ≠	inequality
<	less than
<=, ≤	less than or equal to
>	greater than
>=, ≥	greater than or equal to

```
julia> 1 == 1
true

julia> 1 == 2
false

julia> 1 != 2
true

julia> 1 == 1.0
true

julia> 1 < 2
true

julia> 1.0 > 3
false

julia> 1 >= 1.0
true

julia> -1 <= 1
true

julia> -1 <= -1
true

julia> -1 <= -2
false

julia> 3 < -0.5
false
```

Integers are compared in the standard manner – by comparison of bits. Floating-point numbers are compared according to the IEEE 754 standard:

- Finite numbers are ordered in the usual manner.

- Positive zero is equal but not greater than negative zero.

- `Inf` is equal to itself and greater than everything else except `NaN`.

- `-Inf` is equal to itself and less then everything else except `NaN`.

- `NaN` is not equal to, not less than, and not greater than anything, including itself.

The last point is potentially surprising and thus worth noting:

```julia
julia> NaN == NaN
false

julia> NaN != NaN
true

julia> NaN < NaN
false

julia> NaN > NaN
false
```

and can cause especial headaches with Arrays:

```julia
julia> [1 NaN] == [1 NaN]
false
```

Julia provides additional functions to test numbers for special values, which can be useful in situations like hash key comparisons:

Function	Tests if
isequal(x, y)	x and y are identical
isfinite(x)	x is a finite number
isinf(x)	x is infinite
isnan(x)	x is not a number

isequal() considers NaNs equal to each other:

```julia
julia> isequal(NaN, NaN)
true

julia> isequal([1 NaN], [1 NaN])
true

julia> isequal(NaN, NaN32)
true
```

isequal() can also be used to distinguish signed zeros:

```julia
julia> -0.0 == 0.0
true

julia> isequal(-0.0, 0.0)
false
```

Mixed-type comparisons between signed integers, unsigned integers, and floats can be tricky. A great deal of care has been taken to ensure that Julia does them correctly.

For other types, isequal() defaults to calling ==(), so if you want to define equality for your own types then you only need to add a ==() method. If you define your own equality function, you should probably define a corresponding hash() method to ensure that isequal(x,y) implies hash(x) == hash(y).

Chaining comparisons

Unlike most languages, with the notable exception of Python, comparisons can be arbitrarily chained:

```julia
julia> 1 < 2 <= 2 < 3 == 3 > 2 >= 1 == 1 < 3 != 5
true
```

Chaining comparisons is often quite convenient in numerical code. Chained comparisons use the && operator for scalar comparisons, and the & operator for elementwise comparisons, which allows them to work on arrays. For example, 0 .< A .< 1 gives a boolean array whose entries are true where the corresponding elements of A are between 0 and 1.

Note the evaluation behavior of chained comparisons:

```julia
julia> v(x) = (println(x); x)
v (generic function with 1 method)

julia> v(1) < v(2) <= v(3)
2
1
3
true

julia> v(1) > v(2) <= v(3)
2
1
false
```

The middle expression is only evaluated once, rather than twice as it would be if the expression were written as v(1) < v(2) && v(2) <= v(3). However, the order of evaluations in a chained comparison is undefined. It is strongly recommended not to use expressions with side effects (such as printing) in chained comparisons. If side effects are required, the short-circuit && operator should be used explicitly (see Short-Circuit Evaluation).

Elementary Functions

Julia provides a comprehensive collection of mathematical functions and operators. These mathematical operations are defined over as broad a class of numerical values as permit sensible definitions, including integers, floating-point numbers, rationals, and complex numbers, wherever such definitions make sense.

Moreover, these functions (like any Julia function) can be applied in "vectorized" fashion to arrays and other collections with the dot syntax f.(A), e.g. sin.(A) will compute the sine of each element of an array A.

8.6   Operator Precedence

Julia applies the following order of operations, from highest precedence to lowest:

For a complete list of *every* Julia operator's precedence, see the top of this file: `src/julia-parser.scm`

You can also find the numerical precedence for any given operator via the built-in function `Base.operator_precedence`, where higher numbers take precedence:

```julia
julia> Base.operator_precedence(:+), Base.operator_precedence(:*), Base.operator_precedence(:.)
(9, 11, 15)
```

Category	Operators
Syntax	. followed by ::
Exponentiation	^
Fractions	//
Multiplication	* / % & \
Bitshifts	<< >> >>>
Addition	+ - |
Syntax	: .. followed by |>
Comparisons	> < >= <= == === != !== <:
Control flow	&& followed by || followed by ?
Assignments	= += -= *= /= //= \= ^= ÷= %= |= &= = <<= >>= >>>=

```julia
julia> Base.operator_precedence(:+=), Base.operator_precedence(:(=))  # (Note the necessary parens
↪    on `:(=)`)
(1, 1)
```

8.7   Numerical Conversions

Julia supports three forms of numerical conversion, which differ in their handling of inexact conversions.

- The notation T(x) or convert(T,x) converts x to a value of type T.

  - If T is a floating-point type, the result is the nearest representable value, which could be positive or negative infinity.

  - If T is an integer type, an InexactError is raised if x is not representable by T.

- x % T converts an integer x to a value of integer type T congruent to x modulo 2^n, where n is the number of bits in T. In other words, the binary representation is truncated to fit.

- The Rounding functions take a type T as an optional argument. For example, round(Int,x) is a shorthand for Int(round(x)).

The following examples show the different forms.

```julia
julia> Int8(127)
127

julia> Int8(128)
ERROR: InexactError()
Stacktrace:
 [1] Int8(::Int64) at ./sysimg.jl:24

julia> Int8(127.0)
127

julia> Int8(3.14)
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Int8}, ::Float64) at ./float.jl:658
 [2] Int8(::Float64) at ./sysimg.jl:24
```

```
julia> Int8(128.0)
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Int8}, ::Float64) at ./float.jl:658
 [2] Int8(::Float64) at ./sysimg.jl:24

julia> 127 % Int8
127

julia> 128 % Int8
-128

julia> round(Int8,127.4)
127

julia> round(Int8,127.6)
ERROR: InexactError()
Stacktrace:
 [1] trunc(::Type{Int8}, ::Float64) at ./float.jl:651
 [2] round(::Type{Int8}, ::Float64) at ./float.jl:337
```

See Conversion and Promotion for how to define your own conversions and promotions.

Rounding functions

Function	Description	Return type
round(x)	round x to the nearest integer	typeof(x)
round(T, x)	round x to the nearest integer	T
floor(x)	round x towards -Inf	typeof(x)
floor(T, x)	round x towards -Inf	T
ceil(x)	round x towards +Inf	typeof(x)
ceil(T, x)	round x towards +Inf	T
trunc(x)	round x towards zero	typeof(x)
trunc(T, x)	round x towards zero	T

Division functions

Sign and absolute value functions

Powers, logs and roots

For an overview of why functions like hypot(), expm1(), and log1p() are necessary and useful, see John D. Cook's excellent pair of blog posts on the subject: expm1, log1p, erfc, and hypot.

Trigonometric and hyperbolic functions

All the standard trigonometric and hyperbolic functions are also defined:

```
sin    cos    tan    cot    sec    csc
sinh   cosh   tanh   coth   sech   csch
asin   acos   atan   acot   asec   acsc
asinh  acosh  atanh  acoth  asech  acsch
sinc   cosc   atan2
```

Function	Description
`div(x,y)`	truncated division; quotient rounded towards zero
`fld(x,y)`	floored division; quotient rounded towards `-Inf`
`cld(x,y)`	ceiling division; quotient rounded towards `+Inf`
`rem(x,y)`	remainder; satisfies `x == div(x,y)*y + rem(x,y)`; sign matches x
`mod(x,y)`	modulus; satisfies `x == fld(x,y)*y + mod(x,y)`; sign matches y
`mod1(x,y)`	mod() with offset 1; returns r(0,y] for y>0 or r[y,0) for y<0, where `mod(r, y) == mod(x, y)`
`mod2pi(x)`	modulus with respect to 2pi; `0 <= mod2pi(x)   < 2pi`
`di-vrem(x,y)`	returns `(div(x,y),rem(x,y))`
`fld-mod(x,y)`	returns `(fld(x,y),mod(x,y))`
`gcd(x,y...)`	greatest positive common divisor of x, y,...
`lcm(x,y...)`	least positive common multiple of x, y,...

Function	Description
`abs(x)`	a positive value with the magnitude of x
`abs2(x)`	the squared magnitude of x
`sign(x)`	indicates the sign of x, returning -1, 0, or +1
`signbit(x)`	indicates whether the sign bit is on (true) or off (false)
`copysign(x,y)`	a value with the magnitude of x and the sign of y
`flipsign(x,y)`	a value with the magnitude of x and the sign of x*y

Function	Description
`sqrt(x)`, √x	square root of x
`cbrt(x)`, ∛x	cube root of x
`hypot(x,y)`	hypotenuse of right-angled triangle with other sides of length x and y
`exp(x)`	natural exponential function at x
`expm1(x)`	accurate `exp(x)-1` for x near zero
`ldexp(x,n)`	`x*2^n` computed efficiently for integer values of n
`log(x)`	natural logarithm of x
`log(b,x)`	base b logarithm of x
`log2(x)`	base 2 logarithm of x
`log10(x)`	base 10 logarithm of x
`log1p(x)`	accurate `log(1+x)` for x near zero
`exponent(x)`	binary exponent of x
`significand(x)`	binary significand (a.k.a. mantissa) of a floating-point number x

These are all single-argument functions, with the exception of atan2, which gives the angle in radians between the x-axis and the point specified by its arguments, interpreted as x and y coordinates.

Additionally, `sinpi(x)` and `cospi(x)` are provided for more accurate computations of `sin(pi*x)` and `cos(pi*x)` respectively.

In order to compute trigonometric functions with degrees instead of radians, suffix the function with d. For example, `sind(x)` computes the sine of x where x is specified in degrees. The complete list of trigonometric functions with degree variants is:

```
sind   cosd   tand   cotd   secd   cscd
asind  acosd  atand  acotd  asecd  acscd
```

Special functions

Function	Description
gamma(x)	gamma function at x
lgamma(x)	accurate log(gamma(x)) for large x
lfact(x)	accurate log(factorial(x)) for large x; same as lgamma(x+1) for x > 1, zero otherwise
beta(x,y)	beta function at x, y
lbeta(x,y)	accurate log(beta(x,y)) for large x or y

Chapter 9

Complex and Rational Numbers

Julia ships with predefined types representing both complex and rational numbers, and supports all standard Mathematical Operations and Elementary Functions on them. Conversion and Promotion are defined so that operations on any combination of predefined numeric types, whether primitive or composite, behave as expected.

9.1 Complex Numbers

The global constant `im` is bound to the complex number *i*, representing the principal square root of -1. It was deemed harmful to co-opt the name `i` for a global constant, since it is such a popular index variable name. Since Julia allows numeric literals to be juxtaposed with identifiers as coefficients, this binding suffices to provide convenient syntax for complex numbers, similar to the traditional mathematical notation:

```julia
julia> 1 + 2im
1 + 2im
```

You can perform all the standard arithmetic operations with complex numbers:

```julia
julia> (1 + 2im)*(2 - 3im)
8 + 1im

julia> (1 + 2im)/(1 - 2im)
-0.6 + 0.8im

julia> (1 + 2im) + (1 - 2im)
2 + 0im

julia> (-3 + 2im) - (5 - 1im)
-8 + 3im

julia> (-1 + 2im)^2
-3 - 4im

julia> (-1 + 2im)^2.5
2.7296244647840084 - 6.960664459571898im

julia> (-1 + 2im)^(1 + 1im)
-0.27910381075826657 + 0.087080534141102428im

julia> 3(2 - 5im)
```

```
6 - 15im

julia> 3(2 - 5im)^2
-63 - 60im

julia> 3(2 - 5im)^-1.0
0.20689655172413796 + 0.5172413793103449im
```

The promotion mechanism ensures that combinations of operands of different types just work:

```
julia> 2(1 - 1im)
2 - 2im

julia> (2 + 3im) - 1
1 + 3im

julia> (1 + 2im) + 0.5
1.5 + 2.0im

julia> (2 + 3im) - 0.5im
2.0 + 2.5im

julia> 0.75(1 + 2im)
0.75 + 1.5im

julia> (2 + 3im) / 2
1.0 + 1.5im

julia> (1 - 3im) / (2 + 2im)
-0.5 - 1.0im

julia> 2im^2
-2 + 0im

julia> 1 + 3/4im
1.0 - 0.75im
```

Note that `3/4im == 3/(4*im) == -(3/4*im)`, since a literal coefficient binds more tightly than division.

Standard functions to manipulate complex values are provided:

```
julia> z = 1 + 2im
1 + 2im

julia> real(1 + 2im) # real part of z
1

julia> imag(1 + 2im) # imaginary part of z
2

julia> conj(1 + 2im) # complex conjugate of z
1 - 2im

julia> abs(1 + 2im) # absolute value of z
2.23606797749979
```

```julia
julia> abs2(1 + 2im) # squared absolute value
5

julia> angle(1 + 2im) # phase angle in radians
1.1071487177940904
```

As usual, the absolute value (`abs()`) of a complex number is its distance from zero. `abs2()` gives the square of the absolute value, and is of particular use for complex numbers where it avoids taking a square root. `angle()` returns the phase angle in radians (also known as the *argument* or *arg* function). The full gamut of other Elementary Functions is also defined for complex numbers:

```julia
julia> sqrt(1im)
0.7071067811865476 + 0.7071067811865475im

julia> sqrt(1 + 2im)
1.272019649514069 + 0.7861513777574233im

julia> cos(1 + 2im)
2.0327230070196656 - 3.0518977991518im

julia> exp(1 + 2im)
-1.1312043837568135 + 2.4717266720048188im

julia> sinh(1 + 2im)
-0.4890562590412937 + 1.4031192506220405im
```

Note that mathematical functions typically return real values when applied to real numbers and complex values when applied to complex numbers. For example, `sqrt()` behaves differently when applied to `-1` versus `-1 + 0im` even though `-1 == -1 + 0im`:

```julia
julia> sqrt(-1)
ERROR: DomainError:
sqrt will only return a complex result if called with a complex argument. Try sqrt(complex(x)).
Stacktrace:
 [1] sqrt(::Int64) at ./math.jl:434

julia> sqrt(-1 + 0im)
0.0 + 1.0im
```

The literal numeric coefficient notation does not work when constructing a complex number from variables. Instead, the multiplication must be explicitly written out:

```julia
julia> a = 1; b = 2; a + b*im
1 + 2im
```

However, this is *not* recommended; Use the `complex()` function instead to construct a complex value directly from its real and imaginary parts:

```julia
julia> a = 1; b = 2; complex(a, b)
1 + 2im
```

This construction avoids the multiplication and addition operations.

Inf and NaN propagate through complex numbers in the real and imaginary parts of a complex number as described in the Special floating-point values section:

```julia
julia> 1 + Inf*im
1.0 + Inf*im

julia> 1 + NaN*im
1.0 + NaN*im
```

9.2   Rational Numbers

Julia has a rational number type to represent exact ratios of integers. Rationals are constructed using the `//` operator:

```julia
julia> 2//3
2//3
```

If the numerator and denominator of a rational have common factors, they are reduced to lowest terms such that the denominator is non-negative:

```julia
julia> 6//9
2//3

julia> -4//8
-1//2

julia> 5//-15
-1//3

julia> -4//-12
1//3
```

This normalized form for a ratio of integers is unique, so equality of rational values can be tested by checking for equality of the numerator and denominator. The standardized numerator and denominator of a rational value can be extracted using the `numerator()` and `denominator()` functions:

```julia
julia> numerator(2//3)
2

julia> denominator(2//3)
3
```

Direct comparison of the numerator and denominator is generally not necessary, since the standard arithmetic and comparison operations are defined for rational values:

```julia
julia> 2//3 == 6//9
true

julia> 2//3 == 9//27
false
```

```julia
julia> 3//7 < 1//2
true

julia> 3//4 > 2//3
true

julia> 2//4 + 1//6
2//3

julia> 5//12 - 1//4
1//6

julia> 5//8 * 3//12
5//32

julia> 6//5 / 10//7
21//25
```

Rationals can be easily converted to floating-point numbers:

```julia
julia> float(3//4)
0.75
```

Conversion from rational to floating-point respects the following identity for any integral values of a and b, with the exception of the case a == 0 and b == 0:

```julia
julia> a = 1; b = 2;

julia> isequal(float(a//b), a/b)
true
```

Constructing infinite rational values is acceptable:

```julia
julia> 5//0
1//0

julia> -3//0
-1//0

julia> typeof(ans)
Rational{Int64}
```

Trying to construct a NaN rational value, however, is not:

```julia
julia> 0//0
ERROR: ArgumentError: invalid rational: zero(Int64)//zero(Int64)
Stacktrace:
 [1] Rational{Int64}(::Int64, ::Int64) at ./rational.jl:13
 [2] //(::Int64, ::Int64) at ./rational.jl:40
```

As usual, the promotion system makes interactions with other numeric types effortless:

```julia
julia> 3//5 + 1
8//5

julia> 3//5 - 0.5
0.09999999999999998

julia> 2//7 * (1 + 2im)
2//7 + 4//7*im

julia> 2//7 * (1.5 + 2im)
0.42857142857142855 + 0.5714285714285714im

julia> 3//2 / (1 + 2im)
3//10 - 3//5*im

julia> 1//2 + 2im
1//2 + 2//1*im

julia> 1 + 2//3im
1//1 - 2//3*im

julia> 0.5 == 1//2
true

julia> 0.33 == 1//3
false

julia> 0.33 < 1//3
true

julia> 1//3 - 0.33
0.0033333333333332993
```

Chapter 10

Strings

Strings are finite sequences of characters. Of course, the real trouble comes when one asks what a character is. The characters that English speakers are familiar with are the letters A, B, C, etc., together with numerals and common punctuation symbols. These characters are standardized together with a mapping to integer values between 0 and 127 by the ASCII standard. There are, of course, many other characters used in non-English languages, including variants of the ASCII characters with accents and other modifications, related scripts such as Cyrillic and Greek, and scripts completely unrelated to ASCII and English, including Arabic, Chinese, Hebrew, Hindi, Japanese, and Korean. The Unicode standard tackles the complexities of what exactly a character is, and is generally accepted as the definitive standard addressing this problem. Depending on your needs, you can either ignore these complexities entirely and just pretend that only ASCII characters exist, or you can write code that can handle any of the characters or encodings that one may encounter when handling non-ASCII text. Julia makes dealing with plain ASCII text simple and efficient, and handling Unicode is as simple and efficient as possible. In particular, you can write C-style string code to process ASCII strings, and they will work as expected, both in terms of performance and semantics. If such code encounters non-ASCII text, it will gracefully fail with a clear error message, rather than silently introducing corrupt results. When this happens, modifying the code to handle non-ASCII data is straightforward.

There are a few noteworthy high-level features about Julia's strings:

- The built-in concrete type used for strings (and string literals) in Julia is `String`. This supports the full range of Unicode characters via the UTF-8 encoding. (A `transcode()` function is provided to convert to/from other Unicode encodings.)

- All string types are subtypes of the abstract type `AbstractString`, and external packages define additional `AbstractString` subtypes (e.g. for other encodings). If you define a function expecting a string argument, you should declare the type as `AbstractString` in order to accept any string type.

- Like C and Java, but unlike most dynamic languages, Julia has a first-class type representing a single character, called `Char`. This is just a special kind of 32-bit primitive type whose numeric value represents a Unicode code point.

- As in Java, strings are immutable: the value of an `AbstractString` object cannot be changed. To construct a different string value, you construct a new string from parts of other strings.

- Conceptually, a string is a *partial function* from indices to characters: for some index values, no character value is returned, and instead an exception is thrown. This allows for efficient indexing into strings by the byte index of an encoded representation rather than by a character index, which cannot be implemented both efficiently and simply for variable-width encodings of Unicode strings.

10.1   Characters

A `Char` value represents a single character: it is just a 32-bit primitive type with a special literal representation and appropriate arithmetic behaviors, whose numeric value is interpreted as a Unicode code point. Here is how `Char` values are input and shown:

```
julia> 'x'
'x': ASCII/Unicode U+0078 (category Ll: Letter, lowercase)

julia> typeof(ans)
Char
```

You can convert a `Char` to its integer value, i.e. code point, easily:

```
julia> Int('x')
120

julia> typeof(ans)
Int64
```

On 32-bit architectures, `typeof(ans)` will be `Int32`. You can convert an integer value back to a `Char` just as easily:

```
julia> Char(120)
'x': ASCII/Unicode U+0078 (category Ll: Letter, lowercase)
```

Not all integer values are valid Unicode code points, but for performance, the `Char()` conversion does not check that every character value is valid. If you want to check that each converted value is a valid code point, use the `isvalid()` function:

```
julia> Char(0x110000)
'\U110000': Unicode U+110000 (category Cn: Other, not assigned)

julia> isvalid(Char, 0x110000)
false
```

As of this writing, the valid Unicode code points are `U+00` through `U+d7ff` and `U+e000` through `U+10ffff`. These have not all been assigned intelligible meanings yet, nor are they necessarily interpretable by applications, but all of these values are considered to be valid Unicode characters.

You can input any Unicode character in single quotes using \u followed by up to four hexadecimal digits or \U followed by up to eight hexadecimal digits (the longest valid value only requires six):

```
julia> '\u0'
'\0': ASCII/Unicode U+0000 (category Cc: Other, control)

julia> '\u78'
'x': ASCII/Unicode U+0078 (category Ll: Letter, lowercase)

julia> '\u2200'
'∀': Unicode U+2200 (category Sm: Symbol, math)

julia> '\U10ffff'
'\U10ffff': Unicode U+10ffff (category Cn: Other, not assigned)
```

Julia uses your system's locale and language settings to determine which characters can be printed as-is and which must be output using the generic, escaped \u or \U input forms. In addition to these Unicode escape forms, all of C's traditional escaped input forms can also be used:

```julia
julia> Int('\0')
0

julia> Int('\t')
9

julia> Int('\n')
10

julia> Int('\e')
27

julia> Int('\x7f')
127

julia> Int('\177')
127

julia> Int('\xff')
255
```

You can do comparisons and a limited amount of arithmetic with `Char` values:

```julia
julia> 'A' < 'a'
true

julia> 'A' <= 'a' <= 'Z'
false

julia> 'A' <= 'X' <= 'Z'
true

julia> 'x' - 'a'
23

julia> 'A' + 1
'B': ASCII/Unicode U+0042 (category Lu: Letter, uppercase)
```

10.2 String Basics

String literals are delimited by double quotes or triple double quotes:

```julia
julia> str = "Hello, world.\n"
"Hello, world.\n"

julia> """Contains "quote" characters"""
"Contains \"quote\" characters"
```

If you want to extract a character from a string, you index into it:

```
julia> str[1]
'H': ASCII/Unicode U+0048 (category Lu: Letter, uppercase)

julia> str[6]
',': ASCII/Unicode U+002c (category Po: Punctuation, other)

julia> str[end]
'\n': ASCII/Unicode U+000a (category Cc: Other, control)
```

All indexing in Julia is 1-based: the first element of any integer-indexed object is found at index 1. (As we will see below, this does not necessarily mean that the last element is found at index n, where n is the length of the string.)

In any indexing expression, the keyword end can be used as a shorthand for the last index (computed by endof(str)). You can perform arithmetic and other operations with end, just like a normal value:

```
julia> str[end-1]
'.': ASCII/Unicode U+002e (category Po: Punctuation, other)

julia> str[end÷2]
' ': ASCII/Unicode U+0020 (category Zs: Separator, space)
```

Using an index less than 1 or greater than end raises an error:

```
julia> str[0]
ERROR: BoundsError: attempt to access "Hello, world.\n"
  at index [0]
[...]

julia> str[end+1]
ERROR: BoundsError: attempt to access "Hello, world.\n"
  at index [15]
[...]
```

You can also extract a substring using range indexing:

```
julia> str[4:9]
"lo, wo"
```

Notice that the expressions str[k] and str[k:k] do not give the same result:

```
julia> str[6]
',': ASCII/Unicode U+002c (category Po: Punctuation, other)

julia> str[6:6]
","
```

The former is a single character value of type Char, while the latter is a string value that happens to contain only a single character. In Julia these are very different things.

10.3   Unicode and UTF-8

Julia fully supports Unicode characters and strings. As discussed above, in character literals, Unicode code points can be represented using Unicode \u and \U escape sequences, as well as all the standard C escape sequences. These can likewise be used to write string literals:

```
julia> s = "\u2200 x \u2203 y"
"∀ x ∃ y"
```

Whether these Unicode characters are displayed as escapes or shown as special characters depends on your terminal's locale settings and its support for Unicode. String literals are encoded using the UTF-8 encoding. UTF-8 is a variable-width encoding, meaning that not all characters are encoded in the same number of bytes. In UTF-8, ASCII characters – i.e. those with code points less than 0x80 (128) – are encoded as they are in ASCII, using a single byte, while code points 0x80 and above are encoded using multiple bytes – up to four per character. This means that not every byte index into a UTF-8 string is necessarily a valid index for a character. If you index into a string at such an invalid byte index, an error is thrown:

```
julia> s[1]
'∀': Unicode U+2200 (category Sm: Symbol, math)

julia> s[2]
ERROR: UnicodeError: invalid character index
[...]

julia> s[3]
ERROR: UnicodeError: invalid character index
[...]

julia> s[4]
' ': ASCII/Unicode U+0020 (category Zs: Separator, space)
```

In this case, the character ∀ is a three-byte character, so the indices 2 and 3 are invalid and the next character's index is 4; this next valid index can be computed by `nextind(s,1)`, and the next index after that by `nextind(s,4)` and so on.

Because of variable-length encodings, the number of characters in a string (given by `length(s)`) is not always the same as the last index. If you iterate through the indices 1 through `endof(s)` and index into s, the sequence of characters returned when errors aren't thrown is the sequence of characters comprising the string s. Thus we have the identity that `length(s) <= endof(s)`, since each character in a string must have its own index. The following is an inefficient and verbose way to iterate through the characters of s:

```
julia> for i = 1:endof(s)
           try
               println(s[i])
           catch
               # ignore the index error
           end
       end
∀

x

y
```

The blank lines actually have spaces on them. Fortunately, the above awkward idiom is unnecessary for iterating through the characters in a string, since you can just use the string as an iterable object, no exception handling required:

```
julia> for c in s
           println(c)
       end
∀

x
```

y

Julia uses the UTF-8 encoding by default, and support for new encodings can be added by packages. For example, the LegacyStrings.jl package implements `UTF16String` and `UTF32String` types. Additional discussion of other encodings and how to implement support for them is beyond the scope of this document for the time being. For further discussion of UTF-8 encoding issues, see the section below on byte array literals. The `transcode()` function is provided to convert data between the various UTF-xx encodings, primarily for working with external data and libraries.

10.4 Concatenation

One of the most common and useful string operations is concatenation:

```
julia> greet = "Hello"
"Hello"

julia> whom = "world"
"world"

julia> string(greet, ", ", whom, ".\n")
"Hello, world.\n"
```

Julia also provides `*` for string concatenation:

```
julia> greet * ", " * whom * ".\n"
"Hello, world.\n"
```

While `*` may seem like a surprising choice to users of languages that provide `+` for string concatenation, this use of `*` has precedent in mathematics, particularly in abstract algebra.

In mathematics, `+` usually denotes a *commutative* operation, where the order of the operands does not matter. An example of this is matrix addition, where `A + B == B + A` for any matrices A and B that have the same shape. In contrast, `*` typically denotes a *noncommutative* operation, where the order of the operands *does* matter. An example of this is matrix multiplication, where in general `A * B != B * A`. As with matrix multiplication, string concatenation is noncommutative: `greet * whom != whom * greet`. As such, `*` is a more natural choice for an infix string concatenation operator, consistent with common mathematical use.

More precisely, the set of all finite-length strings S together with the string concatenation operator `*` forms a free monoid (S, `*`). The identity element of this set is the empty string, `""`. Whenever a free monoid is not commutative, the operation is typically represented as `\cdot`, `*`, or a similar symbol, rather than `+`, which as stated usually implies commutativity.

10.5 Interpolation

Constructing strings using concatenation can become a bit cumbersome, however. To reduce the need for these verbose calls to `string()` or repeated multiplications, Julia allows interpolation into string literals using $, as in Perl:

```
julia> "$greet, $whom.\n"
"Hello, world.\n"
```

This is more readable and convenient and equivalent to the above string concatenation – the system rewrites this apparent single string literal into a concatenation of string literals with variables.

The shortest complete expression after the $ is taken as the expression whose value is to be interpolated into the string. Thus, you can interpolate any expression into a string using parentheses:

```julia
julia> "1 + 2 = $(1 + 2)"
"1 + 2 = 3"
```

Both concatenation and string interpolation call `string()` to convert objects into string form. Most non-`AbstractString` objects are converted to strings closely corresponding to how they are entered as literal expressions:

```julia
julia> v = [1,2,3]
3-element Array{Int64,1}:
 1
 2
 3

julia> "v: $v"
"v: [1, 2, 3]"
```

`string()` is the identity for `AbstractString` and `Char` values, so these are interpolated into strings as themselves, unquoted and unescaped:

```julia
julia> c = 'x'
'x': ASCII/Unicode U+0078 (category Ll: Letter, lowercase)

julia> "hi, $c"
"hi, x"
```

To include a literal $ in a string literal, escape it with a backslash:

```julia
julia> print("I have \$100 in my account.\n")
I have $100 in my account.
```

10.6 Triple-Quoted String Literals

When strings are created using triple-quotes (`"""`...`"""`) they have some special behavior that can be useful for creating longer blocks of text. First, if the opening `"""` is followed by a newline, the newline is stripped from the resulting string.

```julia
"""hello"""
```

is equivalent to

```julia
"""
hello"""
```

but

```julia
"""

hello"""
```

will contain a literal newline at the beginning. Trailing whitespace is left unaltered. They can contain " symbols without escaping. Triple-quoted strings are also dedented to the level of the least-indented line. This is useful for defining strings within code that is indented. For example:

```
julia> str = """
           Hello,
           world.
         """
"  Hello,\n  world.\n"
```

In this case the final (empty) line before the closing `"""` sets the indentation level.

Note that line breaks in literal strings, whether single- or triple-quoted, result in a newline (LF) character `\n` in the string, even if your editor uses a carriage return `\r` (CR) or CRLF combination to end lines. To include a CR in a string, use an explicit escape `\r`; for example, you can enter the literal string `"a CRLF line ending\r\n"`.

10.7   Common Operations

You can lexicographically compare strings using the standard comparison operators:

```
julia> "abracadabra" < "xylophone"
true

julia> "abracadabra" == "xylophone"
false

julia> "Hello, world." != "Goodbye, world."
true

julia> "1 + 2 = 3" == "1 + 2 = $(1 + 2)"
true
```

You can search for the index of a particular character using the `search()` function:

```
julia> search("xylophone", 'x')
1

julia> search("xylophone", 'p')
5

julia> search("xylophone", 'z')
0
```

You can start the search for a character at a given offset by providing a third argument:

```
julia> search("xylophone", 'o')
4

julia> search("xylophone", 'o', 5)
7

julia> search("xylophone", 'o', 8)
0
```

You can use the `contains()` function to check if a substring is contained in a string:

```
julia> contains("Hello, world.", "world")
true

julia> contains("Xylophon", "o")
true

julia> contains("Xylophon", "a")
false

julia> contains("Xylophon", 'o')
ERROR: MethodError: no method matching contains(::String, ::Char)
Closest candidates are:
  contains(!Matched::Function, ::Any, !Matched::Any) at reduce.jl:664
  contains(::AbstractString, !Matched::AbstractString) at strings/search.jl:378
```

The last error is because 'o' is a character literal, and contains() is a generic function that looks for subsequences. To look for an element in a sequence, you must use in() instead.

Two other handy string functions are repeat() and join():

```
julia> repeat(".:Z:.", 10)
".:Z:..:Z:..:Z:..:Z:..:Z:..:Z:..:Z:..:Z:..:Z:..:Z:."

julia> join(["apples", "bananas", "pineapples"], ", ", " and ")
"apples, bananas and pineapples"
```

Some other useful functions include:

- endof(str) gives the maximal (byte) index that can be used to index into str.

- length(str) the number of characters in str.

- i = start(str) gives the first valid index at which a character can be found in str (typically 1).

- c, j = next(str,i) returns next character at or after the index i and the next valid character index following that. With start() and endof(), can be used to iterate through the characters in str.

- ind2chr(str,i) gives the number of characters in str up to and including any at index i.

- chr2ind(str,j) gives the index at which the jth character in str occurs.

10.8 Non-Standard String Literals

There are situations when you want to construct a string or use string semantics, but the behavior of the standard string construct is not quite what is needed. For these kinds of situations, Julia provides non-standard string literals. A non-standard string literal looks like a regular double-quoted string literal, but is immediately prefixed by an identifier, and doesn't behave quite like a normal string literal. Regular expressions, byte array literals and version number literals, as described below, are some examples of non-standard string literals. Other examples are given in the Metaprogramming section.

10.9   Regular Expressions

Julia has Perl-compatible regular expressions (regexes), as provided by the PCRE library. Regular expressions are related
to strings in two ways: the obvious connection is that regular expressions are used to find regular patterns in strings;
the other connection is that regular expressions are themselves input as strings, which are parsed into a state machine
that can be used to efficiently search for patterns in strings. In Julia, regular expressions are input using non-standard
string literals prefixed with various identifiers beginning with r. The most basic regular expression literal without any
options turned on just uses r"...":

```julia
julia> r"^\s*(?:#|$)"
r"^\s*(?:#|$)"

julia> typeof(ans)
Regex
```

To check if a regex matches a string, use ismatch():

```julia
julia> ismatch(r"^\s*(?:#|$)", "not a comment")
false

julia> ismatch(r"^\s*(?:#|$)", "# a comment")
true
```

As one can see here, ismatch() simply returns true or false, indicating whether the given regex matches the string or
not. Commonly, however, one wants to know not just whether a string matched, but also *how* it matched. To capture
this information about a match, use the match() function instead:

```julia
julia> match(r"^\s*(?:#|$)", "not a comment")

julia> match(r"^\s*(?:#|$)", "# a comment")
RegexMatch("#")
```

If the regular expression does not match the given string, match() returns nothing – a special value that does not
print anything at the interactive prompt. Other than not printing, it is a completely normal value and you can test for
it programmatically:

```julia
m = match(r"^\s*(?:#|$)", line)
if m === nothing
    println("not a comment")
else
    println("blank or comment")
end
```

If a regular expression does match, the value returned by match() is a RegexMatch object. These objects record how
the expression matches, including the substring that the pattern matches and any captured substrings, if there are any.
This example only captures the portion of the substring that matches, but perhaps we want to capture any non-blank
text after the comment character. We could do the following:

```julia
julia> m = match(r"^\s*(?:#\s*(.*?)\s*$|$)", "# a comment ")
RegexMatch("# a comment ", 1="a comment")
```

When calling `match()`, you have the option to specify an index at which to start the search. For example:

```julia
julia> m = match(r"[0-9]","aaaa1aaaa2aaaa3",1)
RegexMatch("1")

julia> m = match(r"[0-9]","aaaa1aaaa2aaaa3",6)
RegexMatch("2")

julia> m = match(r"[0-9]","aaaa1aaaa2aaaa3",11)
RegexMatch("3")
```

You can extract the following info from a `RegexMatch` object:

- the entire substring matched: `m.match`

- the captured substrings as an array of strings: `m.captures`

- the offset at which the whole match begins: `m.offset`

- the offsets of the captured substrings as a vector: `m.offsets`

For when a capture doesn't match, instead of a substring, `m.captures` contains `nothing` in that position, and `m.off-sets` has a zero offset (recall that indices in Julia are 1-based, so a zero offset into a string is invalid). Here is a pair of somewhat contrived examples:

```julia
julia> m = match(r"(a|b)(c)?(d)", "acd")
RegexMatch("acd", 1="a", 2="c", 3="d")

julia> m.match
"acd"

julia> m.captures
3-element Array{Union{SubString{String}, Void},1}:
 "a"
 "c"
 "d"

julia> m.offset
1

julia> m.offsets
3-element Array{Int64,1}:
 1
 2
 3

julia> m = match(r"(a|b)(c)?(d)", "ad")
RegexMatch("ad", 1="a", 2=nothing, 3="d")

julia> m.match
"ad"

julia> m.captures
3-element Array{Union{SubString{String}, Void},1}:
 "a"
```

```
 nothing
 "d"

julia> m.offset
1

julia> m.offsets
3-element Array{Int64,1}:
 1
 0
 2
```

It is convenient to have captures returned as an array so that one can use destructuring syntax to bind them to local variables:

```
julia> first, second, third = m.captures; first
"a"
```

Captures can also be accessed by indexing the `RegexMatch` object with the number or name of the capture group:

```
julia> m=match(r"(?<hour>\d+):(?<minute>\d+)","12:45")
RegexMatch("12:45", hour="12", minute="45")

julia> m[:minute]
"45"

julia> m[2]
"45"
```

Captures can be referenced in a substitution string when using `replace()` by using \n to refer to the nth capture group and prefixing the subsitution string with s. Capture group 0 refers to the entire match object. Named capture groups can be referenced in the substitution with g<groupname>. For example:

```
julia> replace("first second", r"(\w+) (?<agroup>\w+)", s"\g<agroup> \1")
"second first"
```

Numbered capture groups can also be referenced as \g<n> for disambiguation, as in:

```
julia> replace("a", r".", s"\g<0>1")
"a1"
```

You can modify the behavior of regular expressions by some combination of the flags i, m, s, and x after the closing double quote mark. These flags have the same meaning as they do in Perl, as explained in this excerpt from the perlre manpage:

```
i   Do case-insensitive pattern matching.

    If locale matching rules are in effect, the case map is taken
    from the current locale for code points less than 255, and
    from Unicode rules for larger code points. However, matches
    that would cross the Unicode rules/non-Unicode rules boundary
    (ords 255/256) will not succeed.

m   Treat string as multiple lines.  That is, change "^" and "$"
```

> from matching the start or end of the string to matching the
> start or end of any line anywhere within the string.

> s   Treat string as single line. That is, change "." to match any
> character whatsoever, even a newline, which normally it would
> not match.

> Used together, as r""ms, they let the "." match any character
> whatsoever, while still allowing "^" and "$" to match,
> respectively, just after and just before newlines within the
> string.

> x   Tells the regular expression parser to ignore most whitespace
> that is neither backslashed nor within a character class. You
> can use this to break up your regular expression into
> (slightly) more readable parts. The '#' character is also
> treated as a metacharacter introducing a comment, just as in
> ordinary code.

For example, the following regex has all three flags turned on:

```julia
julia> r"a+.*b+.*?d$"ism
r"a+.*b+.*?d$"ims

julia> match(r"a+.*b+.*?d$"ism, "Goodbye,\nOh, angry,\nBad world\n")
RegexMatch("angry,\nBad world")
```

Triple-quoted regex strings, of the form r"""...""", are also supported (and may be convenient for regular expressions containing quotation marks or newlines).

10.10   Byte Array Literals

Another useful non-standard string literal is the byte-array string literal: b"...". This form lets you use string notation to express literal byte arrays – i.e. arrays of UInt8 values. The rules for byte array literals are the following:

- ASCII characters and ASCII escapes produce a single byte.

- \x and octal escape sequences produce the *byte* corresponding to the escape value.

- Unicode escape sequences produce a sequence of bytes encoding that code point in UTF-8.

There is some overlap between these rules since the behavior of \x and octal escapes less than 0x80 (128) are covered by both of the first two rules, but here these rules agree. Together, these rules allow one to easily use ASCII characters, arbitrary byte values, and UTF-8 sequences to produce arrays of bytes. Here is an example using all three:

```julia
julia> b"DATA\xff\u2200"
8-element Array{UInt8,1}:
 0x44
 0x41
 0x54
 0x41
 0xff
 0xe2
 0x88
 0x80
```

The ASCII string "DATA" corresponds to the bytes 68, 65, 84, 65. `\xff` produces the single byte 255. The Unicode escape `\u2200` is encoded in UTF-8 as the three bytes 226, 136, 128. Note that the resulting byte array does not correspond to a valid UTF-8 string – if you try to use this as a regular string literal, you will get a syntax error:

```julia
julia> "DATA\xff\u2200"
ERROR: syntax: invalid UTF-8 sequence
```

Also observe the significant distinction between `\xff` and `\uff`: the former escape sequence encodes the *byte 255*, whereas the latter escape sequence represents the *code point 255*, which is encoded as two bytes in UTF-8:

```julia
julia> b"\xff"
1-element Array{UInt8,1}:
 0xff

julia> b"\uff"
2-element Array{UInt8,1}:
 0xc3
 0xbf
```

In character literals, this distinction is glossed over and `\xff` is allowed to represent the code point 255, because characters *always* represent code points. In strings, however, `\x` escapes always represent bytes, not code points, whereas `\u` and `\U` escapes always represent code points, which are encoded in one or more bytes. For code points less than `\u80`, it happens that the UTF-8 encoding of each code point is just the single byte produced by the corresponding `\x` escape, so the distinction can safely be ignored. For the escapes `\x80` through `\xff` as compared to `\u80` through `\uff`, however, there is a major difference: the former escapes all encode single bytes, which – unless followed by very specific continuation bytes – do not form valid UTF-8 data, whereas the latter escapes all represent Unicode code points with two-byte encodings.

If this is all extremely confusing, try reading "The Absolute Minimum Every Software Developer Absolutely, Positively Must Know About Unicode and Character Sets". It's an excellent introduction to Unicode and UTF-8, and may help alleviate some confusion regarding the matter.

10.11   Version Number Literals

Version numbers can easily be expressed with non-standard string literals of the form `v"..."`. Version number literals create `VersionNumber` objects which follow the specifications of semantic versioning, and therefore are composed of major, minor and patch numeric values, followed by pre-release and build alpha-numeric annotations. For example, `v"0.2.1-rc1+win64"` is broken into major version 0, minor version 2, patch version 1, pre-release `rc1` and build `win64`. When entering a version literal, everything except the major version number is optional, therefore e.g. `v"0.2"` is equivalent to `v"0.2.0"` (with empty pre-release/build annotations), `v"2"` is equivalent to `v"2.0.0"`, and so on.

`VersionNumber` objects are mostly useful to easily and correctly compare two (or more) versions. For example, the constant `VERSION` holds Julia version number as a `VersionNumber` object, and therefore one can define some version-specific behavior using simple statements as:

```julia
if v"0.2" <= VERSION < v"0.3-"
    # do something specific to 0.2 release series
end
```

Note that in the above example the non-standard version number `v"0.3-"` is used, with a trailing `-`: this notation is a Julia extension of the standard, and it's used to indicate a version which is lower than any `0.3` release, including all of its pre-releases. So in the above example the code would only run with stable `0.2` versions, and exclude such versions

as `v"0.3.0-rc1"`. In order to also allow for unstable (i.e. pre-release) `0.2` versions, the lower bound check should be modified like this: `v"0.2-" <= VERSION`.

Another non-standard version specification extension allows one to use a trailing `+` to express an upper limit on build versions, e.g. `VERSION > v"0.2-rc1+"` can be used to mean any version above `0.2-rc1` and any of its builds: it will return `false` for version `v"0.2-rc1+win64"` and `true` for `v"0.2-rc2"`.

It is good practice to use such special versions in comparisons (particularly, the trailing – should always be used on upper bounds unless there's a good reason not to), but they must not be used as the actual version number of anything, as they are invalid in the semantic versioning scheme.

Besides being used for the `VERSION` constant, `VersionNumber` objects are widely used in the Pkg module, to specify packages versions and their dependencies.

10.12 Raw String Literals

Raw strings without interpolation or unescaping can be expressed with non-standard string literals of the form `raw"..."`. Raw string literals create ordinary `String` objects which contain the enclosed contents exactly as entered with no interpolation or unescaping. This is useful for strings which contain code or markup in other languages which use `$` or `\` as special characters. The exception is quotation marks that still must be escaped, e.g. `raw"\""` is equivalent to `"\""`.

Chapter 11

Functions

In Julia, a function is an object that maps a tuple of argument values to a return value. Julia functions are not pure mathematical functions, in the sense that functions can alter and be affected by the global state of the program. The basic syntax for defining functions in Julia is:

```julia
julia> function f(x,y)
           x + y
       end
f (generic function with 1 method)
```

There is a second, more terse syntax for defining a function in Julia. The traditional function declaration syntax demonstrated above is equivalent to the following compact "assignment form":

```julia
julia> f(x,y) = x + y
f (generic function with 1 method)
```

In the assignment form, the body of the function must be a single expression, although it can be a compound expression (see Compound Expressions). Short, simple function definitions are common in Julia. The short function syntax is accordingly quite idiomatic, considerably reducing both typing and visual noise.

A function is called using the traditional parenthesis syntax:

```julia
julia> f(2,3)
5
```

Without parentheses, the expression f refers to the function object, and can be passed around like any value:

```julia
julia> g = f;

julia> g(2,3)
5
```

As with variables, Unicode can also be used for function names:

```julia
julia> ∑(x,y) = x + y
∑ (generic function with 1 method)

julia> ∑(2, 3)
5
```

11.1   Argument Passing Behavior

Julia function arguments follow a convention sometimes called "pass-by-sharing", which means that values are not copied when they are passed to functions. Function arguments themselves act as new variable *bindings* (new locations that can refer to values), but the values they refer to are identical to the passed values. Modifications to mutable values (such as `Arrays`) made within a function will be visible to the caller. This is the same behavior found in Scheme, most Lisps, Python, Ruby and Perl, among other dynamic languages.

11.2   The `return` Keyword

The value returned by a function is the value of the last expression evaluated, which, by default, is the last expression in the body of the function definition. In the example function, `f`, from the previous section this is the value of the expression `x + y`. As in C and most other imperative or functional languages, the `return` keyword causes a function to return immediately, providing an expression whose value is returned:

```julia
function g(x,y)
    return x * y
    x + y
end
```

Since function definitions can be entered into interactive sessions, it is easy to compare these definitions:

```julia
julia> f(x,y) = x + y
f (generic function with 1 method)

julia> function g(x,y)
           return x * y
           x + y
       end
g (generic function with 1 method)

julia> f(2,3)
5

julia> g(2,3)
6
```

Of course, in a purely linear function body like g, the usage of `return` is pointless since the expression `x + y` is never evaluated and we could simply make `x * y` the last expression in the function and omit the `return`. In conjunction with other control flow, however, `return` is of real use. Here, for example, is a function that computes the hypotenuse length of a right triangle with sides of length x and y, avoiding overflow:

```julia
julia> function hypot(x,y)
           x = abs(x)
           y = abs(y)
           if x > y
               r = y/x
               return x*sqrt(1+r*r)
           end
           if y == 0
               return zero(x)
           end
```

```
          r = x/y
          return y*sqrt(1+r*r)
      end
hypot (generic function with 1 method)

julia> hypot(3, 4)
5.0
```

There are three possible points of return from this function, returning the values of three different expressions, depending on the values of x and y. The `return` on the last line could be omitted since it is the last expression.

11.3   Operators Are Functions

In Julia, most operators are just functions with support for special syntax. (The exceptions are operators with special evaluation semantics like && and ||. These operators cannot be functions since Short-Circuit Evaluation requires that their operands are not evaluated before evaluation of the operator.) Accordingly, you can also apply them using parenthesized argument lists, just as you would any other function:

```
julia> 1 + 2 + 3
6

julia> +(1,2,3)
6
```

The infix form is exactly equivalent to the function application form – in fact the former is parsed to produce the function call internally. This also means that you can assign and pass around operators such as +() and *() just like you would with other function values:

```
julia> f = +;

julia> f(1,2,3)
6
```

Under the name f, the function does not support infix notation, however.

11.4   Operators With Special Names

A few special expressions correspond to calls to functions with non-obvious names. These are:

Expression	Calls
`[A B C ...]`	`hcat()`
`[A; B; C; ...]`	`vcat()`
`[A B; C D; ...]`	`hvcat()`
`A'`	`ctranspose()`
`A.'`	`transpose()`
`1:n`	`colon()`
`A[i]`	`getindex()`
`A[i]=x`	`setindex!()`

These functions are included in the `Base.Operators` module even though they do not have operator-like names.

11.5   Anonymous Functions

Functions in Julia are first-class objects: they can be assigned to variables, and called using the standard function call syntax from the variable they have been assigned to. They can be used as arguments, and they can be returned as values. They can also be created anonymously, without being given a name, using either of these syntaxes:

```julia
julia> x -> x^2 + 2x - 1
(::#1) (generic function with 1 method)

julia> function (x)
           x^2 + 2x - 1
       end
(::#3) (generic function with 1 method)
```

This creates a function taking one argument x and returning the value of the polynomial $x^2 + 2x - 1$ at that value. Notice that the result is a generic function, but with a compiler-generated name based on consecutive numbering.

The primary use for anonymous functions is passing them to functions which take other functions as arguments. A classic example is map(), which applies a function to each value of an array and returns a new array containing the resulting values:

```julia
julia> map(round, [1.2,3.5,1.7])
3-element Array{Float64,1}:
 1.0
 4.0
 2.0
```

This is fine if a named function effecting the transform one wants already exists to pass as the first argument to map(). Often, however, a ready-to-use, named function does not exist. In these situations, the anonymous function construct allows easy creation of a single-use function object without needing a name:

```julia
julia> map(x -> x^2 + 2x - 1, [1,3,-1])
3-element Array{Int64,1}:
  2
 14
 -2
```

An anonymous function accepting multiple arguments can be written using the syntax (x,y,z)->2x+y-z. A zero-argument anonymous function is written as ()->3. The idea of a function with no arguments may seem strange, but is useful for "delaying" a computation. In this usage, a block of code is wrapped in a zero-argument function, which is later invoked by calling it as f().

11.6   Multiple Return Values

In Julia, one returns a tuple of values to simulate returning multiple values. However, tuples can be created and destructured without needing parentheses, thereby providing an illusion that multiple values are being returned, rather than a single tuple value. For example, the following function returns a pair of values:

```julia
julia> function foo(a,b)
           a+b, a*b
       end
foo (generic function with 1 method)
```

If you call it in an interactive session without assigning the return value anywhere, you will see the tuple returned:

```julia
julia> foo(2,3)
(5, 6)
```

A typical usage of such a pair of return values, however, extracts each value into a variable. Julia supports simple tuple "destructuring" that facilitates this:

```julia
julia> x, y = foo(2,3)
(5, 6)

julia> x
5

julia> y
6
```

You can also return multiple values via an explicit usage of the `return` keyword:

```julia
function foo(a,b)
    return a+b, a*b
end
```

This has the exact same effect as the previous definition of `foo`.

11.7   Varargs Functions

It is often convenient to be able to write functions taking an arbitrary number of arguments. Such functions are traditionally known as "varargs" functions, which is short for "variable number of arguments". You can define a varargs function by following the last argument with an ellipsis:

```julia
julia> bar(a,b,x...) = (a,b,x)
bar (generic function with 1 method)
```

The variables a and b are bound to the first two argument values as usual, and the variable x is bound to an iterable collection of the zero or more values passed to `bar` after its first two arguments:

```julia
julia> bar(1,2)
(1, 2, ())

julia> bar(1,2,3)
(1, 2, (3,))

julia> bar(1, 2, 3, 4)
(1, 2, (3, 4))

julia> bar(1,2,3,4,5,6)
(1, 2, (3, 4, 5, 6))
```

In all these cases, x is bound to a tuple of the trailing values passed to `bar`.

It is possible to constrain the number of values passed as a variable argument; this will be discussed later in Parametrically-constrained Varargs methods.

On the flip side, it is often handy to "splice" the values contained in an iterable collection into a function call as individual arguments. To do this, one also uses `...` but in the function call instead:

```
julia> x = (3, 4)
(3, 4)

julia> bar(1,2,x...)
(1, 2, (3, 4))
```

In this case a tuple of values is spliced into a varargs call precisely where the variable number of arguments go. This need not be the case, however:

```
julia> x = (2, 3, 4)
(2, 3, 4)

julia> bar(1,x...)
(1, 2, (3, 4))

julia> x = (1, 2, 3, 4)
(1, 2, 3, 4)

julia> bar(x...)
(1, 2, (3, 4))
```

Furthermore, the iterable object spliced into a function call need not be a tuple:

```
julia> x = [3,4]
2-element Array{Int64,1}:
 3
 4

julia> bar(1,2,x...)
(1, 2, (3, 4))

julia> x = [1,2,3,4]
4-element Array{Int64,1}:
 1
 2
 3
 4

julia> bar(x...)
(1, 2, (3, 4))
```

Also, the function that arguments are spliced into need not be a varargs function (although it often is):

```
julia> baz(a,b) = a + b;

julia> args = [1,2]
2-element Array{Int64,1}:
 1
 2

julia> baz(args...)
3

julia> args = [1,2,3]
3-element Array{Int64,1}:
 1
```

```
  2
  3

julia> baz(args...)
ERROR: MethodError: no method matching baz(::Int64, ::Int64, ::Int64)
Closest candidates are:
  baz(::Any, ::Any) at none:1
```

As you can see, if the wrong number of elements are in the spliced container, then the function call will fail, just as it would if too many arguments were given explicitly.

11.8 Optional Arguments

In many cases, function arguments have sensible default values and therefore might not need to be passed explicitly in every call. For example, the library function `parse(T, num, base)` interprets a string as a number in some base. The `base` argument defaults to `10`. This behavior can be expressed concisely as:

```
function parse(type, num, base=10)
    ###
end
```

With this definition, the function can be called with either two or three arguments, and `10` is automatically passed when a third argument is not specified:

```
julia> parse(Int,"12",10)
12

julia> parse(Int,"12",3)
5

julia> parse(Int,"12")
12
```

Optional arguments are actually just a convenient syntax for writing multiple method definitions with different numbers of arguments (see Note on Optional and keyword Arguments).

11.9 Keyword Arguments

Some functions need a large number of arguments, or have a large number of behaviors. Remembering how to call such functions can be difficult. Keyword arguments can make these complex interfaces easier to use and extend by allowing arguments to be identified by name instead of only by position.

For example, consider a function `plot` that plots a line. This function might have many options, for controlling line style, width, color, and so on. If it accepts keyword arguments, a possible call might look like `plot(x, y, width=2)`, where we have chosen to specify only line width. Notice that this serves two purposes. The call is easier to read, since we can label an argument with its meaning. It also becomes possible to pass any subset of a large number of arguments, in any order.

Functions with keyword arguments are defined using a semicolon in the signature:

```
function plot(x, y; style="solid", width=1, color="black")
    ###
end
```

When the function is called, the semicolon is optional: one can either call `plot(x, y, width=2)` or `plot(x, y; width=2)`, but the former style is more common. An explicit semicolon is required only for passing varargs or computed keywords as described below.

Keyword argument default values are evaluated only when necessary (when a corresponding keyword argument is not passed), and in left-to-right order. Therefore default expressions may refer to prior keyword arguments.

The types of keyword arguments can be made explicit as follows:

```julia
function f(;x::Int64=1)
    ###
end
```

Extra keyword arguments can be collected using `...`, as in varargs functions:

```julia
function f(x; y=0, kwargs...)
    ###
end
```

Inside `f`, `kwargs` will be a collection of `(key,value)` tuples, where each `key` is a symbol. Such collections can be passed as keyword arguments using a semicolon in a call, e.g. `f(x, z=1; kwargs...)`. Dictionaries can also be used for this purpose.

One can also pass `(key,value)` tuples, or any iterable expression (such as a `=>` pair) that can be assigned to such a tuple, explicitly after a semicolon. For example, `plot(x, y; (:width,2))` and `plot(x, y; :width => 2)` are equivalent to `plot(x, y, width=2)`. This is useful in situations where the keyword name is computed at runtime.

The nature of keyword arguments makes it possible to specify the same argument more than once. For example, in the call `plot(x, y; options..., width=2)` it is possible that the `options` structure also contains a value for `width`. In such a case the rightmost occurrence takes precedence; in this example, `width` is certain to have the value 2.

11.10   Evaluation Scope of Default Values

Optional and keyword arguments differ slightly in how their default values are evaluated. When optional argument default expressions are evaluated, only *previous* arguments are in scope. In contrast, *all* the arguments are in scope when keyword arguments default expressions are evaluated. For example, given this definition:

```julia
function f(x, a=b, b=1)
    ###
end
```

the b in a=b refers to a b in an outer scope, not the subsequent argument b. However, if a and b were keyword arguments instead, then both would be created in the same scope and the b in a=b would refer to the subsequent argument b (shadowing any b in an outer scope), which would result in an undefined variable error (since the default expressions are evaluated left-to-right, and b has not been assigned yet).

11.11   Do-Block Syntax for Function Arguments

Passing functions as arguments to other functions is a powerful technique, but the syntax for it is not always convenient. Such calls are especially awkward to write when the function argument requires multiple lines. As an example, consider calling `map()` on a function with several cases:

```
map(x->begin
           if x < 0 && iseven(x)
                return 0
           elseif x == 0
                return 1
           else
                return x
           end
        end,
    [A, B, C])
```

Julia provides a reserved word do for rewriting this code more clearly:

```
map([A, B, C]) do x
    if x < 0 && iseven(x)
        return 0
    elseif x == 0
        return 1
    else
        return x
    end
end
```

The do x syntax creates an anonymous function with argument x and passes it as the first argument to map(). Similarly, do a,b would create a two-argument anonymous function, and a plain do would declare that what follows is an anonymous function of the form () -> ....

How these arguments are initialized depends on the "outer" function; here, map() will sequentially set x to A, B, C, calling the anonymous function on each, just as would happen in the syntax map(func, [A, B, C]).

This syntax makes it easier to use functions to effectively extend the language, since calls look like normal code blocks. There are many possible uses quite different from map(), such as managing system state. For example, there is a version of open() that runs code ensuring that the opened file is eventually closed:

```
open("outfile", "w") do io
    write(io, data)
end
```

This is accomplished by the following definition:

```
function open(f::Function, args...)
    io = open(args...)
    try
        f(io)
    finally
        close(io)
    end
end
```

Here, open() first opens the file for writing and then passes the resulting output stream to the anonymous function you defined in the do ... end block. After your function exits, open() will make sure that the stream is properly closed, regardless of whether your function exited normally or threw an exception. (The try/finally construct will be described in Control Flow.)

With the do block syntax, it helps to check the documentation or implementation to know how the arguments of the user function are initialized.

11.12   Dot Syntax for Vectorizing Functions

In technical-computing languages, it is common to have "vectorized" versions of functions, which simply apply a given function `f(x)` to each element of an array `A` to yield a new array via `f(A)`. This kind of syntax is convenient for data processing, but in other languages vectorization is also often required for performance: if loops are slow, the "vectorized" version of a function can call fast library code written in a low-level language. In Julia, vectorized functions are *not* required for performance, and indeed it is often beneficial to write your own loops (see Performance Tips), but they can still be convenient. Therefore, *any* Julia function f can be applied elementwise to any array (or other collection) with the syntax `f.(A)`. For example `sin` can be applied to all elements in the vector A, like so:

```julia
julia> A = [1.0, 2.0, 3.0]
3-element Array{Float64,1}:
 1.0
 2.0
 3.0

julia> sin.(A)
3-element Array{Float64,1}:
 0.841471
 0.909297
 0.14112
```

Of course, you can omit the dot if you write a specialized "vector" method of f, e.g. via `f(A::AbstractArray) = map(f, A)`, and this is just as efficient as `f.(A)`. But that approach requires you to decide in advance which functions you want to vectorize.

More generally, `f.(args...)` is actually equivalent to `broadcast(f, args...)`, which allows you to operate on multiple arrays (even of different shapes), or a mix of arrays and scalars (see Broadcasting). For example, if you have `f(x,y) = 3x + 4y`, then `f.(pi,A)` will return a new array consisting of `f(pi,a)` for each a in A, and `f.(vec-tor1,vector2)` will return a new vector consisting of `f(vector1[i],vector2[i])` for each index i (throwing an exception if the vectors have different length).

```julia
julia> f(x,y) = 3x + 4y;

julia> A = [1.0, 2.0, 3.0];

julia> B = [4.0, 5.0, 6.0];

julia> f.(pi, A)
3-element Array{Float64,1}:
 13.4248
 17.4248
 21.4248

julia> f.(A, B)
3-element Array{Float64,1}:
 19.0
 26.0
 33.0
```

Moreover, *nested* `f.(args...)` calls are *fused* into a single `broadcast` loop. For example, `sin.(cos.(X))` is equivalent to `broadcast(x -> sin(cos(x)), X)`, similar to `[sin(cos(x)) for x in X]`: there is only a single loop over X, and a single array is allocated for the result. [In contrast, `sin(cos(X))` in a typical "vectorized" language

would first allocate one temporary array for `tmp=cos(X)`, and then compute `sin(tmp)` in a separate loop, allocating a second array.] This loop fusion is not a compiler optimization that may or may not occur, it is a *syntactic guarantee* whenever nested `f.(args...)` calls are encountered. Technically, the fusion stops as soon as a "non-dot" function call is encountered; for example, in `sin.(sort(cos.(X)))` the `sin` and `cos` loops cannot be merged because of the intervening `sort` function.

Finally, the maximum efficiency is typically achieved when the output array of a vectorized operation is *pre-allocated*, so that repeated calls do not allocate new arrays over and over again for the results (Pre-allocating outputs:). A convenient syntax for this is `X .= ...`, which is equivalent to `broadcast!(identity, X, ...)` except that, as above, the `broadcast!` loop is fused with any nested "dot" calls. For example, `X .= sin.(Y)` is equivalent to `broadcast!(sin, X, Y)`, overwriting X with `sin.(Y)` in-place. If the left-hand side is an array-indexing expression, e.g. `X[2:end] .= sin.(Y)`, then it translates to `broadcast!` on a view, e.g. `broadcast!(sin, view(X, 2:endof(X)), Y)`, so that the left-hand side is updated in-place.

Since adding dots to many operations and function calls in an expression can be tedious and lead to code that is difficult to read, the macro `@.` is provided to convert *every* function call, operation, and assignment in an expression into the "dotted" version.

```julia
julia> Y = [1.0, 2.0, 3.0, 4.0];

julia> X = similar(Y); # pre-allocate output array

julia> @. X = sin(cos(Y)) # equivalent to X .= sin.(cos.(Y))
4-element Array{Float64,1}:
  0.514395
 -0.404239
 -0.836022
 -0.608083
```

Binary (or unary) operators like `.+` are handled with the same mechanism: they are equivalent to `broadcast` calls and are fused with other nested "dot" calls. `X .+= Y` etcetera is equivalent to `X .= X .+ Y` and results in a fused in-place assignment; see also dot operators.

11.13 Further Reading

We should mention here that this is far from a complete picture of defining functions. Julia has a sophisticated type system and allows multiple dispatch on argument types. None of the examples given here provide any type annotations on their arguments, meaning that they are applicable to all types of arguments. The type system is described in Types and defining a function in terms of methods chosen by multiple dispatch on run-time argument types is described in Methods.

Chapter 12

Control Flow

Julia provides a variety of control flow constructs:

- Compound Expressions: `begin` and `(;)`.

- Conditional Evaluation: `if-elseif-else` and `?:` (ternary operator).

- Short-Circuit Evaluation: `&&`, `||` and chained comparisons.

- Repeated Evaluation: Loops: `while` and `for`.

- Exception Handling: `try-catch`, `error()` and `throw()`.

- Tasks (aka Coroutines): `yieldto()`.

The first five control flow mechanisms are standard to high-level programming languages. Tasks are not so standard: they provide non-local control flow, making it possible to switch between temporarily-suspended computations. This is a powerful construct: both exception handling and cooperative multitasking are implemented in Julia using tasks. Everyday programming requires no direct usage of tasks, but certain problems can be solved much more easily by using tasks.

12.1 Compound Expressions

Sometimes it is convenient to have a single expression which evaluates several subexpressions in order, returning the value of the last subexpression as its value. There are two Julia constructs that accomplish this: `begin` blocks and `(;)` chains. The value of both compound expression constructs is that of the last subexpression. Here's an example of a `begin` block:

```
julia> z = begin
           x = 1
           y = 2
           x + y
       end
3
```

Since these are fairly small, simple expressions, they could easily be placed onto a single line, which is where the `(;)` chain syntax comes in handy:

```julia
julia> z = (x = 1; y = 2; x + y)
3
```

This syntax is particularly useful with the terse single-line function definition form introduced in Functions. Although it is typical, there is no requirement that `begin` blocks be multiline or that (;) chains be single-line:

```julia
julia> begin x = 1; y = 2; x + y end
3

julia> (x = 1;
        y = 2;
        x + y)
3
```

12.2   Conditional Evaluation

Conditional evaluation allows portions of code to be evaluated or not evaluated depending on the value of a boolean expression. Here is the anatomy of the `if-elseif-else` conditional syntax:

```julia
if x < y
    println("x is less than y")
elseif x > y
    println("x is greater than y")
else
    println("x is equal to y")
end
```

If the condition expression x < y is `true`, then the corresponding block is evaluated; otherwise the condition expression x > y is evaluated, and if it is `true`, the corresponding block is evaluated; if neither expression is true, the `else` block is evaluated. Here it is in action:

```julia
julia> function test(x, y)
           if x < y
               println("x is less than y")
           elseif x > y
               println("x is greater than y")
           else
               println("x is equal to y")
           end
       end
test (generic function with 1 method)

julia> test(1, 2)
x is less than y

julia> test(2, 1)
x is greater than y

julia> test(1, 1)
x is equal to y
```

The `elseif` and `else` blocks are optional, and as many `elseif` blocks as desired can be used. The condition expressions in the `if-elseif-else` construct are evaluated until the first one evaluates to `true`, after which the associated block is evaluated, and no further condition expressions or blocks are evaluated.

`if` blocks are "leaky", i.e. they do not introduce a local scope. This means that new variables defined inside the `if` clauses can be used after the `if` block, even if they weren't defined before. So, we could have defined the `test` function above as

```julia
julia> function test(x,y)
           if x < y
               relation = "less than"
           elseif x == y
               relation = "equal to"
           else
               relation = "greater than"
           end
           println("x is ", relation, " y.")
       end
test (generic function with 1 method)

julia> test(2, 1)
x is greater than y.
```

The variable `relation` is declared inside the `if` block, but used outside. However, when depending on this behavior, make sure all possible code paths define a value for the variable. The following change to the above function results in a runtime error

```julia
julia> function test(x,y)
           if x < y
               relation = "less than"
           elseif x == y
               relation = "equal to"
           end
           println("x is ", relation, " y.")
       end
test (generic function with 1 method)

julia> test(1,2)
x is less than y.

julia> test(2,1)
ERROR: UndefVarError: relation not defined
Stacktrace:
 [1] test(::Int64, ::Int64) at ./none:7
```

`if` blocks also return a value, which may seem unintuitive to users coming from many other languages. This value is simply the return value of the last executed statement in the branch that was chosen, so

```julia
julia> x = 3
3

julia> if x > 0
           "positive!"
       else
```

```
            "negative..."
        end
"positive!"
```

Note that very short conditional statements (one-liners) are frequently expressed using Short-Circuit Evaluation in Julia, as outlined in the next section.

Unlike C, MATLAB, Perl, Python, and Ruby – but like Java, and a few other stricter, typed languages – it is an error if the value of a conditional expression is anything but `true` or `false`:

```
julia> if 1
            println("true")
        end
ERROR: TypeError: non-boolean (Int64) used in boolean context
```

This error indicates that the conditional was of the wrong type: `Int64` rather than the required `Bool`.

The so-called "ternary operator", `? :`, is closely related to the `if-elseif-else` syntax, but is used where a conditional choice between single expression values is required, as opposed to conditional execution of longer blocks of code. It gets its name from being the only operator in most languages taking three operands:

```
a ? b : c
```

The expression a, before the `?`, is a condition expression, and the ternary operation evaluates the expression b, before the `:`, if the condition a is `true` or the expression c, after the `:`, if it is `false`.

The easiest way to understand this behavior is to see an example. In the previous example, the `println` call is shared by all three branches: the only real choice is which literal string to print. This could be written more concisely using the ternary operator. For the sake of clarity, let's try a two-way version first:

```
julia> x = 1; y = 2;

julia> println(x < y ? "less than" : "not less than")
less than

julia> x = 1; y = 0;

julia> println(x < y ? "less than" : "not less than")
not less than
```

If the expression `x < y` is true, the entire ternary operator expression evaluates to the string `"less than"` and otherwise it evaluates to the string `"not less than"`. The original three-way example requires chaining multiple uses of the ternary operator together:

```
julia> test(x, y) = println(x < y ? "x is less than y"    :
                            x > y ? "x is greater than y" : "x is equal to y")
test (generic function with 1 method)

julia> test(1, 2)
x is less than y

julia> test(2, 1)
x is greater than y
```

```
julia> test(1, 1)
x is equal to y
```

To facilitate chaining, the operator associates from right to left.

It is significant that like `if-elseif-else`, the expressions before and after the `:` are only evaluated if the condition expression evaluates to `true` or `false`, respectively:

```
julia> v(x) = (println(x); x)
v (generic function with 1 method)

julia> 1 < 2 ? v("yes") : v("no")
yes
"yes"

julia> 1 > 2 ? v("yes") : v("no")
no
"no"
```

12.3   Short-Circuit Evaluation

Short-circuit evaluation is quite similar to conditional evaluation. The behavior is found in most imperative programming languages having the && and || boolean operators: in a series of boolean expressions connected by these operators, only the minimum number of expressions are evaluated as are necessary to determine the final boolean value of the entire chain. Explicitly, this means that:

- In the expression a && b, the subexpression b is only evaluated if a evaluates to `true`.

- In the expression a || b, the subexpression b is only evaluated if a evaluates to `false`.

The reasoning is that a && b must be `false` if a is `false`, regardless of the value of b, and likewise, the value of a || b must be true if a is `true`, regardless of the value of b. Both && and || associate to the right, but && has higher precedence than || does. It's easy to experiment with this behavior:

```
julia> t(x) = (println(x); true)
t (generic function with 1 method)

julia> f(x) = (println(x); false)
f (generic function with 1 method)

julia> t(1) && t(2)
1
2
true

julia> t(1) && f(2)
1
2
false

julia> f(1) && t(2)
1
false
```

```
julia> f(1) && f(2)
1
false

julia> t(1) || t(2)
1
true

julia> t(1) || f(2)
1
true

julia> f(1) || t(2)
1
2
true

julia> f(1) || f(2)
1
2
false
```

You can easily experiment in the same way with the associativity and precedence of various combinations of && and
|| operators.

This behavior is frequently used in Julia to form an alternative to very short `if` statements. Instead of `if <cond>`
`<statement> end`, one can write `<cond> && <statement>` (which could be read as: `<cond>` *and then* `<state-`
`ment>`). Similarly, instead of `if ! <cond> <statement> end`, one can write `<cond> || <statement>` (which
could be read as: `<cond>` *or else* `<statement>`).

For example, a recursive factorial routine could be defined like this:

```
julia> function fact(n::Int)
           n >= 0 || error("n must be non-negative")
           n == 0 && return 1
           n * fact(n-1)
       end
fact (generic function with 1 method)

julia> fact(5)
120

julia> fact(0)
1

julia> fact(-1)
ERROR: n must be non-negative
Stacktrace:
 [1] fact(::Int64) at ./none:2
```

Boolean operations *without* short-circuit evaluation can be done with the bitwise boolean operators introduced in
Mathematical Operations and Elementary Functions: & and |. These are normal functions, which happen to support
infix operator syntax, but always evaluate their arguments:

```
julia> f(1) & t(2)
```

```
1
2
false

julia> t(1) | t(2)
1
2
true
```

Just like condition expressions used in `if`, `elseif` or the ternary operator, the operands of `&&` or `||` must be boolean values (`true` or `false`). Using a non-boolean value anywhere except for the last entry in a conditional chain is an error:

```
julia> 1 && true
ERROR: TypeError: non-boolean (Int64) used in boolean context
```

On the other hand, any type of expression can be used at the end of a conditional chain. It will be evaluated and returned depending on the preceding conditionals:

```
julia> true && (x = (1, 2, 3))
(1, 2, 3)

julia> false && (x = (1, 2, 3))
false
```

12.4 Repeated Evaluation: Loops

There are two constructs for repeated evaluation of expressions: the `while` loop and the `for` loop. Here is an example of a `while` loop:

```
julia> i = 1;

julia> while i <= 5
           println(i)
           i += 1
       end
1
2
3
4
5
```

The `while` loop evaluates the condition expression (`i <= 5` in this case), and as long it remains `true`, keeps also evaluating the body of the `while` loop. If the condition expression is `false` when the `while` loop is first reached, the body is never evaluated.

The `for` loop makes common repeated evaluation idioms easier to write. Since counting up and down like the above `while` loop does is so common, it can be expressed more concisely with a `for` loop:

```
julia> for i = 1:5
           println(i)
       end
```

```
1
2
3
4
5
```

Here the 1:5 is a Range object, representing the sequence of numbers 1, 2, 3, 4, 5. The for loop iterates through these values, assigning each one in turn to the variable i. One rather important distinction between the previous while loop form and the for loop form is the scope during which the variable is visible. If the variable i has not been introduced in an other scope, in the for loop form, it is visible only inside of the for loop, and not afterwards. You'll either need a new interactive session instance or a different variable name to test this:

```
julia> for j = 1:5
           println(j)
       end
1
2
3
4
5

julia> j
ERROR: UndefVarError: j not defined
```

See Scope of Variables for a detailed explanation of variable scope and how it works in Julia.

In general, the for loop construct can iterate over any container. In these cases, the alternative (but fully equivalent) keyword in or is typically used instead of =, since it makes the code read more clearly:

```
julia> for i in [1,4,0]
           println(i)
       end
1
4
0

julia> for s  ["foo","bar","baz"]
           println(s)
       end
foo
bar
baz
```

Various types of iterable containers will be introduced and discussed in later sections of the manual (see, e.g., Multi-dimensional Arrays).

It is sometimes convenient to terminate the repetition of a while before the test condition is falsified or stop iterating in a for loop before the end of the iterable object is reached. This can be accomplished with the break keyword:

```
julia> i = 1;

julia> while true
           println(i)
```

```
              if i >= 5
                  break
              end
              i += 1
          end
1
2
3
4
5

julia> for i = 1:1000
              println(i)
              if i >= 5
                  break
              end
          end
1
2
3
4
5
```

Without the `break` keyword, the above `while` loop would never terminate on its own, and the `for` loop would iterate up to 1000. These loops are both exited early by using `break`.

In other circumstances, it is handy to be able to stop an iteration and move on to the next one immediately. The `continue` keyword accomplishes this:

```
julia> for i = 1:10
              if i % 3 != 0
                  continue
              end
              println(i)
          end
3
6
9
```

This is a somewhat contrived example since we could produce the same behavior more clearly by negating the condition and placing the `println` call inside the `if` block. In realistic usage there is more code to be evaluated after the `continue`, and often there are multiple points from which one calls `continue`.

Multiple nested `for` loops can be combined into a single outer loop, forming the cartesian product of its iterables:

```
julia> for i = 1:2, j = 3:4
              println((i, j))
          end
(1, 3)
(1, 4)
(2, 3)
(2, 4)
```

A `break` statement inside such a loop exits the entire nest of loops, not just the inner one.

12.5   Exception Handling

When an unexpected condition occurs, a function may be unable to return a reasonable value to its caller. In such cases, it may be best for the exceptional condition to either terminate the program, printing a diagnostic error message, or if the programmer has provided code to handle such exceptional circumstances, allow that code to take the appropriate action.

Built-in `Exceptions`

`Exceptions` are thrown when an unexpected condition has occurred. The built-in `Exceptions` listed below all interrupt the normal flow of control.

Exception
ArgumentError
BoundsError
CompositeException
DivideError
DomainError
EOFError
ErrorException
InexactError
InitError
InterruptException
InvalidStateException
KeyError
LoadError
OutOfMemoryError
ReadOnlyMemoryError
RemoteException
MethodError
OverflowError
ParseError
SystemError
TypeError
UndefRefError
UndefVarError
UnicodeError

For example, the `sqrt()` function throws a `DomainError` if applied to a negative real value:

```julia
julia> sqrt(-1)
ERROR: DomainError:
sqrt will only return a complex result if called with a complex argument. Try sqrt(complex(x)).
Stacktrace:
 [1] sqrt(::Int64) at ./math.jl:434
```

You may define your own exceptions in the following way:

```julia
julia> struct MyCustomException <: Exception end
```

The `throw()` function

Exceptions can be created explicitly with `throw()`. For example, a function defined only for nonnegative numbers could be written to `throw()` a `DomainError` if the argument is negative:

```julia
julia> f(x) = x>=0 ? exp(-x) : throw(DomainError())
f (generic function with 1 method)

julia> f(1)
0.36787944117144233

julia> f(-1)
ERROR: DomainError:
Stacktrace:
 [1] f(::Int64) at ./none:1
```

Note that `DomainError` without parentheses is not an exception, but a type of exception. It needs to be called to obtain an `Exception` object:

```julia
julia> typeof(DomainError()) <: Exception
true

julia> typeof(DomainError) <: Exception
false
```

Additionally, some exception types take one or more arguments that are used for error reporting:

```julia
julia> throw(UndefVarError(:x))
ERROR: UndefVarError: x not defined
```

This mechanism can be implemented easily by custom exception types following the way `UndefVarError` is written:

```julia
julia> struct MyUndefVarError <: Exception
           var::Symbol
       end

julia> Base.showerror(io::IO, e::MyUndefVarError) = print(io, e.var, " not defined")
```

> **Note**
>
> When writing an error message, it is preferred to make the first word lowercase. For example, `size(A) == size(B) || throw(DimensionMismatch("size of A not equal to size of B"))`
>
> is preferred over
>
> `size(A) == size(B) || throw(DimensionMismatch("Size of A not equal to size of B"))`.
>
> However, sometimes it makes sense to keep the uppercase first letter, for instance if an argument to a function is a capital letter: `size(A,1) == size(B,2) || throw(DimensionMismatch("A has first dimension..."))`.

Errors

The `error()` function is used to produce an `ErrorException` that interrupts the normal flow of control.

Suppose we want to stop execution immediately if the square root of a negative number is taken. To do this, we can define a fussy version of the `sqrt()` function that raises an error if its argument is negative:

```
julia> fussy_sqrt(x) = x >= 0 ? sqrt(x) : error("negative x not allowed")
fussy_sqrt (generic function with 1 method)

julia> fussy_sqrt(2)
1.4142135623730951

julia> fussy_sqrt(-1)
ERROR: negative x not allowed
Stacktrace:
 [1] fussy_sqrt(::Int64) at ./none:1
```

If `fussy_sqrt` is called with a negative value from another function, instead of trying to continue execution of the calling function, it returns immediately, displaying the error message in the interactive session:

```
julia> function verbose_fussy_sqrt(x)
           println("before fussy_sqrt")
           r = fussy_sqrt(x)
           println("after fussy_sqrt")
           return r
       end
verbose_fussy_sqrt (generic function with 1 method)

julia> verbose_fussy_sqrt(2)
before fussy_sqrt
after fussy_sqrt
1.4142135623730951

julia> verbose_fussy_sqrt(-1)
before fussy_sqrt
ERROR: negative x not allowed
Stacktrace:
 [1] fussy_sqrt at ./none:1 [inlined]
 [2] verbose_fussy_sqrt(::Int64) at ./none:3
```

Warnings and informational messages

Julia also provides other functions that write messages to the standard error I/O, but do not throw any `Exceptions` and hence do not interrupt execution:

```
julia> info("Hi"); 1+1
INFO: Hi
2

julia> warn("Hi"); 1+1
WARNING: Hi
2

julia> error("Hi"); 1+1
ERROR: Hi
Stacktrace:
 [1] error(::String) at ./error.jl:21
```

The `try/catch` statement

The `try/catch` statement allows for `Exceptions` to be tested for. For example, a customized square root function can be written to automatically call either the real or complex square root method on demand using `Exceptions` :

```julia
julia> f(x) = try
           sqrt(x)
       catch
           sqrt(complex(x, 0))
       end
f (generic function with 1 method)

julia> f(1)
1.0

julia> f(-1)
0.0 + 1.0im
```

It is important to note that in real code computing this function, one would compare x to zero instead of catching an exception. The exception is much slower than simply comparing and branching.

`try/catch` statements also allow the `Exception` to be saved in a variable. In this contrived example, the following example calculates the square root of the second element of x if x is indexable, otherwise assumes x is a real number and returns its square root:

```julia
julia> sqrt_second(x) = try
           sqrt(x[2])
       catch y
           if isa(y, DomainError)
               sqrt(complex(x[2], 0))
           elseif isa(y, BoundsError)
               sqrt(x)
           end
       end
sqrt_second (generic function with 1 method)

julia> sqrt_second([1 4])
2.0

julia> sqrt_second([1 -4])
0.0 + 2.0im

julia> sqrt_second(9)
3.0

julia> sqrt_second(-9)
ERROR: DomainError:
Stacktrace:
 [1] sqrt_second(::Int64) at ./none:7
```

Note that the symbol following `catch` will always be interpreted as a name for the exception, so care is needed when writing `try/catch` expressions on a single line. The following code will *not* work to return the value of x in case of an error:

```julia
try bad() catch x end
```

Instead, use a semicolon or insert a line break after `catch`:

```julia
try bad() catch; x end

try bad()
catch
    x
end
```

The `catch` clause is not strictly necessary; when omitted, the default return value is `nothing`.

```julia
julia> try error() end # Returns nothing
```

The power of the `try/catch` construct lies in the ability to unwind a deeply nested computation immediately to a much higher level in the stack of calling functions. There are situations where no error has occurred, but the ability to unwind the stack and pass a value to a higher level is desirable. Julia provides the `rethrow()`, `backtrace()` and `catch_backtrace()` functions for more advanced error handling.

`finally` Clauses

In code that performs state changes or uses resources like files, there is typically clean-up work (such as closing files) that needs to be done when the code is finished. Exceptions potentially complicate this task, since they can cause a block of code to exit before reaching its normal end. The `finally` keyword provides a way to run some code when a given block of code exits, regardless of how it exits.

For example, here is how we can guarantee that an opened file is closed:

```julia
f = open("file")
try
    # operate on file f
finally
    close(f)
end
```

When control leaves the `try` block (for example due to a `return`, or just finishing normally), `close(f)` will be executed. If the `try` block exits due to an exception, the exception will continue propagating. A `catch` block may be combined with `try` and `finally` as well. In this case the `finally` block will run after `catch` has handled the error.

12.6   Tasks (aka Coroutines)

Tasks are a control flow feature that allows computations to be suspended and resumed in a flexible manner. This feature is sometimes called by other names, such as symmetric coroutines, lightweight threads, cooperative multitasking, or one-shot continuations.

When a piece of computing work (in practice, executing a particular function) is designated as a `Task`, it becomes possible to interrupt it by switching to another `Task`. The original `Task` can later be resumed, at which point it will pick up right where it left off. At first, this may seem similar to a function call. However there are two key differences. First, switching tasks does not use any space, so any number of task switches can occur without consuming the call stack. Second, switching among tasks can occur in any order, unlike function calls, where the called function must finish executing before control returns to the calling function.

This kind of control flow can make it much easier to solve certain problems. In some problems, the various pieces of required work are not naturally related by function calls; there is no obvious "caller" or "callee" among the jobs that

need to be done. An example is the producer-consumer problem, where one complex procedure is generating values and another complex procedure is consuming them. The consumer cannot simply call a producer function to get a value, because the producer may have more values to generate and so might not yet be ready to return. With tasks, the producer and consumer can both run as long as they need to, passing values back and forth as necessary.

Julia provides a `Channel` mechanism for solving this problem. A `Channel` is a waitable first-in first-out queue which can have multiple tasks reading from and writing to it.

Let's define a producer task, which produces values via the `put!` call. To consume values, we need to schedule the producer to run in a new task. A special `Channel` constructor which accepts a 1-arg function as an argument can be used to run a task bound to a channel. We can then `take!()` values repeatedly from the channel object:

```julia
julia> function producer(c::Channel)
           put!(c, "start")
           for n=1:4
               put!(c, 2n)
           end
           put!(c, "stop")
       end;

julia> chnl = Channel(producer);

julia> take!(chnl)
"start"

julia> take!(chnl)
2

julia> take!(chnl)
4

julia> take!(chnl)
6

julia> take!(chnl)
8

julia> take!(chnl)
"stop"
```

One way to think of this behavior is that `producer` was able to return multiple times. Between calls to `put!()`, the producer's execution is suspended and the consumer has control.

The returned `Channel` can be used as an iterable object in a `for` loop, in which case the loop variable takes on all the produced values. The loop is terminated when the channel is closed.

```julia
julia> for x in Channel(producer)
           println(x)
       end
start
2
4
6
8
stop
```

Note that we did not have to explicitly close the channel in the producer. This is because the act of binding a `Channel` to a `Task()` associates the open lifetime of a channel with that of the bound task. The channel object is closed automatically when the task terminates. Multiple channels can be bound to a task, and vice-versa.

While the `Task()` constructor expects a 0-argument function, the `Channel()` method which creates a channel bound task expects a function that accepts a single argument of type `Channel`. A common pattern is for the producer to be parameterized, in which case a partial function application is needed to create a 0 or 1 argument anonymous function.

For `Task()` objects this can be done either directly or by use of a convenience macro:

```
function mytask(myarg)
    ...
end

taskHdl = Task(() -> mytask(7))
# or, equivalently
taskHdl = @task mytask(7)
```

To orchestrate more advanced work distribution patterns, `bind()` and `schedule()` can be used in conjunction with `Task()` and `Channel()` constructors to explicitly link a set of channels with a set of producer/consumer tasks.

Note that currently Julia tasks are not scheduled to run on separate CPU cores. True kernel threads are discussed under the topic of Parallel Computing.

Core task operations

Let us explore the low level construct `yieldto()` to underestand how task switching works. `yieldto(task, value)` suspends the current task, switches to the specified `task`, and causes that task's last `yieldto()` call to return the specified `value`. Notice that `yieldto()` is the only operation required to use task-style control flow; instead of calling and returning we are always just switching to a different task. This is why this feature is also called "symmetric coroutines"; each task is switched to and from using the same mechanism.

`yieldto()` is powerful, but most uses of tasks do not invoke it directly. Consider why this might be. If you switch away from the current task, you will probably want to switch back to it at some point, but knowing when to switch back, and knowing which task has the responsibility of switching back, can require considerable coordination. For example, `put!()` and `take!()` are blocking operations, which, when used in the context of channels maintain state to remember who the consumers are. Not needing to manually keep track of the consuming task is what makes `put!()` easier to use than the low-level `yieldto()`.

In addition to `yieldto()`, a few other basic functions are needed to use tasks effectively.

- `current_task()` gets a reference to the currently-running task.

- `istaskdone()` queries whether a task has exited.

- `istaskstarted()` queries whether a task has run yet.

- `task_local_storage()` manipulates a key-value store specific to the current task.

Tasks and events

Most task switches occur as a result of waiting for events such as I/O requests, and are performed by a scheduler included in the standard library. The scheduler maintains a queue of runnable tasks, and executes an event loop that restarts tasks based on external events such as message arrival.

The basic function for waiting for an event is `wait()`. Several objects implement `wait()`; for example, given a `Process` object, `wait()` will wait for it to exit. `wait()` is often implicit; for example, a `wait()` can happen inside a call to `read()` to wait for data to be available.

In all of these cases, `wait()` ultimately operates on a `Condition` object, which is in charge of queueing and restarting tasks. When a task calls `wait()` on a `Condition`, the task is marked as non-runnable, added to the condition's queue, and switches to the scheduler. The scheduler will then pick another task to run, or block waiting for external events. If all goes well, eventually an event handler will call `notify()` on the condition, which causes tasks waiting for that condition to become runnable again.

A task created explicitly by calling `Task` is initially not known to the scheduler. This allows you to manage tasks manually using `yieldto()` if you wish. However, when such a task waits for an event, it still gets restarted automatically when the event happens, as you would expect. It is also possible to make the scheduler run a task whenever it can, without necessarily waiting for any events. This is done by calling `schedule()`, or using the `@schedule` or `@async` macros (see Parallel Computing for more details).

Task states

Tasks have a `state` field that describes their execution status. A `Task` `state` is one of the following symbols:

Symbol	Meaning
`:runnable`	Currently running, or available to be switched to
`:waiting`	Blocked waiting for a specific event
`:queued`	In the scheduler's run queue about to be restarted
`:done`	Successfully finished executing
`:failed`	Finished with an uncaught exception

Chapter 13

Scope of Variables

The *scope* of a variable is the region of code within which a variable is visible. Variable scoping helps avoid variable naming conflicts. The concept is intuitive: two functions can both have arguments called x without the two x's referring to the same thing. Similarly there are many other cases where different blocks of code can use the same name without referring to the same thing. The rules for when the same variable name does or doesn't refer to the same thing are called scope rules; this section spells them out in detail.

Certain constructs in the language introduce *scope blocks*, which are regions of code that are eligible to be the scope of some set of variables. The scope of a variable cannot be an arbitrary set of source lines; instead, it will always line up with one of these blocks. There are two main types of scopes in Julia, *global scope* and *local scope*, the latter can be nested. The constructs introducing scope blocks are:

Scope name	block/construct introducing this kind of scope
Global Scope	`module`, `baremodule`, at interactive prompt (REPL)
Local Scope	Soft Local Scope: `for`, `while`, comprehensions, try-catch-finally, `let`
Local Scope	Hard Local Scope: functions (either syntax, anonymous & do-blocks), `struct`, `macro`

Notably missing from this table are begin blocks and if blocks, which do *not* introduce new scope blocks. All three types of scopes follow somewhat different rules which will be explained below as well as some extra rules for certain blocks.

Julia uses lexical scoping, meaning that a function's scope does not inherit from its caller's scope, but from the scope in which the function was defined. For example, in the following code the x inside `foo` refers to the x in the global scope of its module `Bar`:

```julia
julia> module Bar
           x = 1
           foo() = x
       end;
```

and not a x in the scope where `foo` is used:

```julia
julia> import .Bar

julia> x = -1;

julia> Bar.foo()
1
```

Thus *lexical scope* means that the scope of variables can be inferred from the source code alone.

13.1   Global Scope

Each module introduces a new global scope, separate from the global scope of all other modules; there is no all-encompassing global scope. Modules can introduce variables of other modules into their scope through the using or import statements or through qualified access using the dot-notation, i.e. each module is a so-called *namespace*. Note that variable bindings can only be changed within their global scope and not from an outside module.

```julia
julia> module A
           a = 1 # a global in A's scope
       end;

julia> module B
           module C
               c = 2
           end
           b = C.c    # can access the namespace of a nested global scope
                      # through a qualified access
           import ..A # makes module A available
           d = A.a
       end;

julia> module D
           b = a # errors as D's global scope is separate from A's
       end;
ERROR: UndefVarError: a not defined

julia> module E
           import ..A # make module A available
           A.a = 2    # throws below error
       end;
ERROR: cannot assign variables in other modules
```

Note that the interactive prompt (aka REPL) is in the global scope of the module `Main`.

13.2   Local Scope

A new local scope is introduced by most code-blocks, see above table for a complete list. A local scope *usually* inherits all the variables from its parent scope, both for reading and writing. There are two subtypes of local scopes, hard and soft, with slightly different rules concerning what variables are inherited. Unlike global scopes, local scopes are not namespaces, thus variables in an inner scope cannot be retrieved from the parent scope through some sort of qualified access.

The following rules and examples pertain to both hard and soft local scopes. A newly introduced variable in a local scope does not back-propagate to its parent scope. For example, here the z is not introduced into the top-level scope:

```julia
julia> for i = 1:10
           z = i
       end

julia> z
ERROR: UndefVarError: z not defined
```

(Note, in this and all following examples it is assumed that their top-level is a global scope with a clean workspace, for instance a newly started REPL.)

Inside a local scope a variable can be forced to be a local variable using the `local` keyword:

```julia
julia> x = 0;

julia> for i = 1:10
           local x
           x = i + 1
       end

julia> x
0
```

Inside a local scope a new global variable can be defined using the keyword `global`:

```julia
julia> for i = 1:10
           global z
           z = i
       end

julia> z
10
```

The location of both the `local` and `global` keywords within the scope block is irrelevant. The following is equivalent to the last example (although stylistically worse):

```julia
julia> for i = 1:10
           z = i
           global z
       end

julia> z
10
```

Soft Local Scope

> In a soft local scope, all variables are inherited from its parent scope unless a variable is specifically marked with the keyword `local`.

Soft local scopes are introduced by for-loops, while-loops, comprehensions, try-catch-finally-blocks, and let-blocks. There are some extra rules for Let Blocks and for For Loops and Comprehensions.

In the following example the x and y refer always to the same variables as the soft local scope inherits both read and write variables:

```julia
julia> x, y = 0, 1;

julia> for i = 1:10
           x = i + y + 1
       end

julia> x
12
```

Within soft scopes, the *global* keyword is never necessary, although allowed. The only case when it would change the semantics is (currently) a syntax error:

```julia
julia> let
           local j = 2
           let
               global j = 3
           end
       end
ERROR: syntax: `global j`: j is local variable in the enclosing scope
```

Hard Local Scope

Hard local scopes are introduced by function definitions (in all their forms), struct type definition blocks, and macro-definitions.

> In a hard local scope, all variables are inherited from its parent scope unless:
>
> - an assignment would result in a modified *global* variable, or
> - a variable is specifically marked with the keyword `local`.

Thus global variables are only inherited for reading but not for writing:

```julia
julia> x, y = 1, 2;

julia> function foo()
           x = 2        # assignment introduces a new local
           return x + y # y refers to the global
       end;

julia> foo()
4

julia> x
1
```

An explicit `global` is needed to assign to a global variable:

```julia
julia> x = 1;

julia> function foobar()
           global x = 2
       end;

julia> foobar();

julia> x
2
```

Note that *nested functions* can behave differently to functions defined in the global scope as they can modify their parent scope's *local* variables:

```
julia> x, y = 1, 2;

julia> function baz()
           x = 2 # introduces a new local
           function bar()
               x = 10      # modifies the parent's x
               return x + y # y is global
           end
           return bar() + x # 12 + 10 (x is modified in call of bar())
       end;

julia> baz()
22

julia> x, y
(1, 2)
```

The distinction between inheriting global and local variables for assignment can lead to some slight differences between functions defined in local vs. global scopes. Consider the modification of the last example by moving `bar` to the global scope:

```
julia> x, y = 1, 2;

julia> function bar()
           x = 10 # local
           return x + y
       end;

julia> function quz()
           x = 2 # local
           return bar() + x # 12 + 2 (x is not modified)
       end;

julia> quz()
14

julia> x, y
(1, 2)
```

Note that above subtlety does not pertain to type and macro definitions as they can only appear at the global scope. There are special scoping rules concerning the evaluation of default and keyword function arguments which are described in the Function section.

An assignment introducing a variable used inside a function, type or macro definition need not come before its inner usage:

```
julia> f = y -> y + a
(::#1) (generic function with 1 method)

julia> f(3)
ERROR: UndefVarError: a not defined
Stacktrace:
 [1] (::##1#2)(::Int64) at ./none:1
```

```julia
julia> a = 1
1

julia> f(3)
4
```

This behavior may seem slightly odd for a normal variable, but allows for named functions – which are just normal variables holding function objects – to be used before they are defined. This allows functions to be defined in whatever order is intuitive and convenient, rather than forcing bottom up ordering or requiring forward declarations, as long as they are defined by the time they are actually called. As an example, here is an inefficient, mutually recursive way to test if positive integers are even or odd:

```julia
julia> even(n) = n == 0 ? true : odd(n-1);

julia> odd(n) = n == 0 ? false : even(n-1);

julia> even(3)
false

julia> odd(3)
true
```

Julia provides built-in, efficient functions to test for oddness and evenness called `iseven()` and `isodd()` so the above definitions should only be taken as examples.

Hard vs. Soft Local Scope

Blocks which introduce a soft local scope, such as loops, are generally used to manipulate the variables in their parent scope. Thus their default is to fully access all variables in their parent scope.

Conversely, the code inside blocks which introduce a hard local scope (function, type, and macro definitions) can be executed at any place in a program. Remotely changing the state of global variables in other modules should be done with care and thus this is an opt-in feature requiring the `global` keyword.

The reason to allow *modifying local* variables of parent scopes in nested functions is to allow constructing closures which have a private state, for instance the `state` variable in the following example:

```julia
julia> let
           state = 0
           global counter
           counter() = state += 1
       end;

julia> counter()
1

julia> counter()
2
```

See also the closures in the examples in the next two sections.

Let Blocks

Unlike assignments to local variables, `let` statements allocate new variable bindings each time they run. An assignment modifies an existing value location, and `let` creates new locations. This difference is usually not important, and is only detectable in the case of variables that outlive their scope via closures. The `let` syntax accepts a comma-separated series of assignments and variable names:

```julia
julia> x, y, z = -1, -1, -1;

julia> let x = 1, z
           println("x: $x, y: $y") # x is local variable, y the global
           println("z: $z") # errors as z has not been assigned yet but is local
       end
x: 1, y: -1
ERROR: UndefVarError: z not defined
```

The assignments are evaluated in order, with each right-hand side evaluated in the scope before the new variable on the left-hand side has been introduced. Therefore it makes sense to write something like `let x = x` since the two x variables are distinct and have separate storage. Here is an example where the behavior of `let` is needed:

```julia
julia> Fs = Array{Any}(2); i = 1;

julia> while i <= 2
           Fs[i] = ()->i
           i += 1
       end

julia> Fs[1]()
3

julia> Fs[2]()
3
```

Here we create and store two closures that return variable `i`. However, it is always the same variable `i`, so the two closures behave identically. We can use `let` to create a new binding for `i`:

```julia
julia> Fs = Array{Any}(2); i = 1;

julia> while i <= 2
           let i = i
               Fs[i] = ()->i
           end
           i += 1
       end

julia> Fs[1]()
1

julia> Fs[2]()
2
```

Since the `begin` construct does not introduce a new scope, it can be useful to use a zero-argument `let` to just introduce a new scope block without creating any new bindings:

```julia
julia> let
           local x = 1
           let
               local x = 2
           end
           x
       end
1
```

Since `let` introduces a new scope block, the inner local x is a different variable than the outer local x.

For Loops and Comprehensions

`for` loops and Comprehensions have the following behavior: any new variables introduced in their body scopes are freshly allocated for each loop iteration. This is in contrast to `while` loops which reuse the variables for all iterations. Therefore these constructs are similar to `while` loops with `let` blocks inside:

```julia
julia> Fs = Array{Any}(2);

julia> for j = 1:2
           Fs[j] = ()->j
       end

julia> Fs[1]()
1

julia> Fs[2]()
2
```

`for` loops will reuse existing variables for its iteration variable:

```julia
julia> i = 0;

julia> for i = 1:3
       end

julia> i
3
```

However, comprehensions do not do this, and always freshly allocate their iteration variables:

```julia
julia> x = 0;

julia> [ x for x = 1:3 ];

julia> x
0
```

13.3   Constants

A common use of variables is giving names to specific, unchanging values. Such variables are only assigned once. This intent can be conveyed to the compiler using the `const` keyword:

```
julia> const e  = 2.71828182845904523536;

julia> const pi = 3.14159265358979323846;
```

The `const` declaration is allowed on both global and local variables, but is especially useful for globals. It is difficult for the compiler to optimize code involving global variables, since their values (or even their types) might change at almost any time. If a global variable will not change, adding a `const` declaration solves this performance problem.

Local constants are quite different. The compiler is able to determine automatically when a local variable is constant, so local constant declarations are not necessary for performance purposes.

Special top-level assignments, such as those performed by the `function` and `struct` keywords, are constant by default.

Note that `const` only affects the variable binding; the variable may be bound to a mutable object (such as an array), and that object may still be modified.

Chapter 14

Types

Type systems have traditionally fallen into two quite different camps: static type systems, where every program expression must have a type computable before the execution of the program, and dynamic type systems, where nothing is known about types until run time, when the actual values manipulated by the program are available. Object orientation allows some flexibility in statically typed languages by letting code be written without the precise types of values being known at compile time. The ability to write code that can operate on different types is called polymorphism. All code in classic dynamically typed languages is polymorphic: only by explicitly checking types, or when objects fail to support operations at run-time, are the types of any values ever restricted.

Julia's type system is dynamic, but gains some of the advantages of static type systems by making it possible to indicate that certain values are of specific types. This can be of great assistance in generating efficient code, but even more significantly, it allows method dispatch on the types of function arguments to be deeply integrated with the language. Method dispatch is explored in detail in Methods, but is rooted in the type system presented here.

The default behavior in Julia when types are omitted is to allow values to be of any type. Thus, one can write many useful Julia programs without ever explicitly using types. When additional expressiveness is needed, however, it is easy to gradually introduce explicit type annotations into previously "untyped" code. Doing so will typically increase both the performance and robustness of these systems, and perhaps somewhat counterintuitively, often significantly simplify them.

Describing Julia in the lingo of type systems, it is: dynamic, nominative and parametric. Generic types can be parameterized, and the hierarchical relationships between types are explicitly declared, rather than implied by compatible structure. One particularly distinctive feature of Julia's type system is that concrete types may not subtype each other: all concrete types are final and may only have abstract types as their supertypes. While this might at first seem unduly restrictive, it has many beneficial consequences with surprisingly few drawbacks. It turns out that being able to inherit behavior is much more important than being able to inherit structure, and inheriting both causes significant difficulties in traditional object-oriented languages. Other high-level aspects of Julia's type system that should be mentioned up front are:

- There is no division between object and non-object values: all values in Julia are true objects having a type that belongs to a single, fully connected type graph, all nodes of which are equally first-class as types.

- There is no meaningful concept of a "compile-time type": the only type a value has is its actual type when the program is running. This is called a "run-time type" in object-oriented languages where the combination of static compilation with polymorphism makes this distinction significant.

- Only values, not variables, have types – variables are simply names bound to values.

- Both abstract and concrete types can be parameterized by other types. They can also be parameterized by symbols, by values of any type for which `isbits()` returns true (essentially, things like numbers and bools that

are stored like C types or structs with no pointers to other objects), and also by tuples thereof. Type parameters may be omitted when they do not need to be referenced or restricted.

Julia's type system is designed to be powerful and expressive, yet clear, intuitive and unobtrusive. Many Julia programmers may never feel the need to write code that explicitly uses types. Some kinds of programming, however, become clearer, simpler, faster and more robust with declared types.

14.1   Type Declarations

The `::` operator can be used to attach type annotations to expressions and variables in programs. There are two primary reasons to do this:

1. As an assertion to help confirm that your program works the way you expect,

2. To provide extra type information to the compiler, which can then improve performance in some cases

When appended to an expression computing a value, the `::` operator is read as "is an instance of". It can be used anywhere to assert that the value of the expression on the left is an instance of the type on the right. When the type on the right is concrete, the value on the left must have that type as its implementation – recall that all concrete types are final, so no implementation is a subtype of any other. When the type is abstract, it suffices for the value to be implemented by a concrete type that is a subtype of the abstract type. If the type assertion is not true, an exception is thrown, otherwise, the left-hand value is returned:

```julia
julia> (1+2)::AbstractFloat
ERROR: TypeError: typeassert: expected AbstractFloat, got Int64

julia> (1+2)::Int
3
```

This allows a type assertion to be attached to any expression in-place.

When appended to a variable on the left-hand side of an assignment, or as part of a `local` declaration, the `::` operator means something a bit different: it declares the variable to always have the specified type, like a type declaration in a statically-typed language such as C. Every value assigned to the variable will be converted to the declared type using `convert()`:

```julia
julia> function foo()
           x::Int8 = 100
           x
       end
foo (generic function with 1 method)

julia> foo()
100

julia> typeof(ans)
Int8
```

This feature is useful for avoiding performance "gotchas" that could occur if one of the assignments to a variable changed its type unexpectedly.

This "declaration" behavior only occurs in specific contexts:

```
local x::Int8   # in a local declaration
x::Int8 = 10    # as the left-hand side of an assignment
```

and applies to the whole current scope, even before the declaration. Currently, type declarations cannot be used in global scope, e.g. in the REPL, since Julia does not yet have constant-type globals.

Declarations can also be attached to function definitions:

```
function sinc(x)::Float64
    if x == 0
        return 1
    end
    return sin(pi*x)/(pi*x)
end
```

Returning from this function behaves just like an assignment to a variable with a declared type: the value is always converted to Float64.

14.2 Abstract Types

Abstract types cannot be instantiated, and serve only as nodes in the type graph, thereby describing sets of related concrete types: those concrete types which are their descendants. We begin with abstract types even though they have no instantiation because they are the backbone of the type system: they form the conceptual hierarchy which makes Julia's type system more than just a collection of object implementations.

Recall that in Integers and Floating-Point Numbers, we introduced a variety of concrete types of numeric values: Int8, UInt8, Int16, UInt16, Int32, UInt32, Int64, UInt64, Int128, UInt128, Float16, Float32, and Float64. Although they have different representation sizes, Int8, Int16, Int32, Int64 and Int128 all have in common that they are signed integer types. Likewise UInt8, UInt16, UInt32, UInt64 and UInt128 are all unsigned integer types, while Float16, Float32 and Float64 are distinct in being floating-point types rather than integers. It is common for a piece of code to make sense, for example, only if its arguments are some kind of integer, but not really depend on what particular *kind* of integer. For example, the greatest common denominator algorithm works for all kinds of integers, but will not work for floating-point numbers. Abstract types allow the construction of a hierarchy of types, providing a context into which concrete types can fit. This allows you, for example, to easily program to any type that is an integer, without restricting an algorithm to a specific type of integer.

Abstract types are declared using the `abstract type` keyword. The general syntaxes for declaring an abstract type are:

```
abstract type «name» end
abstract type «name» <: «supertype» end
```

The `abstract type` keyword introduces a new abstract type, whose name is given by «name». This name can be optionally followed by <: and an already-existing type, indicating that the newly declared abstract type is a subtype of this "parent" type.

When no supertype is given, the default supertype is Any – a predefined abstract type that all objects are instances of and all types are subtypes of. In type theory, Any is commonly called "top" because it is at the apex of the type graph. Julia also has a predefined abstract "bottom" type, at the nadir of the type graph, which is written as Union{}. It is the exact opposite of Any: no object is an instance of Union{} and all types are supertypes of Union{}.

Let's consider some of the abstract types that make up Julia's numerical hierarchy:

```julia
abstract type Number end
abstract type Real      <: Number end
abstract type AbstractFloat <: Real end
abstract type Integer  <: Real end
abstract type Signed   <: Integer end
abstract type Unsigned <: Integer end
```

The Number type is a direct child type of Any, and Real is its child. In turn, Real has two children (it has more, but only two are shown here; we'll get to the others later): Integer and AbstractFloat, separating the world into representations of integers and representations of real numbers. Representations of real numbers include, of course, floating-point types, but also include other types, such as rationals. Hence, AbstractFloat is a proper subtype of Real, including only floating-point representations of real numbers. Integers are further subdivided into Signed and Unsigned varieties.

The <: operator in general means "is a subtype of", and, used in declarations like this, declares the right-hand type to be an immediate supertype of the newly declared type. It can also be used in expressions as a subtype operator which returns true when its left operand is a subtype of its right operand:

```julia
julia> Integer <: Number
true

julia> Integer <: AbstractFloat
false
```

An important use of abstract types is to provide default implementations for concrete types. To give a simple example, consider:

```julia
function myplus(x,y)
    x+y
end
```

The first thing to note is that the above argument declarations are equivalent to x::Any and y::Any. When this function is invoked, say as myplus(2,5), the dispatcher chooses the most specific method named myplus that matches the given arguments. (See Methods for more information on multiple dispatch.)

Assuming no method more specific than the above is found, Julia next internally defines and compiles a method called myplus specifically for two Int arguments based on the generic function given above, i.e., it implicitly defines and compiles:

```julia
function myplus(x::Int,y::Int)
    x+y
end
```

and finally, it invokes this specific method.

Thus, abstract types allow programmers to write generic functions that can later be used as the default method by many combinations of concrete types. Thanks to multiple dispatch, the programmer has full control over whether the default or more specific method is used.

An important point to note is that there is no loss in performance if the programmer relies on a function whose arguments are abstract types, because it is recompiled for each tuple of argument concrete types with which it is invoked. (There may be a performance issue, however, in the case of function arguments that are containers of abstract types; see Performance Tips.)

14.3 Primitive Types

A primitive type is a concrete type whose data consists of plain old bits. Classic examples of primitive types are integers and floating-point values. Unlike most languages, Julia lets you declare your own primitive types, rather than providing only a fixed set of built-in ones. In fact, the standard primitive types are all defined in the language itself:

```
primitive type Float16 <: AbstractFloat 16 end
primitive type Float32 <: AbstractFloat 32 end
primitive type Float64 <: AbstractFloat 64 end

primitive type Bool <: Integer 8 end
primitive type Char 32 end

primitive type Int8    <: Signed    8 end
primitive type UInt8   <: Unsigned 8 end
primitive type Int16   <: Signed   16 end
primitive type UInt16  <: Unsigned 16 end
primitive type Int32   <: Signed   32 end
primitive type UInt32  <: Unsigned 32 end
primitive type Int64   <: Signed   64 end
primitive type UInt64  <: Unsigned 64 end
primitive type Int128  <: Signed  128 end
primitive type UInt128 <: Unsigned 128 end
```

The general syntaxes for declaring a primitive type are:

```
primitive type «name» «bits» end
primitive type «name» <: «supertype» «bits» end
```

The number of bits indicates how much storage the type requires and the name gives the new type a name. A primitive type can optionally be declared to be a subtype of some supertype. If a supertype is omitted, then the type defaults to having Any as its immediate supertype. The declaration of Bool above therefore means that a boolean value takes eight bits to store, and has Integer as its immediate supertype. Currently, only sizes that are multiples of 8 bits are supported. Therefore, boolean values, although they really need just a single bit, cannot be declared to be any smaller than eight bits.

The types Bool, Int8 and UInt8 all have identical representations: they are eight-bit chunks of memory. Since Julia's type system is nominative, however, they are not interchangeable despite having identical structure. A fundamental difference between them is that they have different supertypes: Bool's direct supertype is Integer, Int8's is Signed, and UInt8's is Unsigned. All other differences between Bool, Int8, and UInt8 are matters of behavior – the way functions are defined to act when given objects of these types as arguments. This is why a nominative type system is necessary: if structure determined type, which in turn dictates behavior, then it would be impossible to make Bool behave any differently than Int8 or UInt8.

14.4 Composite Types

Composite types are called records, structs, or objects in various languages. A composite type is a collection of named fields, an instance of which can be treated as a single value. In many languages, composite types are the only kind of user-definable type, and they are by far the most commonly used user-defined type in Julia as well.

In mainstream object oriented languages, such as C++, Java, Python and Ruby, composite types also have named functions associated with them, and the combination is called an "object". In purer object-oriented languages, such as Ruby or Smalltalk, all values are objects whether they are composites or not. In less pure object oriented languages, including C++ and Java, some values, such as integers and floating-point values, are not objects, while instances of

user-defined composite types are true objects with associated methods. In Julia, all values are objects, but functions are not bundled with the objects they operate on. This is necessary since Julia chooses which method of a function to use by multiple dispatch, meaning that the types of *all* of a function's arguments are considered when selecting a method, rather than just the first one (see Methods for more information on methods and dispatch). Thus, it would be inappropriate for functions to "belong" to only their first argument. Organizing methods into function objects rather than having named bags of methods "inside" each object ends up being a highly beneficial aspect of the language design.

Composite types are introduced with the `struct` keyword followed by a block of field names, optionally annotated with types using the `::` operator:

```julia
julia> struct Foo
           bar
           baz::Int
           qux::Float64
       end
```

Fields with no type annotation default to Any, and can accordingly hold any type of value.

New objects of type Foo are created by applying the Foo type object like a function to values for its fields:

```julia
julia> foo = Foo("Hello, world.", 23, 1.5)
Foo("Hello, world.", 23, 1.5)

julia> typeof(foo)
Foo
```

When a type is applied like a function it is called a *constructor*. Two constructors are generated automatically (these are called *default constructors*). One accepts any arguments and calls `convert()` to convert them to the types of the fields, and the other accepts arguments that match the field types exactly. The reason both of these are generated is that this makes it easier to add new definitions without inadvertently replacing a default constructor.

Since the `bar` field is unconstrained in type, any value will do. However, the value for baz must be convertible to `Int`:

```julia
julia> Foo((), 23.5, 1)
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Int64}, ::Float64) at ./float.jl:679
 [2] Foo(::Tuple{}, ::Float64, ::Int64) at ./none:2
```

You may find a list of field names using the `fieldnames` function.

```julia
julia> fieldnames(foo)
3-element Array{Symbol,1}:
 :bar
 :baz
 :qux
```

You can access the field values of a composite object using the traditional `foo.bar` notation:

```julia
julia> foo.bar
"Hello, world."

julia> foo.baz
23

julia> foo.qux
1.5
```

Composite objects declared with `struct` are *immutable*; they cannot be modified after construction. This may seem odd at first, but it has several advantages:

- It can be more efficient. Some structs can be packed efficiently into arrays, and in some cases the compiler is able to avoid allocating immutable objects entirely.

- It is not possible to violate the invariants provided by the type's constructors.

- Code using immutable objects can be easier to reason about.

An immutable object might contain mutable objects, such as arrays, as fields. Those contained objects will remain mutable; only the fields of the immutable object itself cannot be changed to point to different objects.

Where required, mutable composite objects can be declared with the keyword `mutable struct`, to be discussed in the next section.

Composite types with no fields are singletons; there can be only one instance of such types:

```julia
julia> struct NoFields
       end

julia> NoFields() === NoFields()
true
```

The `===` function confirms that the "two" constructed instances of `NoFields` are actually one and the same. Singleton types are described in further detail below.

There is much more to say about how instances of composite types are created, but that discussion depends on both Parametric Types and on Methods, and is sufficiently important to be addressed in its own section: Constructors.

14.5 Mutable Composite Types

If a composite type is declared with `mutable struct` instead of `struct`, then instances of it can be modified:

```julia
julia> mutable struct Bar
           baz
           qux::Float64
       end

julia> bar = Bar("Hello", 1.5);

julia> bar.qux = 2.0
2.0

julia> bar.baz = 1//2
1//2
```

In order to support mutation, such objects are generally allocated on the heap, and have stable memory addresses. A mutable object is like a little container that might hold different values over time, and so can only be reliably identified with its address. In contrast, an instance of an immutable type is associated with specific field values -- the field values alone tell you everything about the object. In deciding whether to make a type mutable, ask whether two instances with the same field values would be considered identical, or if they might need to change independently over time. If they would be considered identical, the type should probably be immutable.

To recap, two essential properties define immutability in Julia:

- An object with an immutable type is passed around (both in assignment statements and in function calls) by copying, whereas a mutable type is passed around by reference.

- It is not permitted to modify the fields of a composite immutable type.

It is instructive, particularly for readers whose background is C/C++, to consider why these two properties go hand in hand. If they were separated, i.e., if the fields of objects passed around by copying could be modified, then it would become more difficult to reason about certain instances of generic code. For example, suppose x is a function argument of an abstract type, and suppose that the function changes a field: `x.isprocessed = true`. Depending on whether x is passed by copying or by reference, this statement may or may not alter the actual argument in the calling routine. Julia sidesteps the possibility of creating functions with unknown effects in this scenario by forbidding modification of fields of objects passed around by copying.

14.6   Declared Types

The three kinds of types discussed in the previous three sections are actually all closely related. They share the same key properties:

- They are explicitly declared.

- They have names.

- They have explicitly declared supertypes.

- They may have parameters.

Because of these shared properties, these types are internally represented as instances of the same concept, `DataType`, which is the type of any of these types:

```julia
julia> typeof(Real)
DataType

julia> typeof(Int)
DataType
```

A `DataType` may be abstract or concrete. If it is concrete, it has a specified size, storage layout, and (optionally) field names. Thus a bits type is a `DataType` with nonzero size, but no field names. A composite type is a `DataType` that has field names or is empty (zero size).

Every concrete value in the system is an instance of some `DataType`.

14.7   Type Unions

A type union is a special abstract type which includes as objects all instances of any of its argument types, constructed using the special `Union` function:

```julia
julia> IntOrString = Union{Int,AbstractString}
Union{AbstractString, Int64}

julia> 1 :: IntOrString
1
```

```
julia> "Hello!" :: IntOrString
"Hello!"

julia> 1.0 :: IntOrString
ERROR: TypeError: typeassert: expected Union{AbstractString, Int64}, got Float64
```

The compilers for many languages have an internal union construct for reasoning about types; Julia simply exposes it to the programmer.

14.8 Parametric Types

An important and powerful feature of Julia's type system is that it is parametric: types can take parameters, so that type declarations actually introduce a whole family of new types – one for each possible combination of parameter values. There are many languages that support some version of generic programming, wherein data structures and algorithms to manipulate them may be specified without specifying the exact types involved. For example, some form of generic programming exists in ML, Haskell, Ada, Eiffel, C++, Java, C#, F#, and Scala, just to name a few. Some of these languages support true parametric polymorphism (e.g. ML, Haskell, Scala), while others support ad-hoc, template-based styles of generic programming (e.g. C++, Java). With so many different varieties of generic programming and parametric types in various languages, we won't even attempt to compare Julia's parametric types to other languages, but will instead focus on explaining Julia's system in its own right. We will note, however, that because Julia is a dynamically typed language and doesn't need to make all type decisions at compile time, many traditional difficulties encountered in static parametric type systems can be relatively easily handled.

All declared types (the `DataType` variety) can be parameterized, with the same syntax in each case. We will discuss them in the following order: first, parametric composite types, then parametric abstract types, and finally parametric bits types.

Parametric Composite Types

Type parameters are introduced immediately after the type name, surrounded by curly braces:

```
julia> struct Point{T}
           x::T
           y::T
       end
```

This declaration defines a new parametric type, `Point{T}`, holding two "coordinates" of type T. What, one may ask, is T? Well, that's precisely the point of parametric types: it can be any type at all (or a value of any bits type, actually, although here it's clearly used as a type). `Point{Float64}` is a concrete type equivalent to the type defined by replacing T in the definition of `Point` with Float64. Thus, this single declaration actually declares an unlimited number of types: `Point{Float64}`, `Point{AbstractString}`, `Point{Int64}`, etc. Each of these is now a usable concrete type:

```
julia> Point{Float64}
Point{Float64}

julia> Point{AbstractString}
Point{AbstractString}
```

The type `Point{Float64}` is a point whose coordinates are 64-bit floating-point values, while the type `Point{AbstractString}` is a "point" whose "coordinates" are string objects (see Strings).

`Point` itself is also a valid type object, containing all instances `Point{Float64}`, `Point{AbstractString}`, etc. as subtypes:

```julia
julia> Point{Float64} <: Point
true

julia> Point{AbstractString} <: Point
true
```

Other types, of course, are not subtypes of it:

```julia
julia> Float64 <: Point
false

julia> AbstractString <: Point
false
```

Concrete Point types with different values of T are never subtypes of each other:

```julia
julia> Point{Float64} <: Point{Int64}
false

julia> Point{Float64} <: Point{Real}
false
```

> **Warning**
>
> This last point is *very* important: even though `Float64 <: Real` we **DO NOT** have `Point{Float64}` `<: Point{Real}`.

In other words, in the parlance of type theory, Julia's type parameters are *invariant*, rather than being covariant (or even contravariant). This is for practical reasons: while any instance of `Point{Float64}` may conceptually be like an instance of `Point{Real}` as well, the two types have different representations in memory:

- An instance of `Point{Float64}` can be represented compactly and efficiently as an immediate pair of 64-bit values;

- An instance of `Point{Real}` must be able to hold any pair of instances of `Real`. Since objects that are instances of `Real` can be of arbitrary size and structure, in practice an instance of `Point{Real}` must be represented as a pair of pointers to individually allocated `Real` objects.

The efficiency gained by being able to store `Point{Float64}` objects with immediate values is magnified enormously in the case of arrays: an `Array{Float64}` can be stored as a contiguous memory block of 64-bit floating-point values, whereas an `Array{Real}` must be an array of pointers to individually allocated `Real` objects – which may well be boxed 64-bit floating-point values, but also might be arbitrarily large, complex objects, which are declared to be implementations of the `Real` abstract type.

Since `Point{Float64}` is not a subtype of `Point{Real}`, the following method can't be applied to arguments of type `Point{Float64}`:

```julia
function norm(p::Point{Real})
    sqrt(p.x^2 + p.y^2)
end
```

A correct way to define a method that accepts all arguments of type `Point{T}` where T is a subtype of `Real` is:

```
function norm(p::Point{<:Real})
    sqrt(p.x^2 + p.y^2)
end
```

(Equivalently, one could define `function norm{T<:Real}(p::Point{T})` or `function norm(p::Point{T} where T<:Real`); see UnionAll Types.)

More examples will be discussed later in Methods.

How does one construct a `Point` object? It is possible to define custom constructors for composite types, which will be discussed in detail in Constructors, but in the absence of any special constructor declarations, there are two default ways of creating new composite objects, one in which the type parameters are explicitly given and the other in which they are implied by the arguments to the object constructor.

Since the type `Point{Float64}` is a concrete type equivalent to `Point` declared with Float64 in place of T, it can be applied as a constructor accordingly:

```
julia> Point{Float64}(1.0, 2.0)
Point{Float64}(1.0, 2.0)

julia> typeof(ans)
Point{Float64}
```

For the default constructor, exactly one argument must be supplied for each field:

```
julia> Point{Float64}(1.0)
ERROR: MethodError: Cannot `convert` an object of type Float64 to an object of type Point{Float64
    }
This may have arisen from a call to the constructor Point{Float64}(...),
since type constructors fall back to convert methods.
Stacktrace:
 [1] Point{Float64}(::Float64) at ./sysimg.jl:24

julia> Point{Float64}(1.0,2.0,3.0)
ERROR: MethodError: no method matching Point{Float64}(::Float64, ::Float64, ::Float64)
```

Only one default constructor is generated for parametric types, since overriding it is not possible. This constructor accepts any arguments and converts them to the field types.

In many cases, it is redundant to provide the type of `Point` object one wants to construct, since the types of arguments to the constructor call already implicitly provide type information. For that reason, you can also apply `Point` itself as a constructor, provided that the implied value of the parameter type T is unambiguous:

```
julia> Point(1.0,2.0)
Point{Float64}(1.0, 2.0)

julia> typeof(ans)
Point{Float64}

julia> Point(1,2)
Point{Int64}(1, 2)

julia> typeof(ans)
Point{Int64}
```

In the case of `Point`, the type of T is unambiguously implied if and only if the two arguments to `Point` have the same type. When this isn't the case, the constructor will fail with a `MethodError`:

```
julia> Point(1,2.5)
ERROR: MethodError: no method matching Point(::Int64, ::Float64)
Closest candidates are:
  Point(::T, !Matched::T) where T at none:2
```

Constructor methods to appropriately handle such mixed cases can be defined, but that will not be discussed until later on in Constructors.

Parametric Abstract Types

Parametric abstract type declarations declare a collection of abstract types, in much the same way:

```
julia> abstract type Pointy{T} end
```

With this declaration, `Pointy{T}` is a distinct abstract type for each type or integer value of T. As with parametric composite types, each such instance is a subtype of `Pointy`:

```
julia> Pointy{Int64} <: Pointy
true

julia> Pointy{1} <: Pointy
true
```

Parametric abstract types are invariant, much as parametric composite types are:

```
julia> Pointy{Float64} <: Pointy{Real}
false

julia> Pointy{Real} <: Pointy{Float64}
false
```

The notation `Pointy{<:Real}` can be used to express the Julia analogue of a *covariant* type, while `Pointy{>:Int}` the analogue of a *contravariant* type, but technically these represent *sets* of types (see UnionAll Types).

```
julia> Pointy{Float64} <: Pointy{<:Real}
true

julia> Pointy{Real} <: Pointy{>:Int}
true
```

Much as plain old abstract types serve to create a useful hierarchy of types over concrete types, parametric abstract types serve the same purpose with respect to parametric composite types. We could, for example, have declared `Point{T}` to be a subtype of `Pointy{T}` as follows:

```
julia> struct Point{T} <: Pointy{T}
           x::T
           y::T
       end
```

Given such a declaration, for each choice of T, we have `Point{T}` as a subtype of `Pointy{T}`:

```
julia> Point{Float64} <: Pointy{Float64}
true

julia> Point{Real} <: Pointy{Real}
true

julia> Point{AbstractString} <: Pointy{AbstractString}
true
```

This relationship is also invariant:

```
julia> Point{Float64} <: Pointy{Real}
false

julia> Point{Float64} <: Pointy{<:Real}
true
```

What purpose do parametric abstract types like `Pointy` serve? Consider if we create a point-like implementation that only requires a single coordinate because the point is on the diagonal line *x = y*:

```
julia> struct DiagPoint{T} <: Pointy{T}
           x::T
       end
```

Now both `Point{Float64}` and `DiagPoint{Float64}` are implementations of the `Pointy{Float64}` abstraction, and similarly for every other possible choice of type T. This allows programming to a common interface shared by all `Pointy` objects, implemented for both `Point` and `DiagPoint`. This cannot be fully demonstrated, however, until we have introduced methods and dispatch in the next section, Methods.

There are situations where it may not make sense for type parameters to range freely over all possible types. In such situations, one can constrain the range of T like so:

```
julia> abstract type Pointy{T<:Real} end
```

With such a declaration, it is acceptable to use any type that is a subtype of Real in place of T, but not types that are not subtypes of `Real`:

```
julia> Pointy{Float64}
Pointy{Float64}

julia> Pointy{Real}
Pointy{Real}

julia> Pointy{AbstractString}
ERROR: TypeError: Pointy: in T, expected T<:Real, got Type{AbstractString}

julia> Pointy{1}
ERROR: TypeError: Pointy: in T, expected T<:Real, got Int64
```

Type parameters for parametric composite types can be restricted in the same manner:

```
struct Point{T<:Real} <: Pointy{T}
    x::T
    y::T
end
```

To give a real-world example of how all this parametric type machinery can be useful, here is the actual definition of Julia's Rational immutable type (except that we omit the constructor here for simplicity), representing an exact ratio of integers:

```
struct Rational{T<:Integer} <: Real
    num::T
    den::T
end
```

It only makes sense to take ratios of integer values, so the parameter type T is restricted to being a subtype of `Integer`, and a ratio of integers represents a value on the real number line, so any `Rational` is an instance of the `Real` abstraction.

Tuple Types

Tuples are an abstraction of the arguments of a function – without the function itself. The salient aspects of a function's arguments are their order and their types. Therefore a tuple type is similar to a parameterized immutable type where each parameter is the type of one field. For example, a 2-element tuple type resembles the following immutable type:

```
struct Tuple2{A,B}
    a::A
    b::B
end
```

However, there are three key differences:

- Tuple types may have any number of parameters.

- Tuple types are *covariant* in their parameters: `Tuple{Int}` is a subtype of `Tuple{Any}`. Therefore `Tuple{Any}` is considered an abstract type, and tuple types are only concrete if their parameters are.

- Tuples do not have field names; fields are only accessed by index.

Tuple values are written with parentheses and commas. When a tuple is constructed, an appropriate tuple type is generated on demand:

```
julia> typeof((1,"foo",2.5))
Tuple{Int64,String,Float64}
```

Note the implications of covariance:

```
julia> Tuple{Int,AbstractString} <: Tuple{Real,Any}
true

julia> Tuple{Int,AbstractString} <: Tuple{Real,Real}
false

julia> Tuple{Int,AbstractString} <: Tuple{Real,}
false
```

Intuitively, this corresponds to the type of a function's arguments being a subtype of the function's signature (when the signature matches).

Vararg Tuple Types

The last parameter of a tuple type can be the special type `Vararg`, which denotes any number of trailing elements:

```julia
julia> mytupletype = Tuple{AbstractString,Vararg{Int}}
Tuple{AbstractString,Vararg{Int64,N} where N}

julia> isa(("1",), mytupletype)
true

julia> isa(("1",1), mytupletype)
true

julia> isa(("1",1,2), mytupletype)
true

julia> isa(("1",1,2,3.0), mytupletype)
false
```

Notice that `Vararg{T}` corresponds to zero or more elements of type T. Vararg tuple types are used to represent the arguments accepted by varargs methods (see Varargs Functions).

The type `Vararg{T,N}` corresponds to exactly N elements of type T. `NTuple{N,T}` is a convenient alias for `Tuple{Vararg{T,N}}`, i.e. a tuple type containing exactly N elements of type T.

Singleton Types

There is a special kind of abstract parametric type that must be mentioned here: singleton types. For each type, T, the "singleton type" `Type{T}` is an abstract type whose only instance is the object T. Since the definition is a little difficult to parse, let's look at some examples:

```julia
julia> isa(Float64, Type{Float64})
true

julia> isa(Real, Type{Float64})
false

julia> isa(Real, Type{Real})
true

julia> isa(Float64, Type{Real})
false
```

In other words, `isa(A, Type{B})` is true if and only if A and B are the same object and that object is a type. Without the parameter, Type is simply an abstract type which has all type objects as its instances, including, of course, singleton types:

```julia
julia> isa(Type{Float64}, Type)
true

julia> isa(Float64, Type)
true

julia> isa(Real, Type)
true
```

Any object that is not a type is not an instance of Type:

```julia
julia> isa(1, Type)
false

julia> isa("foo", Type)
false
```

Until we discuss Parametric Methods and conversions, it is difficult to explain the utility of the singleton type construct, but in short, it allows one to specialize function behavior on specific type *values*. This is useful for writing methods (especially parametric ones) whose behavior depends on a type that is given as an explicit argument rather than implied by the type of one of its arguments.

A few popular languages have singleton types, including Haskell, Scala and Ruby. In general usage, the term "singleton type" refers to a type whose only instance is a single value. This meaning applies to Julia's singleton types, but with that caveat that only type objects have singleton types.

Parametric Primitive Types

Primitive types can also be declared parametrically. For example, pointers are represented as primitive types which would be declared in Julia like this:

```julia
# 32-bit system:
primitive type Ptr{T} 32 end

# 64-bit system:
primitive type Ptr{T} 64 end
```

The slightly odd feature of these declarations as compared to typical parametric composite types, is that the type parameter T is not used in the definition of the type itself – it is just an abstract tag, essentially defining an entire family of types with identical structure, differentiated only by their type parameter. Thus, Ptr{Float64} and Ptr{Int64} are distinct types, even though they have identical representations. And of course, all specific pointer types are subtypes of the umbrella Ptr type:

```julia
julia> Ptr{Float64} <: Ptr
true

julia> Ptr{Int64} <: Ptr
true
```

14.9 UnionAll Types

We have said that a parametric type like Ptr acts as a supertype of all its instances (Ptr{Int64} etc.). How does this work? Ptr itself cannot be a normal data type, since without knowing the type of the referenced data the type clearly cannot be used for memory operations. The answer is that Ptr (or other parametric types like Array) is a different kind of type called a UnionAll type. Such a type expresses the *iterated union* of types for all values of some parameter.

UnionAll types are usually written using the keyword where. For example Ptr could be more accurately written as Ptr{T} where T, meaning all values whose type is Ptr{T} for some value of T. In this context, the parameter T is also often called a "type variable" since it is like a variable that ranges over types. Each where introduces a single type variable, so these expressions are nested for types with multiple parameters, for example Array{T,N} where N where T.

The type application syntax A{B,C} requires A to be a UnionAll type, and first substitutes B for the outermost type variable in A. The result is expected to be another UnionAll type, into which C is then substituted. So A{B,C} is

equivalent to A{B}{C}. This explains why it is possible to partially instantiate a type, as in `Array{Float64}`: the first parameter value has been fixed, but the second still ranges over all possible values. Using explicit `where` syntax, any subset of parameters can be fixed. For example, the type of all 1-dimensional arrays can be written as `Array{T,1} where T`.

Type variables can be restricted with subtype relations. `Array{T} where T<:Integer` refers to all arrays whose element type is some kind of `Integer`. The syntax `Array{<:Integer}` is a convenient shorthand for `Array{T} where T<:Integer`. Type variables can have both lower and upper bounds. `Array{T} where Int<:T<:Number` refers to all arrays of `Number`s that are able to contain `Int`s (since T must be at least as big as Int). The syntax `where T>:Int` also works to specify only the lower bound of a type variable, and `Array{>:Int}` is equivalent to `Array{T} where T>:Int`.

Since `where` expressions nest, type variable bounds can refer to outer type variables. For example `Tuple{T,Array{S}} where S<:AbstractArray{T} where T<:Real` refers to 2-tuples whose first element is some `Real`, and whose second element is an `Array` of any kind of array whose element type contains the type of the first tuple element.

The `where` keyword itself can be nested inside a more complex declaration. For example, consider the two types created by the following declarations:

```julia
julia> const T1 = Array{Array{T,1} where T, 1}
Array{Array{T,1} where T,1}

julia> const T2 = Array{Array{T,1}, 1} where T
Array{Array{T,1},1} where T
```

Type T1 defines a 1-dimensional array of 1-dimensional arrays; each of the inner arrays consists of objects of the same type, but this type may vary from one inner array to the next. On the other hand, type T2 defines a 1-dimensional array of 1-dimensional arrays all of whose inner arrays must have the same type. Note that T2 is an abstract type, e.g., `Array{Array{Int,1},1} <: T2`, whereas T1 is a concrete type. As a consequence, T1 can be constructed with a zero-argument constructor a=T1() but T2 cannot.

There is a convenient syntax for naming such types, similar to the short form of function definition syntax:

```julia
Vector{T} = Array{T,1}
```

This is equivalent to `const Vector = Array{T,1} where T`. Writing `Vector{Float64}` is equivalent to writing `Array{Float64,1}`, and the umbrella type `Vector` has as instances all `Array` objects where the second parameter – the number of array dimensions – is 1, regardless of what the element type is. In languages where parametric types must always be specified in full, this is not especially helpful, but in Julia, this allows one to write just `Vector` for the abstract type including all one-dimensional dense arrays of any element type.

14.10 Type Aliases

Sometimes it is convenient to introduce a new name for an already expressible type. This can be done with a simple assignment statement. For example, UInt is aliased to either `UInt32` or `UInt64` as is appropriate for the size of pointers on the system:

```julia
# 32-bit system:
julia> UInt
UInt32

# 64-bit system:
```

```julia
julia> UInt
UInt64
```

This is accomplished via the following code in `base/boot.jl`:

```julia
if Int === Int64
    const UInt = UInt64
else
    const UInt = UInt32
end
```

Of course, this depends on what `Int` is aliased to – but that is predefined to be the correct type – either `Int32` or `Int64`.

(Note that unlike `Int`, `Float` does not exist as a type alias for a specific sized `AbstractFloat`. Unlike with integer registers, the floating point register sizes are specified by the IEEE-754 standard. Whereas the size of `Int` reflects the size of a native pointer on that machine.)

14.11   Operations on Types

Since types in Julia are themselves objects, ordinary functions can operate on them. Some functions that are particularly useful for working with or exploring types have already been introduced, such as the `<:` operator, which indicates whether its left hand operand is a subtype of its right hand operand.

The `isa` function tests if an object is of a given type and returns true or false:

```julia
julia> isa(1, Int)
true

julia> isa(1, AbstractFloat)
false
```

The `typeof()` function, already used throughout the manual in examples, returns the type of its argument. Since, as noted above, types are objects, they also have types, and we can ask what their types are:

```julia
julia> typeof(Rational{Int})
DataType

julia> typeof(Union{Real,Float64,Rational})
DataType

julia> typeof(Union{Real,String})
Union
```

What if we repeat the process? What is the type of a type of a type? As it happens, types are all composite values and thus all have a type of `DataType`:

```julia
julia> typeof(DataType)
DataType

julia> typeof(Union)
DataType
```

`DataType` is its own type.

Another operation that applies to some types is `supertype()`, which reveals a type's supertype. Only declared types (`DataType`) have unambiguous supertypes:

```julia
julia> supertype(Float64)
AbstractFloat

julia> supertype(Number)
Any

julia> supertype(AbstractString)
Any

julia> supertype(Any)
Any
```

If you apply `supertype()` to other type objects (or non-type objects), a `MethodError` is raised:

```julia
julia> supertype(Union{Float64,Int64})
ERROR: MethodError: no method matching supertype(::Type{Union{Float64, Int64}})
Closest candidates are:
  supertype(!Matched::DataType) at operators.jl:41
  supertype(!Matched::UnionAll) at operators.jl:46
```

14.12 Custom pretty-printing

Often, one wants to customize how instances of a type are displayed. This is accomplished by overloading the `show()` function. For example, suppose we define a type to represent complex numbers in polar form:

```julia
julia> struct Polar{T<:Real} <: Number
           r::T
           Θ::T
       end

julia> Polar(r::RealΘ,::Real) = Polar(promote(rΘ,)...)
Polar
```

Here, we've added a custom constructor function so that it can take arguments of different `Real` types and promote them to a common type (see Constructors and Conversion and Promotion). (Of course, we would have to define lots of other methods, too, to make it act like a `Number`, e.g. `+`, `*`, `one`, `zero`, promotion rules and so on.) By default, instances of this type display rather simply, with information about the type name and the field values, as e.g. `Polar{Float64}(3.0,4.0)`.

If we want it to display instead as `3.0 * exp(4.0im)`, we would define the following method to print the object to a given output object `io` (representing a file, terminal, buffer, etcetera; see Networking and Streams):

```julia
julia> Base.show(io::IO, z::Polar) = print(io, z.r, " * exp(", zΘ., "im)")
```

More fine-grained control over display of `Polar` objects is possible. In particular, sometimes one wants both a verbose multi-line printing format, used for displaying a single object in the REPL and other interactive environments, and also a more compact single-line format used for `print()` or for displaying the object as part of another object (e.g. in an array). Although by default the `show(io, z)` function is called in both cases, you can define a *different* multi-line format for displaying an object by overloading a three-argument form of show that takes the `text/plain` MIME type as its second argument (see Multimedia I/O), for example:

```
julia> Base.show{T}(io::IO, ::MIME"text/plain", z::Polar{T}) =
           print(io, "Polar{$T} complex number:\n   ", z)
```

(Note that `print(..., z)` here will call the 2-argument `show(io, z)` method.) This results in:

```
julia> Polar(3, 4.0)
Polar{Float64} complex number:
   3.0 * exp(4.0im)

julia> [Polar(3, 4.0), Polar(4.0,5.3)]
2-element Array{Polar{Float64},1}:
 3.0 * exp(4.0im)
 4.0 * exp(5.3im)
```

where the single-line `show(io, z)` form is still used for an array of `Polar` values. Technically, the REPL calls `display(z)` to display the result of executing a line, which defaults to `show(STDOUT, MIME("text/plain"), z)`, which in turn defaults to `show(STDOUT, z)`, but you should *not* define new `display()` methods unless you are defining a new multimedia display handler (see Multimedia I/O).

Moreover, you can also define `show` methods for other MIME types in order to enable richer display (HTML, images, etcetera) of objects in environments that support this (e.g. IJulia). For example, we can define formatted HTML display of `Polar` objects, with superscripts and italics, via:

```
julia> Base.show{T}(io::IO, ::MIME"text/html", z::Polar{T}) =
           println(io, "<code>Polar{$T}</code> complex number: ",
                   z.r, " <i>e</i><sup>", z0., " <i>i</i></sup>")
```

A `Polar` object will then display automatically using HTML in an environment that supports HTML display, but you can call `show` manually to get HTML output if you want:

```
julia> show(STDOUT, "text/html", Polar(3.0,4.0))
<code>Polar{Float64}</code> complex number: 3.0 <i>e</i><sup>4.0 <i>i</i></sup>
```

14.13   "Value types"

In Julia, you can't dispatch on a *value* such as `true` or `false`. However, you can dispatch on parametric types, and Julia allows you to include "plain bits" values (Types, Symbols, Integers, floating-point numbers, tuples, etc.) as type parameters. A common example is the dimensionality parameter in `Array{T,N}`, where T is a type (e.g., Float64) but N is just an `Int`.

You can create your own custom types that take values as parameters, and use them to control dispatch of custom types. By way of illustration of this idea, let's introduce a parametric type, `Val{T}`, which serves as a customary way to exploit this technique for cases where you don't need a more elaborate hierarchy.

`Val` is defined as:

```
julia> struct Val{T}
       end
```

There is no more to the implementation of `Val` than this. Some functions in Julia's standard library accept `Val` types as arguments, and you can also use it to write your own functions. For example:

```
julia> firstlast(::Type{Val{true}}) = "First"
firstlast (generic function with 1 method)

julia> firstlast(::Type{Val{false}}) = "Last"
firstlast (generic function with 2 methods)
```

```
julia> firstlast(Val{true})
"First"

julia> firstlast(Val{false})
"Last"
```

For consistency across Julia, the call site should always pass a Val*type* rather than creating an *instance*, i.e., use `foo(Val{:bar})` rather than `foo(Val{:bar}())`.

It's worth noting that it's extremely easy to mis-use parametric "value" types, including `Val`; in unfavorable cases, you can easily end up making the performance of your code much *worse*. In particular, you would never want to write actual code as illustrated above. For more information about the proper (and improper) uses of `Val`, please read the more extensive discussion in the performance tips.

14.14 Nullable Types: Representing Missing Values

In many settings, you need to interact with a value of type T that may or may not exist. To handle these settings, Julia provides a parametric type called `Nullable{T}`, which can be thought of as a specialized container type that can contain either zero or one values. `Nullable{T}` provides a minimal interface designed to ensure that interactions with missing values are safe. At present, the interface consists of several possible interactions:

- Construct a `Nullable` object.

- Check if a `Nullable` object has a missing value.

- Access the value of a `Nullable` object with a guarantee that a `NullException` will be thrown if the object's value is missing.

- Access the value of a `Nullable` object with a guarantee that a default value of type T will be returned if the object's value is missing.

- Perform an operation on the value (if it exists) of a `Nullable` object, getting a `Nullable` result. The result will be missing if the original value was missing.

- Performing a test on the value (if it exists) of a `Nullable` object, getting a result that is missing if either the `Nullable` itself was missing, or the test failed.

- Perform general operations on single `Nullable` objects, propagating the missing data.

Constructing `Nullable` objects

To construct an object representing a missing value of type T, use the `Nullable{T}()` function:

```julia
julia> x1 = Nullable{Int64}()
Nullable{Int64}()

julia> x2 = Nullable{Float64}()
Nullable{Float64}()

julia> x3 = Nullable{Vector{Int64}}()
Nullable{Array{Int64,1}}()
```

To construct an object representing a non-missing value of type T, use the `Nullable(x::T)` function:

```julia
julia> x1 = Nullable(1)
Nullable{Int64}(1)

julia> x2 = Nullable(1.0)
Nullable{Float64}(1.0)

julia> x3 = Nullable([1, 2, 3])
Nullable{Array{Int64,1}}([1, 2, 3])
```

Note the core distinction between these two ways of constructing a `Nullable` object: in one style, you provide a type, T, as a function parameter; in the other style, you provide a single value of type T as an argument.

Checking if a `Nullable` object has a value

You can check if a `Nullable` object has any value using `isnull()`:

```julia
julia> isnull(Nullable{Float64}())
true

julia> isnull(Nullable(0.0))
false
```

Safely accessing the value of a `Nullable` object

You can safely access the value of a `Nullable` object using `get()`:

```julia
julia> get(Nullable{Float64}())
ERROR: NullException()
Stacktrace:
 [1] get(::Nullable{Float64}) at ./nullable.jl:92

julia> get(Nullable(1.0))
1.0
```

If the value is not present, as it would be for `Nullable{Float64}`, a `NullException` error will be thrown. The error-throwing nature of the `get()` function ensures that any attempt to access a missing value immediately fails.

In cases for which a reasonable default value exists that could be used when a `Nullable` object's value turns out to be missing, you can provide this default value as a second argument to `get()`:

```julia
julia> get(Nullable{Float64}(), 0.0)
0.0

julia> get(Nullable(1.0), 0.0)
1.0
```

> **Tip**
>
> Make sure the type of the default value passed to `get()` and that of the `Nullable` object match to avoid type instability, which could hurt performance. Use `convert()` manually if needed.

Performing operations on `Nullable` objects

`Nullable` objects represent values that are possibly missing, and it is possible to write all code using these objects by first testing to see if the value is missing with `isnull()`, and then doing an appropriate action. However, there are some common use cases where the code could be more concise or clear by using a higher-order function.

The `map` function takes as arguments a function f and a `Nullable` value x. It produces a `Nullable`:

- If x is a missing value, then it produces a missing value;

- If x has a value, then it produces a `Nullable` containing `f(get(x))` as value.

This is useful for performing simple operations on values that might be missing if the desired behaviour is to simply propagate the missing values forward.

The `filter` function takes as arguments a predicate function p (that is, a function returning a boolean) and a `Nullable` value x. It produces a `Nullable` value:

- If x is a missing value, then it produces a missing value;

- If `p(get(x))` is true, then it produces the original value x;

- If `p(get(x))` is false, then it produces a missing value.

In this way, `filter` can be thought of as selecting only allowable values, and converting non-allowable values to missing values.

While `map` and `filter` are useful in specific cases, by far the most useful higher-order function is `broadcast`, which can handle a wide variety of cases, including making existing operations work and propagate `Nullables`. An example will motivate the need for `broadcast`. Suppose we have a function that computes the greater of two real roots of a quadratic equation, using the quadratic formula:

```julia
julia> root(a::Real, b::Real, c::Real) = (-b + √(b^2 - 4a*c)) / 2a
root (generic function with 1 method)
```

We may verify that the result of `root(1, -9, 20)` is `5.0`, as we expect, since `5.0` is the greater of two real roots of the quadratic equation.

Suppose now that we want to find the greatest real root of a quadratic equations where the coefficients might be missing values. Having missing values in datasets is a common occurrence in real-world data, and so it is important to be able to deal with them. But we cannot find the roots of an equation if we do not know all the coefficients. The best solution to this will depend on the particular use case; perhaps we should throw an error. However, for this example, we will assume that the best solution is to propagate the missing values forward; that is, if any input is missing, we simply produce a missing output.

The `broadcast()` function makes this task easy; we can simply pass the `root` function we wrote to `broadcast`:

```julia
julia> broadcast(root, Nullable(1), Nullable(-9), Nullable(20))
Nullable{Float64}(5.0)

julia> broadcast(root, Nullable(1), Nullable{Int}(), Nullable{Int}())
Nullable{Float64}()

julia> broadcast(root, Nullable{Int}(), Nullable(-9), Nullable(20))
Nullable{Float64}()
```

If one or more of the inputs is missing, then the output of `broadcast()` will be missing.

There exists special syntactic sugar for the `broadcast()` function using a dot notation:

```
julia> root.(Nullable(1), Nullable(-9), Nullable(20))
Nullable{Float64}(5.0)
```

In particular, the regular arithmetic operators can be `broadcast()` conveniently using `.`-prefixed operators:

```
julia> Nullable(2) ./ Nullable(3) .+ Nullable(1.0)
Nullable{Float64}(1.66667)
```

Chapter 15

Methods

Recall from Functions that a function is an object that maps a tuple of arguments to a return value, or throws an exception if no appropriate value can be returned. It is common for the same conceptual function or operation to be implemented quite differently for different types of arguments: adding two integers is very different from adding two floating-point numbers, both of which are distinct from adding an integer to a floating-point number. Despite their implementation differences, these operations all fall under the general concept of "addition". Accordingly, in Julia, these behaviors all belong to a single object: the + function.

To facilitate using many different implementations of the same concept smoothly, functions need not be defined all at once, but can rather be defined piecewise by providing specific behaviors for certain combinations of argument types and counts. A definition of one possible behavior for a function is called a *method*. Thus far, we have presented only examples of functions defined with a single method, applicable to all types of arguments. However, the signatures of method definitions can be annotated to indicate the types of arguments in addition to their number, and more than a single method definition may be provided. When a function is applied to a particular tuple of arguments, the most specific method applicable to those arguments is applied. Thus, the overall behavior of a function is a patchwork of the behaviors of its various method definitions. If the patchwork is well designed, even though the implementations of the methods may be quite different, the outward behavior of the function will appear seamless and consistent.

The choice of which method to execute when a function is applied is called *dispatch*. Julia allows the dispatch process to choose which of a function's methods to call based on the number of arguments given, and on the types of all of the function's arguments. This is different than traditional object-oriented languages, where dispatch occurs based only on the first argument, which often has a special argument syntax, and is sometimes implied rather than explicitly written as an argument. [1] Using all of a function's arguments to choose which method should be invoked, rather than just the first, is known as multiple dispatch. Multiple dispatch is particularly useful for mathematical code, where it makes little sense to artificially deem the operations to "belong" to one argument more than any of the others: does the addition operation in x + y belong to x any more than it does to y? The implementation of a mathematical operator generally depends on the types of all of its arguments. Even beyond mathematical operations, however, multiple dispatch ends up being a powerful and convenient paradigm for structuring and organizing programs.

15.1 Defining Methods

Until now, we have, in our examples, defined only functions with a single method having unconstrained argument types. Such functions behave just like they would in traditional dynamically typed languages. Nevertheless, we have used multiple dispatch and methods almost continually without being aware of it: all of Julia's standard functions

---

[1]In C++ or Java, for example, in a method call like `obj.meth(arg1,arg2)`, the object obj "receives" the method call and is implicitly passed to the method via the `this` keyword, rather than as an explicit method argument. When the current `this` object is the receiver of a method call, it can be omitted altogether, writing just `meth(arg1,arg2)`, with `this` implied as the receiving object.

and operators, like the aforementioned + function, have many methods defining their behavior over various possible combinations of argument type and count.

When defining a function, one can optionally constrain the types of parameters it is applicable to, using the :: type-assertion operator, introduced in the section on Composite Types:

```julia
julia> f(x::Float64, y::Float64) = 2x + y
f (generic function with 1 method)
```

This function definition applies only to calls where x and y are both values of type Float64:

```julia
julia> f(2.0, 3.0)
7.0
```

Applying it to any other types of arguments will result in a MethodError:

```julia
julia> f(2.0, 3)
ERROR: MethodError: no method matching f(::Float64, ::Int64)
Closest candidates are:
  f(::Float64, !Matched::Float64) at none:1

julia> f(Float32(2.0), 3.0)
ERROR: MethodError: no method matching f(::Float32, ::Float64)
Closest candidates are:
  f(!Matched::Float64, ::Float64) at none:1

julia> f(2.0, "3.0")
ERROR: MethodError: no method matching f(::Float64, ::String)
Closest candidates are:
  f(::Float64, !Matched::Float64) at none:1

julia> f("2.0", "3.0")
ERROR: MethodError: no method matching f(::String, ::String)
```

As you can see, the arguments must be precisely of type Float64. Other numeric types, such as integers or 32-bit floating-point values, are not automatically converted to 64-bit floating-point, nor are strings parsed as numbers. Because Float64 is a concrete type and concrete types cannot be subclassed in Julia, such a definition can only be applied to arguments that are exactly of type Float64. It may often be useful, however, to write more general methods where the declared parameter types are abstract:

```julia
julia> f(x::Number, y::Number) = 2x - y
f (generic function with 2 methods)

julia> f(2.0, 3)
1.0
```

This method definition applies to any pair of arguments that are instances of Number. They need not be of the same type, so long as they are each numeric values. The problem of handling disparate numeric types is delegated to the arithmetic operations in the expression 2x - y.

To define a function with multiple methods, one simply defines the function multiple times, with different numbers and types of arguments. The first method definition for a function creates the function object, and subsequent method definitions add new methods to the existing function object. The most specific method definition matching the number and types of the arguments will be executed when the function is applied. Thus, the two method definitions above, taken together, define the behavior for f over all pairs of instances of the abstract type Number – but with a different behavior specific to pairs of Float64 values. If one of the arguments is a 64-bit float but the other one is not, then the f(Float64,Float64) method cannot be called and the more general f(Number,Number) method must be used:

```
julia> f(2.0, 3.0)
7.0

julia> f(2, 3.0)
1.0

julia> f(2.0, 3)
1.0

julia> f(2, 3)
1
```

The 2x + y definition is only used in the first case, while the 2x – y definition is used in the others. No automatic casting or conversion of function arguments is ever performed: all conversion in Julia is non-magical and completely explicit. Conversion and Promotion, however, shows how clever application of sufficiently advanced technology can be indistinguishable from magic. [2]

For non-numeric values, and for fewer or more than two arguments, the function f remains undefined, and applying it will still result in a MethodError:

```
julia> f("foo", 3)
ERROR: MethodError: no method matching f(::String, ::Int64)
Closest candidates are:
  f(!Matched::Number, ::Number) at none:1

julia> f()
ERROR: MethodError: no method matching f()
Closest candidates are:
  f(!Matched::Float64, !Matched::Float64) at none:1
  f(!Matched::Number, !Matched::Number) at none:1
```

You can easily see which methods exist for a function by entering the function object itself in an interactive session:

```
julia> f
f (generic function with 2 methods)
```

This output tells us that f is a function object with two methods. To find out what the signatures of those methods are, use the methods() function:

```
julia> methods(f)
# 2 methods for generic function "f":
f(x::Float64, y::Float64) in Main at none:1
f(x::Number, y::Number) in Main at none:1
```

which shows that f has two methods, one taking two Float64 arguments and one taking arguments of type Number. It also indicates the file and line number where the methods were defined: because these methods were defined at the REPL, we get the apparent line number none:1.

In the absence of a type declaration with ::, the type of a method parameter is Any by default, meaning that it is unconstrained since all values in Julia are instances of the abstract type Any. Thus, we can define a catch-all method for f like so:

```
julia> f(x,y) = println("Whoa there, Nelly.")
f (generic function with 3 methods)

julia> f("foo", 1)
Whoa there, Nelly.
```

This catch-all is less specific than any other possible method definition for a pair of parameter values, so it will only be called on pairs of arguments to which no other method definition applies.

Although it seems a simple concept, multiple dispatch on the types of values is perhaps the single most powerful and central feature of the Julia language. Core operations typically have dozens of methods:

```julia
julia> methods(+)
# 180 methods for generic function "+":
+(x::Bool, z::Complex{Bool}) in Base at complex.jl:224
+(x::Bool, y::Bool) in Base at bool.jl:89
+(x::Bool) in Base at bool.jl:86
+(x::Bool, y::T) where T<:AbstractFloat in Base at bool.jl:96
+(x::Bool, z::Complex) in Base at complex.jl:231
+(a::Float16, b::Float16) in Base at float.jl:372
+(x::Float32, y::Float32) in Base at float.jl:374
+(x::Float64, y::Float64) in Base at float.jl:375
+(z::Complex{Bool}, x::Bool) in Base at complex.jl:225
+(z::Complex{Bool}, x::Real) in Base at complex.jl:239
+(x::Char, y::Integer) in Base at char.jl:40
+(c::BigInt, x::BigFloat) in Base.MPFR at mpfr.jl:303
+(a::BigInt, b::BigInt, c::BigInt, d::BigInt, e::BigInt) in Base.GMP at gmp.jl:303
+(a::BigInt, b::BigInt, c::BigInt, d::BigInt) in Base.GMP at gmp.jl:296
+(a::BigInt, b::BigInt, c::BigInt) in Base.GMP at gmp.jl:290
+(x::BigInt, y::BigInt) in Base.GMP at gmp.jl:258
+(x::BigInt, c::Union{UInt16, UInt32, UInt64, UInt8}) in Base.GMP at gmp.jl:315
...
+(a, b, c, xs...) at operators.jl:119
```

Multiple dispatch together with the flexible parametric type system give Julia its ability to abstractly express high-level algorithms decoupled from implementation details, yet generate efficient, specialized code to handle each case at run time.

15.2   Method Ambiguities

It is possible to define a set of function methods such that there is no unique most specific method applicable to some combinations of arguments:

```julia
julia> g(x::Float64, y) = 2x + y
g (generic function with 1 method)

julia> g(x, y::Float64) = x + 2y
g (generic function with 2 methods)

julia> g(2.0, 3)
7.0

julia> g(2, 3.0)
8.0

julia> g(2.0, 3.0)
ERROR: MethodError: g(::Float64, ::Float64) is ambiguous.
[...]
```

Here the call g(2.0, 3.0) could be handled by either the g(Float64, Any) or the g(Any, Float64) method, and neither is more specific than the other. In such cases, Julia raises a MethodError rather than arbitrarily picking a method. You can avoid method ambiguities by specifying an appropriate method for the intersection case:

```
julia> g(x::Float64, y::Float64) = 2x + 2y
g (generic function with 3 methods)

julia> g(2.0, 3)
7.0

julia> g(2, 3.0)
8.0

julia> g(2.0, 3.0)
10.0
```

It is recommended that the disambiguating method be defined first, since otherwise the ambiguity exists, if transiently, until the more specific method is defined.

In more complex cases, resolving method ambiguities involves a certain element of design; this topic is explored further below.

15.3 Parametric Methods

Method definitions can optionally have type parameters qualifying the signature:

```
julia> same_type(x::T, y::T) where {T} = true
same_type (generic function with 1 method)

julia> same_type(x,y) = false
same_type (generic function with 2 methods)
```

The first method applies whenever both arguments are of the same concrete type, regardless of what type that is, while the second method acts as a catch-all, covering all other cases. Thus, overall, this defines a boolean function that checks whether its two arguments are of the same type:

```
julia> same_type(1, 2)
true

julia> same_type(1, 2.0)
false

julia> same_type(1.0, 2.0)
true

julia> same_type("foo", 2.0)
false

julia> same_type("foo", "bar")
true

julia> same_type(Int32(1), Int64(2))
false
```

Such definitions correspond to methods whose type signatures are UnionAll types (see UnionAll Types).

This kind of definition of function behavior by dispatch is quite common – idiomatic, even – in Julia. Method type parameters are not restricted to being used as the types of arguments: they can be used anywhere a value would be in the signature of the function or body of the function. Here's an example where the method type parameter T is used as the type parameter to the parametric type Vector{T} in the method signature:

```julia
julia> myappend(v::Vector{T}, x::T) where {T} = [v..., x]
myappend (generic function with 1 method)

julia> myappend([1,2,3],4)
4-element Array{Int64,1}:
 1
 2
 3
 4

julia> myappend([1,2,3],2.5)
ERROR: MethodError: no method matching myappend(::Array{Int64,1}, ::Float64)
Closest candidates are:
  myappend(::Array{T,1}, !Matched::T) where T at none:1

julia> myappend([1.0,2.0,3.0],4.0)
4-element Array{Float64,1}:
 1.0
 2.0
 3.0
 4.0

julia> myappend([1.0,2.0,3.0],4)
ERROR: MethodError: no method matching myappend(::Array{Float64,1}, ::Int64)
Closest candidates are:
  myappend(::Array{T,1}, !Matched::T) where T at none:1
```

As you can see, the type of the appended element must match the element type of the vector it is appended to, or else a `MethodError` is raised. In the following example, the method type parameter T is used as the return value:

```julia
julia> mytypeof(x::T) where {T} = T
mytypeof (generic function with 1 method)

julia> mytypeof(1)
Int64

julia> mytypeof(1.0)
Float64
```

Just as you can put subtype constraints on type parameters in type declarations (see Parametric Types), you can also constrain type parameters of methods:

```julia
julia> same_type_numeric(x::T, y::T) where {T<:Number} = true
same_type_numeric (generic function with 1 method)

julia> same_type_numeric(x::Number, y::Number) = false
same_type_numeric (generic function with 2 methods)

julia> same_type_numeric(1, 2)
true

julia> same_type_numeric(1, 2.0)
false
```

```
julia> same_type_numeric(1.0, 2.0)
true

julia> same_type_numeric("foo", 2.0)
ERROR: MethodError: no method matching same_type_numeric(::String, ::Float64)
Closest candidates are:
  same_type_numeric(!Matched::T<:Number, ::T<:Number) where T<:Number at none:1
  same_type_numeric(!Matched::Number, ::Number) at none:1

julia> same_type_numeric("foo", "bar")
ERROR: MethodError: no method matching same_type_numeric(::String, ::String)

julia> same_type_numeric(Int32(1), Int64(2))
false
```

The `same_type_numeric` function behaves much like the `same_type` function defined above, but is only defined for pairs of numbers.

Parametric methods allow the same syntax as `where` expressions used to write types (see UnionAll Types). If there is only a single parameter, the enclosing curly braces (in `where {T}`) can be omitted, but are often preferred for clarity. Multiple parameters can be separated with commas, e.g. `where {T, S<:Real}`, or written using nested `where`, e.g. `where S<:Real where T`.

15.4 Redefining Methods

When redefining a method or adding new methods, it is important to realize that these changes don't take effect immediately. This is key to Julia's ability to statically infer and compile code to run fast, without the usual JIT tricks and overhead. Indeed, any new method definition won't be visible to the current runtime environment, including Tasks and Threads (and any previously defined `@generated` functions). Let's start with an example to see what this means:

```
julia> function tryeval()
           @eval newfun() = 1
           newfun()
       end
tryeval (generic function with 1 method)

julia> tryeval()
ERROR: MethodError: no method matching newfun()
The applicable method may be too new: running in world age xxxx1, while current world is xxxx2.
Closest candidates are:
  newfun() at none:1 (method too new to be called from this world context.)
 in tryeval() at none:1
 ...

julia> newfun()
1
```

In this example, observe that the new definition for `newfun` has been created, but can't be immediately called. The new global is immediately visible to the `tryeval` function, so you could write `return newfun` (without parentheses). But neither you, nor any of your callers, nor the functions they call, or etc. can call this new method definition!

But there's an exception: future calls to `newfun` *from the REPL* work as expected, being able to both see and call the new definition of `newfun`.

However, future calls to `tryeval` will continue to see the definition of `newfun` as it was *at the previous statement at the REPL*, and thus before that call to `tryeval`.

You may want to try this for yourself to see how it works.

The implementation of this behavior is a "world age counter". This monotonically increasing value tracks each method definition operation. This allows describing "the set of method definitions visible to a given runtime environment" as a single number, or "world age". It also allows comparing the methods available in two worlds just by comparing their ordinal value. In the example above, we see that the "current world" (in which the method `newfun()` exists), is one greater than the task-local "runtime world" that was fixed when the execution of `tryeval` started.

Sometimes it is necessary to get around this (for example, if you are implementing the above REPL). Fortunately, there is an easy solution: call the function using `Base.invokelatest`:

```julia
julia> function tryeval2()
           @eval newfun2() = 2
           Base.invokelatest(newfun2)
       end
tryeval2 (generic function with 1 method)

julia> tryeval2()
2
```

Finally, let's take a look at some more complex examples where this rule comes into play. Define a function `f(x)`, which initially has one method:

```julia
julia> f(x) = "original definition"
f (generic function with 1 method)
```

Start some other operations that use `f(x)`:

```julia
julia> g(x) = f(x)
g (generic function with 1 method)

julia> t = @async f(wait()); yield();
```

Now we add some new methods to `f(x)`:

```julia
julia> f(x::Int) = "definition for Int"
f (generic function with 2 methods)

julia> f(x::Type{Int}) = "definition for Type{Int}"
f (generic function with 3 methods)
```

Compare how these results differ:

```julia
julia> f(1)
"definition for Int"

julia> g(1)
"definition for Int"

julia> wait(schedule(t, 1))
"original definition"

julia> t = @async f(wait()); yield();

julia> wait(schedule(t, 1))
"definition for Int"
```

15.5   Parametrically-constrained Varargs methods

Function parameters can also be used to constrain the number of arguments that may be supplied to a "varargs" function (Varargs Functions). The notation `Vararg{T,N}` is used to indicate such a constraint. For example:

```julia
julia> bar(a,b,x::Vararg{Any,2}) = (a,b,x)
bar (generic function with 1 method)

julia> bar(1,2,3)
ERROR: MethodError: no method matching bar(::Int64, ::Int64, ::Int64)
Closest candidates are:
  bar(::Any, ::Any, ::Any, !Matched::Any) at none:1

julia> bar(1,2,3,4)
(1, 2, (3, 4))

julia> bar(1,2,3,4,5)
ERROR: MethodError: no method matching bar(::Int64, ::Int64, ::Int64, ::Int64, ::Int64)
Closest candidates are:
  bar(::Any, ::Any, ::Any, ::Any) at none:1
```

More usefully, it is possible to constrain varargs methods by a parameter. For example:

```julia
function getindex(A::AbstractArray{T,N}, indexes::Vararg{Number,N}) where {T,N}
```

would be called only when the number of `indexes` matches the dimensionality of the array.

15.6   Note on Optional and keyword Arguments

As mentioned briefly in Functions, optional arguments are implemented as syntax for multiple method definitions. For example, this definition:

```julia
f(a=1,b=2) = a+2b
```

translates to the following three methods:

```julia
f(a,b) = a+2b
f(a) = f(a,2)
f() = f(1,2)
```

This means that calling `f()` is equivalent to calling `f(1,2)`. In this case the result is 5, because `f(1,2)` invokes the first method of `f` above. However, this need not always be the case. If you define a fourth method that is more specialized for integers:

```julia
f(a::Int,b::Int) = a-2b
```

then the result of both `f()` and `f(1,2)` is -3. In other words, optional arguments are tied to a function, not to any specific method of that function. It depends on the types of the optional arguments which method is invoked. When optional arguments are defined in terms of a global variable, the type of the optional argument may even change at run-time.

Keyword arguments behave quite differently from ordinary positional arguments. In particular, they do not participate in method dispatch. Methods are dispatched based only on positional arguments, with keyword arguments processed after the matching method is identified.

15.7   Function-like objects

Methods are associated with types, so it is possible to make any arbitrary Julia object "callable" by adding methods to its type. (Such "callable" objects are sometimes called "functors.")

For example, you can define a type that stores the coefficients of a polynomial, but behaves like a function evaluating the polynomial:

```julia
julia> struct Polynomial{R}
           coeffs::Vector{R}
       end

julia> function (p::Polynomial)(x)
           v = p.coeffs[end]
           for i = (length(p.coeffs)-1):-1:1
               v = v*x + p.coeffs[i]
           end
           return v
       end
```

Notice that the function is specified by type instead of by name. In the function body, p will refer to the object that was called. A `Polynomial` can be used as follows:

```julia
julia> p = Polynomial([1,10,100])
Polynomial{Int64}([1, 10, 100])

julia> p(3)
931
```

This mechanism is also the key to how type constructors and closures (inner functions that refer to their surrounding environment) work in Julia, discussed later in the manual.

15.8   Empty generic functions

Occasionally it is useful to introduce a generic function without yet adding methods. This can be used to separate interface definitions from implementations. It might also be done for the purpose of documentation or code readability. The syntax for this is an empty `function` block without a tuple of arguments:

```julia
function emptyfunc
end
```

15.9   Method design and the avoidance of ambiguities

Julia's method polymorphism is one of its most powerful features, yet exploiting this power can pose design challenges. In particular, in more complex method hierarchies it is not uncommon for ambiguities to arise.

Above, it was pointed out that one can resolve ambiguities like

```julia
f(x, y::Int) = 1
f(x::Int, y) = 2
```

by defining a method

```julia
f(x::Int, y::Int) = 3
```

This is often the right strategy; however, there are circumstances where following this advice blindly can be counter-productive. In particular, the more methods a generic function has, the more possibilities there are for ambiguities. When your method hierarchies get more complicated than this simple example, it can be worth your while to think carefully about alternative strategies.

Below we discuss particular challenges and some alternative ways to resolve such issues.

Tuple and NTuple arguments

Tuple (and NTuple) arguments present special challenges. For example,

```julia
f(x::NTuple{N,Int}) where {N} = 1
f(x::NTuple{N,Float64}) where {N} = 2
```

are ambiguous because of the possibility that N == 0: there are no elements to determine whether the Int or Float64 variant should be called. To resolve the ambiguity, one approach is define a method for the empty tuple:

```julia
f(x::Tuple{}) = 3
```

Alternatively, for all methods but one you can insist that there is at least one element in the tuple:

```julia
f(x::NTuple{N,Int}) where {N} = 1          # this is the fallback
f(x::Tuple{Float64, Vararg{Float64}}) = 2  # this requires at least one Float64
```

Orthogonalize your design

When you might be tempted to dispatch on two or more arguments, consider whether a "wrapper" function might make for a simpler design. For example, instead of writing multiple variants:

```julia
f(x::A, y::A) = ...
f(x::A, y::B) = ...
f(x::B, y::A) = ...
f(x::B, y::B) = ...
```

you might consider defining

```julia
f(x::A, y::A) = ...
f(x, y) = f(g(x), g(y))
```

where g converts the argument to type A. This is a very specific example of the more general principle of orthogonal design, in which separate concepts are assigned to separate methods. Here, g will most likely need a fallback definition

```julia
g(x::A) = x
```

A related strategy exploits promote to bring x and y to a common type:

```julia
f(x::T, y::T) where {T} = ...
f(x, y) = f(promote(x, y)...)
```

One risk with this design is the possibility that if there is no suitable promotion method converting x and y to the same type, the second method will recurse on itself infinitely and trigger a stack overflow. The non-exported function Base.promote_noncircular can be used as an alternative; when promotion fails it will still throw an error, but one that fails faster with a more specific error message.

Dispatch on one argument at a time

If you need to dispatch on multiple arguments, and there are many fallbacks with too many combinations to make it practical to define all possible variants, then consider introducing a "name cascade" where (for example) you dispatch on the first argument and then call an internal method:

```
f(x::A, y) = _fA(x, y)
f(x::B, y) = _fB(x, y)
```

Then the internal methods `_fA` and `_fB` can dispatch on `y` without concern about ambiguities with each other with respect to `x`.

Be aware that this strategy has at least one major disadvantage: in many cases, it is not possible for users to further customize the behavior of `f` by defining further specializations of your exported function `f`. Instead, they have to define specializations for your internal methods `_fA` and `_fB`, and this blurs the lines between exported and internal methods.

Abstract containers and element types

Where possible, try to avoid defining methods that dispatch on specific element types of abstract containers. For example,

```
-(A::AbstractArray{T}, b::Date) where {T<:Date}
```

generates ambiguities for anyone who defines a method

```
-(A::MyArrayType{T}, b::T) where {T}
```

The best approach is to avoid defining *either* of these methods: instead, rely on a generic method `-(A::AbstractArray, b)` and make sure this method is implemented with generic calls (like `similar` and `-`) that do the right thing for each container type and element type *separately*. This is just a more complex variant of the advice to orthogonalize your methods.

When this approach is not possible, it may be worth starting a discussion with other developers about resolving the ambiguity; just because one method was defined first does not necessarily mean that it can't be modified or eliminated. As a last resort, one developer can define the "band-aid" method

```
-(A::MyArrayType{T}, b::Date) where {T<:Date} = ...
```

that resolves the ambiguity by brute force.

Complex method "cascades" with default arguments

If you are defining a method "cascade" that supplies defaults, be careful about dropping any arguments that correspond to potential defaults. For example, suppose you're writing a digital filtering algorithm and you have a method that handles the edges of the signal by applying padding:

```
function myfilter(A, kernel, ::Replicate)
    Apadded = replicate_edges(A, size(kernel))
    myfilter(Apadded, kernel)  # now perform the "real" computation
end
```

This will run afoul of a method that supplies default padding:

```
myfilter(A, kernel) = myfilter(A, kernel, Replicate()) # replicate the edge by default
```

Together, these two methods generate an infinite recursion with A constantly growing bigger.

The better design would be to define your call hierarchy like this:

```
struct NoPad end   # indicate that no padding is desired, or that it's already applied

myfilter(A, kernel) = myfilter(A, kernel, Replicate())  # default boundary conditions

function myfilter(A, kernel, ::Replicate)
    Apadded = replicate_edges(A, size(kernel))
    myfilter(Apadded, kernel, NoPad())  # indicate the new boundary conditions
end

# other padding methods go here

function myfilter(A, kernel, ::NoPad)
    # Here's the "real" implementation of the core computation
end
```

NoPad is supplied in the same argument position as any other kind of padding, so it keeps the dispatch hierarchy well organized and with reduced likelihood of ambiguities. Moreover, it extends the "public" `myfilter` interface: a user who wants to control the padding explicitly can call the NoPad variant directly.

---

[2]Arthur C. Clarke, *Profiles of the Future* (1961): Clarke's Third Law.

Chapter 16

Constructors

Constructors [1] are functions that create new objects – specifically, instances of Composite Types. In Julia, type objects also serve as constructor functions: they create new instances of themselves when applied to an argument tuple as a function. This much was already mentioned briefly when composite types were introduced. For example:

```
julia> struct Foo
           bar
           baz
       end

julia> foo = Foo(1, 2)
Foo(1, 2)

julia> foo.bar
1

julia> foo.baz
2
```

For many types, forming new objects by binding their field values together is all that is ever needed to create instances. There are, however, cases where more functionality is required when creating composite objects. Sometimes invariants must be enforced, either by checking arguments or by transforming them. Recursive data structures, especially those that may be self-referential, often cannot be constructed cleanly without first being created in an incomplete state and then altered programmatically to be made whole, as a separate step from object creation. Sometimes, it's just convenient to be able to construct objects with fewer or different types of parameters than they have fields. Julia's system for object construction addresses all of these cases and more.

16.1 Outer Constructor Methods

A constructor is just like any other function in Julia in that its overall behavior is defined by the combined behavior of its methods. Accordingly, you can add functionality to a constructor by simply defining new methods. For example, let's say you want to add a constructor method for Foo objects that takes only one argument and uses the given value for both the bar and baz fields. This is simple:

```
julia> Foo(x) = Foo(x,x)
```

---

[1]Nomenclature: while the term "constructor" generally refers to the entire function which constructs objects of a type, it is common to abuse terminology slightly and refer to specific constructor methods as "constructors". In such situations, it is generally clear from context that the term is used to mean "constructor method" rather than "constructor function", especially as it is often used in the sense of singling out a particular method of the constructor from all of the others.

```
Foo

julia> Foo(1)
Foo(1, 1)
```

You could also add a zero-argument Foo constructor method that supplies default values for both of the bar and baz fields:

```
julia> Foo() = Foo(0)
Foo

julia> Foo()
Foo(0, 0)
```

Here the zero-argument constructor method calls the single-argument constructor method, which in turn calls the automatically provided two-argument constructor method. For reasons that will become clear very shortly, additional constructor methods declared as normal methods like this are called *outer* constructor methods. Outer constructor methods can only ever create a new instance by calling another constructor method, such as the automatically provided default ones.

16.2   Inner Constructor Methods

While outer constructor methods succeed in addressing the problem of providing additional convenience methods for constructing objects, they fail to address the other two use cases mentioned in the introduction of this chapter: enforcing invariants, and allowing construction of self-referential objects. For these problems, one needs *inner* constructor methods. An inner constructor method is much like an outer constructor method, with two differences:

1.   It is declared inside the block of a type declaration, rather than outside of it like normal methods.

2.   It has access to a special locally existent function called new that creates objects of the block's type.

For example, suppose one wants to declare a type that holds a pair of real numbers, subject to the constraint that the first number is not greater than the second one. One could declare it like this:

```
julia> struct OrderedPair
           x::Real
           y::Real
           OrderedPair(x,y) = x > y ? error("out of order") : new(x,y)
       end
```

Now OrderedPair objects can only be constructed such that x <= y:

```
julia> OrderedPair(1, 2)
OrderedPair(1, 2)

julia> OrderedPair(2,1)
ERROR: out of order
Stacktrace:
 [1] OrderedPair(::Int64, ::Int64) at ./none:4
```

If the type were declared mutable, you could reach in and directly change the field values to violate this invariant, but messing around with an object's internals uninvited is considered poor form. You (or someone else) can also provide additional outer constructor methods at any later point, but once a type is declared, there is no way to add more inner constructor methods. Since outer constructor methods can only create objects by calling other constructor methods,

ultimately, some inner constructor must be called to create an object. This guarantees that all objects of the declared type must come into existence by a call to one of the inner constructor methods provided with the type, thereby giving some degree of enforcement of a type's invariants.

If any inner constructor method is defined, no default constructor method is provided: it is presumed that you have supplied yourself with all the inner constructors you need. The default constructor is equivalent to writing your own inner constructor method that takes all of the object's fields as parameters (constrained to be of the correct type, if the corresponding field has a type), and passes them to new, returning the resulting object:

```julia
julia> struct Foo
           bar
           baz
           Foo(bar,baz) = new(bar,baz)
       end
```

This declaration has the same effect as the earlier definition of the Foo type without an explicit inner constructor method. The following two types are equivalent – one with a default constructor, the other with an explicit constructor:

```julia
julia> struct T1
           x::Int64
       end

julia> struct T2
           x::Int64
           T2(x) = new(x)
       end

julia> T1(1)
T1(1)

julia> T2(1)
T2(1)

julia> T1(1.0)
T1(1)

julia> T2(1.0)
T2(1)
```

It is considered good form to provide as few inner constructor methods as possible: only those taking all arguments explicitly and enforcing essential error checking and transformation. Additional convenience constructor methods, supplying default values or auxiliary transformations, should be provided as outer constructors that call the inner constructors to do the heavy lifting. This separation is typically quite natural.

16.3   Incomplete Initialization

The final problem which has still not been addressed is construction of self-referential objects, or more generally, recursive data structures. Since the fundamental difficulty may not be immediately obvious, let us briefly explain it. Consider the following recursive type declaration:

```julia
julia> mutable struct SelfReferential
           obj::SelfReferential
       end
```

This type may appear innocuous enough, until one considers how to construct an instance of it. If a is an instance of `SelfReferential`, then a second instance can be created by the call:

```julia
julia> b = SelfReferential(a)
```

But how does one construct the first instance when no instance exists to provide as a valid value for its `obj` field? The only solution is to allow creating an incompletely initialized instance of `SelfReferential` with an unassigned `obj` field, and using that incomplete instance as a valid value for the `obj` field of another instance, such as, for example, itself.

To allow for the creation of incompletely initialized objects, Julia allows the new function to be called with fewer than the number of fields that the type has, returning an object with the unspecified fields uninitialized. The inner constructor method can then use the incomplete object, finishing its initialization before returning it. Here, for example, we take another crack at defining the `SelfReferential` type, with a zero-argument inner constructor returning instances having `obj` fields pointing to themselves:

```julia
julia> mutable struct SelfReferential
           obj::SelfReferential
           SelfReferential() = (x = new(); x.obj = x)
       end
```

We can verify that this constructor works and constructs objects that are, in fact, self-referential:

```julia
julia> x = SelfReferential();

julia> x === x
true

julia> x === x.obj
true

julia> x === x.obj.obj
true
```

Although it is generally a good idea to return a fully initialized object from an inner constructor, incompletely initialized objects can be returned:

```julia
julia> mutable struct Incomplete
           xx
           Incomplete() = new()
       end

julia> z = Incomplete();
```

While you are allowed to create objects with uninitialized fields, any access to an uninitialized reference is an immediate error:

```julia
julia> z.xx
ERROR: UndefRefError: access to undefined reference
```

This avoids the need to continually check for `null` values. However, not all object fields are references. Julia considers some types to be "plain data", meaning all of their data is self-contained and does not reference other objects. The plain data types consist of primitive types (e.g. `Int`) and immutable structs of other plain data types. The initial contents of a plain data type is undefined:

```julia
julia> struct HasPlain
           n::Int
           HasPlain() = new()
       end

julia> HasPlain()
HasPlain(438103441441)
```

Arrays of plain data types exhibit the same behavior.

You can pass incomplete objects to other functions from inner constructors to delegate their completion:

```julia
julia> mutable struct Lazy
           xx
           Lazy(v) = complete_me(new(), v)
       end
```

As with incomplete objects returned from constructors, if `complete_me` or any of its callees try to access the `xx` field of the `Lazy` object before it has been initialized, an error will be thrown immediately.

16.4 Parametric Constructors

Parametric types add a few wrinkles to the constructor story. Recall from Parametric Types that, by default, instances of parametric composite types can be constructed either with explicitly given type parameters or with type parameters implied by the types of the arguments given to the constructor. Here are some examples:

```julia
julia> struct Point{T<:Real}
           x::T
           y::T
       end

julia> Point(1,2) ## implicit T ##
Point{Int64}(1, 2)

julia> Point(1.0,2.5) ## implicit T ##
Point{Float64}(1.0, 2.5)

julia> Point(1,2.5) ## implicit T ##
ERROR: MethodError: no method matching Point(::Int64, ::Float64)
Closest candidates are:
  Point(::T<:Real, !Matched::T<:Real) where T<:Real at none:2

julia> Point{Int64}(1, 2) ## explicit T ##
Point{Int64}(1, 2)

julia> Point{Int64}(1.0,2.5) ## explicit T ##
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Int64}, ::Float64) at ./float.jl:679
 [2] Point{Int64}(::Float64, ::Float64) at ./none:2

julia> Point{Float64}(1.0, 2.5) ## explicit T ##
Point{Float64}(1.0, 2.5)

julia> Point{Float64}(1,2) ## explicit T ##
Point{Float64}(1.0, 2.0)
```

As you can see, for constructor calls with explicit type parameters, the arguments are converted to the implied field types: `Point{Int64}(1,2)` works, but `Point{Int64}(1.0,2.5)` raises an `InexactError` when converting `2.5` to `Int64`. When the type is implied by the arguments to the constructor call, as in `Point(1,2)`, then the types of the arguments must agree – otherwise the T cannot be determined – but any pair of real arguments with matching type may be given to the generic `Point` constructor.

What's really going on here is that `Point`, `Point{Float64}` and `Point{Int64}` are all different constructor functions. In fact, `Point{T}` is a distinct constructor function for each type T. Without any explicitly provided inner constructors, the declaration of the composite type `Point{T<:Real}` automatically provides an inner constructor, `Point{T}`, for each possible type `T<:Real`, that behaves just like non-parametric default inner constructors do. It also provides a single general outer `Point` constructor that takes pairs of real arguments, which must be of the same type. This automatic provision of constructors is equivalent to the following explicit declaration:

```julia
julia> struct Point{T<:Real}
           x::T
           y::T
           Point{T}(x,y) where {T<:Real} = new(x,y)
       end

julia> Point(x::T, y::T) where {T<:Real} = Point{T}(x,y);
```

Notice that each definition looks like the form of constructor call that it handles. The call `Point{Int64}(1,2)` will invoke the definition `Point{T}(x,y)` inside the `type` block. The outer constructor declaration, on the other hand, defines a method for the general `Point` constructor which only applies to pairs of values of the same real type. This declaration makes constructor calls without explicit type parameters, like `Point(1,2)` and `Point(1.0,2.5)`, work. Since the method declaration restricts the arguments to being of the same type, calls like `Point(1,2.5)`, with arguments of different types, result in "no method" errors.

Suppose we wanted to make the constructor call `Point(1,2.5)` work by "promoting" the integer value 1 to the floating-point value `1.0`. The simplest way to achieve this is to define the following additional outer constructor method:

```julia
julia> Point(x::Int64, y::Float64) = Point(convert(Float64,x),y);
```

This method uses the `convert()` function to explicitly convert x to `Float64` and then delegates construction to the general constructor for the case where both arguments are `Float64`. With this method definition what was previously a `MethodError` now successfully creates a point of type `Point{Float64}`:

```julia
julia> Point(1,2.5)
Point{Float64}(1.0, 2.5)

julia> typeof(ans)
Point{Float64}
```

However, other similar calls still don't work:

```julia
julia> Point(1.5,2)
ERROR: MethodError: no method matching Point(::Float64, ::Int64)
Closest candidates are:
  Point(::T<:Real, !Matched::T<:Real) where T<:Real at none:1
```

For a more general way to make all such calls work sensibly, see Conversion and Promotion. At the risk of spoiling the suspense, we can reveal here that all it takes is the following outer method definition to make all calls to the general `Point` constructor work as one would expect:

```julia
julia> Point(x::Real, y::Real) = Point(promote(x,y)...);
```

The `promote` function converts all its arguments to a common type – in this case `Float64`. With this method definition, the `Point` constructor promotes its arguments the same way that numeric operators like `+` do, and works for all kinds of real numbers:

```julia
julia> Point(1.5,2)
Point{Float64}(1.5, 2.0)

julia> Point(1,1//2)
Point{Rational{Int64}}(1//1, 1//2)

julia> Point(1.0,1//2)
Point{Float64}(1.0, 0.5)
```

Thus, while the implicit type parameter constructors provided by default in Julia are fairly strict, it is possible to make them behave in a more relaxed but sensible manner quite easily. Moreover, since constructors can leverage all of the power of the type system, methods, and multiple dispatch, defining sophisticated behavior is typically quite simple.

16.5   Case Study: Rational

Perhaps the best way to tie all these pieces together is to present a real world example of a parametric composite type and its constructor methods. To that end, here is the (slightly modified) beginning of `rational.jl`, which implements Julia's Rational Numbers:

```julia
julia> struct OurRational{T<:Integer} <: Real
           num::T
           den::T
           function OurRational{T}(num::T, den::T) where T<:Integer
               if num == 0 && den == 0
                   error("invalid rational: 0//0")
               end
               g = gcd(den, num)
               num = div(num, g)
               den = div(den, g)
               new(num, den)
           end
       end

julia> OurRational(n::T, d::T) where {T<:Integer} = OurRational{T}(n,d)
OurRational

julia> OurRational(n::Integer, d::Integer) = OurRational(promote(n,d)...)
OurRational

julia> OurRational(n::Integer) = OurRational(n,one(n))
OurRational

julia> //(n::Integer, d::Integer) = OurRational(n,d)
// (generic function with 1 method)

julia> //(x::OurRational, y::Integer) = x.num // (x.den*y)
// (generic function with 2 methods)

julia> //(x::Integer, y::OurRational) = (x*y.den) // y.num
// (generic function with 3 methods)

julia> //(x::Complex, y::Real) = complex(real(x)//y, imag(x)//y)
```

```
// (generic function with 4 methods)

julia> //(x::Real, y::Complex) = x*y'//real(y*y')
// (generic function with 5 methods)

julia> function //(x::Complex, y::Complex)
           xy = x*y'
           yy = real(y*y')
           complex(real(xy)//yy, imag(xy)//yy)
       end
// (generic function with 6 methods)
```

The first line – struct OurRational{T<:Integer} <: Real – declares that OurRational takes one type parameter of an integer type, and is itself a real type. The field declarations num::T and den::T indicate that the data held in a OurRational{T} object are a pair of integers of type T, one representing the rational value's numerator and the other representing its denominator.

Now things get interesting. OurRational has a single inner constructor method which checks that both of num and den aren't zero and ensures that every rational is constructed in "lowest terms" with a non-negative denominator. This is accomplished by dividing the given numerator and denominator values by their greatest common divisor, computed using the gcd function. Since gcd returns the greatest common divisor of its arguments with sign matching the first argument (den here), after this division the new value of den is guaranteed to be non-negative. Because this is the only inner constructor for OurRational, we can be certain that OurRational objects are always constructed in this normalized form.

OurRational also provides several outer constructor methods for convenience. The first is the "standard" general constructor that infers the type parameter T from the type of the numerator and denominator when they have the same type. The second applies when the given numerator and denominator values have different types: it promotes them to a common type and then delegates construction to the outer constructor for arguments of matching type. The third outer constructor turns integer values into rationals by supplying a value of 1 as the denominator.

Following the outer constructor definitions, we have a number of methods for the // operator, which provides a syntax for writing rationals. Before these definitions, // is a completely undefined operator with only syntax and no meaning. Afterwards, it behaves just as described in Rational Numbers – its entire behavior is defined in these few lines. The first and most basic definition just makes a//b construct a OurRational by applying the OurRational constructor to a and b when they are integers. When one of the operands of // is already a rational number, we construct a new rational for the resulting ratio slightly differently; this behavior is actually identical to division of a rational with an integer. Finally, applying // to complex integral values creates an instance of Complex{OurRational} – a complex number whose real and imaginary parts are rationals:

```
julia> ans = (1 + 2im)//(1 - 2im);

julia> typeof(ans)
Complex{OurRational{Int64}}

julia> ans <: Complex{OurRational}
false
```

Thus, although the // operator usually returns an instance of OurRational, if either of its arguments are complex integers, it will return an instance of Complex{OurRational} instead. The interested reader should consider perusing the rest of rational.jl: it is short, self-contained, and implements an entire basic Julia type.

16.6   Constructors and Conversion

Constructors T(args...) in Julia are implemented like other callable objects: methods are added to their types. The type of a type is Type, so all constructor methods are stored in the method table for the Type type. This means that

you can declare more flexible constructors, e.g. constructors for abstract types, by explicitly defining methods for the appropriate types.

However, in some cases you could consider adding methods to `Base.convert` *instead* of defining a constructor, because Julia falls back to calling `convert()` if no matching constructor is found. For example, if no constructor `T(args...) = ...` exists `Base.convert(::Type{T}, args...) = ...` is called.

`convert` is used extensively throughout Julia whenever one type needs to be converted to another (e.g. in assignment, `ccall`, etcetera), and should generally only be defined (or successful) if the conversion is lossless. For example, `convert(Int, 3.0)` produces 3, but `convert(Int, 3.2)` throws an `InexactError`. If you want to define a constructor for a lossless conversion from one type to another, you should probably define a `convert` method instead.

On the other hand, if your constructor does not represent a lossless conversion, or doesn't represent "conversion" at all, it is better to leave it as a constructor rather than a `convert` method. For example, the `Array{Int}()` constructor creates a zero-dimensional `Array` of the type `Int`, but is not really a "conversion" from `Int` to an `Array`.

16.7 Outer-only constructors

As we have seen, a typical parametric type has inner constructors that are called when type parameters are known; e.g. they apply to `Point{Int}` but not to `Point`. Optionally, outer constructors that determine type parameters automatically can be added, for example constructing a `Point{Int}` from the call `Point(1,2)`. Outer constructors call inner constructors to do the core work of making an instance. However, in some cases one would rather not provide inner constructors, so that specific type parameters cannot be requested manually.

For example, say we define a type that stores a vector along with an accurate representation of its sum:

```julia
julia> struct SummedArray{T<:Number,S<:Number}
           data::Vector{T}
           sum::S
       end

julia> SummedArray(Int32[1; 2; 3], Int32(6))
SummedArray{Int32,Int32}(Int32[1, 2, 3], 6)
```

The problem is that we want S to be a larger type than T, so that we can sum many elements with less information loss. For example, when T is `Int32`, we would like S to be `Int64`. Therefore we want to avoid an interface that allows the user to construct instances of the type `SummedArray{Int32,Int32}`. One way to do this is to provide a constructor only for `SummedArray`, but inside the `type` definition block to suppress generation of default constructors:

```julia
julia> struct SummedArray{T<:Number,S<:Number}
           data::Vector{T}
           sum::S
           function SummedArray(a::Vector{T}) where T
               S = widen(T)
               new{T,S}(a, sum(S, a))
           end
       end

julia> SummedArray(Int32[1; 2; 3], Int32(6))
ERROR: MethodError: no method matching SummedArray(::Array{Int32,1}, ::Int32)
Closest candidates are:
  SummedArray(::Array{T,1}) where T at none:5
```

This constructor will be invoked by the syntax `SummedArray(a)`. The syntax `new{T,S}` allows specifying parameters for the type to be constructed, i.e. this call will return a `SummedArray{T,S}`. `new{T,S}` can be used in any constructor

definition, but for convenience the parameters to new{} are automatically derived from the type being constructed when possible.

Chapter 17

Conversion and Promotion

Julia has a system for promoting arguments of mathematical operators to a common type, which has been mentioned in various other sections, including Integers and Floating-Point Numbers, Mathematical Operations and Elementary Functions, Types, and Methods. In this section, we explain how this promotion system works, as well as how to extend it to new types and apply it to functions besides built-in mathematical operators. Traditionally, programming languages fall into two camps with respect to promotion of arithmetic arguments:

- **Automatic promotion for built-in arithmetic types and operators.** In most languages, built-in numeric types, when used as operands to arithmetic operators with infix syntax, such as `+`, `-`, `*`, and `/`, are automatically promoted to a common type to produce the expected results. C, Java, Perl, and Python, to name a few, all correctly compute the sum `1 + 1.5` as the floating-point value `2.5`, even though one of the operands to `+` is an integer. These systems are convenient and designed carefully enough that they are generally all-but-invisible to the programmer: hardly anyone consciously thinks of this promotion taking place when writing such an expression, but compilers and interpreters must perform conversion before addition since integers and floating-point values cannot be added as-is. Complex rules for such automatic conversions are thus inevitably part of specifications and implementations for such languages.

- **No automatic promotion.** This camp includes Ada and ML – very "strict" statically typed languages. In these languages, every conversion must be explicitly specified by the programmer. Thus, the example expression `1 + 1.5` would be a compilation error in both Ada and ML. Instead one must write `real(1) + 1.5`, explicitly converting the integer 1 to a floating-point value before performing addition. Explicit conversion everywhere is so inconvenient, however, that even Ada has some degree of automatic conversion: integer literals are promoted to the expected integer type automatically, and floating-point literals are similarly promoted to appropriate floating-point types.

In a sense, Julia falls into the "no automatic promotion" category: mathematical operators are just functions with special syntax, and the arguments of functions are never automatically converted. However, one may observe that applying mathematical operations to a wide variety of mixed argument types is just an extreme case of polymorphic multiple dispatch – something which Julia's dispatch and type systems are particularly well-suited to handle. "Automatic" promotion of mathematical operands simply emerges as a special application: Julia comes with pre-defined catch-all dispatch rules for mathematical operators, invoked when no specific implementation exists for some combination of operand types. These catch-all rules first promote all operands to a common type using user-definable promotion rules, and then invoke a specialized implementation of the operator in question for the resulting values, now of the same type. User-defined types can easily participate in this promotion system by defining methods for conversion to and from other types, and providing a handful of promotion rules defining what types they should promote to when mixed with other types.

17.1 Conversion

Conversion of values to various types is performed by the `convert` function. The `convert` function generally takes two arguments: the first is a type object while the second is a value to convert to that type; the returned value is the value converted to an instance of given type. The simplest way to understand this function is to see it in action:

```julia
julia> x = 12
12

julia> typeof(x)
Int64

julia> convert(UInt8, x)
0x0c

julia> typeof(ans)
UInt8

julia> convert(AbstractFloat, x)
12.0

julia> typeof(ans)
Float64

julia> a = Any[1 2 3; 4 5 6]
2×3 Array{Any,2}:
 1  2  3
 4  5  6

julia> convert(Array{Float64}, a)
2×3 Array{Float64,2}:
 1.0  2.0  3.0
 4.0  5.0  6.0
```

Conversion isn't always possible, in which case a no method error is thrown indicating that `convert` doesn't know how to perform the requested conversion:

```julia
julia> convert(AbstractFloat, "foo")
ERROR: MethodError: Cannot `convert` an object of type String to an object of type AbstractFloat
This may have arisen from a call to the constructor AbstractFloat(...),
since type constructors fall back to convert methods.
```

Some languages consider parsing strings as numbers or formatting numbers as strings to be conversions (many dynamic languages will even perform conversion for you automatically), however Julia does not: even though some strings can be parsed as numbers, most strings are not valid representations of numbers, and only a very limited subset of them are. Therefore in Julia the dedicated `parse()` function must be used to perform this operation, making it more explicit.

Defining New Conversions

To define a new conversion, simply provide a new method for `convert()`. That's really all there is to it. For example, the method to convert a real number to a boolean is this:

```julia
convert(::Type{Bool}, x::Real) = x==0 ? false : x==1 ? true : throw(InexactError())
```

The type of the first argument of this method is a singleton type, Type{Bool}, the only instance of which is Bool. Thus, this method is only invoked when the first argument is the type value Bool. Notice the syntax used for the first argument: the argument name is omitted prior to the :: symbol, and only the type is given. This is the syntax in Julia for a function argument whose type is specified but whose value is never used in the function body. In this example, since the type is a singleton, there would never be any reason to use its value within the body. When invoked, the method determines whether a numeric value is true or false as a boolean, by comparing it to one and zero:

```julia
julia> convert(Bool, 1)
true

julia> convert(Bool, 0)
false

julia> convert(Bool, 1im)
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Bool}, ::Complex{Int64}) at ./complex.jl:31

julia> convert(Bool, 0im)
false
```

The method signatures for conversion methods are often quite a bit more involved than this example, especially for parametric types. The example above is meant to be pedagogical, and is not the actual Julia behaviour. This is the actual implementation in Julia:

```julia
convert(::Type{T}, z::Complex) where {T<:Real} =
    (imag(z) == 0 ? convert(T, real(z)) : throw(InexactError()))
```

Case Study: Rational Conversions

To continue our case study of Julia's Rational type, here are the conversions declared in rational.jl, right after the declaration of the type and its constructors:

```julia
convert(::Type{Rational{T}}, x::Rational) where {T<:Integer} =
↪   Rational(convert(T,x.num),convert(T,x.den))
convert(::Type{Rational{T}}, x::Integer) where {T<:Integer} = Rational(convert(T,x),
↪   convert(T,1))

function convert(::Type{Rational{T}}, x::AbstractFloat, tol::Real) where T<:Integer
    if isnan(x); return zero(T)//zero(T); end
    if isinf(x); return sign(x)//zero(T); end
    y = x
    a = d = one(T)
    b = c = zero(T)
    while true
        f = convert(T,round(y)); y -= f
        a, b, c, d = f*a+c, f*b+d, a, b
        if y == 0 || abs(a/b-x) <= tol
            return a//b
        end
        y = 1/y
    end
end
convert(rt::Type{Rational{T}}, x::AbstractFloat) where {T<:Integer} = convert(rt,x,eps(x))
```

```
convert(::Type{T}, x::Rational) where {T<:AbstractFloat} = convert(T,x.num)/convert(T,x.den)
convert(::Type{T}, x::Rational) where {T<:Integer} = div(convert(T,x.num),convert(T,x.den))
```

The initial four convert methods provide conversions to rational types. The first method converts one type of rational to another type of rational by converting the numerator and denominator to the appropriate integer type. The second method does the same conversion for integers by taking the denominator to be 1. The third method implements a standard algorithm for approximating a floating-point number by a ratio of integers to within a given tolerance, and the fourth method applies it, using machine epsilon at the given value as the threshold. In general, one should have `a//b == convert(Rational{Int64}, a/b)`.

The last two convert methods provide conversions from rational types to floating-point and integer types. To convert to floating point, one simply converts both numerator and denominator to that floating point type and then divides. To convert to integer, one can use the `div` operator for truncated integer division (rounded towards zero).

17.2 Promotion

Promotion refers to converting values of mixed types to a single common type. Although it is not strictly necessary, it is generally implied that the common type to which the values are converted can faithfully represent all of the original values. In this sense, the term "promotion" is appropriate since the values are converted to a "greater" type – i.e. one which can represent all of the input values in a single common type. It is important, however, not to confuse this with object-oriented (structural) super-typing, or Julia's notion of abstract super-types: promotion has nothing to do with the type hierarchy, and everything to do with converting between alternate representations. For instance, although every Int32 value can also be represented as a Float64 value, Int32 is not a subtype of Float64.

Promotion to a common "greater" type is performed in Julia by the `promote` function, which takes any number of arguments, and returns a tuple of the same number of values, converted to a common type, or throws an exception if promotion is not possible. The most common use case for promotion is to convert numeric arguments to a common type:

```julia
julia> promote(1, 2.5)
(1.0, 2.5)

julia> promote(1, 2.5, 3)
(1.0, 2.5, 3.0)

julia> promote(2, 3//4)
(2//1, 3//4)

julia> promote(1, 2.5, 3, 3//4)
(1.0, 2.5, 3.0, 0.75)

julia> promote(1.5, im)
(1.5 + 0.0im, 0.0 + 1.0im)

julia> promote(1 + 2im, 3//4)
(1//1 + 2//1*im, 3//4 + 0//1*im)
```

Floating-point values are promoted to the largest of the floating-point argument types. Integer values are promoted to the larger of either the native machine word size or the largest integer argument type. Mixtures of integers and floating-point values are promoted to a floating-point type big enough to hold all the values. Integers mixed with rationals are promoted to rationals. Rationals mixed with floats are promoted to floats. Complex values mixed with real values are promoted to the appropriate kind of complex value.

That is really all there is to using promotions. The rest is just a matter of clever application, the most typical "clever" application being the definition of catch-all methods for numeric operations like the arithmetic operators `+`, `-`, `*` and `/`. Here are some of the catch-all method definitions given in `promotion.jl`:

```
+(x::Number, y::Number) = +(promote(x,y)...)
-(x::Number, y::Number) = -(promote(x,y)...)
*(x::Number, y::Number) = *(promote(x,y)...)
/(x::Number, y::Number) = /(promote(x,y)...)
```

These method definitions say that in the absence of more specific rules for adding, subtracting, multiplying and dividing pairs of numeric values, promote the values to a common type and then try again. That's all there is to it: nowhere else does one ever need to worry about promotion to a common numeric type for arithmetic operations – it just happens automatically. There are definitions of catch-all promotion methods for a number of other arithmetic and mathematical functions in `promotion.jl`, but beyond that, there are hardly any calls to `promote` required in the Julia standard library. The most common usages of `promote` occur in outer constructors methods, provided for convenience, to allow constructor calls with mixed types to delegate to an inner type with fields promoted to an appropriate common type. For example, recall that `rational.jl` provides the following outer constructor method:

```
Rational(n::Integer, d::Integer) = Rational(promote(n,d)...)
```

This allows calls like the following to work:

```
julia> Rational(Int8(15),Int32(-5))
-3//1

julia> typeof(ans)
Rational{Int32}
```

For most user-defined types, it is better practice to require programmers to supply the expected types to constructor functions explicitly, but sometimes, especially for numeric problems, it can be convenient to do promotion automatically.

Defining Promotion Rules

Although one could, in principle, define methods for the `promote` function directly, this would require many redundant definitions for all possible permutations of argument types. Instead, the behavior of `promote` is defined in terms of an auxiliary function called `promote_rule`, which one can provide methods for. The `promote_rule` function takes a pair of type objects and returns another type object, such that instances of the argument types will be promoted to the returned type. Thus, by defining the rule:

```
promote_rule(::Type{Float64}, ::Type{Float32}) = Float64
```

one declares that when 64-bit and 32-bit floating-point values are promoted together, they should be promoted to 64-bit floating-point. The promotion type does not need to be one of the argument types, however; the following promotion rules both occur in Julia's standard library:

```
promote_rule(::Type{UInt8}, ::Type{Int8}) = Int
promote_rule(::Type{BigInt}, ::Type{Int8}) = BigInt
```

In the latter case, the result type is `BigInt` since `BigInt` is the only type large enough to hold integers for arbitrary-precision integer arithmetic. Also note that one does not need to define both `promote_rule(::Type{A}, ::Type{B})` and `promote_rule(::Type{B}, ::Type{A})` – the symmetry is implied by the way `promote_rule` is used in the promotion process.

The `promote_rule` function is used as a building block to define a second function called `promote_type`, which, given any number of type objects, returns the common type to which those values, as arguments to `promote` should be promoted. Thus, if one wants to know, in absence of actual values, what type a collection of values of certain types would promote to, one can use `promote_type`:

```julia
julia> promote_type(Int8, UInt16)
Int64
```

Internally, `promote_type` is used inside of `promote` to determine what type argument values should be converted to for promotion. It can, however, be useful in its own right. The curious reader can read the code in `promotion.jl`, which defines the complete promotion mechanism in about 35 lines.

Case Study: Rational Promotions

Finally, we finish off our ongoing case study of Julia's rational number type, which makes relatively sophisticated use of the promotion mechanism with the following promotion rules:

```julia
promote_rule(::Type{Rational{T}}, ::Type{S}) where {T<:Integer,S<:Integer} =
↪   Rational{promote_type(T,S)}
promote_rule(::Type{Rational{T}}, ::Type{Rational{S}}) where {T<:Integer,S<:Integer} =
↪   Rational{promote_type(T,S)}
promote_rule(::Type{Rational{T}}, ::Type{S}) where {T<:Integer,S<:AbstractFloat} =
↪   promote_type(T,S)
```

The first rule says that promoting a rational number with any other integer type promotes to a rational type whose numerator/denominator type is the result of promotion of its numerator/denominator type with the other integer type. The second rule applies the same logic to two different types of rational numbers, resulting in a rational of the promotion of their respective numerator/denominator types. The third and final rule dictates that promoting a rational with a float results in the same type as promoting the numerator/denominator type with the float.

This small handful of promotion rules, together with the conversion methods discussed above, are sufficient to make rational numbers interoperate completely naturally with all of Julia's other numeric types – integers, floating-point numbers, and complex numbers. By providing appropriate conversion methods and promotion rules in the same manner, any user-defined numeric type can interoperate just as naturally with Julia's predefined numerics.

Chapter 18

Interfaces

A lot of the power and extensibility in Julia comes from a collection of informal interfaces. By extending a few specific methods to work for a custom type, objects of that type not only receive those functionalities, but they are also able to be used in other methods that are written to generically build upon those behaviors.

18.1   Iteration

Required methods		Brief description
`start(iter)`		Returns the initial iteration state
`next(iter, state)`		Returns the current item and the next state
`done(iter, state)`		Tests if there are any items remaining
Important optional methods	**Default definition**	**Brief description**
`iteratorsize(Iter-Type)`	`HasLength()`	One of `HasLength()`, `HasShape()`, `IsInfinite()`, or `SizeUnknown()` as appropriate
`iteratorel-type(IterType)`	`HasEl-type()`	Either `EltypeUnknown()` or `HasEltype()` as appropriate
`eltype(IterType)`	`Any`	The type the items returned by `next()`
`length(iter)`	*(undefined)*	The number of items, if known
`size(iter, [dim...])`	*(undefined)*	The number of items in each dimension, if known

Value returned by `iteratorsize(IterType)`	Required Methods
`HasLength()`	`length(iter)`
`HasShape()`	`length(iter)` and `size(iter, [dim...])`
`IsInfinite()`	*(none)*
`SizeUnknown()`	*(none)*

Value returned by `iteratoreltype(IterType)`	Required Methods
`HasEltype()`	`eltype(IterType)`
`EltypeUnknown()`	*(none)*

Sequential iteration is implemented by the methods `start()`, `done()`, and `next()`. Instead of mutating objects as they are iterated over, Julia provides these three methods to keep track of the iteration state externally from the object. The `start(iter)` method returns the initial state for the iterable object `iter`. That state gets passed along to `done(iter, state)`, which tests if there are any elements remaining, and `next(iter, state)`, which returns

161

a tuple containing the current element and an updated `state`. The `state` object can be anything, and is generally considered to be an implementation detail private to the iterable object.

Any object defines these three methods is iterable and can be used in the many functions that rely upon iteration. It can also be used directly in a `for` loop since the syntax:

```julia
for i in iter   # or  "for i = iter"
    # body
end
```

is translated into:

```julia
state = start(iter)
while !done(iter, state)
    (i, state) = next(iter, state)
    # body
end
```

A simple example is an iterable sequence of square numbers with a defined length:

```julia
julia> struct Squares
            count::Int
        end

julia> Base.start(::Squares) = 1

julia> Base.next(S::Squares, state) = (state*state, state+1)

julia> Base.done(S::Squares, state) = state > S.count

julia> Base.eltype(::Type{Squares}) = Int # Note that this is defined for the type

julia> Base.length(S::Squares) = S.count
```

With only start, next, and done definitions, the `Squares` type is already pretty powerful. We can iterate over all the elements:

```julia
julia> for i in Squares(7)
            println(i)
        end
1
4
9
16
25
36
49
```

We can use many of the builtin methods that work with iterables, like in(), mean() and std():

```julia
julia> 25 in Squares(10)
true

julia> mean(Squares(100))
3383.5
```

```
julia> std(Squares(100))
3024.355854282583
```

There are a few more methods we can extend to give Julia more information about this iterable collection. We know that the elements in a `Squares` sequence will always be `Int`. By extending the `eltype()` method, we can give that information to Julia and help it make more specialized code in the more complicated methods. We also know the number of elements in our sequence, so we can extend `length()`, too.

Now, when we ask Julia to `collect()` all the elements into an array it can preallocate a `Vector{Int}` of the right size instead of blindly `push!`ing each element into a `Vector{Any}`:

```
julia> collect(Squares(10))' # transposed to save space
1×10 RowVector{Int64,Array{Int64,1}}:
 1  4  9  16  25  36  49  64  81  100
```

While we can rely upon generic implementations, we can also extend specific methods where we know there is a simpler algorithm. For example, there's a formula to compute the sum of squares, so we can override the generic iterative version with a more performant solution:

```
julia> Base.sum(S::Squares) = (n = S.count; return n*(n+1)*(2n+1)÷6)

julia> sum(Squares(1803))
1955361914
```

This is a very common pattern throughout the Julia standard library: a small set of required methods define an informal interface that enable many fancier behaviors. In some cases, types will want to additionally specialize those extra behaviors when they know a more efficient algorithm can be used in their specific case.

18.2 Indexing

Methods to implement	Brief description
`getindex(X, i)`	`X[i]`, indexed element access
`setindex!(X, v, i)`	`X[i] = v`, indexed assignment
`endof(X)`	The last index, used in `X[end]`

For the `Squares` iterable above, we can easily compute the `i`th element of the sequence by squaring it. We can expose this as an indexing expression `S[i]`. To opt into this behavior, `Squares` simply needs to define `getindex()`:

```
julia> function Base.getindex(S::Squares, i::Int)
           1 <= i <= S.count || throw(BoundsError(S, i))
           return i*i
       end

julia> Squares(100)[23]
529
```

Additionally, to support the syntax `S[end]`, we must define `endof()` to specify the last valid index:

```
julia> Base.endof(S::Squares) = length(S)

julia> Squares(23)[end]
529
```

Note, though, that the above *only* defines `getindex()` with one integer index. Indexing with anything other than an `Int` will throw a `MethodError` saying that there was no matching method. In order to support indexing with ranges or vectors of `Int`s, separate methods must be written:

```
julia> Base.getindex(S::Squares, i::Number) = S[convert(Int, i)]

julia> Base.getindex(S::Squares, I) = [S[i] for i in I]

julia> Squares(10)[[3,4.,5]]
3-element Array{Int64,1}:
  9
 16
 25
```

While this is starting to support more of the indexing operations supported by some of the builtin types, there's still quite a number of behaviors missing. This `Squares` sequence is starting to look more and more like a vector as we've added behaviors to it. Instead of defining all these behaviors ourselves, we can officially define it as a subtype of an `AbstractArray`.

18.3   Abstract Arrays

If a type is defined as a subtype of `AbstractArray`, it inherits a very large set of rich behaviors including iteration and multidimensional indexing built on top of single-element access. See the arrays manual page and standard library section for more supported methods.

A key part in defining an `AbstractArray` subtype is `IndexStyle`. Since indexing is such an important part of an array and often occurs in hot loops, it's important to make both indexing and indexed assignment as efficient as possible. Array data structures are typically defined in one of two ways: either it most efficiently accesses its elements using just one index (linear indexing) or it intrinsically accesses the elements with indices specified for every dimension. These two modalities are identified by Julia as `IndexLinear()` and `IndexCartesian()`. Converting a linear index to multiple indexing subscripts is typically very expensive, so this provides a traits-based mechanism to enable efficient generic code for all array types.

This distinction determines which scalar indexing methods the type must define. `IndexLinear()` arrays are simple: just define `getindex(A::ArrayType, i::Int)`. When the array is subsequently indexed with a multidimensional set of indices, the fallback `getindex(A::AbstractArray, I...)()` efficiently converts the indices into one linear index and then calls the above method. `IndexCartesian()` arrays, on the other hand, require methods to be defined for each supported dimensionality with `ndims(A)Int` indices. For example, the builtin `SparseMatrixCSC` type only supports two dimensions, so it just defines `getindex(A::SparseMatrixCSC, i::Int, j::Int)()`. The same holds for `setindex!()`.

Returning to the sequence of squares from above, we could instead define it as a subtype of an `AbstractArray{Int, 1}`:

```
julia> struct SquaresVector <: AbstractArray{Int, 1}
           count::Int
       end

julia> Base.size(S::SquaresVector) = (S.count,)

julia> Base.IndexStyle(::Type{<:SquaresVector}) = IndexLinear()

julia> Base.getindex(S::SquaresVector, i::Int) = i*i
```

Note that it's very important to specify the two parameters of the `AbstractArray`; the first defines the `eltype()`, and the second defines the `ndims()`. That supertype and those three methods are all it takes for `SquaresVector` to be an iterable, indexable, and completely functional array:

```
julia> s = SquaresVector(7)
7-element SquaresVector:
```

Methods to implement		Brief description
`size(A)`		Returns a tuple containing the dimensions of `A`
`getindex(A, i::Int)`		(if `IndexLinear`) Linear scalar indexing
`getindex(A, I::Vararg{Int, N})`		(if `IndexCartesian`, where `N = ndims(A)`) N-dimensional scalar indexing
`setindex!(A, v, i::Int)`		(if `IndexLinear`) Scalar indexed assignment
`setindex!(A, v, I::Vararg{Int, N})`		(if `IndexCartesian`, where `N = ndims(A)`) N-dimensional scalar indexed assignment
Optional methods	**Default definition**	**Brief description**
`IndexStyle(::Type)`	`IndexCartesian()`	Returns either `IndexLinear()` or `IndexCartesian()`. See the description below.
`getindex(A, I...)`	defined in terms of scalar `getindex()`	Multidimensional and nonscalar indexing
`setindex!(A, I...)`	defined in terms of scalar `setindex!()`	Multidimensional and nonscalar indexed assignment
`start()/next()/done()`	defined in terms of scalar `getindex()`	Iteration
`length(A)`	`prod(size(A))`	Number of elements
`similar(A)`	`similar(A, eltype(A), size(A))`	Return a mutable array with the same shape and element type
`similar(A, ::Type{S})`	`similar(A, S, size(A))`	Return a mutable array with the same shape and the specified element type
`similar(A, dims::NTuple{Int})`	`similar(A, eltype(A), dims)`	Return a mutable array with the same element type and size *dims*
`similar(A, ::Type{S}, dims::NTuple{Int})`	`Array{S}(dims)`	Return a mutable array with the specified element type and size
Non-traditional indices	**Default definition**	**Brief description**
`indices(A)`	`map(OneTo, size(A))`	Return the `AbstractUnitRange` of valid indices
`Base.similar(A, ::Type{S}, inds::NTuple{Ind})`	`similar(A, S, Base.to_shape(inds))`	Return a mutable array with the specified indices inds (see below)
`Base.similar(T::Union{Type,Function}, inds)`	`T(Base.to_shape(inds))`	Return an array similar to T with the specified indices inds (see below)

```
  1
  4
  9
 16
 25
 36
 49

julia> s[s .> 20]
3-element Array{Int64,1}:
 25
 36
```

```
 49

julia> s \ [1 2; 3 4; 5 6; 7 8; 9 10; 11 12; 13 14]
1×2 Array{Float64,2}:
 0.305389  0.335329

julia> s  s # dot(s, s)
4676
```

As a more complicated example, let's define our own toy N-dimensional sparse-like array type built on top of `Dict`:

```
julia> struct SparseArray{T,N} <: AbstractArray{T,N}
           data::Dict{NTuple{N,Int}, T}
           dims::NTuple{N,Int}
       end

julia> SparseArray{T}(::Type{T}, dims::Int...) = SparseArray(T, dims);

julia> SparseArray{T,N}(::Type{T}, dims::NTuple{N,Int}) = SparseArray{T,N}(Dict{NTuple{N,Int}, T
    }(), dims);

julia> Base.size(A::SparseArray) = A.dims

julia> Base.similar(A::SparseArray, ::Type{T}, dims::Dims) where {T} = SparseArray(T, dims)

julia> Base.getindex(A::SparseArray{T,N}, I::Vararg{Int,N}) where {T,N} = get(A.data, I, zero(T))

julia> Base.setindex!(A::SparseArray{T,N}, v, I::Vararg{Int,N}) where {T,N} = (A.data[I] = v)
```

Notice that this is an `IndexCartesian` array, so we must manually define `getindex()` and `setindex!()` at the dimensionality of the array. Unlike the `SquaresVector`, we are able to define `setindex!()`, and so we can mutate the array:

```
julia> A = SparseArray(Float64, 3, 3)
3×3 SparseArray{Float64,2}:
 0.0  0.0  0.0
 0.0  0.0  0.0
 0.0  0.0  0.0

julia> fill!(A, 2)
3×3 SparseArray{Float64,2}:
 2.0  2.0  2.0
 2.0  2.0  2.0
 2.0  2.0  2.0

julia> A[:] = 1:length(A); A
3×3 SparseArray{Float64,2}:
 1.0  4.0  7.0
 2.0  5.0  8.0
 3.0  6.0  9.0
```

The result of indexing an `AbstractArray` can itself be an array (for instance when indexing by a Range). The AbstractArray fallback methods use `similar()` to allocate an `Array` of the appropriate size and element type, which is filled in using the basic indexing method described above. However, when implementing an array wrapper you often want the result to be wrapped as well:

```
julia> A[1:2,:]
```

```
2×3 SparseArray{Float64,2}:
 1.0  4.0  7.0
 2.0  5.0  8.0
```

In this example it is accomplished by defining `Base.similar{T}(A::SparseArray, ::Type{T}, dims::Dims)` to create the appropriate wrapped array. (Note that while `similar` supports 1- and 2-argument forms, in most case you only need to specialize the 3-argument form.) For this to work it's important that `SparseArray` is mutable (supports `setindex!`). Defining `similar()`, `getindex()` and `setindex!()` for `SparseArray` also makes it possible to copy() the array:

```
julia> copy(A)
3×3 SparseArray{Float64,2}:
 1.0  4.0  7.0
 2.0  5.0  8.0
 3.0  6.0  9.0
```

In addition to all the iterable and indexable methods from above, these types can also interact with each other and use most of the methods defined in the standard library for `AbstractArrays`:

```
julia> A[SquaresVector(3)]
3-element SparseArray{Float64,1}:
 1.0
 4.0
 9.0

julia> dot(A[:,1],A[:,2])
32.0
```

If you are defining an array type that allows non-traditional indexing (indices that start at something other than 1), you should specialize `indices`. You should also specialize similar so that the `dims` argument (ordinarily a `Dims` size-tuple) can accept `AbstractUnitRange` objects, perhaps range-types `Ind` of your own design. For more information, see Arrays with custom indices.

Chapter 19

Modules

Modules in Julia are separate variable workspaces, i.e. they introduce a new global scope. They are delimited syntactically, inside `module Name ... end`. Modules allow you to create top-level definitions (aka global variables) without worrying about name conflicts when your code is used together with somebody else's. Within a module, you can control which names from other modules are visible (via importing), and specify which of your names are intended to be public (via exporting).

The following example demonstrates the major features of modules. It is not meant to be run, but is shown for illustrative purposes:

```
module MyModule
using Lib

using BigLib: thing1, thing2

import Base.show

importall OtherLib

export MyType, foo

struct MyType
    x
end

bar(x) = 2x
foo(a::MyType) = bar(a.x) + 1

show(io::IO, a::MyType) = print(io, "MyType $(a.x)")
end
```

Note that the style is not to indent the body of the module, since that would typically lead to whole files being indented.

This module defines a type `MyType`, and two functions. Function `foo` and type `MyType` are exported, and so will be available for importing into other modules. Function `bar` is private to `MyModule`.

The statement `using Lib` means that a module called `Lib` will be available for resolving names as needed. When a global variable is encountered that has no definition in the current module, the system will search for it among variables exported by `Lib` and import it if it is found there. This means that all uses of that global within the current module will resolve to the definition of that variable in `Lib`.

The statement `using BigLib: thing1, thing2` is a syntactic shortcut for `using BigLib.thing1, BigLib.thing2`.

The `import` keyword supports all the same syntax as `using`, but only operates on a single name at a time. It does not add modules to be searched the way `using` does. `import` also differs from `using` in that functions must be imported using `import` to be extended with new methods.

In `MyModule` above we wanted to add a method to the standard show function, so we had to write `import Base.show`. Functions whose names are only visible via `using` cannot be extended.

The keyword `importall` explicitly imports all names exported by the specified module, as if `import` were individually used on all of them.

Once a variable is made visible via `using` or `import`, a module may not create its own variable with the same name. Imported variables are read-only; assigning to a global variable always affects a variable owned by the current module, or else raises an error.

19.1   Summary of module usage

To load a module, two main keywords can be used: `using` and `import`. To understand their differences, consider the following example:

```
module MyModule

export x, y

x() = "x"
y() = "y"
p() = "p"

end
```

In this module we export the `x` and `y` functions (with the keyword `export`), and also have the non-exported function p. There are several different ways to load the Module and its inner functions into the current workspace:

Import Command	What is brought into scope	Available for method extension
`using MyModule`	All exported names (x and y), `MyModule.x`, `MyModule.y` and `MyModule.p`	`MyModule.x`, `MyModule.y` and `MyModule.p`
`using MyModule.x, MyModule.p`	x and p	
`using MyModule: x, p`	x and p	
`import MyModule`	`MyModule.x`, `MyModule.y` and `MyModule.p`	`MyModule.x`, `MyModule.y` and `MyModule.p`
`import MyModule.x, MyModule.p`	x and p	x and p
`import MyModule: x, p`	x and p	x and p
`importall MyModule`	All exported names (x and y)	x and y

Modules and files

Files and file names are mostly unrelated to modules; modules are associated only with module expressions. One can have multiple files per module, and multiple modules per file:

```julia
module Foo

include("file1.jl")
include("file2.jl")

end
```

Including the same code in different modules provides mixin-like behavior. One could use this to run the same code with different base definitions, for example testing code by running it with "safe" versions of some operators:

```julia
module Normal
include("mycode.jl")
end

module Testing
include("safe_operators.jl")
include("mycode.jl")
end
```

Standard modules

There are three important standard modules: Main, Core, and Base.

Main is the top-level module, and Julia starts with Main set as the current module. Variables defined at the prompt go in Main, and whos() lists variables in Main.

Core contains all identifiers considered "built in" to the language, i.e. part of the core language and not libraries. Every module implicitly specifies using Core, since you can't do anything without those definitions.

Base is the standard library (the contents of base/). All modules implicitly contain using Base, since this is needed in the vast majority of cases.

Default top-level definitions and bare modules

In addition to using Base, modules also automatically contain a definition of the eval function, which evaluates expressions within the context of that module.

If these default definitions are not wanted, modules can be defined using the keyword baremodule instead (note: Core is still imported, as per above). In terms of baremodule, a standard module looks like this:

```julia
baremodule Mod

using Base

eval(x) = Core.eval(Mod, x)
eval(m,x) = Core.eval(m, x)

...

end
```

Relative and absolute module paths

Given the statement using Foo, the system looks for Foo within Main. If the module does not exist, the system attempts to require("Foo"), which typically results in loading code from an installed package.

However, some modules contain submodules, which means you sometimes need to access a module that is not directly available in `Main`. There are two ways to do this. The first is to use an absolute path, for example `using  Base.Sort`. The second is to use a relative path, which makes it easier to import submodules of the current module or any of its enclosing modules:

```
module Parent

module Utils
...
end

using .Utils

...
end
```

Here module `Parent` contains a submodule `Utils`, and code in `Parent` wants the contents of `Utils` to be visible. This is done by starting the `using` path with a period. Adding more leading periods moves up additional levels in the module hierarchy. For example `using  ..Utils` would look for `Utils` in `Parent`'s enclosing module rather than in `Parent` itself.

Note that relative-import qualifiers are only valid in `using` and `import` statements.

Module file paths

The global variable `LOAD_PATH` contains the directories Julia searches for modules when calling `require`. It can be extended using `push!`:

```
push!(LOAD_PATH, "/Path/To/My/Module/")
```

Putting this statement in the file `~/.juliarc.jl` will extend `LOAD_PATH` on every Julia startup. Alternatively, the module load path can be extended by defining the environment variable `JULIA_LOAD_PATH`.

Namespace miscellanea

If a name is qualified (e.g. `Base.sin`), then it can be accessed even if it is not exported. This is often useful when debugging. It can also have methods added to it by using the qualified name as the function name. However, due to syntactic ambiguities that arise, if you wish to add methods to a function in a different module whose name contains only symbols, such as an operator, `Base.+` for example, you must use `Base.:+` to refer to it. If the operator is more than one character in length you must surround it in brackets, such as: `Base.:(==)`.

Macro names are written with `@` in import and export statements, e.g. `import Mod.@mac`. Macros in other modules can be invoked as `Mod.@mac` or `@Mod.mac`.

The syntax `M.x  =  y` does not work to assign a global in another module; global assignment is always module-local.

A variable can be "reserved" for the current module without assigning to it by declaring it as `global  x` at the top level. This can be used to prevent name conflicts for globals initialized after load time.

Module initialization and precompilation

Large modules can take several seconds to load because executing all of the statements in a module often involves compiling a large amount of code. Julia provides the ability to create precompiled versions of modules to reduce this time.

To create an incremental precompiled module file, add `__precompile__()` at the top of your module file (before the `module` starts). This will cause it to be automatically compiled the first time it is imported. Alternatively, you can manually call `Base.compilecache(modulename)`. The resulting cache files will be stored in `Base.LOAD_CACHE_PATH[1]`. Subsequently, the module is automatically recompiled upon `import` whenever any of its dependencies change; dependencies are modules it imports, the Julia build, files it includes, or explicit dependencies declared by `include_dependency(path)` in the module file(s).

For file dependencies, a change is determined by examining whether the modification time (mtime) of each file loaded by `include` or added explicitly by `include_dependency` is unchanged, or equal to the modification time truncated to the nearest second (to accommodate systems that can't copy mtime with sub-second accuracy). It also takes into account whether the path to the file chosen by the search logic in `require` matches the path that had created the precompile file.

It also takes into account the set of dependencies already loaded into the current process and won't recompile those modules, even if their files change or disappear, in order to avoid creating incompatibilities between the running system and the precompile cache. If you want to have changes to the source reflected in the running system, you should call `reload("Module")` on the module you changed, and any module that depended on it in which you want to see the change reflected.

Precompiling a module also recursively precompiles any modules that are imported therein. If you know that it is *not* safe to precompile your module (for the reasons described below), you should put `__precompile__(false)` in the module file to cause `Base.compilecache` to throw an error (and thereby prevent the module from being imported by any other precompiled module).

`__precompile__()` should *not* be used in a module unless all of its dependencies are also using `__precompile__()`. Failure to do so can result in a runtime error when loading the module.

In order to make your module work with precompilation, however, you may need to change your module to explicitly separate any initialization steps that must occur at *runtime* from steps that can occur at *compile time*. For this purpose, Julia allows you to define an `__init__()` function in your module that executes any initialization steps that must occur at runtime. This function will not be called during compilation (`--output-*` or `__precompile__()`). You may, of course, call it manually if necessary, but the default is to assume this function deals with computing state for the local machine, which does not need to be – or even should not be – captured in the compiled image. It will be called after the module is loaded into a process, including if it is being loaded into an incremental compile (`--output-incremental=yes`), but not if it is being loaded into a full-compilation process.

In particular, if you define a `function __init__()` in a module, then Julia will call `__init__()` immediately *after* the module is loaded (e.g., by `import`, `using`, or `require`) at runtime for the *first* time (i.e., `__init__` is only called once, and only after all statements in the module have been executed). Because it is called after the module is fully imported, any submodules or other imported modules have their `__init__` functions called *before* the `__init__` of the enclosing module.

Two typical uses of `__init__` are calling runtime initialization functions of external C libraries and initializing global constants that involve pointers returned by external libraries. For example, suppose that we are calling a C library `libfoo` that requires us to call a `foo_init()` initialization function at runtime. Suppose that we also want to define a global constant `foo_data_ptr` that holds the return value of a `void *foo_data()` function defined by `libfoo` – this constant must be initialized at runtime (not at compile time) because the pointer address will change from run to run. You could accomplish this by defining the following `__init__` function in your module:

```
const foo_data_ptr = Ref{Ptr{Void}}(0)
function __init__()
    ccall((:foo_init, :libfoo), Void, ())
    foo_data_ptr[] = ccall((:foo_data, :libfoo), Ptr{Void}, ())
end
```

Notice that it is perfectly possible to define a global inside a function like `__init__`; this is one of the advantages of using a dynamic language. But by making it a constant at global scope, we can ensure that the type is known to the compiler and allow it to generate better optimized code. Obviously, any other globals in your module that depends on `foo_data_ptr` would also have to be initialized in `__init__`.

Constants involving most Julia objects that are not produced by `ccall` do not need to be placed in `__init__`: their definitions can be precompiled and loaded from the cached module image. This includes complicated heap-allocated objects like arrays. However, any routine that returns a raw pointer value must be called at runtime for precompilation to work (Ptr objects will turn into null pointers unless they are hidden inside an isbits object). This includes the return values of the Julia functions `cfunction` and `pointer`.

Dictionary and set types, or in general anything that depends on the output of a `hash(key)` method, are a trickier case. In the common case where the keys are numbers, strings, symbols, ranges, `Expr`, or compositions of these types (via arrays, tuples, sets, pairs, etc.) they are safe to precompile. However, for a few other key types, such as `Function` or `DataType` and generic user-defined types where you haven't defined a `hash` method, the fallback `hash` method depends on the memory address of the object (via its `object_id`) and hence may change from run to run. If you have one of these key types, or if you aren't sure, to be safe you can initialize this dictionary from within your `__init__` function. Alternatively, you can use the `ObjectIdDict` dictionary type, which is specially handled by precompilation so that it is safe to initialize at compile-time.

When using precompilation, it is important to keep a clear sense of the distinction between the compilation phase and the execution phase. In this mode, it will often be much more clearly apparent that Julia is a compiler which allows execution of arbitrary Julia code, not a standalone interpreter that also generates compiled code.

Other known potential failure scenarios include:

1. Global counters (for example, for attempting to uniquely identify objects) Consider the following code snippet:

```
mutable struct UniquedById
    myid::Int
    let counter = 0
        UniquedById() = new(counter += 1)
    end
end
```

   while the intent of this code was to give every instance a unique id, the counter value is recorded at the end of compilation. All subsequent usages of this incrementally compiled module will start from that same counter value.

   Note that `object_id` (which works by hashing the memory pointer) has similar issues (see notes on `Dict` usage below).

   One alternative is to store both `current_module()` and the current `counter` value, however, it may be better to redesign the code to not depend on this global state.

2. Associative collections (such as `Dict` and `Set`) need to be re-hashed in `__init__`. (In the future, a mechanism may be provided to register an initializer function.)

3. Depending on compile-time side-effects persisting through load-time. Example include: modifying arrays or other variables in other Julia modules; maintaining handles to open files or devices; storing pointers to other system resources (including memory);

4. Creating accidental "copies" of global state from another module, by referencing it directly instead of via its lookup path. For example, (in global scope):

```
#mystdout = Base.STDOUT #= will not work correctly, since this will copy Base.STDOUT into
↪    this module =#
# instead use accessor functions:
```

```julia
getstdout() = Base.STDOUT #= best option =#
# or move the assignment into the runtime:
__init__() = global mystdout = Base.STDOUT #= also works =#
```

Several additional restrictions are placed on the operations that can be done while precompiling code to help the user avoid other wrong-behavior situations:

1. Calling `eval` to cause a side-effect in another module. This will also cause a warning to be emitted when the incremental precompile flag is set.

2. `global const` statements from local scope after `__init__()` has been started (see issue #12010 for plans to add an error for this)

3. Replacing a module (or calling `workspace()`) is a runtime error while doing an incremental precompile.

A few other points to be aware of:

1. No code reload / cache invalidation is performed after changes are made to the source files themselves, (including by `Pkg.update`), and no cleanup is done after `Pkg.rm`

2. The memory sharing behavior of a reshaped array is disregarded by precompilation (each view gets its own copy)

3. Expecting the filesystem to be unchanged between compile-time and runtime e.g. `@__FILE__`/`source_path()` to find resources at runtime, or the BinDeps `@checked_lib` macro. Sometimes this is unavoidable. However, when possible, it can be good practice to copy resources into the module at compile-time so they won't need to be found at runtime.

4. `WeakRef` objects and finalizers are not currently handled properly by the serializer (this will be fixed in an upcoming release).

5. It is usually best to avoid capturing references to instances of internal metadata objects such as `Method`, `MethodInstance`, `MethodTable`, `TypeMapLevel`, `TypeMapEntry` and fields of those objects, as this can confuse the serializer and may not lead to the outcome you desire. It is not necessarily an error to do this, but you simply need to be prepared that the system will try to copy some of these and to create a single unique instance of others.

It is sometimes helpful during module development to turn off incremental precompilation. The command line flag `--compilecache={yes|no}` enables you to toggle module precompilation on and off. When Julia is started with `--compilecache=no` the serialized modules in the compile cache are ignored when loading modules and module dependencies. `Base.compilecache()` can still be called manually and it will respect `__precompile__()` directives for the module. The state of this command line flag is passed to `Pkg.build()` to disable automatic precompilation triggering when installing, updating, and explicitly building packages.

Chapter 20

Documentation

Julia enables package developers and users to document functions, types and other objects easily via a built-in documentation system since Julia 0.4.

The basic syntax is very simple: any string appearing at the top-level right before an object (function, macro, type or instance) will be interpreted as documenting it (these are called *docstrings*). Here is a very simple example:

```
"Tell whether there are too foo items in the array."
foo(xs::Array) = ...
```

Documentation is interpreted as Markdown, so you can use indentation and code fences to delimit code examples from text. Technically, any object can be associated with any other as metadata; Markdown happens to be the default, but one can construct other string macros and pass them to the @doc macro just as well.

Here is a more complex example, still using Markdown:

```
"""
    bar(x[, y])

Compute the Bar index between `x` and `y`. If `y` is missing, compute
the Bar index between all pairs of columns of `x`.

# Examples
```julia-repl
julia> bar([1, 2], [1, 2])
1
```
"""
function bar(x, y) ...
```

As in the example above, we recommend following some simple conventions when writing documentation:

1. Always show the signature of a function at the top of the documentation, with a four-space indent so that it is printed as Julia code.

   This can be identical to the signature present in the Julia code (like mean(x::AbstractArray)), or a simplified form. Optional arguments should be represented with their default values (i.e. f(x, y=1)) when possible, following the actual Julia syntax. Optional arguments which do not have a default value should be put in brackets (i.e. f(x[, y]) and f(x[, y[, z]])). An alternative solution is to use several lines: one without optional

arguments, the other(s) with them. This solution can also be used to document several related methods of a given function. When a function accepts many keyword arguments, only include a `<keyword arguments>` placeholder in the signature (i.e. `f(x; <keyword arguments>)`), and give the complete list under an `# Arguments` section (see point 4 below).

2. Include a single one-line sentence describing what the function does or what the object represents after the simplified signature block. If needed, provide more details in a second paragraph, after a blank line.

    The one-line sentence should use the imperative form ("Do this", "Return that") instead of the third person (do not write "Returns the length...") when documenting functions. It should end with a period. If the meaning of a function cannot be summarized easily, splitting it into separate composable parts could be beneficial (this should not be taken as an absolute requirement for every single case though).

3. Do not repeat yourself.

    Since the function name is given by the signature, there is no need to start the documentation with "The function `bar`...": go straight to the point. Similarly, if the signature specifies the types of the arguments, mentioning them in the description is redundant.

4. Only provide an argument list when really necessary.

    For simple functions, it is often clearer to mention the role of the arguments directly in the description of the function's purpose. An argument list would only repeat information already provided elsewhere. However, providing an argument list can be a good idea for complex functions with many arguments (in particular keyword arguments). In that case, insert it after the general description of the function, under an `# Arguments` header, with one – bullet for each argument. The list should mention the types and default values (if any) of the arguments:

    ```
    """
    ...
    # Arguments
    - `n::Integer`: the number of elements to compute.
    - `dim::Integer=1`: the dimensions along which to perform the computation.
    ...
    """
    ```

5. Include any code examples in an `# Examples` section.

    Examples should, whenever possible, be written as *doctests*. A *doctest* is a fenced code block (see Code blocks) starting with ```` ```jldoctest ```` and contains any number of `julia>` prompts together with inputs and expected outputs that mimic the Julia REPL.

    For example in the following docstring a variable a is defined and the expected result, as printed in a Julia REPL, appears afterwards:

    ```
    """
    Some nice documentation here.

    # Examples

    ```jldoctest
    julia> a = [1 2; 3 4]
    2×2 Array{Int64,2}:
     1  2
     3  4
    ```
    """
    ```

> **Warning**
>
> Calling `rand` and other RNG-related functions should be avoided in doctests since they will not produce consistent outputs during different Julia sessions.
>
> Operating system word size (`Int32` or `Int64`) as well as path separator differences (`/` or `\`) will also affect the reproducibility of some doctests.
>
> Note that whitespace in your doctest is significant! The doctest will fail if you misalign the output of pretty-printing an array, for example.

You can then run `make -C doc doctest` to run all the doctests in the Julia Manual, which will ensure that your example works.

Examples that are untestable should be written within fenced code blocks starting with ```` ```julia ```` so that they are highlighted correctly in the generated documentation.

> **Tip**
>
> Wherever possible examples should be **self-contained** and **runnable** so that readers are able to try them out without having to include any dependencies.

6. Use backticks to identify code and equations.

   Julia identifiers and code excerpts should always appear between backticks `` ` `` to enable highlighting. Equations in the LaTeX syntax can be inserted between double backticks ``` `` ```. Use Unicode characters rather than their LaTeX escape sequence, i.e. ``` ``α = 1`` ``` rather than ``` ``\alpha = 1`` ```.

7. Place the starting and ending `"""` characters on lines by themselves.

   That is, write:

   ```
   """
   ...

   ...
   """
   f(x, y) = ...
   ```

   rather than:

   ```
   """...

   ..."""
   f(x, y) = ...
   ```

   This makes it more clear where docstrings start and end.

8. Respect the line length limit used in the surrounding code.

   Docstrings are edited using the same tools as code. Therefore, the same conventions should apply. It it advised to add line breaks after 92 characters.

20.1 Accessing Documentation

Documentation can be accessed at the REPL or in IJulia by typing ? followed by the name of a function or macro, and pressing `Enter`. For example,

```
?fft
?@time
?r""
```

will bring up docs for the relevant function, macro or string macro respectively. In Juno using `Ctrl-J, Ctrl-D` will bring up documentation for the object under the cursor.

20.2   Functions & Methods

Functions in Julia may have multiple implementations, known as methods. While it's good practice for generic functions to have a single purpose, Julia allows methods to be documented individually if necessary. In general, only the most generic method should be documented, or even the function itself (i.e. the object created without any methods by `function bar end`). Specific methods should only be documented if their behaviour differs from the more generic ones. In any case, they should not repeat the information provided elsewhere. For example:

```
"""
    *(x, y, z...)

Multiplication operator. `x * y * z *...` calls this function with multiple
arguments, i.e. `*(x, y, z...)`.
"""
function *(x, y, z...)
    # ... [implementation sold separately] ...
end

"""
    *(x::AbstractString, y::AbstractString, z::AbstractString...)

When applied to strings, concatenates them.
"""
function *(x::AbstractString, y::AbstractString, z::AbstractString...)
    # ... [insert secret sauce here] ...
end

help?> *
search: * .*

  *(x, y, z...)

  Multiplication operator. x * y * z *... calls this function with multiple
  arguments, i.e. *(x,y,z...).

  *(x::AbstractString, y::AbstractString, z::AbstractString...)

  When applied to strings, concatenates them.
```

When retrieving documentation for a generic function, the metadata for each method is concatenated with the `catdoc` function, which can of course be overridden for custom types.

20.3   Advanced Usage

The `@doc` macro associates its first argument with its second in a per-module dictionary called `META`. By default, documentation is expected to be written in Markdown, and the doc`""` string macro simply creates an object representing the Markdown content. In the future it is likely to do more advanced things such as allowing for relative image or link paths.

When used for retrieving documentation, the `@doc` macro (or equally, the doc function) will search all `META` dictionaries for metadata relevant to the given object and return it. The returned object (some Markdown content, for example)

will by default display itself intelligently. This design also makes it easy to use the doc system in a programmatic way; for example, to re-use documentation between different versions of a function:

```
@doc "..." foo!
@doc (@doc foo!) foo
```

Or for use with Julia's metaprogramming functionality:

```
for (f, op) in ((:add, :+), (:subtract, :-), (:multiply, :*), (:divide, :/))
    @eval begin
        $f(a,b) = $op(a,b)
    end
end
@doc "`add(a,b)` adds `a` and `b` together" add
@doc "`subtract(a,b)` subtracts `b` from `a`" subtract
```

Documentation written in non-toplevel blocks, such as `begin`, `if`, `for`, and `let`, is added to the documentation system as blocks are evaluated. For example:

```
if VERSION > v"0.5"
    "..."
    f(x) = x
end
```

will add documentation to `f(x)` when the condition is `true`. Note that even if `f(x)` goes out of scope at the end of the block, its documentation will remain.

Dynamic documentation

Sometimes the appropriate documentation for an instance of a type depends on the field values of that instance, rather than just on the type itself. In these cases, you can add a method to `Docs.getdoc` for your custom type that returns the documentation on a per-instance basis. For instance,

```
struct MyType
    value::String
end

Docs.getdoc(t::MyType) = "Documentation for MyType with value $(t.value)"

x = MyType("x")
y = MyType("y")
```

?x will display "Documentation for MyType with value x" while ?y will display "Documentation for MyType with value y".

20.4   Syntax Guide

A comprehensive overview of all documentable Julia syntax.

In the following examples " . . . " is used to illustrate an arbitrary docstring which may be one of the follow four variants and contain arbitrary text:

```
"..."

doc"..."

"""
...
"""

doc"""
...
"""
```

`@doc_str` should only be used when the docstring contains `$` or `\` characters that should not be parsed by Julia such as LaTeX syntax or Julia source code examples containing interpolation.

Functions and Methods

```
"..."
function f end

"..."
f
```

Adds docstring `"..."` to `Function f`. The first version is the preferred syntax, however both are equivalent.

```
"..."
f(x) = x

"..."
function f(x)
    x
end

"..."
f(x)
```

Adds docstring `"..."` to `Method f(::Any)`.

```
"..."
f(x, y = 1) = x + y
```

Adds docstring `"..."` to two `Method`s, namely `f(::Any)` and `f(::Any, ::Any)`.

Macros

```
"..."
macro m(x) end
```

Adds docstring `"..."` to the `@m(::Any)` macro definition.

```
"..."
:(@m)
```

Adds docstring `"..."` to the macro named `@m`.

Types

```
"..."
abstract type T1 end

"..."
mutable struct T2
    ...
end

"..."
struct T3
    ...
end
```

Adds the docstring `"..."` to types T1, T2, and T3.

```
"..."
struct T
    "x"
    x
    "y"
    y
end
```

Adds docstring `"..."` to type T, `"x"` to field `T.x` and `"y"` to field `T.y`. Also applicable to `mutable struct` types.

Modules

```
"..."
module M end

module M

"..."
M

end
```

Adds docstring `"..."` to the `ModuleM`. Adding the docstring above the `Module` is the preferred syntax, however both are equivalent.

```
"..."
baremodule M
# ...
end

baremodule M

import Base: @doc

"..."
f(x) = x

end
```

Documenting a `baremodule` by placing a docstring above the expression automatically imports `@doc` into the module. These imports must be done manually when the module expression is not documented. Empty `baremodules` cannot be documented.

Global Variables

```
"..."
const a = 1

"..."
b = 2

"..."
global c = 3
```

Adds docstring `"..."` to the `Bindings` a, b, and c.

`Bindings` are used to store a reference to a particular `Symbol` in a `Module` without storing the referenced value itself.

> **Note**
>
> When a `const` definition is only used to define an alias of another definition, such as is the case with the function `div` and its alias ÷ in `Base`, do not document the alias and instead document the actual function.
>
> If the alias is documented and not the real definition then the docsystem (? mode) will not return the docstring attached to the alias when the real definition is searched for.
>
> For example you should write
>
> ```
> "..."
> f(x) = x + 1
> const alias = f
> ```
>
> rather than
>
> ```
> f(x) = x + 1
> "..."
> const alias = f
> ```

```
"..."
sym
```

Adds docstring `"..."` to the value associated with `sym`. Users should prefer documenting `sym` at it's definition.

Multiple Objects

```
"..."
a, b
```

Adds docstring `"..."` to a and b each of which should be a documentable expression. This syntax is equivalent to

```
"..."
a

"..."
b
```

Any number of expressions many be documented together in this way. This syntax can be useful when two functions are related, such as non-mutating and mutating versions `f` and `f!`.

Macro-generated code

```
"..."
@m expression
```

Adds docstring `"..."` to expression generated by expanding `@m expression`. This allows for expressions decorated with `@inline`, `@noinline`, `@generated`, or any other macro to be documented in the same way as undecorated expressions.

Macro authors should take note that only macros that generate a single expression will automatically support docstrings. If a macro returns a block containing multiple subexpressions then the subexpression that should be documented must be marked using the `@__doc__` macro.

The @enum macro makes use of @__doc__ to allow for documenting Enums. Examining it's definition should serve as an example of how to use `@__doc__` correctly.

`Core.@__doc__` – Macro.

```
@__doc__(ex)
```

Low-level macro used to mark expressions returned by a macro that should be documented. If more than one expression is marked then the same docstring is applied to each expression.

```
macro example(f)
    quote
        $(f)() = 0
        @__doc__ $(f)(x) = 1
        $(f)(x, y) = 2
    end |> esc
end
```

`@__doc__` has no effect when a macro that uses it is not documented.

source

20.5 Markdown syntax

The following markdown syntax is supported in Julia.

Inline elements

Here "inline" refers to elements that can be found within blocks of text, i.e. paragraphs. These include the following elements.

Bold

Surround words with two asterisks, `**`, to display the enclosed text in boldface.

```
A paragraph containing a **bold** word.
```

Italics

Surround words with one asterisk, `*`, to display the enclosed text in italics.

```
A paragraph containing an *emphasised* word.
```

Literals

Surround text that should be displayed exactly as written with single backticks, ` .

```
A paragraph containing a `literal` word.
```

Literals should be used when writing text that refers to names of variables, functions, or other parts of a Julia program.

> **Tip**
>
> To include a backtick character within literal text use three backticks rather than one to enclose the text.
>
> ```
> A paragraph containing a ``` `backtick` character ```.
> ```
>
> By extension any odd number of backticks may be used to enclose a lesser number of backticks.

LaTeX

Surround text that should be displayed as mathematics using LaTeXsyntax with double backticks, `` .

```
A paragraph containing some ``\LaTeX`` markup.
```

> **Tip**
>
> As with literals in the previous section, if literal backticks need to be written within double backticks use an even number greater than two. Note that if a single literal backtick needs to be included within LaTeXmarkup then two enclosing backticks is sufficient.

Links

Links to either external or internal addresses can be written using the following syntax, where the text enclosed in square brackets, [   ], is the name of the link and the text enclosed in parentheses, (   ), is the URL.

```
A paragraph containing a link to [Julia](http://www.julialang.org).
```

It's also possible to add cross-references to other documented functions/methods/variables within the Julia documentation itself. For example:

```
"""
    eigvals!(A,[irange,][vl,][vu]) -> values

Same as [`eigvals`](@ref), but saves space by overwriting the input `A`, instead of creating a
↪  copy.
"""
```

This will create a link in the generated docs to the `eigvals` documentation (which has more information about what this function actually does). It's good to include cross references to mutating/non-mutating versions of a function, or to highlight a difference between two similar-seeming functions.

> **Note**
>
> The above cross referencing is *not* a Markdown feature, and relies on Documenter.jl, which is used to build base Julia's documentation.

Footnote references

Named and numbered footnote references can be written using the following syntax. A footnote name must be a single alphanumeric word containing no punctuation.

```
A paragraph containing a numbered footnote [^1] and a named one [^named].
```

> **Note**
>
> The text associated with a footnote can be written anywhere within the same page as the footnote reference. The syntax used to define the footnote text is discussed in the Footnotes section below.

Toplevel elements

The following elements can be written either at the "toplevel" of a document or within another "toplevel" element.

Paragraphs

A paragraph is a block of plain text, possibly containing any number of inline elements defined in the Inline elements section above, with one or more blank lines above and below it.

```
This is a paragraph.

And this is *another* one containing some emphasised text.
A new line, but still part of the same paragraph.
```

Headers

A document can be split up into different sections using headers. Headers use the following syntax:

```
# Level One
## Level Two
### Level Three
#### Level Four
##### Level Five
###### Level Six
```

A header line can contain any inline syntax in the same way as a paragraph can.

> **Tip**
>
> Try to avoid using too many levels of header within a single document. A heavily nested document may be indicative of a need to restructure it or split it into several pages covering separate topics.

Code blocks

Source code can be displayed as a literal block using an indent of four spaces as shown in the following example.

```
This is a paragraph.

    function func(x)
        # ...
    end

Another paragraph.
```

Additionally, code blocks can be enclosed using triple backticks with an optional "language" to specify how a block of code should be highlighted.

```
A code block without a "language":

```
function func(x)
    # ...
end
```

and another one with the "language" specified as `julia`:

```julia
function func(x)
    # ...
end
```
```

> **Note**
>
> "Fenced" code blocks, as shown in the last example, should be prefered over indented code blocks since there is no way to specify what language an indented code block is written in.

Block quotes

Text from external sources, such as quotations from books or websites, can be quoted using > characters prepended to each line of the quote as follows.

```
Here's a quote:

> Julia is a high-level, high-performance dynamic programming language for
> technical computing, with syntax that is familiar to users of other
> technical computing environments.
```

Note that a single space must appear after the > character on each line. Quoted blocks may themselves contain other toplevel or inline elements.

Images

The syntax for images is similar to the link syntax mentioned above. Prepending a ! character to a link will display an image from the specified URL rather than a link to it.

```
![alternative text](link/to/image.png)
```

Lists

Unordered lists can be written by prepending each item in a list with either *, +, or -.

```
A list of items:

  * item one
  * item two
  * item three
```

Note the two spaces before each * and the single space after each one.

Lists can contain other nested toplevel elements such as lists, code blocks, or quoteblocks. A blank line should be left between each list item when including any toplevel elements within a list.

```
Another list:

  * item one

  * item two

    ```
    f(x) = x
    ```

  * And a sublist:

      + sub-item one
      + sub-item two
```

> **Note**
>
> The contents of each item in the list must line up with the first line of the item. In the above example the fenced code block must be indented by four spaces to align with the `i` in `item two`.

Ordered lists are written by replacing the "bullet" character, either *, +, or -, with a positive integer followed by either . or).

```
Two ordered lists:

 1. item one
 2. item two
 3. item three

 5) item five
 6) item six
 7) item seven
```

An ordered list may start from a number other than one, as in the second list of the above example, where it is numbered from five. As with unordered lists, ordered lists can contain nested toplevel elements.

Display equations

Large LaTeX equations that do not fit inline within a paragraph may be written as display equations using a fenced code block with the "language" `math` as in the example below.

```math
f(a) = \frac{1}{2\pi}\int_{0}^{2\pi} (\alpha+R\cos(\theta))d\theta
```

Footnotes

This syntax is paired with the inline syntax for Footnote references. Make sure to read that section as well.

Footnote text is defined using the following syntax, which is similar to footnote reference syntax, aside from the `:` character that is appended to the footnote label.

```
[^1]: Numbered footnote text.

[^note]:

    Named footnote text containing several toplevel elements.

      * item one
      * item two
      * item three

    ```julia
    function func(x)
        # ...
    end
    ```
```

> **Note**
>
> No checks are done during parsing to make sure that all footnote references have matching footnotes.

Horizontal rules

The equivalent of an `<hr>` HTML tag can be written using the following syntax:

```
Text above the line.

---

And text below the line.
```

Tables

Basic tables can be written using the syntax described below. Note that markdown tables have limited features and cannot contain nested toplevel elements unlike other elements discussed above – only inline elements are allowed. Tables must always contain a header row with column names. Cells cannot span multiple rows or columns of the table.

```
| Column One | Column Two | Column Three |
|:---------- | ---------- |:------------:|
| Row `1`    | Column `2` |              |
| *Row* 2    | **Row** 2  | Column ``3`` |
```

> **Note**
>
> As illustrated in the above example each column of `|` characters must be aligned vertically.
>
> A `:` character on either end of a column's header separator (the row containing – characters) specifies whether the row is left-aligned, right-aligned, or (when `:` appears on both ends) center-aligned. Providing no `:` characters will default to right-aligning the column.

Admonitions

Specially formatted blocks with titles such as "Notes", "Warning", or "Tips" are known as admonitions and are used when some part of a document needs special attention. They can be defined using the following `!!!` syntax:

```
!!! note

    This is the content of the note.

!!! warning "Beware!"

    And this is another one.

    This warning admonition has a custom title: `"Beware!"`.
```

Admonitions, like most other toplevel elements, can contain other toplevel elements. When no title text, specified after the admonition type in double quotes, is included then the title used will be the type of the block, i.e. `"Note"` in the case of the `note` admonition.

20.6 Markdown Syntax Extensions

Julia's markdown supports interpolation in a very similar way to basic string literals, with the difference that it will store the object itself in the Markdown tree (as opposed to converting it to a string). When the Markdown content is rendered the usual show methods will be called, and these can be overridden as usual. This design allows the Markdown to be extended with arbitrarily complex features (such as references) without cluttering the basic syntax.

In principle, the Markdown parser itself can also be arbitrarily extended by packages, or an entirely custom flavour of Markdown can be used, but this should generally be unnecessary.

Chapter 21

Metaprogramming

The strongest legacy of Lisp in the Julia language is its metaprogramming support. Like Lisp, Julia represents its own code as a data structure of the language itself. Since code is represented by objects that can be created and manipulated from within the language, it is possible for a program to transform and generate its own code. This allows sophisticated code generation without extra build steps, and also allows true Lisp-style macros operating at the level of abstract syntax trees. In contrast, preprocessor "macro" systems, like that of C and C++, perform textual manipulation and substitution before any actual parsing or interpretation occurs. Because all data types and code in Julia are represented by Julia data structures, powerful reflection capabilities are available to explore the internals of a program and its types just like any other data.

21.1 Program representation

Every Julia program starts life as a string:

```
julia> prog = "1 + 1"
"1 + 1"
```

What happens next?

The next step is to parse each string into an object called an expression, represented by the Julia type Expr:

```
julia> ex1 = parse(prog)
:(1 + 1)

julia> typeof(ex1)
Expr
```

Expr objects contain three parts:

- a Symbol identifying the kind of expression. A symbol is an interned string identifier (more discussion below).

```
julia> ex1.head
:call
```

- the expression arguments, which may be symbols, other expressions, or literal values:

```
julia> ex1.args
3-element Array{Any,1}:
   :+
  1
  1
```

- finally, the expression result type, which may be annotated by the user or inferred by the compiler (and may be ignored completely for the purposes of this chapter):

```
julia> ex1.typ
Any
```

Expressions may also be constructed directly in prefix notation:

```
julia> ex2 = Expr(:call, :+, 1, 1)
:(1 + 1)
```

The two expressions constructed above – by parsing and by direct construction – are equivalent:

```
julia> ex1 == ex2
true
```

The key point here is that Julia code is internally represented as a data structure that is accessible from the language itself.

The dump() function provides indented and annotated display of Expr objects:

```
julia> dump(ex2)
Expr
  head: Symbol call
  args: Array{Any}((3,))
    1: Symbol +
    2: Int64 1
    3: Int64 1
  typ: Any
```

Expr objects may also be nested:

```
julia> ex3 = parse("(4 + 4) / 2")
:((4 + 4) / 2)
```

Another way to view expressions is with Meta.show_sexpr, which displays the S-expression form of a given Expr, which may look very familiar to users of Lisp. Here's an example illustrating the display on a nested Expr:

```
julia> Meta.show_sexpr(ex3)
(:call, :/, (:call, :+, 4, 4), 2)
```

Symbols

The : character has two syntactic purposes in Julia. The first form creates a Symbol, an interned string used as one building-block of expressions:

```
julia> :foo
:foo

julia> typeof(ans)
Symbol
```

The Symbol constructor takes any number of arguments and creates a new symbol by concatenating their string representations together:

```julia
julia> :foo == Symbol("foo")
true

julia> Symbol("func",10)
:func10

julia> Symbol(:var,'_',"sym")
:var_sym
```

In the context of an expression, symbols are used to indicate access to variables; when an expression is evaluated, a symbol is replaced with the value bound to that symbol in the appropriate scope.

Sometimes extra parentheses around the argument to : are needed to avoid ambiguity in parsing.:

```julia
julia> :(:)
:(:)

julia> :(::)
:(::)
```

21.2 Expressions and evaluation

Quoting

The second syntactic purpose of the : character is to create expression objects without using the explicit Expr constructor. This is referred to as *quoting*. The : character, followed by paired parentheses around a single statement of Julia code, produces an Expr object based on the enclosed code. Here is example of the short form used to quote an arithmetic expression:

```julia
julia> ex = :(a+b*c+1)
:(a + b * c + 1)

julia> typeof(ex)
Expr
```

(to view the structure of this expression, try ex.head and ex.args, or use dump() as above)

Note that equivalent expressions may be constructed using parse() or the direct Expr form:

```julia
julia>       :(a + b*c + 1)   ==
        parse("a + b*c + 1") ==
        Expr(:call, :+, :a, Expr(:call, :*, :b, :c), 1)
true
```

Expressions provided by the parser generally only have symbols, other expressions, and literal values as their args, whereas expressions constructed by Julia code can have arbitrary run-time values without literal forms as args. In this specific example, + and a are symbols, *(b,c) is a subexpression, and 1 is a literal 64-bit signed integer.

There is a second syntactic form of quoting for multiple expressions: blocks of code enclosed in quote ... end. Note that this form introduces QuoteNode elements to the expression tree, which must be considered when directly manipulating an expression tree generated from quote blocks. For other purposes, :(...) and quote .. end blocks are treated identically.

```julia
julia> ex = quote
           x = 1
           y = 2
           x + y
       end
quote  # none, line 2:
    x = 1 # none, line 3:
    y = 2 # none, line 4:
    x + y
end

julia> typeof(ex)
Expr
```

Interpolation

Direct construction of `Expr` objects with value arguments is powerful, but `Expr` constructors can be tedious compared to "normal" Julia syntax. As an alternative, Julia allows "splicing" or interpolation of literals or expressions into quoted expressions. Interpolation is indicated by the $ prefix.

In this example, the literal value of a is interpolated:

```julia
julia> a = 1;

julia> ex = :($a + b)
:(1 + b)
```

Interpolating into an unquoted expression is not supported and will cause a compile-time error:

```julia
julia> $a + b
ERROR: unsupported or misplaced expression $
 ...
```

In this example, the tuple `(1,2,3)` is interpolated as an expression into a conditional test:

```julia
julia> ex = :(a in $:((1,2,3)) )
:(a in (1, 2, 3))
```

Interpolating symbols into a nested expression requires enclosing each symbol in an enclosing quote block:

```julia
julia> :( :a in $( :(:a + :b) ) )
                  ^^^^^^^^^^
                 quoted inner expression
```

The use of $ for expression interpolation is intentionally reminiscent of string interpolation and command interpolation. Expression interpolation allows convenient, readable programmatic construction of complex Julia expressions.

`eval()` and effects

Given an expression object, one can cause Julia to evaluate (execute) it at global scope using `eval()`:

```julia
julia> :(1 + 2)
:(1 + 2)

julia> eval(ans)
3
```

```
julia> ex = :(a + b)
:(a + b)

julia> eval(ex)
ERROR: UndefVarError: b not defined
[...]

julia> a = 1; b = 2;

julia> eval(ex)
3
```

Every module has its own eval() function that evaluates expressions in its global scope. Expressions passed to eval() are not limited to returning values – they can also have side-effects that alter the state of the enclosing module's environment:

```
julia> ex = :(x = 1)
:(x = 1)

julia> x
ERROR: UndefVarError: x not defined

julia> eval(ex)
1

julia> x
1
```

Here, the evaluation of an expression object causes a value to be assigned to the global variable x.

Since expressions are just Expr objects which can be constructed programmatically and then evaluated, it is possible to dynamically generate arbitrary code which can then be run using eval(). Here is a simple example:

```
julia> a = 1;

julia> ex = Expr(:call, :+, a, :b)
:(1 + b)

julia> a = 0; b = 2;

julia> eval(ex)
3
```

The value of a is used to construct the expression ex which applies the + function to the value 1 and the variable b. Note the important distinction between the way a and b are used:

- The value of the *variable* a at expression construction time is used as an immediate value in the expression. Thus, the value of a when the expression is evaluated no longer matters: the value in the expression is already 1, independent of whatever the value of a might be.

- On the other hand, the *symbol* :b is used in the expression construction, so the value of the variable b at that time is irrelevant – :b is just a symbol and the variable b need not even be defined. At expression evaluation time, however, the value of the symbol :b is resolved by looking up the value of the variable b.

Functions on Expressions

As hinted above, one extremely useful feature of Julia is the capability to generate and manipulate Julia code within Julia itself. We have already seen one example of a function returning `Expr` objects: the `parse()` function, which takes a string of Julia code and returns the corresponding `Expr`. A function can also take one or more `Expr` objects as arguments, and return another `Expr`. Here is a simple, motivating example:

```julia
julia> function math_expr(op, op1, op2)
           expr = Expr(:call, op, op1, op2)
           return expr
       end
math_expr (generic function with 1 method)

julia>  ex = math_expr(:+, 1, Expr(:call, :*, 4, 5))
:(1 + 4 * 5)

julia> eval(ex)
21
```

As another example, here is a function that doubles any numeric argument, but leaves expressions alone:

```julia
julia> function make_expr2(op, opr1, opr2)
           opr1f, opr2f = map(x -> isa(x, Number) ? 2*x : x, (opr1, opr2))
           retexpr = Expr(:call, op, opr1f, opr2f)
           return retexpr
       end
make_expr2 (generic function with 1 method)

julia> make_expr2(:+, 1, 2)
:(2 + 4)

julia> ex = make_expr2(:+, 1, Expr(:call, :*, 5, 8))
:(2 + 5 * 8)

julia> eval(ex)
42
```

21.3   Macros

Macros provide a method to include generated code in the final body of a program. A macro maps a tuple of arguments to a returned *expression*, and the resulting expression is compiled directly rather than requiring a runtime `eval()` call. Macro arguments may include expressions, literal values, and symbols.

Basics

Here is an extraordinarily simple macro:

```julia
julia> macro sayhello()
           return :( println("Hello, world!") )
       end
@sayhello (macro with 1 method)
```

Macros have a dedicated character in Julia's syntax: the @ (at-sign), followed by the unique name declared in a `macro NAME ... end` block. In this example, the compiler will replace all instances of `@sayhello` with:

```
:( println("Hello, world!") )
```

When @sayhello is entered in the REPL, the expression executes immediately, thus we only see the evaluation result:

```
julia> @sayhello()
Hello, world!
```

Now, consider a slightly more complex macro:

```
julia> macro sayhello(name)
           return :( println("Hello, ", $name) )
       end
@sayhello (macro with 1 method)
```

This macro takes one argument: name. When @sayhello is encountered, the quoted expression is *expanded* to interpolate the value of the argument into the final expression:

```
julia> @sayhello("human")
Hello, human
```

We can view the quoted return expression using the function `macroexpand()` (**important note:** this is an extremely useful tool for debugging macros):

```
julia> ex = macroexpand( :(@sayhello("human")) )
:((println)("Hello, ", "human"))

julia> typeof(ex)
Expr
```

We can see that the "human" literal has been interpolated into the expression.

There also exists a macro `@macroexpand` that is perhaps a bit more convenient than the `macroexpand` function:

```
julia> @macroexpand @sayhello "human"
:((println)("Hello, ", "human"))
```

Hold up: why macros?

We have already seen a function `f(::Expr...) -> Expr` in a previous section. In fact, `macroexpand()` is also such a function. So, why do macros exist?

Macros are necessary because they execute when code is parsed, therefore, macros allow the programmer to generate and include fragments of customized code *before* the full program is run. To illustrate the difference, consider the following example:

```
julia> macro twostep(arg)
           println("I execute at parse time. The argument is: ", arg)
           return :(println("I execute at runtime. The argument is: ", $arg))
       end
@twostep (macro with 1 method)

julia> ex = macroexpand( :(@twostep :(1, 2, 3)) );
I execute at parse time. The argument is: $(Expr(:quote, :((1, 2, 3))))
```

The first call to `println()` is executed when `macroexpand()` is called. The resulting expression contains *only* the second `println`:

```
julia> typeof(ex)
Expr

julia> ex
:((println)("I execute at runtime. The argument is: ", $(Expr(:copyast, :($(QuoteNode(:((1, 2, 3)
    )))))))))

julia> eval(ex)
I execute at runtime. The argument is: (1, 2, 3)
```

Macro invocation

Macros are invoked with the following general syntax:

```
@name expr1 expr2 ...
@name(expr1, expr2, ...)
```

Note the distinguishing @ before the macro name and the lack of commas between the argument expressions in the
first form, and the lack of whitespace after @name in the second form. The two styles should not be mixed. For
example, the following syntax is different from the examples above; it passes the tuple (expr1, expr2, ...) as
one argument to the macro:

```
@name (expr1, expr2, ...)
```

It is important to emphasize that macros receive their arguments as expressions, literals, or symbols. One way to
explore macro arguments is to call the show() function within the macro body:

```
julia> macro showarg(x)
           show(x)
           # ... remainder of macro, returning an expression
       end
@showarg (macro with 1 method)

julia> @showarg(a)
:a

julia> @showarg(1+1)
:(1 + 1)

julia> @showarg(println("Yo!"))
:(println("Yo!"))
```

Building an advanced macro

Here is a simplified definition of Julia's @assert macro:

```
julia> macro assert(ex)
           return :( $ex ? nothing : throw(AssertionError($(string(ex)))) )
       end
@assert (macro with 1 method)
```

This macro can be used like this:

```
julia> @assert 1 == 1.0

julia> @assert 1 == 0
ERROR: AssertionError: 1 == 0
```

In place of the written syntax, the macro call is expanded at parse time to its returned result. This is equivalent to writing:

```
1 == 1.0 ? nothing : throw(AssertionError("1 == 1.0"))
1 == 0 ? nothing : throw(AssertionError("1 == 0"))
```

That is, in the first call, the expression `:(1 == 1.0)` is spliced into the test condition slot, while the value of `string(:(1 == 1.0))` is spliced into the assertion message slot. The entire expression, thus constructed, is placed into the syntax tree where the `@assert` macro call occurs. Then at execution time, if the test expression evaluates to true, then `nothing` is returned, whereas if the test is false, an error is raised indicating the asserted expression that was false. Notice that it would not be possible to write this as a function, since only the *value* of the condition is available and it would be impossible to display the expression that computed it in the error message.

The actual definition of `@assert` in the standard library is more complicated. It allows the user to optionally specify their own error message, instead of just printing the failed expression. Just like in functions with a variable number of arguments, this is specified with an ellipses following the last argument:

```
julia> macro assert(ex, msgs...)
           msg_body = isempty(msgs) ? ex : msgs[1]
           msg = string(msg_body)
           return :($ex ? nothing : throw(AssertionError($msg)))
       end
@assert (macro with 1 method)
```

Now `@assert` has two modes of operation, depending upon the number of arguments it receives! If there's only one argument, the tuple of expressions captured by `msgs` will be empty and it will behave the same as the simpler definition above. But now if the user specifies a second argument, it is printed in the message body instead of the failing expression. You can inspect the result of a macro expansion with the aptly named `macroexpand()` function:

```
julia> macroexpand(:(@assert a == b))
:(if a == b
        nothing
    else
        (throw)((AssertionError)("a == b"))
    end)

julia> macroexpand(:(@assert a==b "a should equal b!"))
:(if a == b
        nothing
    else
        (throw)((AssertionError)("a should equal b!"))
    end)
```

There is yet another case that the actual `@assert` macro handles: what if, in addition to printing "a should equal b," we wanted to print their values? One might naively try to use string interpolation in the custom message, e.g., `@assert a==b "a ($a) should equal b ($b)!"`, but this won't work as expected with the above macro. Can you see why? Recall from string interpolation that an interpolated string is rewritten to a call to `string()`. Compare:

```
julia> typeof(:("a should equal b"))
String
```

```
julia> typeof(:("a ($a) should equal b ($b)!"))
Expr

julia> dump(:("a ($a) should equal b ($b)!"))
Expr
  head: Symbol string
  args: Array{Any}((5,))
    1: String "a ("
    2: Symbol a
    3: String ") should equal b ("
    4: Symbol b
    5: String ")!"
  typ: Any
```

So now instead of getting a plain string in `msg_body`, the macro is receiving a full expression that will need to be evaluated in order to display as expected. This can be spliced directly into the returned expression as an argument to the `string()` call; see `error.jl` for the complete implementation.

The `@assert` macro makes great use of splicing into quoted expressions to simplify the manipulation of expressions inside the macro body.

Hygiene

An issue that arises in more complex macros is that of hygiene. In short, macros must ensure that the variables they introduce in their returned expressions do not accidentally clash with existing variables in the surrounding code they expand into. Conversely, the expressions that are passed into a macro as arguments are often *expected* to evaluate in the context of the surrounding code, interacting with and modifying the existing variables. Another concern arises from the fact that a macro may be called in a different module from where it was defined. In this case we need to ensure that all global variables are resolved to the correct module. Julia already has a major advantage over languages with textual macro expansion (like C) in that it only needs to consider the returned expression. All the other variables (such as `msg` in `@assert` above) follow the normal scoping block behavior.

To demonstrate these issues, let us consider writing a `@time` macro that takes an expression as its argument, records the time, evaluates the expression, records the time again, prints the difference between the before and after times, and then has the value of the expression as its final value. The macro might look like this:

```
macro time(ex)
    return quote
        local t0 = time()
        local val = $ex
        local t1 = time()
        println("elapsed time: ", t1-t0, " seconds")
        val
    end
end
```

Here, we want `t0`, `t1`, and `val` to be private temporary variables, and we want `time` to refer to the `time()` function in the standard library, not to any `time` variable the user might have (the same applies to `println`). Imagine the problems that could occur if the user expression `ex` also contained assignments to a variable called `t0`, or defined its own `time` variable. We might get errors, or mysteriously incorrect behavior.

Julia's macro expander solves these problems in the following way. First, variables within a macro result are classified as either local or global. A variable is considered local if it is assigned to (and not declared global), declared local, or

used as a function argument name. Otherwise, it is considered global. Local variables are then renamed to be unique (using the `gensym()` function, which generates new symbols), and global variables are resolved within the macro definition environment. Therefore both of the above concerns are handled; the macro's locals will not conflict with any user variables, and `time` and `println` will refer to the standard library definitions.

One problem remains however. Consider the following use of this macro:

```
module MyModule
import Base.@time

time() = ... # compute something

@time time()
end
```

Here the user expression `ex` is a call to `time`, but not the same `time` function that the macro uses. It clearly refers to `MyModule.time`. Therefore we must arrange for the code in `ex` to be resolved in the macro call environment. This is done by "escaping" the expression with `esc()`:

```
macro time(ex)
    ...
    local val = $(esc(ex))
    ...
end
```

An expression wrapped in this manner is left alone by the macro expander and simply pasted into the output verbatim. Therefore it will be resolved in the macro call environment.

This escaping mechanism can be used to "violate" hygiene when necessary, in order to introduce or manipulate user variables. For example, the following macro sets `x` to zero in the call environment:

```
julia> macro zerox()
           return esc(:(x = 0))
       end
@zerox (macro with 1 method)

julia> function foo()
           x = 1
           @zerox
           return x # is zero
       end
foo (generic function with 1 method)

julia> foo()
0
```

This kind of manipulation of variables should be used judiciously, but is occasionally quite handy.

21.4 Code Generation

When a significant amount of repetitive boilerplate code is required, it is common to generate it programmatically to avoid redundancy. In most languages, this requires an extra build step, and a separate program to generate the repetitive code. In Julia, expression interpolation and `eval()` allow such code generation to take place in the normal

course of program execution. For example, the following code defines a series of operators on three arguments in terms of their 2-argument forms:

```
for op = (:+, :*, :&, :|, :$)
    eval(quote
        ($op)(a,b,c) = ($op)(($op)(a,b),c)
    end)
end
```

In this manner, Julia acts as its own preprocessor, and allows code generation from inside the language. The above code could be written slightly more tersely using the `:` prefix quoting form:

```
for op = (:+, :*, :&, :|, :$)
    eval(:(($op)(a,b,c) = ($op)(($op)(a,b),c)))
end
```

This sort of in-language code generation, however, using the `eval(quote(...))` pattern, is common enough that Julia comes with a macro to abbreviate this pattern:

```
for op = (:+, :*, :&, :|, :$)
    @eval ($op)(a,b,c) = ($op)(($op)(a,b),c)
end
```

The `@eval` macro rewrites this call to be precisely equivalent to the above longer versions. For longer blocks of generated code, the expression argument given to `@eval` can be a block:

```
@eval begin
    # multiple lines
end
```

21.5   Non-Standard String Literals

Recall from Strings that string literals prefixed by an identifier are called non-standard string literals, and can have different semantics than un-prefixed string literals. For example:

- `r"^\s*(?:#|$)"` produces a regular expression object rather than a string

- `b"DATA\xff\u2200"` is a byte array literal for $[68, 65, 84, 65, 255, 226, 136, 128]$.

Perhaps surprisingly, these behaviors are not hard-coded into the Julia parser or compiler. Instead, they are custom behaviors provided by a general mechanism that anyone can use: prefixed string literals are parsed as calls to specially-named macros. For example, the regular expression macro is just the following:

```
macro r_str(p)
    Regex(p)
end
```

That's all. This macro says that the literal contents of the string literal `r"^\s*(?:#|$)"` should be passed to the `@r_str` macro and the result of that expansion should be placed in the syntax tree where the string literal occurs. In other words, the expression `r"^\s*(?:#|$)"` is equivalent to placing the following object directly into the syntax tree:

```
Regex("^\\s*(?:#|\$)")
```

Not only is the string literal form shorter and far more convenient, but it is also more efficient: since the regular expression is compiled and the Regex object is actually created *when the code is compiled*, the compilation occurs only once, rather than every time the code is executed. Consider if the regular expression occurs in a loop:

```
for line = lines
    m = match(r"^\s*(?:#|$)", line)
    if m === nothing
        # non-comment
    else
        # comment
    end
end
```

Since the regular expression r"^\s*(?:#|$)" is compiled and inserted into the syntax tree when this code is parsed, the expression is only compiled once instead of each time the loop is executed. In order to accomplish this without macros, one would have to write this loop like this:

```
re = Regex("^\\s*(?:#|\$)")
for line = lines
    m = match(re, line)
    if m === nothing
        # non-comment
    else
        # comment
    end
end
```

Moreover, if the compiler could not determine that the regex object was constant over all loops, certain optimizations might not be possible, making this version still less efficient than the more convenient literal form above. Of course, there are still situations where the non-literal form is more convenient: if one needs to interpolate a variable into the regular expression, one must take this more verbose approach; in cases where the regular expression pattern itself is dynamic, potentially changing upon each loop iteration, a new regular expression object must be constructed on each iteration. In the vast majority of use cases, however, regular expressions are not constructed based on run-time data. In this majority of cases, the ability to write regular expressions as compile-time values is invaluable.

Like non-standard string literals, non-standard command literals exist using a prefixed variant of the command literal syntax. The command literal custom`literal` is parsed as @custom_cmd "literal". Julia itself does not contain any non-standard command literals, but packages can make use of this syntax. Aside from the different syntax and the _cmd suffix instead of the _str suffix, non-standard command literals behave exactly like non-standard string literals.

In the event that two modules provide non-standard string or command literals with the same name, it is possible to qualify the string or command literal with a module name. For instance, if both Foo and Bar provide non-standard string literal @x_str, then one can write Foo.x"literal" or Bar.x"literal" to disambiguate between the two.

The mechanism for user-defined string literals is deeply, profoundly powerful. Not only are Julia's non-standard literals implemented using it, but also the command literal syntax (`echo "Hello, $person"`) is implemented with the following innocuous-looking macro:

```
macro cmd(str)
    :(cmd_gen($(shell_parse(str)[1])))
end
```

Of course, a large amount of complexity is hidden in the functions used in this macro definition, but they are just functions, written entirely in Julia. You can read their source and see precisely what they do – and all they do is construct expression objects to be inserted into your program's syntax tree.

21.6   Generated functions

A very special macro is `@generated`, which allows you to define so-called *generated functions*. These have the capability to generate specialized code depending on the types of their arguments with more flexibility and/or less code than what can be achieved with multiple dispatch. While macros work with expressions at parsing-time and cannot access the types of their inputs, a generated function gets expanded at a time when the types of the arguments are known, but the function is not yet compiled.

Instead of performing some calculation or action, a generated function declaration returns a quoted expression which then forms the body for the method corresponding to the types of the arguments. When called, the body expression is first evaluated and compiled, then the returned expression is compiled and run. To make this efficient, the result is often cached. And to make this inferable, only a limited subset of the language is usable. Thus, generated functions provide a flexible framework to move work from run-time to compile-time, at the expense of greater restrictions on the allowable constructs.

When defining generated functions, there are four main differences to ordinary functions:

1. You annotate the function declaration with the `@generated` macro. This adds some information to the AST that lets the compiler know that this is a generated function.

2. In the body of the generated function you only have access to the *types* of the arguments – not their values – and any function that was defined *before* the definition of the generated function.

3. Instead of calculating something or performing some action, you return a *quoted expression* which, when evaluated, does what you want.

4. Generated functions must not *mutate* or *observe* any non-constant global state (including, for example, IO, locks, non-local dictionaries, or using `method_exists`). This means they can only read global constants, and cannot have any side effects. In other words, they must be completely pure. Due to an implementation limitation, this also means that they currently cannot define a closure or untyped generator.

It's easiest to illustrate this with an example. We can declare a generated function `foo` as

```
julia> @generated function foo(x)
           Core.println(x)
           return :(x * x)
       end
foo (generic function with 1 method)
```

Note that the body returns a quoted expression, namely `:(x * x)`, rather than just the value of `x * x`.

From the caller's perspective, they are very similar to regular functions; in fact, you don't have to know if you're calling a regular or generated function - the syntax and result of the call is just the same. Let's see how `foo` behaves:

```
julia> x = foo(2); # note: output is from println() statement in the body
Int64

julia> x           # now we print x
4

julia> y = foo("bar");
```

```
String

julia> y
"barbar"
```

So, we see that in the body of the generated function, x is the *type* of the passed argument, and the value returned by the generated function, is the result of evaluating the quoted expression we returned from the definition, now with the *value* of x.

What happens if we evaluate foo again with a type that we have already used?

```
julia> foo(4)
16
```

Note that there is no printout of Int64. We can see that the body of the generated function was only executed once here, for the specific set of argument types, and the result was cached. After that, for this example, the expression returned from the generated function on the first invocation was re-used as the method body. However, the actual caching behavior is an implementation-defined performance optimization, so it is invalid to depend too closely on this behavior.

The number of times a generated function is generated *might* be only once, but it *might* also be more often, or appear to not happen at all. As a consequence, you should *never* write a generated function with side effects - when, and how often, the side effects occur is undefined. (This is true for macros too - and just like for macros, the use of eval() in a generated function is a sign that you're doing something the wrong way.) However, unlike macros, the runtime system cannot correctly handle a call to eval(), so it is disallowed.

It is also important to see how @generated functions interact with method redefinition. Following the principle that a correct @generated function must not observe any mutable state or cause any mutation of global state, we see the following behavior. Observe that the generated function *cannot* call any method that was not defined prior to the *definition* of the generated function itself.

Initially f(x) has one definition

```
julia> f(x) = "original definition";
```

Define other operations that use f(x):

```
julia> g(x) = f(x);

julia> @generated gen1(x) = f(x);

julia> @generated gen2(x) = :(f(x));
```

We now add some new definitions for f(x):

```
julia> f(x::Int) = "definition for Int";

julia> f(x::Type{Int}) = "definition for Type{Int}";
```

and compare how these results differ:

```
julia> f(1)
"definition for Int"

julia> g(1)
"definition for Int"

julia> gen1(1)
```

```
"original definition"

julia> gen2(1)
"definition for Int"
```

Each method of a generated function has its own view of defined functions:

```
julia> @generated gen1(x::Real) = f(x);

julia> gen1(1)
"definition for Type{Int}"
```

The example generated function `foo` above did not do anything a normal function `foo(x) = x * x` could not do (except printing the type on the first invocation, and incurring higher overhead). However, the power of a generated function lies in its ability to compute different quoted expressions depending on the types passed to it:

```
julia> @generated function bar(x)
           if x <: Integer
               return :(x ^ 2)
           else
               return :(x)
           end
       end
bar (generic function with 1 method)

julia> bar(4)
16

julia> bar("baz")
"baz"
```

(although of course this contrived example would be more easily implemented using multiple dispatch...)

Abusing this will corrupt the runtime system and cause undefined behavior:

```
julia> @generated function baz(x)
           if rand() < .9
               return :(x^2)
           else
               return :("boo!")
           end
       end
baz (generic function with 1 method)
```

Since the body of the generated function is non-deterministic, its behavior, *and the behavior of all subsequent code* is undefined.

Don't copy these examples!

These examples are hopefully helpful to illustrate how generated functions work, both in the definition end and at the call site; however, *don't copy them*, for the following reasons:

- the `foo` function has side-effects (the call to `Core.println`), and it is undefined exactly when, how often or how many times these side-effects will occur

- the `bar` function solves a problem that is better solved with multiple dispatch - defining `bar(x) = x` and `bar(x::Integer) = x ^ 2` will do the same thing, but it is both simpler and faster.

- the `baz` function is pathologically insane

Note that the set of operations that should not be attempted in a generated function is unbounded, and the runtime system can currently only detect a subset of the invalid operations. There are many other operations that will simply corrupt the runtime system without notification, usually in subtle ways not obviously connected to the bad definition. Because the function generator is run during inference, it must respect all of the limitations of that code.

Some operations that should not be attempted include:

1. Caching of native pointers.

2. Interacting with the contents or methods of Core.Inference in any way.

3. Observing any mutable state.

    – Inference on the generated function may be run at *any* time, including while your code is attempting to observe or mutate this state.

4. Taking any locks: C code you call out to may use locks internally, (for example, it is not problematic to call `malloc`, even though most implementations require locks internally) but don't attempt to hold or acquire any while executing Julia code.

5. Calling any function that is defined after the body of the generated function. This condition is relaxed for incrementally-loaded precompiled modules to allow calling any function in the module.

Alright, now that we have a better understanding of how generated functions work, let's use them to build some more advanced (and valid) functionality...

An advanced example

Julia's base library has a `sub2ind()` function to calculate a linear index into an n-dimensional array, based on a set of n multilinear indices - in other words, to calculate the index `i` that can be used to index into an array A using `A[i]`, instead of `A[x,y,z,...]`. One possible implementation is the following:

```julia
julia> function sub2ind_loop(dims::NTuple{N}, I::Integer...) where N
           ind = I[N] - 1
           for i = N-1:-1:1
               ind = I[i]-1 + dims[i]*ind
           end
           return ind + 1
       end
sub2ind_loop (generic function with 1 method)

julia> sub2ind_loop((3, 5), 1, 2)
4
```

The same thing can be done using recursion:

```julia
julia> sub2ind_rec(dims::Tuple{}) = 1;

julia> sub2ind_rec(dims::Tuple{}, i1::Integer, I::Integer...) =
           i1 == 1 ? sub2ind_rec(dims, I...) : throw(BoundsError());
```

```
julia> sub2ind_rec(dims::Tuple{Integer, Vararg{Integer}}, i1::Integer) = i1;

julia> sub2ind_rec(dims::Tuple{Integer, Vararg{Integer}}, i1::Integer, I::Integer...) =
           i1 + dims[1] * (sub2ind_rec(Base.tail(dims), I...) - 1);

julia> sub2ind_rec((3, 5), 1, 2)
4
```

Both these implementations, although different, do essentially the same thing: a runtime loop over the dimensions of the array, collecting the offset in each dimension into the final index.

However, all the information we need for the loop is embedded in the type information of the arguments. Thus, we can utilize generated functions to move the iteration to compile-time; in compiler parlance, we use generated functions to manually unroll the loop. The body becomes almost identical, but instead of calculating the linear index, we build up an *expression* that calculates the index:

```
julia> @generated function sub2ind_gen(dims::NTuple{N}, I::Integer...) where N
           ex = :(I[$N] - 1)
           for i = (N - 1):-1:1
               ex = :(I[$i] - 1 + dims[$i] * $ex)
           end
           return :($ex + 1)
       end
sub2ind_gen (generic function with 1 method)

julia> sub2ind_gen((3, 5), 1, 2)
4
```

What code will this generate?

An easy way to find out is to extract the body into another (regular) function:

```
julia> @generated function sub2ind_gen(dims::NTuple{N}, I::Integer...) where N
           return sub2ind_gen_impl(dims, I...)
       end
sub2ind_gen (generic function with 1 method)

julia> function sub2ind_gen_impl(dims::Type{T}, I...) where T <: NTuple{N,Any} where N
           length(I) == N || return :(error("partial indexing is unsupported"))
           ex = :(I[$N] - 1)
           for i = (N - 1):-1:1
               ex = :(I[$i] - 1 + dims[$i] * $ex)
           end
           return :($ex + 1)
       end
sub2ind_gen_impl (generic function with 1 method)
```

We can now execute `sub2ind_gen_impl` and examine the expression it returns:

```
julia> sub2ind_gen_impl(Tuple{Int,Int}, Int, Int)
:(((I[1] - 1) + dims[1] * (I[2] - 1)) + 1)
```

So, the method body that will be used here doesn't include a loop at all - just indexing into the two tuples, multiplication and addition/subtraction. All the looping is performed compile-time, and we avoid looping during execution entirely. Thus, we only loop *once per type*, in this case once per N (except in edge cases where the function is generated more than once - see disclaimer above).

Chapter 22

Multi-dimensional Arrays

Julia, like most technical computing languages, provides a first-class array implementation. Most technical comput-
ing languages pay a lot of attention to their array implementation at the expense of other containers. Julia does
not treat arrays in any special way. The array library is implemented almost completely in Julia itself, and derives its
performance from the compiler, just like any other code written in Julia. As such, it's also possible to define custom
array types by inheriting from `AbstractArray`. See the manual section on the AbstractArray interface for more
details on implementing a custom array type.

An array is a collection of objects stored in a multi-dimensional grid. In the most general case, an array may contain
objects of type Any. For most computational purposes, arrays should contain objects of a more specific type, such as
`Float64` or `Int32`.

In general, unlike many other technical computing languages, Julia does not expect programs to be written in a vector-
ized style for performance. Julia's compiler uses type inference and generates optimized code for scalar array indexing,
allowing programs to be written in a style that is convenient and readable, without sacrificing performance, and using
less memory at times.

In Julia, all arguments to functions are passed by reference. Some technical computing languages pass arrays by value,
and this is convenient in many cases. In Julia, modifications made to input arrays within a function will be visible in the
parent function. The entire Julia array library ensures that inputs are not modified by library functions. User code, if it
needs to exhibit similar behavior, should take care to create a copy of inputs that it may modify.

22.1   Arrays

Basic Functions

Function	Description
`eltype(A)`	the type of the elements contained in A
`length(A)`	the number of elements in A
`ndims(A)`	the number of dimensions of A
`size(A)`	a tuple containing the dimensions of A
`size(A,n)`	the size of A along dimension n
`indices(A)`	a tuple containing the valid indices of A
`indices(A,n)`	a range expressing the valid indices along dimension n
`eachindex(A)`	an efficient iterator for visiting each position in A
`stride(A,k)`	the stride (linear index distance between adjacent elements) along dimension k
`strides(A)`	a tuple of the strides in each dimension

Construction and Initialization

Many functions for constructing and initializing arrays are provided. In the following list of such functions, calls with a `dims...` argument can either take a single tuple of dimension sizes or a series of dimension sizes passed as a variable number of arguments. Most of these functions also accept a first input T, which is the element type of the array. If the type T is omitted it will default to `Float64`.

Function	Description
`Array{T}(dims...)`	an uninitialized dense `Array`
`zeros(T, dims...)`	an `Array` of all zeros
`zeros(A)`	an array of all zeros with the same type, element type and shape as A
`ones(T, dims...)`	an `Array` of all ones
`ones(A)`	an array of all ones with the same type, element type and shape as A
`trues(dims...)`	a `BitArray` with all values `true`
`trues(A)`	a `BitArray` with all values `true` and the same shape as A
`falses(dims...)`	a `BitArray` with all values `false`
`falses(A)`	a `BitArray` with all values `false` and the same shape as A
`reshape(A, dims...)`	an array containing the same data as A, but with different dimensions
`copy(A)`	copy A
`deepcopy(A)`	copy A, recursively copying its elements
`similar(A, T, dims...)`	an uninitialized array of the same type as A (dense, sparse, etc.), but with the specified element type and dimensions. The second and third arguments are both optional, defaulting to the element type and dimensions of A if omitted.
`reinterpret(T, A)`	an array with the same binary data as A, but with element type T
`rand(T, dims...)`	an `Array` with random, iid [1] and uniformly distributed values in the half-open interval $[0, 1)$
`randn(T, dims...)`	an `Array` with random, iid and standard normally distributed values
`eye(T, n)`	n-by-n identity matrix
`eye(T, m, n)`	m-by-n identity matrix
`linspace(start, stop, n)`	range of n linearly spaced elements from `start` to `stop`
`fill!(A, x)`	fill the array A with the value x
`fill(x, dims...)`	an `Array` filled with the value x

The syntax `[A, B, C, ...]` constructs a 1-d array (vector) of its arguments. If all arguments have a common promotion type then they get converted to that type using `convert()`.

Concatenation

Arrays can be constructed and also concatenated using the following functions:

Scalar values passed to these functions are treated as 1-element arrays.

---
[1] *iid*, independently and identically distributed.

Function	Description
`cat(k, A...)`	concatenate input n-d arrays along the dimension k
`vcat(A...)`	shorthand for `cat(1, A...)`
`hcat(A...)`	shorthand for `cat(2, A...)`

The concatenation functions are used so often that they have special syntax:

Expression	Calls
`[A; B; C; ...]`	`vcat()`
`[A B C ...]`	`hcat()`
`[A B; C D; ...]`	`hvcat()`

`hvcat()` concatenates in both dimension 1 (with semicolons) and dimension 2 (with spaces).

Typed array initializers

An array with a specific element type can be constructed using the syntax `T[A, B, C, ...]`. This will construct a 1-d array with element type T, initialized to contain elements A, B, C, etc. For example `Any[x, y, z]` constructs a heterogeneous array that can contain any values.

Concatenation syntax can similarly be prefixed with a type to specify the element type of the result.

```
julia> [[1 2] [3 4]]
1×4 Array{Int64,2}:
 1  2  3  4

julia> Int8[[1 2] [3 4]]
1×4 Array{Int8,2}:
 1  2  3  4
```

Comprehensions

Comprehensions provide a general and powerful way to construct arrays. Comprehension syntax is similar to set construction notation in mathematics:

```
A = [ F(x,y,...) for x=rx, y=ry, ... ]
```

The meaning of this form is that `F(x,y,...)` is evaluated with the variables x, y, etc. taking on each value in their given list of values. Values can be specified as any iterable object, but will commonly be ranges like `1:n` or `2:(n-1)`, or explicit arrays of values like `[1.2, 3.4, 5.7]`. The result is an N-d dense array with dimensions that are the concatenation of the dimensions of the variable ranges rx, ry, etc. and each `F(x,y,...)` evaluation returns a scalar.

The following example computes a weighted average of the current element and its left and right neighbor along a 1-d grid. :

```
julia> x = rand(8)
8-element Array{Float64,1}:
 0.843025
 0.869052
 0.365105
 0.699456
 0.977653
 0.994953
```

```
 0.41084
 0.809411

julia> [ 0.25*x[i-1] + 0.5*x[i] + 0.25*x[i+1] for i=2:length(x)-1 ]
6-element Array{Float64,1}:
 0.736559
 0.57468
 0.685417
 0.912429
 0.8446
 0.656511
```

The resulting array type depends on the types of the computed elements. In order to control the type explicitly, a type can be prepended to the comprehension. For example, we could have requested the result in single precision by writing:

```
Float32[ 0.25*x[i-1] + 0.5*x[i] + 0.25*x[i+1] for i=2:length(x)-1 ]
```

Generator Expressions

Comprehensions can also be written without the enclosing square brackets, producing an object known as a generator. This object can be iterated to produce values on demand, instead of allocating an array and storing them in advance (see Iteration). For example, the following expression sums a series without allocating memory:

```
julia> sum(1/n^2 for n=1:1000)
1.6439345666815615
```

When writing a generator expression with multiple dimensions inside an argument list, parentheses are needed to separate the generator from subsequent arguments:

```
julia> map(tuple, 1/(i+j) for i=1:2, j=1:2, [1:4;])
ERROR: syntax: invalid iteration specification
```

All comma-separated expressions after `for` are interpreted as ranges. Adding parentheses lets us add a third argument to `map`:

```
julia> map(tuple, (1/(i+j) for i=1:2, j=1:2), [1 3; 2 4])
2×2 Array{Tuple{Float64,Int64},2}:
 (0.5, 1)       (0.333333, 3)
 (0.333333, 2)  (0.25, 4)
```

Ranges in generators and comprehensions can depend on previous ranges by writing multiple `for` keywords:

```
julia> [(i,j) for i=1:3 for j=1:i]
6-element Array{Tuple{Int64,Int64},1}:
 (1, 1)
 (2, 1)
 (2, 2)
 (3, 1)
 (3, 2)
 (3, 3)
```

In such cases, the result is always 1-d.

Generated values can be filtered using the `if` keyword:

```julia
julia> [(i,j) for i=1:3 for j=1:i if i+j == 4]
2-element Array{Tuple{Int64,Int64},1}:
 (2, 2)
 (3, 1)
```

Indexing

The general syntax for indexing into an n-dimensional array A is:

```julia
X = A[I_1, I_2, ..., I_n]
```

where each `I_k` may be a scalar integer, an array of integers, or any other supported index. This includes Colon (`:`) to select all indices within the entire dimension, ranges of the form `a:c` or `a:b:c` to select contiguous or strided subsections, and arrays of booleans to select elements at their `true` indices.

If all the indices are scalars, then the result X is a single element from the array A. Otherwise, X is an array with the same number of dimensions as the sum of the dimensionalities of all the indices.

If all indices are vectors, for example, then the shape of X would be (`length(I_1)`, `length(I_2)`, ..., `length(I_n)`), with location (`i_1`, `i_2`, ..., `i_n`) of X containing the value `A[I_1[i_1], I_2[i_2], ..., I_n[i_n]]`. If `I_1` is changed to a two-dimensional matrix, then X becomes an n+1-dimensional array of shape (`size(I_1, 1)`, `size(I_1, 2)`, `length(I_2)`, ..., `length(I_n)`). The matrix adds a dimension. The location (`i_1`, `i_2`, `i_3`, ..., `i_{n+1}`) contains the value at `A[I_1[i_1, i_2], I_2[i_3], ..., I_n[i_{n+1}]]`. All dimensions indexed with scalars are dropped. For example, the result of `A[2, I, 3]` is an array with size `size(I)`. Its ith element is populated by `A[2, I[i], 3]`.

As a special part of this syntax, the end keyword may be used to represent the last index of each dimension within the indexing brackets, as determined by the size of the innermost array being indexed. Indexing syntax without the end keyword is equivalent to a call to `getindex`:

```julia
X = getindex(A, I_1, I_2, ..., I_n)
```

Example:

```julia
julia> x = reshape(1:16, 4, 4)
4×4 Base.ReshapedArray{Int64,2,UnitRange{Int64},Tuple{}}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> x[2:3, 2:end-1]
2×2 Array{Int64,2}:
 6  10
 7  11

julia> x[1, [2 3; 4 1]]
2×2 Array{Int64,2}:
  5  9
 13  1
```

Empty ranges of the form n:n-1 are sometimes used to indicate the inter-index location between n-1 and n. For example, the searchsorted() function uses this convention to indicate the insertion point of a value not found in a sorted array:

```julia
julia> a = [1,2,5,6,7];

julia> searchsorted(a, 3)
3:2
```

Assignment

The general syntax for assigning values in an n-dimensional array A is:

```
A[I_1, I_2, ..., I_n] = X
```

where each I_k may be a scalar integer, an array of integers, or any other supported index. This includes Colon (:) to select all indices within the entire dimension, ranges of the form a:c or a:b:c to select contiguous or strided subsections, and arrays of booleans to select elements at their true indices.

If X is an array, it must have the same number of elements as the product of the lengths of the indices: prod(length(I_1), length(I_2), ..., length(I_n)). The value in location I_1[i_1], I_2[i_2], ..., I_n[i_n] of A is overwritten with the value X[i_1, i_2, ..., i_n]. If X is not an array, its value is written to all referenced locations of A.

Just as in Indexing, the end keyword may be used to represent the last index of each dimension within the indexing brackets, as determined by the size of the array being assigned into. Indexed assignment syntax without the end keyword is equivalent to a call to setindex!():

```
setindex!(A, X, I_1, I_2, ..., I_n)
```

Example:

```julia
julia> x = collect(reshape(1:9, 3, 3))
3×3 Array{Int64,2}:
 1  4  7
 2  5  8
 3  6  9

julia> x[1:2, 2:3] = -1
-1

julia> x
3×3 Array{Int64,2}:
 1  -1  -1
 2  -1  -1
 3   6   9
```

Supported index types

In the expression A[I_1, I_2, ..., I_n], each I_k may be a scalar index, an array of scalar indices, or an object that represents an array of scalar indices and can be converted to such by to_indices:

1. A scalar index. By default this includes:

    – Non-boolean integers

– `CartesianIndex{N}`s, which behave like an N-tuple of integers spanning multiple dimensions (see below for more details)

2. An array of scalar indices. This includes:

– Vectors and multidimensional arrays of integers

– Empty arrays like `[ ]`, which select no elements

– Ranges of the form `a:c` or `a:b:c`, which select contiguous or strided subsections from a to c (inclusive)

– Any custom array of scalar indices that is a subtype of `AbstractArray`

– Arrays of `CartesianIndex{N}` (see below for more details)

3. An object that represents an array of scalar indices and can be converted to such by `to_indices`. By default this includes:

– `Colon()` (`:`), which represents all indices within an entire dimension or across the entire array

– Arrays of booleans, which select elements at their `true` indices (see below for more details)

Cartesian indices

The special `CartesianIndex{N}` object represents a scalar index that behaves like an N-tuple of integers spanning multiple dimensions. For example:

```julia
julia> A = reshape(1:32, 4, 4, 2);

julia> A[3, 2, 1]
7

julia> A[CartesianIndex(3, 2, 1)] == A[3, 2, 1] == 7
true
```

Considered alone, this may seem relatively trivial; `CartesianIndex` simply gathers multiple integers together into one object that represents a single multidimensional index. When combined with other indexing forms and iterators that yield `CartesianIndexes`, however, this can lead directly to very elegant and efficient code. See Iteration below, and for some more advanced examples, see this blog post on multidimensional algorithms and iteration.

Arrays of `CartesianIndex{N}` are also supported. They represent a collection of scalar indices that each span N dimensions, enabling a form of indexing that is sometimes referred to as pointwise indexing. For example, it enables accessing the diagonal elements from the first "page" of A from above:

```julia
julia> page = A[:,:,1]
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> page[[CartesianIndex(1,1),
             CartesianIndex(2,2),
             CartesianIndex(3,3),
             CartesianIndex(4,4)]]
4-element Array{Int64,1}:
  1
  6
 11
 16
```

This can be expressed much more simply with dot broadcasting and by combining it with a normal integer index (instead of extracting the first page from A as a separate step). It can even be combined with a : to extract both diagonals from the two pages at the same time:

```
julia> A[CartesianIndex.(indices(A, 1), indices(A, 2)), 1]
4-element Array{Int64,1}:
  1
  6
 11
 16

julia> A[CartesianIndex.(indices(A, 1), indices(A, 2)), :]
4×2 Array{Int64,2}:
  1  17
  6  22
 11  27
 16  32
```

> **Warning**
>
> CartesianIndex and arrays of CartesianIndex are not compatible with the end keyword to rep-
> resent the last index of a dimension. Do not use end in indexing expressions that may contain either
> CartesianIndex or arrays thereof.

Logical indexing

Often referred to as logical indexing or indexing with a logical mask, indexing by a boolean array selects elements at the indices where its values are true. Indexing by a boolean vector B is effectively the same as indexing by the vector of integers that is returned by find(B). Similarly, indexing by a N-dimensional boolean array is effectively the same as indexing by the vector of CartesianIndex{N}s where its values are true. A logical index must be a vector of the same length as the dimension it indexes into, or it must be the only index provided and match the size and dimensionality of the array it indexes into. It is generally more efficient to use boolean arrays as indices directly instead of first calling find().

```
julia> x = reshape(1:16, 4, 4)
4×4 Base.ReshapedArray{Int64,2,UnitRange{Int64},Tuple{}}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> x[[false, true, true, false], :]
2×4 Array{Int64,2}:
 2  6  10  14
 3  7  11  15

julia> mask = map(ispow2, x)
4×4 Array{Bool,2}:
  true  false  false  false
  true  false  false  false
 false  false  false  false
  true   true  false   true

julia> x[mask]
5-element Array{Int64,1}:
```

```
  1
  2
  4
  8
 16
```

Iteration

The recommended ways to iterate over a whole array are

```julia
for a in A
    # Do something with the element a
end

for i in eachindex(A)
    # Do something with i and/or A[i]
end
```

The first construct is used when you need the value, but not index, of each element. In the second construct, `i` will be an `Int` if `A` is an array type with fast linear indexing; otherwise, it will be a `CartesianIndex`:

```julia
julia> A = rand(4,3);

julia> B = view(A, 1:3, 2:3);

julia> for i in eachindex(B)
           @show i
       end
i = CartesianIndex{2}((1, 1))
i = CartesianIndex{2}((2, 1))
i = CartesianIndex{2}((3, 1))
i = CartesianIndex{2}((1, 2))
i = CartesianIndex{2}((2, 2))
i = CartesianIndex{2}((3, 2))
```

In contrast with `for i = 1:length(A)`, iterating with `eachindex` provides an efficient way to iterate over any array type.

Array traits

If you write a custom `AbstractArray` type, you can specify that it has fast linear indexing using

```julia
Base.IndexStyle(::Type{<:MyArray}) = IndexLinear()
```

This setting will cause `eachindex` iteration over a `MyArray` to use integers. If you don't specify this trait, the default value `IndexCartesian()` is used.

Array and Vectorized Operators and Functions

The following operators are supported for arrays:

1. Unary arithmetic – `-`, `+`

2. Binary arithmetic – -, +, *, /, \, ^

3. Comparison – ==, !=, ≈ (`isapprox`),

Most of the binary arithmetic operators listed above also operate elementwise when one argument is scalar: -, +, and * when either argument is scalar, and / and \ when the denominator is scalar. For example, `[1, 2] + 3 == [4, 5]` and `[6, 4] / 2 == [3, 2]`.

Additionally, to enable convenient vectorization of mathematical and other operations, Julia provides the dot syntax `f.(args...)`, e.g. `sin.(x)` or `min.(x,y)`, for elementwise operations over arrays or mixtures of arrays and scalars (a Broadcasting operation); these have the additional advantage of "fusing" into a single loop when combined with other dot calls, e.g. `sin.(cos.(x))`.

Also, *every* binary operator supports a dot version that can be applied to arrays (and combinations of arrays and scalars) in such fused broadcasting operations, e.g. `z .== sin.(x .* y)`.

Note that comparisons such as == operate on whole arrays, giving a single boolean answer. Use dot operators like `.==` for elementwise comparisons. (For comparison operations like <, *only* the elementwise `.<` version is applicable to arrays.)

Also notice the difference between `max.(a,b)`, which broadcasts `max()` elementwise over a and b, and `maximum(a)`, which finds the largest value within a. The same relationship holds for `min.(a,b)` and `minimum(a)`.

Broadcasting

It is sometimes useful to perform element-by-element binary operations on arrays of different sizes, such as adding a vector to each column of a matrix. An inefficient way to do this would be to replicate the vector to the size of the matrix:

```julia
julia> a = rand(2,1); A = rand(2,3);

julia> repmat(a,1,3)+A
2×3 Array{Float64,2}:
 1.20813  1.82068  1.25387
 1.56851  1.86401  1.67846
```

This is wasteful when dimensions get large, so Julia offers `broadcast()`, which expands singleton dimensions in array arguments to match the corresponding dimension in the other array without using extra memory, and applies the given function elementwise:

```julia
julia> broadcast(+, a, A)
2×3 Array{Float64,2}:
 1.20813  1.82068  1.25387
 1.56851  1.86401  1.67846

julia> b = rand(1,2)
1×2 Array{Float64,2}:
 0.867535  0.00457906

julia> broadcast(+, a, b)
2×2 Array{Float64,2}:
 1.71056  0.847604
 1.73659  0.873631
```

Dotted operators such as `.+` and `.*` are equivalent to `broadcast` calls (except that they fuse, as described below). There is also a `broadcast!()` function to specify an explicit destination (which can also be accessed in a fusing fashion by `.=` assignment), and functions `broadcast_getindex()` and `broadcast_setindex!()` that broadcast the indices before indexing. Moreover, `f.(args...)` is equivalent to `broadcast(f, args...)`, providing a convenient syntax to broadcast any function (dot syntax). Nested "dot calls" `f.(...)` (including calls to `.+` etcetera) automatically fuse into a single `broadcast` call.

Additionally, `broadcast()` is not limited to arrays (see the function documentation), it also handles tuples and treats any argument that is not an array, tuple or `Ref` (except for `Ptr`) as a "scalar".

```julia
julia> convert.(Float32, [1, 2])
2-element Array{Float32,1}:
 1.0
 2.0

julia> ceil.((UInt8,), [1.2 3.4; 5.6 6.7])
2×2 Array{UInt8,2}:
 0x02  0x04
 0x06  0x07

julia> string.(1:3, ". ", ["First", "Second", "Third"])
3-element Array{String,1}:
 "1. First"
 "2. Second"
 "3. Third"
```

Implementation

The base array type in Julia is the abstract type `AbstractArray{T,N}`. It is parametrized by the number of dimensions `N` and the element type `T`. `AbstractVector` and `AbstractMatrix` are aliases for the 1-d and 2-d cases. Operations on `AbstractArray` objects are defined using higher level operators and functions, in a way that is independent of the underlying storage. These operations generally work correctly as a fallback for any specific array implementation.

The `AbstractArray` type includes anything vaguely array-like, and implementations of it might be quite different from conventional arrays. For example, elements might be computed on request rather than stored. However, any concrete `AbstractArray{T,N}` type should generally implement at least `size(A)` (returning an `Int` tuple), `getindex(A,i)` and `getindex(A,i1,...,iN)`; mutable arrays should also implement `setindex!()`. It is recommended that these operations have nearly constant time complexity, or technically $\tilde{O}(1)$ complexity, as otherwise some array functions may be unexpectedly slow. Concrete types should also typically provide a `similar(A,T=eltype(A),dims=size(A))` method, which is used to allocate a similar array for `copy()` and other out-of-place operations. No matter how an `AbstractArray{T,N}` is represented internally, `T` is the type of object returned by *integer* indexing (`A[1, ..., 1]`, when `A` is not empty) and `N` should be the length of the tuple returned by `size()`.

`DenseArray` is an abstract subtype of `AbstractArray` intended to include all arrays that are laid out at regular offsets in memory, and which can therefore be passed to external C and Fortran functions expecting this memory layout. Subtypes should provide a method `stride(A,k)` that returns the "stride" of dimension `k`: increasing the index of dimension `k` by 1 should increase the index `i` of `getindex(A,i)` by `stride(A,k)`. If a pointer conversion method `Base.unsafe_convert(Ptr{T}, A)` is provided, the memory layout should correspond in the same way to these strides.

The `Array` type is a specific instance of `DenseArray` where elements are stored in column-major order (see additional notes in Performance Tips). `Vector` and `Matrix` are aliases for the 1-d and 2-d cases. Specific operations such as scalar indexing, assignment, and a few other basic storage-specific operations are all that have to be implemented for `Array`, so that the rest of the array library can be implemented in a generic manner.

SubArray is a specialization of AbstractArray that performs indexing by reference rather than by copying. A Sub-Array is created with the `view()` function, which is called the same way as `getindex()` (with an array and a series of index arguments). The result of `view()` looks the same as the result of `getindex()`, except the data is left in place. `view()` stores the input index vectors in a SubArray object, which can later be used to index the original array indirectly. By putting the `@views` macro in front of an expression or block of code, any `array[...]` slice in that expression will be converted to create a SubArray view instead.

StridedVector and StridedMatrix are convenient aliases defined to make it possible for Julia to call a wider range of BLAS and LAPACK functions by passing them either `Array` or SubArray objects, and thus saving inefficiencies from memory allocation and copying.

The following example computes the QR decomposition of a small section of a larger array, without creating any temporaries, and by calling the appropriate LAPACK function with the right leading dimension size and stride parameters.

```julia
julia> a = rand(10,10)
10×10 Array{Float64,2}:
 0.561255   0.226678   0.203391  0.308912   …  0.750307  0.235023   0.217964
 0.718915   0.537192   0.556946  0.996234      0.666232  0.509423   0.660788
 0.493501   0.0565622  0.118392  0.493498      0.262048  0.940693   0.252965
 0.0470779  0.736979   0.264822  0.228787      0.161441  0.897023   0.567641
 0.343935   0.32327    0.795673  0.452242      0.468819  0.628507   0.511528
 0.935597   0.991511   0.571297  0.74485    …  0.84589   0.178834   0.284413
 0.160706   0.672252   0.133158  0.65554       0.371826  0.770628   0.0531208
 0.306617   0.836126   0.301198  0.0224702     0.39344   0.0370205  0.536062
 0.890947   0.168877   0.32002   0.486136      0.096078  0.172048   0.77672
 0.507762   0.573567   0.220124  0.165816      0.211049  0.433277   0.539476

julia> b = view(a, 2:2:8,2:2:4)
4×2 SubAr-
↪   ray{Float64,2,Array{Float64,2},Tuple{StepRange{Int64,Int64},StepRange{Int64,Int64}},false}:
 0.537192  0.996234
 0.736979  0.228787
 0.991511  0.74485
 0.836126  0.0224702

julia> (q,r) = qr(b);

julia> q
4×2 Array{Float64,2}:
 -0.338809   0.78934
 -0.464815  -0.230274
 -0.625349   0.194538
 -0.527347  -0.534856

julia> r
2×2 Array{Float64,2}:
 -1.58553  -0.921517
  0.0       0.866567
```

22.2   Sparse Matrices

Sparse matrices are matrices that contain enough zeros that storing them in a special data structure leads to savings in space and execution time. Sparse matrices may be used when operations on the sparse representation of a matrix lead to considerable gains in either time or space when compared to performing the same operations on a dense matrix.

Compressed Sparse Column (CSC) Storage

In Julia, sparse matrices are stored in the Compressed Sparse Column (CSC) format. Julia sparse matrices have the type `SparseMatrixCSC{Tv,Ti}`, where `Tv` is the type of the stored values, and `Ti` is the integer type for storing column pointers and row indices.:

```julia
struct SparseMatrixCSC{Tv,Ti<:Integer} <: AbstractSparseMatrix{Tv,Ti}
    m::Int                  # Number of rows
    n::Int                  # Number of columns
    colptr::Vector{Ti}      # Column i is in colptr[i]:(colptr[i+1]-1)
    rowval::Vector{Ti}      # Row indices of stored values
    nzval::Vector{Tv}       # Stored values, typically nonzeros
end
```

The compressed sparse column storage makes it easy and quick to access the elements in the column of a sparse matrix, whereas accessing the sparse matrix by rows is considerably slower. Operations such as insertion of previously unstored entries one at a time in the CSC structure tend to be slow. This is because all elements of the sparse matrix that are beyond the point of insertion have to be moved one place over.

All operations on sparse matrices are carefully implemented to exploit the CSC data structure for performance, and to avoid expensive operations.

If you have data in CSC format from a different application or library, and wish to import it in Julia, make sure that you use 1-based indexing. The row indices in every column need to be sorted. If your `SparseMatrixCSC` object contains unsorted row indices, one quick way to sort them is by doing a double transpose.

In some applications, it is convenient to store explicit zero values in a `SparseMatrixCSC`. These *are* accepted by functions in `Base` (but there is no guarantee that they will be preserved in mutating operations). Such explicitly stored zeros are treated as structural nonzeros by many routines. The `nnz()` function returns the number of elements explicitly stored in the sparse data structure, including structural nonzeros. In order to count the exact number of actual values that are nonzero, use `countnz()`, which inspects every stored element of a sparse matrix.

Sparse matrix constructors

The simplest way to create sparse matrices is to use functions equivalent to the `zeros()` and `eye()` functions that Julia provides for working with dense matrices. To produce sparse matrices instead, you can use the same names with an `sp` prefix:

```julia
julia> spzeros(3,5)
3×5 SparseMatrixCSC{Float64,Int64} with 0 stored entries

julia> speye(3,5)
3×5 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0
```

The `sparse()` function is often a handy way to construct sparse matrices. It takes as its input a vector `I` of row indices, a vector `J` of column indices, and a vector `V` of stored values. `sparse(I,J,V)` constructs a sparse matrix such that `S[I[k], J[k]] = V[k]`.

```julia
julia> I = [1, 4, 3, 5]; J = [4, 7, 18, 9]; V = [1, 2, -5, 3];

julia> S = sparse(I,J,V)
5×18 SparseMatrixCSC{Int64,Int64} with 4 stored entries:
```

```
  [1 ,   4]   =   1
  [4 ,   7]   =   2
  [5 ,   9]   =   3
  [3 , 18]   =  -5
```

The inverse of the `sparse()` function is `findn()`, which retrieves the inputs used to create the sparse matrix.

```
julia> findn(S)
([1, 4, 5, 3], [4, 7, 9, 18])

julia> findnz(S)
([1, 4, 5, 3], [4, 7, 9, 18], [1, 2, 3, -5])
```

Another way to create sparse matrices is to convert a dense matrix into a sparse matrix using the `sparse()` function:

```
julia> sparse(eye(5))
5×5 SparseMatrixCSC{Float64,Int64} with 5 stored entries:
  [1, 1]   =   1.0
  [2, 2]   =   1.0
  [3, 3]   =   1.0
  [4, 4]   =   1.0
  [5, 5]   =   1.0
```

You can go in the other direction using the `full()` function. The `issparse()` function can be used to query if a matrix is sparse.

```
julia> issparse(speye(5))
true
```

Sparse matrix operations

Arithmetic operations on sparse matrices also work as they do on dense matrices. Indexing of, assignment into, and concatenation of sparse matrices work in the same way as dense matrices. Indexing operations, especially assignment, are expensive, when carried out one element at a time. In many cases it may be better to convert the sparse matrix into (I,J,V) format using `findnz()`, manipulate the values or the structure in the dense vectors (I,J,V), and then reconstruct the sparse matrix.

Correspondence of dense and sparse methods

The following table gives a correspondence between built-in methods on sparse matrices and their corresponding methods on dense matrix types. In general, methods that generate sparse matrices differ from their dense counterparts in that the resulting matrix follows the same sparsity pattern as a given sparse matrix S, or that the resulting sparse matrix has density d, i.e. each matrix element has a probability d of being non-zero.

Details can be found in the Sparse Vectors and Matrices section of the standard library reference.

Sparse	Dense	Description
`spze-ros(m,n)`	`ze-ros(m,n)`	Creates a *m*-by-*n* matrix of zeros. (`spzeros(m,n)` is empty.)
`spones(S)`	`ones(m,n)`	Creates a matrix filled with ones. Unlike the dense version, `spones()` has the same sparsity pattern as *S*.
`speye(n)`	`eye(n)`	Creates a *n*-by-*n* identity matrix.
`full(S)`	`sparse(A)`	Interconverts between dense and sparse formats.
`sprand(m,n,d)`	`rand(m,n)`	Creates a *m*-by-*n* random matrix (of density *d*) with iid non-zero elements distributed uniformly on the half-open interval $[0, 1)$.
`sprandn(m,n,d)`	`randn(m,n)`	Creates a *m*-by-*n* random matrix (of density *d*) with iid non-zero elements distributed according to the standard normal (Gaussian) distribution.
`sprandn(m,n,d,X)`	`randn(m,n)`	Creates a *m*-by-*n* random matrix (of density *d*) with iid non-zero elements distributed according to the *X* distribution. (Requires the `Distributions` package.)

Chapter 23

Linear algebra

In addition to (and as part of) its support for multi-dimensional arrays, Julia provides native implementations of many common and useful linear algebra operations. Basic operations, such as `trace`, `det`, and `inv` are all supported:

```julia
julia> A = [1 2 3; 4 1 6; 7 8 1]
3×3 Array{Int64,2}:
 1  2  3
 4  1  6
 7  8  1

julia> trace(A)
3

julia> det(A)
104.0

julia> inv(A)
3×3 Array{Float64,2}:
 -0.451923   0.211538    0.0865385
  0.365385  -0.192308    0.0576923
  0.240385   0.0576923  -0.0673077
```

As well as other useful operations, such as finding eigenvalues or eigenvectors:

```julia
julia> A = [1.5 2 -4; 3 -1 -6; -10 2.3 4]
3×3 Array{Float64,2}:
   1.5   2.0  -4.0
   3.0  -1.0  -6.0
 -10.0   2.3   4.0

julia> eigvals(A)
3-element Array{Complex{Float64},1}:
  9.31908+0.0im
 -2.40954+2.72095im
 -2.40954-2.72095im

julia> eigvecs(A)
3×3 Array{Complex{Float64},2}:
 -0.488645+0.0im  0.182546-0.39813im   0.182546+0.39813im
 -0.540358+0.0im  0.692926+0.0im       0.692926-0.0im
  0.68501+0.0im   0.254058-0.513301im  0.254058+0.513301im
```

In addition, Julia provides many factorizations which can be used to speed up problems such as linear solve or matrix exponentiation by pre-factorizing a matrix into a form more amenable (for performance or memory reasons) to the problem. See the documentation on `factorize` for more information. As an example:

```julia
julia> A = [1.5 2 -4; 3 -1 -6; -10 2.3 4]
3×3 Array{Float64,2}:
   1.5   2.0  -4.0
   3.0  -1.0  -6.0
 -10.0   2.3   4.0

julia> factorize(A)
Base.LinAlg.LU{Float64,Array{Float64,2}} with factors L and U:
[1.0 0.0 0.0; -0.15 1.0 0.0; -0.3 -0.132196 1.0]
[-10.0 2.3 4.0; 0.0 2.345 -3.4; 0.0 0.0 -5.24947]
```

Since A is not Hermitian, symmetric, triangular, tridiagonal, or bidiagonal, an LU factorization may be the best we can do. Compare with:

```julia
julia> B = [1.5 2 -4; 2 -1 -3; -4 -3 5]
3×3 Array{Float64,2}:
  1.5   2.0  -4.0
  2.0  -1.0  -3.0
 -4.0  -3.0   5.0

julia> factorize(B)
Base.LinAlg.BunchKaufman{Float64,Array{Float64,2}}([-1.64286 0.142857 -0.8; 2.0 -2.8 -0.6; -4.0
↪   -3.0 5.0], [1, 2, 3], 'U', true, false, 0)
```

Here, Julia was able to detect that B is in fact symmetric, and used a more appropriate factorization. Often it's possible to write more efficient code for a matrix that is known to have certain properties e.g. it is symmetric, or tridiagonal. Julia provides some special types so that you can "tag" matrices as having these properties. For instance:

```julia
julia> B = [1.5 2 -4; 2 -1 -3; -4 -3 5]
3×3 Array{Float64,2}:
  1.5   2.0  -4.0
  2.0  -1.0  -3.0
 -4.0  -3.0   5.0

julia> sB = Symmetric(B)
3×3 Symmetric{Float64,Array{Float64,2}}:
  1.5   2.0  -4.0
  2.0  -1.0  -3.0
 -4.0  -3.0   5.0
```

sB has been tagged as a matrix that's (real) symmetric, so for later operations we might perform on it, such as eigenfactorization or computing matrix-vector products, efficiencies can be found by only referencing half of it. For example:

```julia
julia> B = [1.5 2 -4; 2 -1 -3; -4 -3 5]
3×3 Array{Float64,2}:
  1.5   2.0  -4.0
  2.0  -1.0  -3.0
 -4.0  -3.0   5.0
```

```
julia> sB = Symmetric(B)
3×3 Symmetric{Float64,Array{Float64,2}}:
  1.5   2.0  -4.0
  2.0  -1.0  -3.0
 -4.0  -3.0   5.0

julia> x = [1; 2; 3]
3-element Array{Int64,1}:
 1
 2
 3

julia> sB\x
3-element Array{Float64,1}:
 -1.73913
 -1.1087
 -1.45652
```

The \ operation here performs the linear solution. Julia's parser provides convenient dispatch to specialized methods for the *transpose* of a matrix left-divided by a vector, or for the various combinations of transpose operations in matrix-matrix solutions. Many of these are further specialized for certain special matrix types. For example, A\B will end up calling `Base.LinAlg.A_ldiv_B!` while A'\B will end up calling `Base.LinAlg.Ac_ldiv_B`, even though we used the same left-division operator. This works for matrices too: `A.'\B.'` would call `Base.LinAlg.At_ldiv_Bt`. The left-division operator is pretty powerful and it's easy to write compact, readable code that is flexible enough to solve all sorts of systems of linear equations.

23.1 Special matrices

Matrices with special symmetries and structures arise often in linear algebra and are frequently associated with various matrix factorizations. Julia features a rich collection of special matrix types, which allow for fast computation with specialized routines that are specially developed for particular matrix types.

The following tables summarize the types of special matrices that have been implemented in Julia, as well as whether hooks to various optimized methods for them in LAPACK are available.

Type	Description
Hermitian	Hermitian matrix
UpperTriangular	Upper triangular matrix
LowerTriangular	Lower triangular matrix
Tridiagonal	Tridiagonal matrix
SymTridiagonal	Symmetric tridiagonal matrix
Bidiagonal	Upper/lower bidiagonal matrix
Diagonal	Diagonal matrix
UniformScaling	Uniform scaling operator

Elementary operations

Legend:

Matrix factorizations

Legend:

Matrix type	+	−	*	\	Other functions with optimized methods
Hermitian				MV	`inv()`, `sqrtm()`, `expm()`
UpperTriangular			MV	MV	`inv()`, `det()`
LowerTriangular			MV	MV	`inv()`, `det()`
SymTridiagonal	M	M	MS	MV	`eigmax()`, `eigmin()`
Tridiagonal	M	M	MS	MV	
Bidiagonal	M	M	MS	MV	
Diagonal	M	M	MV	MV	`inv()`, `det()`, `logdet()`, `/()`
UniformScaling	M	M	MVS	MVS	`/()`

Key	Description
M (matrix)	An optimized method for matrix-matrix operations is available
V (vector)	An optimized method for matrix-vector operations is available
S (scalar)	An optimized method for matrix-scalar operations is available

Matrix type	LAPACK	`eig()`	`eigvals()`	`eigvecs()`	`svd()`	`svdvals()`
Hermitian	HE		ARI			
UpperTriangular	TR	A	A	A		
LowerTriangular	TR	A	A	A		
SymTridiagonal	ST	A	ARI	AV		
Tridiagonal	GT					
Bidiagonal	BD				A	A
Diagonal	DI		A			

The uniform scaling operator

A `UniformScaling` operator represents a scalar times the identity operator, $\lambda*I$. The identity operator `I` is defined as a constant and is an instance of `UniformScaling`. The size of these operators are generic and match the other matrix in the binary operations `+`, `−`, `*` and `\`. For `A+I` and `A−I` this means that `A` must be square. Multiplication with the identity operator `I` is a noop (except for checking that the scaling factor is one) and therefore almost without overhead.

23.2 Matrix factorizations

Matrix factorizations (a.k.a. matrix decompositions) compute the factorization of a matrix into a product of matrices, and are one of the central concepts in linear algebra.

The following table summarizes the types of matrix factorizations that have been implemented in Julia. Details of their associated methods can be found in the Linear Algebra section of the standard library documentation.

Key	Description	Example
A (all)	An optimized method to find all the characteristic values and/or vectors is available	e.g. `eigvals(M)`
R (range)	An optimized method to find the `il`th through the `ih`th characteristic values are available	`eigvals(M, il, ih)`
I (interval)	An optimized method to find the characteristic values in the interval [vl, vh] is available	`eigvals(M, vl, vh)`
V (vectors)	An optimized method to find the characteristic vectors corresponding to the characteristic values x=[x1, x2,...] is available	`eigvecs(M, x)`

Type	Description
`Cholesky`	Cholesky factorization
`CholeskyPivoted`	Pivoted Cholesky factorization
`LU`	LU factorization
`LUTridiagonal`	LU factorization for `Tridiagonal` matrices
`UmfpackLU`	LU factorization for sparse matrices (computed by UMFPack)
`QR`	QR factorization
`QRCompactWY`	Compact WY form of the QR factorization
`QRPivoted`	Pivoted QR factorization
`Hessenberg`	Hessenberg decomposition
`Eigen`	Spectral decomposition
`SVD`	Singular value decomposition
`GeneralizedSVD`	Generalized SVD

Chapter 24

Networking and Streams

Julia provides a rich interface to deal with streaming I/O objects such as terminals, pipes and TCP sockets. This interface, though asynchronous at the system level, is presented in a synchronous manner to the programmer and it is usually unnecessary to think about the underlying asynchronous operation. This is achieved by making heavy use of Julia cooperative threading (coroutine) functionality.

24.1   Basic Stream I/O

All Julia streams expose at least a `read()` and a `write()` method, taking the stream as their first argument, e.g.:

```julia
julia> write(STDOUT,"Hello World");  # suppress return value 11 with ;
Hello World
julia> read(STDIN,Char)

'\n': ASCII/Unicode U+000a (category Cc: Other, control)
```

Note that `write()` returns 11, the number of bytes (in "Hello World") written to `STDOUT`, but this return value is suppressed with the `;`.

Here Enter was pressed again so that Julia would read the newline. Now, as you can see from this example, `write()` takes the data to write as its second argument, while `read()` takes the type of the data to be read as the second argument.

For example, to read a simple byte array, we could do:

```julia
julia> x = zeros(UInt8, 4)
4-element Array{UInt8,1}:
 0x00
 0x00
 0x00
 0x00

julia> read!(STDIN, x)
abcd
4-element Array{UInt8,1}:
 0x61
 0x62
 0x63
 0x64
```

However, since this is slightly cumbersome, there are several convenience methods provided. For example, we could have written the above as:

```julia
julia> read(STDIN,4)
abcd
4-element Array{UInt8,1}:
 0x61
 0x62
 0x63
 0x64
```

or if we had wanted to read the entire line instead:

```julia
julia> readline(STDIN)
abcd
"abcd"
```

Note that depending on your terminal settings, your TTY may be line buffered and might thus require an additional enter before the data is sent to Julia.

To read every line from `STDIN` you can use `eachline()`:

```julia
for line in eachline(STDIN)
    print("Found $line")
end
```

or `read()` if you wanted to read by character instead:

```julia
while !eof(STDIN)
    x = read(STDIN, Char)
    println("Found: $x")
end
```

24.2   Text I/O

Note that the `write()` method mentioned above operates on binary streams. In particular, values do not get converted to any canonical text representation but are written out as is:

```julia
julia> write(STDOUT,0x61);  # suppress return value 1 with ;
a
```

Note that a is written to `STDOUT` by the `write()` function and that the returned value is 1 (since `0x61` is one byte).

For text I/O, use the `print()` or `show()` methods, depending on your needs (see the standard library reference for a detailed discussion of the difference between the two):

```julia
julia> print(STDOUT, 0x61)
97
```

24.3　IO Output Contextual Properties

Sometimes IO output can benefit from the ability to pass contextual information into show methods. The `IOContext` object provides this framework for associating arbitrary metadata with an IO object. For example, `showcompact` adds a hinting parameter to the IO object that the invoked show method should print a shorter output (if applicable).

24.4　Working with Files

Like many other environments, Julia has an `open()` function, which takes a filename and returns an `IOStream` object that you can use to read and write things from the file. For example if we have a file, `hello.txt`, whose contents are `Hello, World!`:

```julia
julia> f = open("hello.txt")
IOStream(<file hello.txt>)

julia> readlines(f)
1-element Array{String,1}:
 "Hello, World!"
```

If you want to write to a file, you can open it with the write (`"w"`) flag:

```julia
julia> f = open("hello.txt","w")
IOStream(<file hello.txt>)

julia> write(f,"Hello again.")
12
```

If you examine the contents of `hello.txt` at this point, you will notice that it is empty; nothing has actually been written to disk yet. This is because the `IOStream` must be closed before the write is actually flushed to disk:

```julia
julia> close(f)
```

Examining `hello.txt` again will show its contents have been changed.

Opening a file, doing something to its contents, and closing it again is a very common pattern. To make this easier, there exists another invocation of `open()` which takes a function as its first argument and filename as its second, opens the file, calls the function with the file as an argument, and then closes it again. For example, given a function:

```julia
function read_and_capitalize(f::IOStream)
    return uppercase(readstring(f))
end
```

You can call:

```julia
julia> open(read_and_capitalize, "hello.txt")
"HELLO AGAIN."
```

to open `hello.txt`, call `read_and_capitalize on it`, close `hello.txt` and return the capitalized contents.

To avoid even having to define a named function, you can use the do syntax, which creates an anonymous function on the fly:

```julia
julia> open("hello.txt") do f
           uppercase(readstring(f))
       end
"HELLO AGAIN."
```

24.5   A simple TCP example

Let's jump right in with a simple example involving TCP sockets. Let's first create a simple server:

```julia
julia> @async begin
           server = listen(2000)
           while true
               sock = accept(server)
               println("Hello World\n")
           end
       end
Task (runnable) @0x00007fd31dc11ae0
```

To those familiar with the Unix socket API, the method names will feel familiar, though their usage is somewhat simpler than the raw Unix socket API. The first call to `listen()` will create a server waiting for incoming connections on the specified port (2000) in this case. The same function may also be used to create various other kinds of servers:

```julia
julia> listen(2000) # Listens on localhost:2000 (IPv4)
TCPServer(active)

julia> listen(ip"127.0.0.1",2000) # Equivalent to the first
TCPServer(active)

julia> listen(ip"::1",2000) # Listens on localhost:2000 (IPv6)
TCPServer(active)

julia> listen(IPv4(0),2001) # Listens on port 2001 on all IPv4 interfaces
TCPServer(active)

julia> listen(IPv6(0),2001) # Listens on port 2001 on all IPv6 interfaces
TCPServer(active)

julia> listen("testsocket") # Listens on a UNIX domain socket/named pipe
PipeServer(active)
```

Note that the return type of the last invocation is different. This is because this server does not listen on TCP, but rather on a named pipe (Windows) or UNIX domain socket. The difference is subtle and has to do with the `accept()` and `connect()` methods. The `accept()` method retrieves a connection to the client that is connecting on the server we just created, while the `connect()` function connects to a server using the specified method. The `connect()` function takes the same arguments as `listen()`, so, assuming the environment (i.e. host, cwd, etc.) is the same you should be able to pass the same arguments to `connect()` as you did to listen to establish the connection. So let's try that out (after having created the server above):

```julia
julia> connect(2000)
TCPSocket(open, 0 bytes waiting)

julia> Hello World
```

As expected we saw "Hello World" printed. So, let's actually analyze what happened behind the scenes. When we called `connect()`, we connect to the server we had just created. Meanwhile, the accept function returns a server-side connection to the newly created socket and prints "Hello World" to indicate that the connection was successful.

A great strength of Julia is that since the API is exposed synchronously even though the I/O is actually happening asynchronously, we didn't have to worry callbacks or even making sure that the server gets to run. When we called `connect()` the current task waited for the connection to be established and only continued executing after that was done. In this pause, the server task resumed execution (because a connection request was now available), accepted the connection, printed the message and waited for the next client. Reading and writing works in the same way. To see this, consider the following simple echo server:

```julia
julia> @async begin
           server = listen(2001)
           while true
               sock = accept(server)
               @async while isopen(sock)
                   write(sock,readline(sock))
               end
           end
       end
Task (runnable) @0x00007fd31dc12e60

julia> clientside = connect(2001)
TCPSocket(RawFD(28) open, 0 bytes waiting)

julia> @async while true
           write(STDOUT,readline(clientside))
       end
Task (runnable) @0x00007fd31dc11870

julia> println(clientside,"Hello World from the Echo Server")
Hello World from the Echo Server
```

As with other streams, use `close()` to disconnect the socket:

```julia
julia> close(clientside)
```

24.6 Resolving IP Addresses

One of the `connect()` methods that does not follow the `listen()` methods is `connect(host::String,port)`, which will attempt to connect to the host given by the `host` parameter on the port given by the port parameter. It allows you to do things like:

```julia
julia> connect("google.com",80)
TCPSocket(RawFD(30) open, 0 bytes waiting)
```

At the base of this functionality is `getaddrinfo()`, which will do the appropriate address resolution:

```julia
julia> getaddrinfo("google.com")
ip"74.125.226.225"
```

Chapter 25

Parallel Computing

Most modern computers possess more than one CPU, and several computers can be combined together in a cluster. Harnessing the power of these multiple CPUs allows many computations to be completed more quickly. There are two major factors that influence performance: the speed of the CPUs themselves, and the speed of their access to memory. In a cluster, it's fairly obvious that a given CPU will have fastest access to the RAM within the same computer (node). Perhaps more surprisingly, similar issues are relevant on a typical multicore laptop, due to differences in the speed of main memory and the cache. Consequently, a good multiprocessing environment should allow control over the "ownership" of a chunk of memory by a particular CPU. Julia provides a multiprocessing environment based on message passing to allow programs to run on multiple processes in separate memory domains at once.

Julia's implementation of message passing is different from other environments such as MPI [1]. Communication in Julia is generally "one-sided", meaning that the programmer needs to explicitly manage only one process in a two-process operation. Furthermore, these operations typically do not look like "message send" and "message receive" but rather resemble higher-level operations like calls to user functions.

Parallel programming in Julia is built on two primitives: *remote references* and *remote calls*. A remote reference is an object that can be used from any process to refer to an object stored on a particular process. A remote call is a request by one process to call a certain function on certain arguments on another (possibly the same) process.

Remote references come in two flavors: `Future` and `RemoteChannel`.

A remote call returns a `Future` to its result. Remote calls return immediately; the process that made the call proceeds to its next operation while the remote call happens somewhere else. You can wait for a remote call to finish by calling `wait()` on the returned `Future`, and you can obtain the full value of the result using `fetch()`.

On the other hand, `RemoteChannel`s are rewritable. For example, multiple processes can co-ordinate their processing by referencing the same remote `Channel`.

Each process has an associated identifier. The process providing the interactive Julia prompt always has an `id` equal to 1. The processes used by default for parallel operations are referred to as "workers". When there is only one process, process 1 is considered a worker. Otherwise, workers are considered to be all processes other than process 1.

Let's try this out. Starting with `julia -p n` provides n worker processes on the local machine. Generally it makes sense for n to equal the number of CPU cores on the machine.

```
$ ./julia -p 2

julia> r = remotecall(rand, 2, 2, 2)
Future(2, 1, 4, Nullable{Any}())

julia> s = @spawnat 2 1 .+ fetch(r)
Future(2, 1, 5, Nullable{Any}())
```

```
julia> fetch(s)
2×2 Array{Float64,2}:
 1.18526  1.50912
 1.16296  1.60607
```

The first argument to `remotecall()` is the function to call. Most parallel programming in Julia does not reference specific processes or the number of processes available, but `remotecall()` is considered a low-level interface providing finer control. The second argument to `remotecall()` is the id of the process that will do the work, and the remaining arguments will be passed to the function being called.

As you can see, in the first line we asked process 2 to construct a 2-by-2 random matrix, and in the second line we asked it to add 1 to it. The result of both calculations is available in the two futures, r and s. The `@spawnat` macro evaluates the expression in the second argument on the process specified by the first argument.

Occasionally you might want a remotely-computed value immediately. This typically happens when you read from a remote object to obtain data needed by the next local operation. The function `remotecall_fetch()` exists for this purpose. It is equivalent to `fetch(remotecall(...))` but is more efficient.

```
julia> remotecall_fetch(getindex, 2, r, 1, 1)
0.18526337335308085
```

Remember that `getindex(r,1,1)` is equivalent to `r[1,1]`, so this call fetches the first element of the future r.

The syntax of `remotecall()` is not especially convenient. The macro `@spawn` makes things easier. It operates on an expression rather than a function, and picks where to do the operation for you:

```
julia> r = @spawn rand(2,2)
Future(2, 1, 4, Nullable{Any}())

julia> s = @spawn 1 .+ fetch(r)
Future(3, 1, 5, Nullable{Any}())

julia> fetch(s)
2×2 Array{Float64,2}:
 1.38854  1.9098
 1.20939  1.57158
```

Note that we used `1 .+ fetch(r)` instead of `1 .+ r`. This is because we do not know where the code will run, so in general a `fetch()` might be required to move r to the process doing the addition. In this case, `@spawn` is smart enough to perform the computation on the process that owns r, so the `fetch()` will be a no-op (no work is done).

(It is worth noting that `@spawn` is not built-in but defined in Julia as a `macro`. It is possible to define your own such constructs.)

An important thing to remember is that, once fetched, a `Future` will cache its value locally. Further `fetch()` calls do not entail a network hop. Once all referencing `Future`s have fetched, the remote stored value is deleted.

25.1 Code Availability and Loading Packages

Your code must be available on any process that runs it. For example, type the following into the Julia prompt:

```julia
julia> function rand2(dims...)
           return 2*rand(dims...)
       end

julia> rand2(2,2)
2×2 Array{Float64,2}:
 0.153756  0.368514
 1.15119   0.918912

julia> fetch(@spawn rand2(2,2))
ERROR: RemoteException(2, CapturedException(UndefVarError(Symbol("#rand2"))
[...]
```

Process 1 knew about the function `rand2`, but process 2 did not.

Most commonly you'll be loading code from files or packages, and you have a considerable amount of flexibility in controlling which processes load code. Consider a file, `DummyModule.jl`, containing the following code:

```julia
module DummyModule

export MyType, f

mutable struct MyType
    a::Int
end

f(x) = x^2+1

println("loaded")

end
```

Starting Julia with `julia -p 2`, you can use this to verify the following:

- `include("DummyModule.jl")` loads the file on just a single process (whichever one executes the statement).

- `using DummyModule` causes the module to be loaded on all processes; however, the module is brought into scope only on the one executing the statement.

- As long as `DummyModule` is loaded on process 2, commands like

  ```julia
  rr = RemoteChannel(2)
  put!(rr, MyType(7))
  ```

  allow you to store an object of type `MyType` on process 2 even if `DummyModule` is not in scope on process 2.

You can force a command to run on all processes using the `@everywhere` macro. For example, `@everywhere` can also be used to directly define a function on all processes:

```julia
julia> @everywhere id = myid()

julia> remotecall_fetch(()->id, 2)
2
```

A file can also be preloaded on multiple processes at startup, and a driver script can be used to drive the computation:

```
julia -p <n> -L file1.jl -L file2.jl driver.jl
```

The Julia process running the driver script in the example above has an id equal to 1, just like a process providing an interactive prompt.

The base Julia installation has in-built support for two types of clusters:

- A local cluster specified with the -p option as shown above.

- A cluster spanning machines using the --machinefile option. This uses a passwordless ssh login to start Julia worker processes (from the same path as the current host) on the specified machines.

Functions addprocs(), rmprocs(), workers(), and others are available as a programmatic means of adding, removing and querying the processes in a cluster.

Note that workers do not run a .juliarc.jl startup script, nor do they synchronize their global state (such as global variables, new method definitions, and loaded modules) with any of the other running processes.

Other types of clusters can be supported by writing your own custom ClusterManager, as described below in the ClusterManagers section.

25.2   Data Movement

Sending messages and moving data constitute most of the overhead in a parallel program. Reducing the number of messages and the amount of data sent is critical to achieving performance and scalability. To this end, it is important to understand the data movement performed by Julia's various parallel programming constructs.

fetch() can be considered an explicit data movement operation, since it directly asks that an object be moved to the local machine. @spawn (and a few related constructs) also moves data, but this is not as obvious, hence it can be called an implicit data movement operation. Consider these two approaches to constructing and squaring a random matrix:

Method 1:

```
julia> A = rand(1000,1000);

julia> Bref = @spawn A^2;

[...]

julia> fetch(Bref);
```

Method 2:

```
julia> Bref = @spawn rand(1000,1000)^2;

[...]

julia> fetch(Bref);
```

The difference seems trivial, but in fact is quite significant due to the behavior of @spawn. In the first method, a random matrix is constructed locally, then sent to another process where it is squared. In the second method, a random matrix

is both constructed and squared on another process. Therefore the second method sends much less data than the first.

In this toy example, the two methods are easy to distinguish and choose from. However, in a real program designing data movement might require more thought and likely some measurement. For example, if the first process needs matrix A then the first method might be better. Or, if computing A is expensive and only the current process has it, then moving it to another process might be unavoidable. Or, if the current process has very little to do between the @spawn and `fetch(Bref)`, it might be better to eliminate the parallelism altogether. Or imagine `rand(1000,1000)` is replaced with a more expensive operation. Then it might make sense to add another @spawn statement just for this step.

Chapter 26

Global variables

Expressions executed remotely via @spawn, or closures specified for remote execution using `remotecall` may refer to global variables. Global bindings under module `Main` are treated a little differently compared to global bindings in other modules. Consider the following code snippet:

```
A = rand(10,10)
remotecall_fetch(()->foo(A), 2)
```

Note that A is a global variable defined in the local workspace. Worker 2 does not have a variable called A under `Main`. The act of shipping the closure `()->foo(A)` to worker 2 results in `Main.A` being defined on 2. `Main.A` continues to exist on worker 2 even after the call `remotecall_fetch` returns. Remote calls with embedded global references (under `Main` module only) manage globals as follows:

- New global bindings are created on destination workers if they are referenced as part of a remote call.

- Global constants are declared as constants on remote nodes too.

- Globals are re-sent to a destination worker only in the context of a remote call, and then only if its value has changed. Also, the cluster does not synchronize global bindings across nodes. For example:

  ```
  A = rand(10,10)
  remotecall_fetch(()->foo(A), 2) # worker 2
  A = rand(10,10)
  remotecall_fetch(()->foo(A), 3) # worker 3
  A = nothing
  ```

  Executing the above snippet results in `Main.A` on worker 2 having a different value from `Main.A` on worker 3, while the value of `Main.A` on node 1 is set to `nothing`.

As you may have realized, while memory associated with globals may be collected when they are reassigned on the master, no such action is taken on the workers as the bindings continue to be valid. `clear!` can be used to manually reassign specific globals on remote nodes to `nothing` once they are no longer required. This will release any memory associated with them as part of a regular garbage collection cycle.

Thus programs should be careful referencing globals in remote calls. In fact, it is preferable to avoid them altogether if possible. If you must reference globals, consider using `let` blocks to localize global variables.

For example:

```julia
julia> A = rand(10,10);

julia> remotecall_fetch(()->A, 2);

julia> B = rand(10,10);

julia> let B = B
           remotecall_fetch(()->B, 2)
       end;

julia> @spawnat 2 whos();

julia>   From worker 2:                          A    800 bytes  10×10 Array{Float64,2}
         From worker 2:                          Base            Module
         From worker 2:                          Core            Module
         From worker 2:                          Main            Module
```

As can be seen, global variable A is defined on worker 2, but B is captured as a local variable and hence a binding for B does not exist on worker 2.

26.1   Parallel Map and Loops

Fortunately, many useful parallel computations do not require data movement. A common example is a Monte Carlo simulation, where multiple processes can handle independent simulation trials simultaneously. We can use @spawn to flip coins on two processes. First, write the following function in count_heads.jl:

```julia
function count_heads(n)
    c::Int = 0
    for i = 1:n
        c += rand(Bool)
    end
    c
end
```

The function count_heads simply adds together n random bits. Here is how we can perform some trials on two machines, and add together the results:

```julia
julia> @everywhere include("count_heads.jl")

julia> a = @spawn count_heads(100000000)
Future(2, 1, 6, Nullable{Any}())

julia> b = @spawn count_heads(100000000)
Future(3, 1, 7, Nullable{Any}())

julia> fetch(a)+fetch(b)
100001564
```

This example demonstrates a powerful and often-used parallel programming pattern. Many iterations run independently over several processes, and then their results are combined using some function. The combination process is called a *reduction*, since it is generally tensor-rank-reducing: a vector of numbers is reduced to a single number, or a matrix is reduced to a single row or column, etc. In code, this typically looks like the pattern x = f(x,v[i]), where

x is the accumulator, f is the reduction function, and the v[i] are the elements being reduced. It is desirable for f to be associative, so that it does not matter what order the operations are performed in.

Notice that our use of this pattern with count_heads can be generalized. We used two explicit @spawn statements, which limits the parallelism to two processes. To run on any number of processes, we can use a *parallel for loop*, which can be written in Julia using @parallel like this:

```julia
nheads = @parallel (+) for i = 1:200000000
    Int(rand(Bool))
end
```

This construct implements the pattern of assigning iterations to multiple processes, and combining them with a specified reduction (in this case (+)). The result of each iteration is taken as the value of the last expression inside the loop. The whole parallel loop expression itself evaluates to the final answer.

Note that although parallel for loops look like serial for loops, their behavior is dramatically different. In particular, the iterations do not happen in a specified order, and writes to variables or arrays will not be globally visible since iterations run on different processes. Any variables used inside the parallel loop will be copied and broadcast to each process.

For example, the following code will not work as intended:

```julia
a = zeros(100000)
@parallel for i = 1:100000
    a[i] = i
end
```

This code will not initialize all of a, since each process will have a separate copy of it. Parallel for loops like these must be avoided. Fortunately, Shared Arrays can be used to get around this limitation:

```julia
a = SharedArray{Float64}(10)
@parallel for i = 1:10
    a[i] = i
end
```

Using "outside" variables in parallel loops is perfectly reasonable if the variables are read-only:

```julia
a = randn(1000)
@parallel (+) for i = 1:100000
    f(a[rand(1:end)])
end
```

Here each iteration applies f to a randomly-chosen sample from a vector a shared by all processes.

As you could see, the reduction operator can be omitted if it is not needed. In that case, the loop executes asynchronously, i.e. it spawns independent tasks on all available workers and returns an array of Future immediately without waiting for completion. The caller can wait for the Future completions at a later point by calling fetch() on them, or wait for completion at the end of the loop by prefixing it with @sync, like @sync @parallel for.

In some cases no reduction operator is needed, and we merely wish to apply a function to all integers in some range (or, more generally, to all elements in some collection). This is another useful operation called *parallel map*, implemented in Julia as the pmap() function. For example, we could compute the singular values of several large random matrices in parallel as follows:

```julia
julia> M = Matrix{Float64}[rand(1000,1000) for i = 1:10];

julia> pmap(svd, M);
```

Julia's pmap() is designed for the case where each function call does a large amount of work. In contrast, @parallel for can handle situations where each iteration is tiny, perhaps merely summing two numbers. Only worker processes are used by both pmap() and @parallel for for the parallel computation. In case of @parallel for, the final reduction is done on the calling process.

26.2   Synchronization With Remote References

26.3   Scheduling

Julia's parallel programming platform uses Tasks (aka Coroutines) to switch among multiple computations. Whenever code performs a communication operation like fetch() or wait(), the current task is suspended and a scheduler picks another task to run. A task is restarted when the event it is waiting for completes.

For many problems, it is not necessary to think about tasks directly. However, they can be used to wait for multiple events at the same time, which provides for *dynamic scheduling*. In dynamic scheduling, a program decides what to compute or where to compute it based on when other jobs finish. This is needed for unpredictable or unbalanced workloads, where we want to assign more work to processes only when they finish their current tasks.

As an example, consider computing the singular values of matrices of different sizes:

```julia
julia> M = Matrix{Float64}[rand(800,800), rand(600,600), rand(800,800), rand(600,600)];

julia> pmap(svd, M);
```

If one process handles both 800×800 matrices and another handles both 600×600 matrices, we will not get as much scalability as we could. The solution is to make a local task to "feed" work to each process when it completes its current task. For example, consider a simple pmap() implementation:

```julia
function pmap(f, lst)
    np = nprocs()  # determine the number of processes available
    n = length(lst)
    results = Vector{Any}(n)
    i = 1
    # function to produce the next work item from the queue.
    # in this case it's just an index.
    nextidx() = (idx=i; i+=1; idx)
    @sync begin
        for p=1:np
            if p != myid() || np == 1
                @async begin
                    while true
                        idx = nextidx()
                        if idx > n
                            break
                        end
                        results[idx] = remotecall_fetch(f, p, lst[idx])
                    end
                end
            end
        end
```

```
        end
    end
    results
end
```

@async is similar to @spawn, but only runs tasks on the local process. We use it to create a "feeder" task for each process. Each task picks the next index that needs to be computed, then waits for its process to finish, then repeats until we run out of indexes. Note that the feeder tasks do not begin to execute until the main task reaches the end of the @sync block, at which point it surrenders control and waits for all the local tasks to complete before returning from the function. The feeder tasks are able to share state via nextidx() because they all run on the same process. No locking is required, since the threads are scheduled cooperatively and not preemptively. This means context switches only occur at well-defined points: in this case, when remotecall_fetch() is called.

26.4   Channels

The section on Tasks in Control Flow discussed the execution of multiple functions in a co-operative manner. Channels can be quite useful to pass data between running tasks, particularly those involving I/O operations.

Examples of operations involving I/O include reading/writing to files, accessing web services, executing external programs, etc. In all these cases, overall execution time can be improved if other tasks can be run while a file is being read, or while waiting for an external service/program to complete.

A channel can be visualized as a pipe, i.e., it has a write end and read end.

- Multiple writers in different tasks can write to the same channel concurrently via put!() calls.

- Multiple readers in different tasks can read data concurrently via take!() calls.

- As an example:

```
# Given Channels c1 and c2,
c1 = Channel(32)
c2 = Channel(32)

# and a function `foo()` which reads items from from c1, processes the item read
# and writes a result to c2,
function foo()
    while true
        data = take!(c1)
        [...]                  # process data
        put!(c2, result)    # write out result
    end
end

# we can schedule `n` instances of `foo()` to be active concurrently.
for _ in 1:n
    @schedule foo()
end
```

- Channels are created via the Channel{T}(sz) constructor. The channel will only hold objects of type T. If the type is not specified, the channel can hold objects of any type. sz refers to the maximum number of elements that can be held in the channel at any time. For example, Channel(32) creates a channel that can hold a maximum of 32 objects of any type. A Channel{MyType}(64) can hold up to 64 objects of MyType at any time.

- If a `Channel` is empty, readers (on a `take!()` call) will block until data is available.

- If a `Channel` is full, writers (on a `put!()` call) will block until space becomes available.

- `isready()` tests for the presence of any object in the channel, while `wait()` waits for an object to become available.

- A `Channel` is in an open state initially. This means that it can be read from and written to freely via `take!()` and `put!()` calls. `close()` closes a `Channel`. On a closed `Channel`, `put!()` will fail. For example:

```julia
julia> c = Channel(2);

julia> put!(c, 1) # `put!` on an open channel succeeds
1

julia> close(c);

julia> put!(c, 2) # `put!` on a closed channel throws an exception.
ERROR: InvalidStateException("Channel is closed.",:closed)
[...]
```

- `take!()` and `fetch()` (which retrieves but does not remove the value) on a closed channel successfully return any existing values until it is emptied. Continuing the above example:

```julia
julia> fetch(c) # Any number of `fetch` calls succeed.
1

julia> fetch(c)
1

julia> take!(c) # The first `take!` removes the value.
1

julia> take!(c) # No more data available on a closed channel.
ERROR: InvalidStateException("Channel is closed.",:closed)
[...]
```

A `Channel` can be used as an iterable object in a `for` loop, in which case the loop runs as long as the `Channel` has data or is open. The loop variable takes on all values added to the `Channel`. The `for` loop is terminated once the `Channel` is closed and emptied.

For example, the following would cause the `for` loop to wait for more data:

```julia
julia> c = Channel{Int}(10);

julia> foreach(i->put!(c, i), 1:3) # add a few entries

julia> data = [i for i in c]
```

while this will return after reading all data:

```julia
julia> c = Channel{Int}(10);

julia> foreach(i->put!(c, i), 1:3); # add a few entries
```

```julia
julia> close(c);                      # `for` loops can exit

julia> data = [i for i in c]
3-element Array{Int64,1}:
 1
 2
 3
```

Consider a simple example using channels for inter-task communication. We start 4 tasks to process data from a single `jobs` channel. Jobs, identified by an id (`job_id`), are written to the channel. Each task in this simulation reads a `job_id`, waits for a random amout of time and writes back a tuple of `job_id` and the simulated time to the results channel. Finally all the `results` are printed out.

```julia
julia> const jobs = Channel{Int}(32);

julia> const results = Channel{Tuple}(32);

julia> function do_work()
           for job_id in jobs
               exec_time = rand()
               sleep(exec_time)                  # simulates elapsed time doing actual work
                                                 # typically performed externally.
               put!(results, (job_id, exec_time))
           end
       end;

julia> function make_jobs(n)
           for i in 1:n
               put!(jobs, i)
           end
       end;

julia> n = 12;

julia> @schedule make_jobs(n); # feed the jobs channel with "n" jobs

julia> for i in 1:4 # start 4 tasks to process requests in parallel
           @schedule do_work()
       end

julia> @elapsed while n > 0 # print out results
           job_id, exec_time = take!(results)
           println("$job_id finished in $(round(exec_time,2)) seconds")
           n = n - 1
       end
4 finished in 0.22 seconds
3 finished in 0.45 seconds
1 finished in 0.5 seconds
7 finished in 0.14 seconds
2 finished in 0.78 seconds
5 finished in 0.9 seconds
9 finished in 0.36 seconds
6 finished in 0.87 seconds
8 finished in 0.79 seconds
```

```
10 finished in 0.64 seconds
12 finished in 0.5 seconds
11 finished in 0.97 seconds
0.029772311
```

The current version of Julia multiplexes all tasks onto a single OS thread. Thus, while tasks involving I/O operations benefit from parallel execution, compute bound tasks are effectively executed sequentially on a single OS thread. Future versions of Julia may support scheduling of tasks on multiple threads, in which case compute bound tasks will see benefits of parallel execution too.

26.5   Remote References and AbstractChannels

Remote references always refer to an implementation of an `AbstractChannel`.

A concrete implementation of an `AbstractChannel` (like Channel), is required to implement `put!()`, `take!()`, `fetch()`, `isready()` and `wait()`. The remote object referred to by a `Future` is stored in a `Channel{Any}(1)`, i.e., a `Channel` of size 1 capable of holding objects of Any type.

`RemoteChannel`, which is rewritable, can point to any type and size of channels, or any other implementation of an `AbstractChannel`.

The constructor `RemoteChannel(f::Function, pid)()` allows us to construct references to channels holding more than one value of a specific type. `f()` is a function executed on `pid` and it must return an `AbstractChannel`.

For example, `RemoteChannel(()->Channel{Int}(10), pid)`, will return a reference to a channel of type `Int` and size 10. The channel exists on worker `pid`.

Methods `put!()`, `take!()`, `fetch()`, `isready()` and `wait()` on a `RemoteChannel` are proxied onto the backing store on the remote process.

`RemoteChannel` can thus be used to refer to user implemented `AbstractChannel` objects. A simple example of this is provided in `examples/dictchannel.jl` which uses a dictionary as its remote store.

26.6   Channels and RemoteChannels

- A `Channel` is local to a process. Worker 2 cannot directly refer to a `Channel` on worker 3 and vice-versa. A `RemoteChannel`, however, can put and take values across workers.

- A `RemoteChannel` can be thought of as a *handle* to a `Channel`.

- The process id, `pid`, associated with a `RemoteChannel` identifies the process where the backing store, i.e., the backing `Channel` exists.

- Any process with a reference to a `RemoteChannel` can put and take items from the channel. Data is automatically sent to (or retrieved from) the process a `RemoteChannel` is associated with.

- Serializing a `Channel` also serializes any data present in the channel. Deserializing it therefore effectively makes a copy of the original object.

- On the other hand, serializing a `RemoteChannel` only involves the serialization of an identifier that identifies the location and instance of `Channel` referred to by the handle. A deserialized `RemoteChannel` object (on any worker), therefore also points to the same backing store as the original.

The channels example from above can be modified for interprocess communication, as shown below.

We start 4 workers to process a single jobs remote channel. Jobs, identified by an id (job_id), are written to the channel. Each remotely executing task in this simulation reads a job_id, waits for a random amount of time and writes back a tuple of job_id, time taken and its own pid to the results channel. Finally all the results are printed out on the master process.

```julia
julia> addprocs(4); # add worker processes

julia> const jobs = RemoteChannel(()->Channel{Int}(32));

julia> const results = RemoteChannel(()->Channel{Tuple}(32));

julia> @everywhere function do_work(jobs, results) # define work function everywhere
           while true
               job_id = take!(jobs)
               exec_time = rand()
               sleep(exec_time) # simulates elapsed time doing actual work
               put!(results, (job_id, exec_time, myid()))
           end
       end

julia> function make_jobs(n)
           for i in 1:n
               put!(jobs, i)
           end
       end;

julia> n = 12;

julia> @schedule make_jobs(n); # feed the jobs channel with "n" jobs

julia> for p in workers() # start tasks on the workers to process requests in parallel
           @async remote_do(do_work, p, jobs, results)
       end

julia> @elapsed while n > 0 # print out results
           job_id, exec_time, where = take!(results)
           println("$job_id finished in $(round(exec_time,2)) seconds on worker $where")
           n = n - 1
       end
1 finished in 0.18 seconds on worker 4
2 finished in 0.26 seconds on worker 5
6 finished in 0.12 seconds on worker 4
7 finished in 0.18 seconds on worker 4
5 finished in 0.35 seconds on worker 5
4 finished in 0.68 seconds on worker 2
3 finished in 0.73 seconds on worker 3
11 finished in 0.01 seconds on worker 3
12 finished in 0.02 seconds on worker 3
9 finished in 0.26 seconds on worker 5
8 finished in 0.57 seconds on worker 4
10 finished in 0.58 seconds on worker 2
0.055971741
```

26.7   Remote References and Distributed Garbage Collection

Objects referred to by remote references can be freed only when *all* held references in the cluster are deleted.

The node where the value is stored keeps track of which of the workers have a reference to it. Every time a `RemoteChannel` or a (unfetched) `Future` is serialized to a worker, the node pointed to by the reference is notified. And every time a `RemoteChannel` or a (unfetched) `Future` is garbage collected locally, the node owning the value is again notified.

The notifications are done via sending of "tracking" messages–an "add reference" message when a reference is serialized to a different process and a "delete reference" message when a reference is locally garbage collected.

Since `Future`s are write-once and cached locally, the act of `fetch()`ing a `Future` also updates reference tracking information on the node owning the value.

The node which owns the value frees it once all references to it are cleared.

With `Future`s, serializing an already fetched `Future` to a different node also sends the value since the original remote store may have collected the value by this time.

It is important to note that *when* an object is locally garbage collected depends on the size of the object and the current memory pressure in the system.

In case of remote references, the size of the local reference object is quite small, while the value stored on the remote node may be quite large. Since the local object may not be collected immediately, it is a good practice to explicitly call `finalize()` on local instances of a `RemoteChannel`, or on unfetched `Future`s. Since calling `fetch()` on a `Future` also removes its reference from the remote store, this is not required on fetched `Future`s. Explicitly calling `finalize()` results in an immediate message sent to the remote node to go ahead and remove its reference to the value.

Once finalized, a reference becomes invalid and cannot be used in any further calls.

26.8   Shared Arrays

Shared Arrays use system shared memory to map the same array across many processes. While there are some similarities to a `DArray`, the behavior of a `SharedArray` is quite different. In a `DArray`, each process has local access to just a chunk of the data, and no two processes share the same chunk; in contrast, in a `SharedArray` each "participating" process has access to the entire array. A `SharedArray` is a good choice when you want to have a large amount of data jointly accessible to two or more processes on the same machine.

`SharedArray` indexing (assignment and accessing values) works just as with regular arrays, and is efficient because the underlying memory is available to the local process. Therefore, most algorithms work naturally on `SharedArray`s, albeit in single-process mode. In cases where an algorithm insists on an `Array` input, the underlying array can be retrieved from a `SharedArray` by calling `sdata()`. For other `AbstractArray` types, `sdata()` just returns the object itself, so it's safe to use `sdata()` on any `Array`-type object.

The constructor for a shared array is of the form:

```
SharedArray{T,N}(dims::NTuple; init=false, pids=Int[])
```

which creates an N-dimensional shared array of a bits type T and size `dims` across the processes specified by `pids`. Unlike distributed arrays, a shared array is accessible only from those participating workers specified by the `pids` named argument (and the creating process too, if it is on the same host).

If an `init` function, of signature `initfn(S::SharedArray)`, is specified, it is called on all the participating workers. You can specify that each worker runs the `init` function on a distinct portion of the array, thereby parallelizing initialization.

Here's a brief example:

```julia
julia> addprocs(3)
3-element Array{Int64,1}:
 2
 3
 4

julia> S = SharedArray{Int,2}((3,4), init = S -> S[Base.localindexes(S)] = myid())
3×4 SharedArray{Int64,2}:
 2  2  3  4
 2  3  3  4
 2  3  4  4

julia> S[3,2] = 7
7

julia> S
3×4 SharedArray{Int64,2}:
 2  2  3  4
 2  3  3  4
 2  7  4  4
```

`Base.localindexes()` provides disjoint one-dimensional ranges of indexes, and is sometimes convenient for splitting up tasks among processes. You can, of course, divide the work any way you wish:

```julia
julia> S = SharedArray{Int,2}((3,4), init = S -> S[indexpids(S):length(procs(S)):length(S)] =
↪ myid())
3×4 SharedArray{Int64,2}:
 2  2  2  2
 3  3  3  3
 4  4  4  4
```

Since all processes have access to the underlying data, you do have to be careful not to set up conflicts. For example:

```julia
@sync begin
    for p in procs(S)
        @async begin
            remotecall_wait(fill!, p, S, p)
        end
    end
end
```

would result in undefined behavior. Because each process fills the *entire* array with its own `pid`, whichever process is the last to execute (for any particular element of S) will have its `pid` retained.

As a more extended and complex example, consider running the following "kernel" in parallel:

```julia
q[i,j,t+1] = q[i,j,t] + u[i,j,t]
```

In this case, if we try to split up the work using a one-dimensional index, we are likely to run into trouble: if `q[i,j,t]` is near the end of the block assigned to one worker and `q[i,j,t+1]` is near the beginning of the block assigned to another, it's very likely that `q[i,j,t]` will not be ready at the time it's needed for computing `q[i,j,t+1]`. In such cases, one is better off chunking the array manually. Let's split along the second dimension. Define a function that returns the (`irange`, `jrange`) indexes assigned to this worker:

```julia
julia> @everywhere function myrange(q::SharedArray)
           idx = indexpids(q)
           if idx == 0 # This worker is not assigned a piece
               return 1:0, 1:0
           end
           nchunks = length(procs(q))
           splits = [round(Int, s) for s in linspace(0,size(q,2),nchunks+1)]
           1:size(q,1), splits[idx]+1:splits[idx+1]
       end
```

Next, define the kernel:

```julia
julia> @everywhere function advection_chunk!(q, u, irange, jrange, trange)
           @show (irange, jrange, trange)  # display so we can see what's happening
           for t in trange, j in jrange, i in irange
               q[i,j,t+1] = q[i,j,t] + u[i,j,t]
           end
           q
       end
```

We also define a convenience wrapper for a `SharedArray` implementation

```julia
julia> @everywhere advection_shared_chunk!(q, u) =
           advection_chunk!(q, u, myrange(q)..., 1:size(q,3)-1)
```

Now let's compare three different versions, one that runs in a single process:

```julia
julia> advection_serial!(q, u) = advection_chunk!(q, u, 1:size(q,1), 1:size(q,2), 1:size(q,3)-1);
```

one that uses `@parallel`:

```julia
julia> function advection_parallel!(q, u)
           for t = 1:size(q,3)-1
               @sync @parallel for j = 1:size(q,2)
                   for i = 1:size(q,1)
                       q[i,j,t+1]= q[i,j,t] + u[i,j,t]
                   end
               end
           end
           q
       end;
```

and one that delegates in chunks:

```julia
julia> function advection_shared!(q, u)
           @sync begin
               for p in procs(q)
                   @async remotecall_wait(advection_shared_chunk!, p, q, u)
               end
           end
           q
       end;
```

If we create `SharedArrays` and time these functions, we get the following results (with `julia -p 4`):

```julia
julia> q = SharedArray{Float64,3}((500,500,500));

julia> u = SharedArray{Float64,3}((500,500,500));
```

Run the functions once to JIT-compile and `@time` them on the second run:

```julia
julia> @time advection_serial!(q, u);
(irange,jrange,trange) = (1:500,1:500,1:499)
 830.220 milliseconds (216 allocations: 13820 bytes)

julia> @time advection_parallel!(q, u);
  2.495 seconds       (3999 k allocations: 289 MB, 2.09% gc time)

julia> @time advection_shared!(q,u);
        From worker 2:        (irange,jrange,trange) = (1:500,1:125,1:499)
        From worker 4:        (irange,jrange,trange) = (1:500,251:375,1:499)
        From worker 3:        (irange,jrange,trange) = (1:500,126:250,1:499)
        From worker 5:        (irange,jrange,trange) = (1:500,376:500,1:499)
 238.119 milliseconds (2264 allocations: 169 KB)
```

The biggest advantage of `advection_shared!` is that it minimizes traffic among the workers, allowing each to compute for an extended time on the assigned piece.

26.9 Shared Arrays and Distributed Garbage Collection

Like remote references, shared arrays are also dependent on garbage collection on the creating node to release references from all participating workers. Code which creates many short lived shared array objects would benefit from explicitly finalizing these objects as soon as possible. This results in both memory and file handles mapping the shared segment being released sooner.

26.10 ClusterManagers

The launching, management and networking of Julia processes into a logical cluster is done via cluster managers. A `ClusterManager` is responsible for

- launching worker processes in a cluster environment

- managing events during the lifetime of each worker

- optionally, providing data transport

A Julia cluster has the following characteristics:

- The initial Julia process, also called the `master`, is special and has an `id` of 1.

- Only the `master` process can add or remove worker processes.

- All processes can directly communicate with each other.

Connections between workers (using the in-built TCP/IP transport) is established in the following manner:

- `addprocs()` is called on the master process with a `ClusterManager` object.

- `addprocs()` calls the appropriate `launch()` method which spawns required number of worker processes on appropriate machines.

- Each worker starts listening on a free port and writes out its host and port information to `STDOUT`.

- The cluster manager captures the `STDOUT` of each worker and makes it available to the master process.

- The master process parses this information and sets up TCP/IP connections to each worker.

- Every worker is also notified of other workers in the cluster.

- Each worker connects to all workers whose `id` is less than the worker's own `id`.

- In this way a mesh network is established, wherein every worker is directly connected with every other worker.

While the default transport layer uses plain `TCPSocket`, it is possible for a Julia cluster to provide its own transport.

Julia provides two in-built cluster managers:

- `LocalManager`, used when `addprocs()` or `addprocs(np::Integer)` are called

- `SSHManager`, used when `addprocs(hostnames::Array)` is called with a list of hostnames

`LocalManager` is used to launch additional workers on the same host, thereby leveraging multi-core and multi-processor hardware.

Thus, a minimal cluster manager would need to:

- be a subtype of the abstract `ClusterManager`

- implement `launch()`, a method responsible for launching new workers

- implement `manage()`, which is called at various events during a worker's lifetime (for example, sending an interrupt signal)

`addprocs(manager::FooManager)` requires FooManager to implement:

```julia
function launch(manager::FooManager, params::Dict, launched::Array, c::Condition)
    [...]
end

function manage(manager::FooManager, id::Integer, config::WorkerConfig, op::Symbol)
    [...]
end
```

As an example let us see how the `LocalManager`, the manager responsible for starting workers on the same host, is implemented:

```julia
struct LocalManager <: ClusterManager
    np::Integer
end

function launch(manager::LocalManager, params::Dict, launched::Array, c::Condition)
    [...]
end

function manage(manager::LocalManager, id::Integer, config::WorkerConfig, op::Symbol)
    [...]
end
```

The `launch()` method takes the following arguments:

- `manager::ClusterManager`: the cluster manager that `addprocs()` is called with

- `params::Dict`: all the keyword arguments passed to `addprocs()`

- `launched::Array`: the array to append one or more `WorkerConfig` objects to

- `c::Condition`: the condition variable to be notified as and when workers are launched

The `launch()` method is called asynchronously in a separate task. The termination of this task signals that all requested workers have been launched. Hence the `launch()` function MUST exit as soon as all the requested workers have been launched.

Newly launched workers are connected to each other, and the master process, in an all-to-all manner. Specifying the command argument `--worker <cookie>` results in the launched processes initializing themselves as workers and connections being set up via TCP/IP sockets. Optionally, `--bind-to bind_addr[:port]` may also be specified to enable other workers to connect to it at the specified `bind_addr` and `port`. This is useful for multi-homed hosts.

As an example of a non-TCP/IP transport, an implementation may choose to use MPI, in which case `--worker` must NOT be specified. Instead, newly launched workers should call `init_worker(cookie)` before using any of the parallel constructs.

For every worker launched, the `launch()` method must add a `WorkerConfig` object (with appropriate fields initialized) to `launched`

```julia
mutable struct WorkerConfig
    # Common fields relevant to all cluster managers
    io::Nullable{IO}
    host::Nullable{AbstractString}
    port::Nullable{Integer}

    # Used when launching additional workers at a host
    count::Nullable{Union{Int, Symbol}}
    exename::Nullable{AbstractString}
    exeflags::Nullable{Cmd}

    # External cluster managers can use this to store information at a per-worker level
    # Can be a dict if multiple fields need to be stored.
    userdata::Nullable{Any}

    # SSHManager / SSH tunnel connections to workers
    tunnel::Nullable{Bool}
```

```
    bind_addr::Nullable{AbstractString}
    sshflags::Nullable{Cmd}
    max_parallel::Nullable{Integer}

    connect_at::Nullable{Any}

    [...]
end
```

Most of the fields in `WorkerConfig` are used by the inbuilt managers. Custom cluster managers would typically specify only `io` or `host` / `port`:

- If `io` is specified, it is used to read host/port information. A Julia worker prints out its bind address and port at startup. This allows Julia workers to listen on any free port available instead of requiring worker ports to be configured manually.

- If `io` is not specified, `host` and `port` are used to connect.

- `count`, `exename` and `exeflags` are relevant for launching additional workers from a worker. For example, a cluster manager may launch a single worker per node, and use that to launch additional workers.

    - `count` with an integer value n will launch a total of n workers.
    - `count` with a value of `:auto` will launch as many workers as the number of cores on that machine.
    - `exename` is the name of the `julia` executable including the full path.
    - `exeflags` should be set to the required command line arguments for new workers.

- `tunnel`, `bind_addr`, `sshflags` and `max_parallel` are used when a ssh tunnel is required to connect to the workers from the master process.

- `userdata` is provided for custom cluster managers to store their own worker-specific information.

`manage(manager::FooManager, id::Integer, config::WorkerConfig, op::Symbol)` is called at different times during the worker's lifetime with appropriate op values:

- with `:register`/`:deregister` when a worker is added / removed from the Julia worker pool.

- with `:interrupt` when `interrupt(workers)` is called. The `ClusterManager` should signal the appropriate worker with an interrupt signal.

- with `:finalize` for cleanup purposes.

26.11   Cluster Managers with Custom Transports

Replacing the default TCP/IP all-to-all socket connections with a custom transport layer is a little more involved. Each Julia process has as many communication tasks as the workers it is connected to. For example, consider a Julia cluster of 32 processes in an all-to-all mesh network:

- Each Julia process thus has 31 communication tasks.

- Each task handles all incoming messages from a single remote worker in a message-processing loop.

- The message-processing loop waits on an `IO` object (for example, a `TCPSocket` in the default implementation), reads an entire message, processes it and waits for the next one.

- Sending messages to a process is done directly from any Julia task–not just communication tasks–again, via the appropriate `IO` object.

Replacing the default transport requires the new implementation to set up connections to remote workers and to provide appropriate `IO` objects that the message-processing loops can wait on. The manager-specific callbacks to be implemented are:

```
connect(manager::FooManager, pid::Integer, config::WorkerConfig)
kill(manager::FooManager, pid::Int, config::WorkerConfig)
```

The default implementation (which uses TCP/IP sockets) is implemented as `connect(manager::ClusterManager, pid::Integer, config::WorkerConfig)`.

`connect` should return a pair of `IO` objects, one for reading data sent from worker `pid`, and the other to write data that needs to be sent to worker `pid`. Custom cluster managers can use an in-memory `BufferStream` as the plumbing to proxy data between the custom, possibly non-`IO` transport and Julia's in-built parallel infrastructure.

A `BufferStream` is an in-memory `IOBuffer` which behaves like an `IO`–it is a stream which can be handled asynchronously.

Folder `examples/clustermanager/0mq` contains an example of using ZeroMQ to connect Julia workers in a star topology with a 0MQ broker in the middle. Note: The Julia processes are still all *logically* connected to each other–any worker can message any other worker directly without any awareness of 0MQ being used as the transport layer.

When using custom transports:

- Julia workers must NOT be started with `--worker`. Starting with `--worker` will result in the newly launched workers defaulting to the TCP/IP socket transport implementation.

- For every incoming logical connection with a worker, `Base.process_messages(rd::IO, wr::IO)()` must be called. This launches a new task that handles reading and writing of messages from/to the worker represented by the `IO` objects.

- `init_worker(cookie, manager::FooManager)` MUST be called as part of worker process initialization.

- Field `connect_at::Any` in `WorkerConfig` can be set by the cluster manager when `launch()` is called. The value of this field is passed in all `connect()` callbacks. Typically, it carries information on *how to connect* to a worker. For example, the TCP/IP socket transport uses this field to specify the (`host, port`) tuple at which to connect to a worker.

`kill(manager, pid, config)` is called to remove a worker from the cluster. On the master process, the corresponding `IO` objects must be closed by the implementation to ensure proper cleanup. The default implementation simply executes an `exit()` call on the specified remote worker.

`examples/clustermanager/simple` is an example that shows a simple implementation using UNIX domain sockets for cluster setup.

26.12 Network Requirements for LocalManager and SSHManager

Julia clusters are designed to be executed on already secured environments on infrastructure such as local laptops, departmental clusters, or even the cloud. This section covers network security requirements for the inbuilt `LocalManager` and `SSHManager`:

- The master process does not listen on any port. It only connects out to the workers.

- Each worker binds to only one of the local interfaces and listens on the first free port starting from `9009`.

- `LocalManager`, used by `addprocs(N)`, by default binds only to the loopback interface. This means that workers started later on remote hosts (or by anyone with malicious intentions) are unable to connect to the cluster. An `addprocs(4)` followed by an `addprocs(["remote_host"])` will fail. Some users may need to create a cluster comprising their local system and a few remote systems. This can be done by explicitly requesting `LocalManager` to bind to an external network interface via the `restrict` keyword argument: `addprocs(4; restrict=false)`.

- `SSHManager`, used by `addprocs(list_of_remote_hosts)`, launches workers on remote hosts via SSH. By default SSH is only used to launch Julia workers. Subsequent master-worker and worker-worker connections use plain, unencrypted TCP/IP sockets. The remote hosts must have passwordless login enabled. Additional SSH flags or credentials may be specified via keyword argument `sshflags`.

- `addprocs(list_of_remote_hosts; tunnel=true, sshflags=<ssh keys and other flags>)` is useful when we wish to use SSH connections for master-worker too. A typical scenario for this is a local laptop running the Julia REPL (i.e., the master) with the rest of the cluster on the cloud, say on Amazon EC2. In this case only port 22 needs to be opened at the remote cluster coupled with SSH client authenticated via public key infrastructure (PKI). Authentication credentials can be supplied via `sshflags`, for example `sshflags=`-e <keyfile>``.

  Note that worker-worker connections are still plain TCP and the local security policy on the remote cluster must allow for free connections between worker nodes, at least for ports 9009 and above.

  Securing and encrypting all worker-worker traffic (via SSH) or encrypting individual messages can be done via a custom ClusterManager.

26.13   Cluster Cookie

All processes in a cluster share the same cookie which, by default, is a randomly generated string on the master process:

- `Base.cluster_cookie()` returns the cookie, while `Base.cluster_cookie(cookie)()` sets it and returns the new cookie.

- All connections are authenticated on both sides to ensure that only workers started by the master are allowed to connect to each other.

- The cookie must be passed to the workers at startup via argument `--worker <cookie>`. Custom ClusterManagers can retrieve the cookie on the master by calling `Base.cluster_cookie()`. Cluster managers not using the default TCP/IP transport (and hence not specifying `--worker`) must call `init_worker(cookie, manager)` with the same cookie as on the master.

Note that environments requiring higher levels of security can implement this via a custom `ClusterManager`. For example, cookies can be pre-shared and hence not specified as a startup argument.

26.14   Specifying Network Topology (Experimental)

The keyword argument `topology` passed to `addprocs` is used to specify how the workers must be connected to each other:

- `:all_to_all`, the default: all workers are connected to each other.

- `:master_slave`: only the driver process, i.e. `pid 1`, has connections to the workers.

- `:custom`: the `launch` method of the cluster manager specifies the connection topology via the fields `ident` and `connect_idents` in `WorkerConfig`. A worker with a cluster-manager-provided identity `ident` will connect to all workers specified in `connect_idents`.

Currently, sending a message between unconnected workers results in an error. This behaviour, as with the functionality and interface, should be considered experimental in nature and may change in future releases.

26.15 Multi-Threading (Experimental)

In addition to tasks, remote calls, and remote references, Julia from v0.5 forwards will natively support multi-threading. Note that this section is experimental and the interfaces may change in the future.

Setup

By default, Julia starts up with a single thread of execution. This can be verified by using the command `Threads.nthreads()`:

```
julia> Threads.nthreads()
1
```

The number of threads Julia starts up with is controlled by an environment variable called JULIA_NUM_THREADS. Now, let's start up Julia with 4 threads:

```
export JULIA_NUM_THREADS=4
```

(The above command works on bourne shells on Linux and OSX. Note that if you're using a C shell on these platforms, you should use the keyword `set` instead of `export`. If you're on Windows, start up the command line in the location of `julia.exe` and use `set` instead of `export`.)

Let's verify there are 4 threads at our disposal.

```
julia> Threads.nthreads()
4
```

But we are currently on the master thread. To check, we use the command `Threads.threadid()`

```
julia> Threads.threadid()
1
```

The `@threads` Macro

Let's work a simple example using our native threads. Let us create an array of zeros:

```
julia> a = zeros(10)
10-element Array{Float64,1}:
 0.0
 0.0
 0.0
 0.0
 0.0
 0.0
 0.0
 0.0
 0.0
 0.0
```

Let us operate on this array simultaneously using 4 threads. We'll have each thread write its thread ID into each location.

Julia supports parallel loops using the `Threads.@threads` macro. This macro is affixed in front of a `for` loop to indicate to Julia that the loop is a multi-threaded region:

```julia
julia> Threads.@threads for i = 1:10
           a[i] = Threads.threadid()
       end
```

The iteration space is split amongst the threads, after which each thread writes its thread ID to its assigned locations:

```julia
julia> a
10-element Array{Float64,1}:
 1.0
 1.0
 1.0
 2.0
 2.0
 2.0
 3.0
 3.0
 4.0
 4.0
```

Note that `Threads.@threads` does not have an optional reduction parameter like `@parallel`.

26.16   @threadcall (Experimental)

All I/O tasks, timers, REPL commands, etc are multiplexed onto a single OS thread via an event loop. A patched version of libuv (http://docs.libuv.org/en/v1.x/) provides this functionality. Yield points provide for co-operatively scheduling multiple tasks onto the same OS thread. I/O tasks and timers yield implicitly while waiting for the event to occur. Calling `yield()` explicitly allows for other tasks to be scheduled.

Thus, a task executing a `ccall` effectively prevents the Julia scheduler from executing any other tasks till the call returns. This is true for all calls into external libraries. Exceptions are calls into custom C code that call back into Julia (which may then yield) or C code that calls `jl_yield()` (C equivalent of `yield()`).

Note that while Julia code runs on a single thread (by default), libraries used by Julia may launch their own internal threads. For example, the BLAS library may start as many threads as there are cores on a machine.

The `@threadcall` macro addresses scenarios where we do not want a `ccall` to block the main Julia event loop. It schedules a C function for execution in a separate thread. A threadpool with a default size of 4 is used for this. The size of the threadpool is controlled via environment variable `UV_THREADPOOL_SIZE`. While waiting for a free thread, and during function execution once a thread is available, the requesting task (on the main Julia event loop) yields to other tasks. Note that `@threadcall` does not return till the execution is complete. From a user point of view, it is therefore a blocking call like other Julia APIs.

It is very important that the called function does not call back into Julia.

`@threadcall` may be removed/changed in future versions of Julia.

---

[1] In this context, MPI refers to the MPI-1 standard. Beginning with MPI-2, the MPI standards committee introduced a new set of communication mechanisms, collectively referred to as Remote Memory Access (RMA). The motivation for adding RMA to the MPI standard was to facilitate one-sided communication patterns. For additional information on the latest MPI standard, see http://mpi-forum.org/docs.

Chapter 27

Date and DateTime

The `Dates` module provides two types for working with dates: `Date` and `DateTime`, representing day and millisecond precision, respectively; both are subtypes of the abstract `TimeType`. The motivation for distinct types is simple: some operations are much simpler, both in terms of code and mental reasoning, when the complexities of greater precision don't have to be dealt with. For example, since the `Date` type only resolves to the precision of a single date (i.e. no hours, minutes, or seconds), normal considerations for time zones, daylight savings/summer time, and leap seconds are unnecessary and avoided.

Both `Date` and `DateTime` are basically immutable `Int64` wrappers. The single `instant` field of either type is actually a `UTInstant{P}` type, which represents a continuously increasing machine timeline based on the UT second [1]. The `DateTime` type is not aware of time zones (*naive*, in Python parlance), analogous to a *LocalDateTime* in Java 8. Additional time zone functionality can be added through the TimeZones.jl package, which compiles the IANA time zone database. Both `Date` and `DateTime` are based on the ISO 8601 standard, which follows the proleptic Gregorian calendar. One note is that the ISO 8601 standard is particular about BC/BCE dates. In general, the last day of the BC/BCE era, 1-12-31 BC/BCE, was followed by 1-1-1 AD/CE, thus no year zero exists. The ISO standard, however, states that 1 BC/BCE is year zero, so `0000-12-31` is the day before `0001-01-01`, and year `-0001` (yes, negative one for the year) is 2 BC/BCE, year `-0002` is 3 BC/BCE, etc.

27.1 Constructors

`Date` and `DateTime` types can be constructed by integer or `Period` types, by parsing, or through adjusters (more on those later):

```julia
julia> DateTime(2013)
2013-01-01T00:00:00

julia> DateTime(2013,7)
2013-07-01T00:00:00

julia> DateTime(2013,7,1)
2013-07-01T00:00:00
```

---

[1] The notion of the UT second is actually quite fundamental. There are basically two different notions of time generally accepted, one based on the physical rotation of the earth (one full rotation = 1 day), the other based on the SI second (a fixed, constant value). These are radically different! Think about it, a "UT second", as defined relative to the rotation of the earth, may have a different absolute length depending on the day! Anyway, the fact that `Date` and `DateTime` are based on UT seconds is a simplifying, yet honest assumption so that things like leap seconds and all their complexity can be avoided. This basis of time is formally called UT or UT1. Basing types on the UT second basically means that every minute has 60 seconds and every day has 24 hours and leads to more natural calculations when working with calendar dates.

```julia
julia> DateTime(2013,7,1,12)
2013-07-01T12:00:00

julia> DateTime(2013,7,1,12,30)
2013-07-01T12:30:00

julia> DateTime(2013,7,1,12,30,59)
2013-07-01T12:30:59

julia> DateTime(2013,7,1,12,30,59,1)
2013-07-01T12:30:59.001

julia> Date(2013)
2013-01-01

julia> Date(2013,7)
2013-07-01

julia> Date(2013,7,1)
2013-07-01

julia> Date(Dates.Year(2013),Dates.Month(7),Dates.Day(1))
2013-07-01

julia> Date(Dates.Month(7),Dates.Year(2013))
2013-07-01
```

Date or DateTime parsing is accomplished by the use of format strings. Format strings work by the notion of defining *delimited* or *fixed-width* "slots" that contain a period to parse and passing the text to parse and format string to a Date or DateTime constructor, of the form Date("2015-01-01","y-m-d") or DateTime("20150101","yyyymmdd").

Delimited slots are marked by specifying the delimiter the parser should expect between two subsequent periods; so "y-m-d" lets the parser know that between the first and second slots in a date string like "2014-07-16", it should find the – character. The y, m, and d characters let the parser know which periods to parse in each slot.

Fixed-width slots are specified by repeating the period character the number of times corresponding to the width with no delimiter between characters. So "yyyymmdd" would correspond to a date string like "20140716". The parser distinguishes a fixed-width slot by the absence of a delimiter, noting the transition "yyyymm" from one period character to the next.

Support for text-form month parsing is also supported through the u and U characters, for abbreviated and full-length month names, respectively. By default, only English month names are supported, so u corresponds to "Jan", "Feb", "Mar", etc. And U corresponds to "January", "February", "March", etc. Similar to other name=>value mapping functions dayname() and monthname(), custom locales can be loaded by passing in the locale=>Dict{String,Int} mapping to the MONTHTOVALUEABBR and MONTHTOVALUE dicts for abbreviated and full-name month names, respectively.

One note on parsing performance: using the Date(date_string, format_string) function is fine if only called a few times. If there are many similarly formatted date strings to parse however, it is much more efficient to first create a Dates.DateFormat, and pass it instead of a raw format string.

```julia
julia> df = DateFormat("y-m-d");

julia> dt = Date("2015-01-01",df)
2015-01-01
```

```
julia> dt2 = Date("2015-01-02",df)
2015-01-02
```

You can also use the `dateformat""` string macro. This macro creates the `DateFormat` object once when the macro is expanded and uses the same `DateFormat` object even if a code snippet is run multiple times.

```
julia> for i = 1:10^5
           Date("2015-01-01", dateformat"y-m-d")
       end
```

A full suite of parsing and formatting tests and examples is available in `tests/dates/io.jl`.

27.2 Durations/Comparisons

Finding the length of time between two `Date` or `DateTime` is straightforward given their underlying representation as `UTInstant{Day}` and `UTInstant{Millisecond}`, respectively. The difference between `Date` is returned in the number of `Day`, and `DateTime` in the number of `Millisecond`. Similarly, comparing `TimeType` is a simple matter of comparing the underlying machine instants (which in turn compares the internal `Int64` values).

```
julia> dt = Date(2012,2,29)
2012-02-29

julia> dt2 = Date(2000,2,1)
2000-02-01

julia> dump(dt)
Date
  instant: Base.Dates.UTInstant{Base.Dates.Day}
    periods: Base.Dates.Day
      value: Int64 734562

julia> dump(dt2)
Date
  instant: Base.Dates.UTInstant{Base.Dates.Day}
    periods: Base.Dates.Day
      value: Int64 730151

julia> dt > dt2
true

julia> dt != dt2
true

julia> dt + dt2
ERROR: MethodError: no method matching +(::Date, ::Date)
[...]

julia> dt * dt2
ERROR: MethodError: no method matching *(::Date, ::Date)
[...]

julia> dt / dt2
ERROR: MethodError: no method matching /(::Date, ::Date)
```

```
[...]

julia> dt - dt2
4411 days

julia> dt2 - dt
-4411 days

julia> dt = DateTime(2012,2,29)
2012-02-29T00:00:00

julia> dt2 = DateTime(2000,2,1)
2000-02-01T00:00:00

julia> dt - dt2
381110400000 milliseconds
```

27.3   Accessor Functions

Because the Date and DateTime types are stored as single Int64 values, date parts or fields can be retrieved through accessor functions. The lowercase accessors return the field as an integer:

```
julia> t = Date(2014, 1, 31)
2014-01-31

julia> Dates.year(t)
2014

julia> Dates.month(t)
1

julia> Dates.week(t)
5

julia> Dates.day(t)
31
```

While propercase return the same value in the corresponding Period type:

```
julia> Dates.Year(t)
2014 years

julia> Dates.Day(t)
31 days
```

Compound methods are provided, as they provide a measure of efficiency if multiple fields are needed at the same time:

```
julia> Dates.yearmonth(t)
(2014, 1)

julia> Dates.monthday(t)
(1, 31)

julia> Dates.yearmonthday(t)
(2014, 1, 31)
```

One may also access the underlying UTInstant or integer value:

```
julia> dump(t)
Date
  instant: Base.Dates.UTInstant{Base.Dates.Day}
    periods: Base.Dates.Day
      value: Int64 735264

julia> t.instant
Base.Dates.UTInstant{Base.Dates.Day}(735264 days)

julia> Dates.value(t)
735264
```

27.4 Query Functions

Query functions provide calendrical information about a TimeType. They include information about the day of the week:

```
julia> t = Date(2014, 1, 31)
2014-01-31

julia> Dates.dayofweek(t)
5

julia> Dates.dayname(t)
"Friday"

julia> Dates.dayofweekofmonth(t) # 5th Friday of January
5
```

Month of the year:

```
julia> Dates.monthname(t)
"January"

julia> Dates.daysinmonth(t)
31
```

As well as information about the TimeType's year and quarter:

```
julia> Dates.isleapyear(t)
false

julia> Dates.dayofyear(t)
31

julia> Dates.quarterofyear(t)
1

julia> Dates.dayofquarter(t)
31
```

The dayname() and monthname() methods can also take an optional locale keyword that can be used to return the name of the day or month of the year for other languages/locales. There are also versions of these functions returning the abbreviated names, namely dayabbr() and monthabbr(). First the mapping is loaded into the LOCALES variable:

```julia
julia> french_months = ["janvier", "février", "mars", "avril", "mai", "juin",
                        "juillet", "août", "septembre", "octobre", "novembre", "décembre"];

julia> french_monts_abbrev = ["janv","févr","mars","avril","mai","juin",
                             "juil","août","sept","oct","nov","déc"];

julia> french_days = ["lundi","mardi","mercredi","jeudi","vendredi","samedi","dimanche"];

julia> Dates.LOCALES["french"] = Dates.DateLocale(french_months, french_monts_abbrev, french_days
    , [""]);
```

The above mentioned functions can then be used to perform the queries:

```julia
julia> Dates.dayname(t;locale="french")
"vendredi"

julia> Dates.monthname(t;locale="french")
"janvier"

julia> Dates.monthabbr(t;locale="french")
"janv"
```

Since the abbreviated versions of the days are not loaded, trying to use the function dayabbr() will error.

```julia
julia> Dates.dayabbr(t;locale="french")
ERROR: BoundsError: attempt to access 1-element Array{String,1} at index [5]
Stacktrace:
 [1] #dayabbr#6(::String, ::Function, ::Int64) at ./dates/query.jl:114
 [2] (::Base.Dates.#kw##dayabbr)(::Array{Any,1}, ::Base.Dates.#dayabbr, ::Int64) at ./<missing>:0
     (repeats 2 times)
```

27.5   TimeType-Period Arithmetic

It's good practice when using any language/date framework to be familiar with how date-period arithmetic is handled as there are some tricky issues to deal with (though much less so for day-precision types).

The Dates module approach tries to follow the simple principle of trying to change as little as possible when doing Period arithmetic. This approach is also often known as *calendrical* arithmetic or what you would probably guess if someone were to ask you the same calculation in a conversation. Why all the fuss about this? Let's take a classic example: add 1 month to January 31st, 2014. What's the answer? Javascript will say March 3 (assumes 31 days). PHP says March 2 (assumes 30 days). The fact is, there is no right answer. In the Dates module, it gives the result of February 28th. How does it figure that out? I like to think of the classic 7-7-7 gambling game in casinos.

Now just imagine that instead of 7-7-7, the slots are Year-Month-Day, or in our example, 2014-01-31. When you ask to add 1 month to this date, the month slot is incremented, so now we have 2014-02-31. Then the day number is checked if it is greater than the last valid day of the new month; if it is (as in the case above), the day number is adjusted down to the last valid day (28). What are the ramifications with this approach? Go ahead and add another month to our date, 2014-02-28 + Month(1) == 2014-03-28. What? Were you expecting the last day of March? Nope, sorry, remember the 7-7-7 slots. As few slots as possible are going to change, so we first increment the month slot by 1, 2014-03-28, and boom, we're done because that's a valid date. On the other hand, if we were to add 2 months to our original date, 2014-01-31, then we end up with 2014-03-31, as expected. The other ramification of this approach is a loss in associativity when a specific ordering is forced (i.e. adding things in different orders results in different outcomes). For example:

```julia
julia> (Date(2014,1,29)+Dates.Day(1)) + Dates.Month(1)
2014-02-28
```

```
julia> (Date(2014,1,29)+Dates.Month(1)) + Dates.Day(1)
2014-03-01
```

What's going on there? In the first line, we're adding 1 day to January 29th, which results in 2014-01-30; then we add 1 month, so we get 2014-02-30, which then adjusts down to 2014-02-28. In the second example, we add 1 month *first*, where we get 2014-02-29, which adjusts down to 2014-02-28, and *then* add 1 day, which results in 2014-03-01. One design principle that helps in this case is that, in the presence of multiple Periods, the operations will be ordered by the Periods' *types*, not their value or positional order; this means Year will always be added first, then Month, then Week, etc. Hence the following *does* result in associativity and Just Works:

```
julia> Date(2014,1,29) + Dates.Day(1) + Dates.Month(1)
2014-03-01

julia> Date(2014,1,29) + Dates.Month(1) + Dates.Day(1)
2014-03-01
```

Tricky? Perhaps. What is an innocent Dates user to do? The bottom line is to be aware that explicitly forcing a certain associativity, when dealing with months, may lead to some unexpected results, but otherwise, everything should work as expected. Thankfully, that's pretty much the extent of the odd cases in date-period arithmetic when dealing with time in UT (avoiding the "joys" of dealing with daylight savings, leap seconds, etc.).

As a bonus, all period arithmetic objects work directly with ranges:

```
julia> dr = Date(2014,1,29):Date(2014,2,3)
2014-01-29:1 day:2014-02-03

julia> collect(dr)
6-element Array{Date,1}:
 2014-01-29
 2014-01-30
 2014-01-31
 2014-02-01
 2014-02-02
 2014-02-03

julia> dr = Date(2014,1,29):Dates.Month(1):Date(2014,07,29)
2014-01-29:1 month:2014-07-29

julia> collect(dr)
7-element Array{Date,1}:
 2014-01-29
 2014-02-28
 2014-03-29
 2014-04-29
 2014-05-29
 2014-06-29
 2014-07-29
```

27.6   Adjuster Functions

As convenient as date-period arithmetics are, often the kinds of calculations needed on dates take on a *calendrical* or *temporal* nature rather than a fixed number of periods. Holidays are a perfect example; most follow rules such

as "Memorial Day = Last Monday of May", or "Thanksgiving = 4th Thursday of November". These kinds of temporal expressions deal with rules relative to the calendar, like first or last of the month, next Tuesday, or the first and third Wednesdays, etc.

The `Dates` module provides the *adjuster* API through several convenient methods that aid in simply and succinctly expressing temporal rules. The first group of adjuster methods deal with the first and last of weeks, months, quarters, and years. They each take a single `TimeType` as input and return or *adjust to* the first or last of the desired period relative to the input.

```julia
julia> Dates.firstdayofweek(Date(2014,7,16)) # Adjusts the input to the Monday of the input's week
2014-07-14

julia> Dates.lastdayofmonth(Date(2014,7,16)) # Adjusts to the last day of the input's month
2014-07-31

julia> Dates.lastdayofquarter(Date(2014,7,16)) # Adjusts to the last day of the input's quarter
2014-09-30
```

The next two higher-order methods, `tonext()`, and `toprev()`, generalize working with temporal expressions by taking a `DateFunction` as first argument, along with a starting `TimeType`. A `DateFunction` is just a function, usually anonymous, that takes a single `TimeType` as input and returns a `Bool`, `true` indicating a satisfied adjustment criterion. For example:

```julia
julia> istuesday = x->Dates.dayofweek(x) == Dates.Tuesday # Returns true if the day of the week of
↪   x is Tuesday
(::#1) (generic function with 1 method)

julia> Dates.tonext(istuesday, Date(2014,7,13)) # 2014-07-13 is a Sunday
2014-07-15

julia> Dates.tonext(Date(2014,7,13), Dates.Tuesday) # Convenience method provided for day of the
↪   week adjustments
2014-07-15
```

This is useful with the do-block syntax for more complex temporal expressions:

```julia
julia> Dates.tonext(Date(2014,7,13)) do x
           # Return true on the 4th Thursday of November (Thanksgiving)
           Dates.dayofweek(x) == Dates.Thursday &&
           Dates.dayofweekofmonth(x) == 4 &&
           Dates.month(x) == Dates.November
       end
2014-11-27
```

The `Base.filter()` method can be used to obtain all valid dates/moments in a specified range:

```julia
# Pittsburgh street cleaning; Every 2nd Tuesday from April to November
# Date range from January 1st, 2014 to January 1st, 2015
julia> dr = Dates.Date(2014):Dates.Date(2015);

julia> filter(dr) do x
           Dates.dayofweek(x) == Dates.Tue &&
           Dates.April <= Dates.month(x) <= Dates.Nov &&
```

```
        Dates.dayofweekofmonth(x) == 2
    end
8-element Array{Date,1}:
 2014-04-08
 2014-05-13
 2014-06-10
 2014-07-08
 2014-08-12
 2014-09-09
 2014-10-14
 2014-11-11
```

Additional examples and tests are available in `test/dates/adjusters.jl`.

27.7 Period Types

Periods are a human view of discrete, sometimes irregular durations of time. Consider 1 month; it could represent, in days, a value of 28, 29, 30, or 31 depending on the year and month context. Or a year could represent 365 or 366 days in the case of a leap year. `Period` types are simple `Int64` wrappers and are constructed by wrapping any `Int64` convertible type, i.e. `Year(1)` or `Month(3.0)`. Arithmetic between `Period` of the same type behave like integers, and limited `Period-Real` arithmetic is available.

```
julia> y1 = Dates.Year(1)
1 year

julia> y2 = Dates.Year(2)
2 years

julia> y3 = Dates.Year(10)
10 years

julia> y1 + y2
3 years

julia> div(y3,y2)
5

julia> y3 - y2
8 years

julia> y3 % y2
0 years

julia> div(y3,3) # mirrors integer division
3 years
```

27.8 Rounding

`Date` and `DateTime` values can be rounded to a specified resolution (e.g., 1 month or 15 minutes) with `floor()`, `ceil()`, or `round()`:

```
julia> floor(Date(1985, 8, 16), Dates.Month)
1985-08-01
```

```
julia> ceil(DateTime(2013, 2, 13, 0, 31, 20), Dates.Minute(15))
2013-02-13T00:45:00

julia> round(DateTime(2016, 8, 6, 20, 15), Dates.Day)
2016-08-07T00:00:00
```

Unlike the numeric `round()` method, which breaks ties toward the even number by default, the `TimeTyperound()` method uses the `RoundNearestTiesUp` rounding mode. (It's difficult to guess what breaking ties to nearest "even" `TimeType` would entail.) Further details on the available `RoundingMode`s can be found in the API reference.

Rounding should generally behave as expected, but there are a few cases in which the expected behaviour is not obvious.

Rounding Epoch

In many cases, the resolution specified for rounding (e.g., `Dates.Second(30)`) divides evenly into the next largest period (in this case, `Dates.Minute(1)`). But rounding behaviour in cases in which this is not true may lead to confusion. What is the expected result of rounding a `DateTime` to the nearest 10 hours?

```
julia> round(DateTime(2016, 7, 17, 11, 55), Dates.Hour(10))
2016-07-17T12:00:00
```

That may seem confusing, given that the hour (12) is not divisible by 10. The reason that `2016-07-17T12:00:00` was chosen is that it is 17,676,660 hours after `0000-01-01T00:00:00`, and 17,676,660 is divisible by 10.

As Julia `Date` and `DateTime` values are represented according to the ISO 8601 standard, `0000-01-01T00:00:00` was chosen as base (or "rounding epoch") from which to begin the count of days (and milliseconds) used in rounding calculations. (Note that this differs slightly from Julia's internal representation of `Date`s using Rata Die notation; but since the ISO 8601 standard is most visible to the end user, `0000-01-01T00:00:00` was chosen as the rounding epoch instead of the `0000-12-31T00:00:00` used internally to minimize confusion.)

The only exception to the use of `0000-01-01T00:00:00` as the rounding epoch is when rounding to weeks. Rounding to the nearest week will always return a Monday (the first day of the week as specified by ISO 8601). For this reason, we use `0000-01-03T00:00:00` (the first day of the first week of year 0000, as defined by ISO 8601) as the base when rounding to a number of weeks.

Here is a related case in which the expected behaviour is not necessarily obvious: What happens when we round to the nearest P(2), where P is a `Period` type? In some cases (specifically, when `P <: Dates.TimePeriod`) the answer is clear:

```
julia> round(DateTime(2016, 7, 17, 8, 55, 30), Dates.Hour(2))
2016-07-17T08:00:00

julia> round(DateTime(2016, 7, 17, 8, 55, 30), Dates.Minute(2))
2016-07-17T08:56:00
```

This seems obvious, because two of each of these periods still divides evenly into the next larger order period. But in the case of two months (which still divides evenly into one year), the answer may be surprising:

```
julia> round(DateTime(2016, 7, 17, 8, 55, 30), Dates.Month(2))
2016-07-01T00:00:00
```

Why round to the first day in July, even though it is month 7 (an odd number)? The key is that months are 1-indexed (the first month is assigned 1), unlike hours, minutes, seconds, and milliseconds (the first of which are assigned 0).

This means that rounding a `DateTime` to an even multiple of seconds, minutes, hours, or years (because the ISO 8601 specification includes a year zero) will result in a `DateTime` with an even value in that field, while rounding a `DateTime` to an even multiple of months will result in the months field having an odd value. Because both months and years may contain an irregular number of days, whether rounding to an even number of days will result in an even value in the days field is uncertain.

See the API reference for additional information on methods exported from the `Dates` module.

Chapter 28

Interacting With Julia

Julia comes with a full-featured interactive command-line REPL (read-eval-print loop) built into the `julia` executable. In addition to allowing quick and easy evaluation of Julia statements, it has a searchable history, tab-completion, many helpful keybindings, and dedicated help and shell modes. The REPL can be started by simply calling `julia` with no arguments or double-clicking on the executable:

```
$ julia
               _
   _       _ _(_)_     |  A fresh approach to technical computing
  (_)     | (_) (_)    |  Documentation: https://docs.julialang.org
   _ _   _| |_  __ _   |  Type "?help" for help.
  | | | | | | |/ _` |  |
  | | |_| | | | (_| |  |  Version 0.6.0-dev.2493 (2017-01-31 18:53 UTC)
 _/ |\__'_|_|_|\__'_|  |  Commit c99e12c* (0 days old master)
|__/                   |  x86_64-linux-gnu

julia>
```

To exit the interactive session, type `^D` – the control key together with the d key on a blank line – or type `quit()` followed by the return or enter key. The REPL greets you with a banner and a `julia>` prompt.

28.1 The different prompt modes

The Julian mode

The REPL has four main modes of operation. The first and most common is the Julian prompt. It is the default mode of operation; each new line initially starts with `julia>`. It is here that you can enter Julia expressions. Hitting return or enter after a complete expression has been entered will evaluate the entry and show the result of the last expression.

```
julia> string(1 + 2)
"3"
```

There are a number useful features unique to interactive work. In addition to showing the result, the REPL also binds the result to the variable `ans`. A trailing semicolon on the line can be used as a flag to suppress showing the result.

```
julia> string(3 * 4);

julia> ans
"12"
```

In Julia mode, the REPL supports something called *prompt pasting*. This activates when pasting text that starts with `julia>` into the REPL. In that case, only expressions starting with `julia>` are parsed, others are removed. This makes it is possible to paste a chunk of code that has been copied from a REPL session without having to scrub away prompts and outputs. This feature is enabled by default but can be disabled or enabled at will with `Base.REPL.en-able_promptpaste(::Bool)`. If it is enabled, you can try it out by pasting the code block above this paragraph straight into the REPL. This feature does not work on the standard Windows command prompt due to its limitation at detecting when a paste occurs.

Help mode

When the cursor is at the beginning of the line, the prompt can be changed to a help mode by typing ?. Julia will attempt to print help or documentation for anything entered in help mode:

```
julia> ? # upon typing ?, the prompt changes (in place) to: help?>

help?> string
search: string String stringmime Cstring Cwstring RevString readstring randstring bytestring
↪    SubString

  string(xs...)

  Create a string from any values using the print function.
```

Macros, types and variables can also be queried:

```
help?> @time
  @time

  A macro to execute an expression, printing the time it took to execute, the number of
    allocations,
  and the total number of bytes its execution caused to be allocated, before returning the value
    of the
  expression.

  See also @timev, @timed, @elapsed, and @allocated.

help?> AbstractString
search: AbstractString AbstractSparseMatrix AbstractSparseVector AbstractSet

  No documentation found.

  Summary:

  abstract AbstractString <: Any

  Subtypes:

  Base.Test.GenericString
  DirectIndexString
  String
```

Help mode can be exited by pressing backspace at the beginning of the line.

Shell mode

Just as help mode is useful for quick access to documentation, another common task is to use the system shell to execute system commands. Just as ? entered help mode when at the beginning of the line, a semicolon (;) will enter the shell mode. And it can be exited by pressing backspace at the beginning of the line.

```julia
julia> ; # upon typing ;, the prompt changes (in place) to: shell>

shell> echo hello
hello
```

Search modes

In all of the above modes, the executed lines get saved to a history file, which can be searched. To initiate an incremental search through the previous history, type ^R – the control key together with the r key. The prompt will change to (reverse-i-search)`'`:, and as you type the search query will appear in the quotes. The most recent result that matches the query will dynamically update to the right of the colon as more is typed. To find an older result using the same query, simply type ^R again.

Just as ^R is a reverse search, ^S is a forward search, with the prompt (i-search)`'`:. The two may be used in conjunction with each other to move through the previous or next matching results, respectively.

28.2   Key bindings

The Julia REPL makes great use of key bindings. Several control-key bindings were already introduced above (^D to exit, ^R and ^S for searching), but there are many more. In addition to the control-key, there are also meta-key bindings. These vary more by platform, but most terminals default to using alt- or option- held down with a key to send the meta-key (or can be configured to do so).

Customizing keybindings

Julia's REPL keybindings may be fully customized to a user's preferences by passing a dictionary to REPL.setup_interface(). The keys of this dictionary may be characters or strings. The key '*' refers to the default action. Control plus character x bindings are indicated with "^x". Meta plus x can be written "\\Mx". The values of the custom keymap must be nothing (indicating that the input should be ignored) or functions that accept the signature (PromptState, AbstractREPL, Char). The REPL.setup_interface() function must be called before the REPL is initialized, by registering the operation with atreplinit(). For example, to bind the up and down arrow keys to move through history without prefix search, one could put the following code in .juliarc.jl:

```julia
import Base: LineEdit, REPL

const mykeys = Dict{Any,Any}(
    # Up Arrow
    "\e[A" => (s,o...)->(LineEdit.edit_move_up(s) || LineEdit.history_prev(s,
↪  LineEdit.mode(s).hist)),
    # Down Arrow
    "\e[B" => (s,o...)->(LineEdit.edit_move_up(s) || LineEdit.history_next(s,
↪  LineEdit.mode(s).hist))
)

function customize_keys(repl)
    repl.interface = REPL.setup_interface(repl; extra_repl_keymap = mykeys)
end
```

Keybinding	Description
Program control	
^D	Exit (when buffer is empty)
^C	Interrupt or cancel
^L	Clear console screen
Return/Enter, ^J	New line, executing if it is complete
meta-Return/Enter	Insert new line without executing it
? or ;	Enter help or shell mode (when at start of a line)
^R, ^S	Incremental history search, described above
Cursor movement	
Right arrow, ^F	Move right one character
Left arrow, ^B	Move left one character
Home, ^A	Move to beginning of line
End, ^E	Move to end of line
^P	Change to the previous or next history entry
^N	Change to the next history entry
Up arrow	Move up one line (or to the previous history entry)
Down arrow	Move down one line (or to the next history entry)
Page-up	Change to the previous history entry that matches the text before the cursor
Page-down	Change to the next history entry that matches the text before the cursor
meta-F	Move right one word
meta-B	Move left one word
Editing	
Backspace, ^H	Delete the previous character
Delete, ^D	Forward delete one character (when buffer has text)
meta-Backspace	Delete the previous word
meta-D	Forward delete the next word
^W	Delete previous text up to the nearest whitespace
^K	"Kill" to end of line, placing the text in a buffer
^Y	"Yank" insert the text from the kill buffer
^T	Transpose the characters about the cursor
^Q	Write a number in REPL and press ^Q to open editor at corresponding stackframe

```
atreplinit(customize_keys)
```

Users should refer to `base/LineEdit.jl` to discover the available actions on key input.

28.3   Tab completion

In both the Julian and help modes of the REPL, one can enter the first few characters of a function or type and then press the tab key to get a list all matches:

```
julia> stri[TAB]
stride      strides      string       stringmime  strip

julia> Stri[TAB]
StridedArray     StridedMatrix     StridedVecOrMat  StridedVector     String
```

The tab key can also be used to substitute LaTeX math symbols with their Unicode equivalents, and get a list of LaTeX matches as well:

```
julia> \pi[TAB]
julia> π
π = 3.1415926535897...

julia> e\_1[TAB] = [1,0]
julia> e = [1,0]
2-element Array{Int64,1}:
 1
 0

julia> e\^1[TAB] = [1 0]
julia> e¹ = [1 0]
1×2 Array{Int64,2}:
 1  0

julia> \sqrt[TAB]2     # √ is equivalent to the sqrt() function
julia> √2
1.4142135623730951

julia> \hbar[TAB](h) = h / 2\pi[TAB]
julia> ℏ(h) = h / 2π
ℏ (generic function with 1 method)

julia> \h[TAB]
\hat              \hermitconjmatrix  \hkswarow         \hrectangle
\hatapprox        \hexagon           \hookleftarrow    \hrectangleblack
\hbar             \hexagonblack      \hookrightarrow   \hslash
\heartsuit        \hksearow          \house            \hspace

julia> α="\alpha[TAB]"   # LaTeX completion also works in strings
julia> α="α"
```

A full list of tab-completions can be found in the Unicode Input section of the manual.

Completion of paths works for strings and julia's shell mode:

```
julia> path="/[TAB]"
.dockerenv  .juliabox/  boot/       etc/        lib/        media/      opt/        root/
↪   sbin/        sys/         usr/
.dockerinit bin/        dev/        home/       lib64/      mnt/        proc/       run/
↪   srv/         tmp/         var/
shell> /[TAB]
.dockerenv  .juliabox/  boot/       etc/        lib/        media/      opt/        root/
↪   sbin/        sys/         usr/
.dockerinit bin/        dev/        home/       lib64/      mnt/        proc/       run/
↪   srv/         tmp/         var/
```

Tab completion can help with investigation of the available methods matching the input arguments:

```
julia> max([TAB] # All methods are displayed, not shown here due to size of the list
```

```
julia> max([1, 2], [TAB] # All methods where `Vector{Int}` matches as first argument
max(x, y) in Base at operators.jl:215
max(a, b, c, xs...) in Base at operators.jl:281

julia> max([1, 2], max(1, 2), [TAB] # All methods matching the arguments.
max(x, y) in Base at operators.jl:215
max(a, b, c, xs...) in Base at operators.jl:281
```

Keywords are also displayed in the suggested methods, see second line after `;` where `limit` and `keep` are keyword arguments:

```
julia> split("1 1 1", [TAB]
split(str::AbstractString) in Base at strings/util.jl:278
split{T<:AbstractString}(str::T, splitter; limit, keep) in Base at strings/util.jl:254
```

The completion of the methods uses type inference and can therefore see if the arguments match even if the arguments are output from functions. The function needs to be type stable for the completion to be able to remove non-matching methods.

Tab completion can also help completing fields:

```
julia> Pkg.a[TAB]
add      available
```

Fields for output from functions can also be completed:

```
julia> split("","")[1].[TAB]
endof  offset  string
```

The completion of fields for output from functions uses type inference, and it can only suggest fields if the function is type stable.

28.4   Customizing Colors

The colors used by Julia and the REPL can be customized, as well. To change the color of the Julia prompt you can add something like the following to your `.juliarc.jl` file, which is to be placed inside your home directory:

```
function customize_colors(repl)
    repl.prompt_color = Base.text_colors[:cyan]
end

atreplinit(customize_colors)
```

The available color keys can be seen by typing `Base.text_colors` in the help mode of the REPL. In addition, the integers 0 to 255 can be used as color keys for terminals with 256 color support.

You can also change the colors for the help and shell prompts and input and answer text by setting the appropriate field of `repl` in the `customize_colors` function above (respectively, `help_color`, `shell_color`, `input_color`, and `answer_color`). For the latter two, be sure that the `envcolors` field is also set to false.

It is also possible to apply boldface formatting by using `Base.text_colors[:bold]` as a color. For instance, to print answers in boldface font, one can use the following as a `.juliarc.jl`:

```julia
function customize_colors(repl)
    repl.envcolors = false
    repl.answer_color = Base.text_colors[:bold]
end

atreplinit(customize_colors)
```

You can also customize the color used to render warning and informational messages by setting the appropriate environment variables. For instance, to render error, warning, and informational messages respectively in magenta, yellow, and cyan you can add the following to your `.juliarc.jl` file:

```julia
ENV["JULIA_ERROR_COLOR"] = :magenta
ENV["JULIA_WARN_COLOR"] = :yellow
ENV["JULIA_INFO_COLOR"] = :cyan
```

Chapter 29

Running External Programs

Julia borrows backtick notation for commands from the shell, Perl, and Ruby. However, in Julia, writing

```julia
julia> `echo hello`
`echo hello`
```

differs in several aspects from the behavior in various shells, Perl, or Ruby:

- Instead of immediately running the command, backticks create a `Cmd` object to represent the command. You can use this object to connect the command to others via pipes, run it, and read or write to it.

- When the command is run, Julia does not capture its output unless you specifically arrange for it to. Instead, the output of the command by default goes to STDOUT as it would using `libc`'s `system` call.

- The command is never run with a shell. Instead, Julia parses the command syntax directly, appropriately interpolating variables and splitting on words as the shell would, respecting shell quoting syntax. The command is run as `julia`'s immediate child process, using `fork` and `exec` calls.

Here's a simple example of running an external program:

```julia
julia> mycommand = `echo hello`
`echo hello`

julia> typeof(mycommand)
Cmd

julia> run(mycommand)
hello
```

The `hello` is the output of the `echo` command, sent to STDOUT. The run method itself returns `nothing`, and throws an `ErrorException` if the external command fails to run successfully.

If you want to read the output of the external command, `readstring()` can be used instead:

```julia
julia> a = readstring(`echo hello`)
"hello\n"

julia> chomp(a) == "hello"
true
```

More generally, you can use open() to read from or write to an external command.

```
julia> open(`less`, "w", STDOUT) do io
           for i = 1:3
               println(io, i)
           end
       end
1
2
3
```

29.1   Interpolation

Suppose you want to do something a bit more complicated and use the name of a file in the variable file as an argument to a command. You can use $ for interpolation much as you would in a string literal (see Strings):

```
julia> file = "/etc/passwd"
"/etc/passwd"

julia> `sort $file`
`sort /etc/passwd`
```

A common pitfall when running external programs via a shell is that if a file name contains characters that are special to the shell, they may cause undesirable behavior. Suppose, for example, rather than /etc/passwd, we wanted to sort the contents of the file /Volumes/External HD/data.csv. Let's try it:

```
julia> file = "/Volumes/External HD/data.csv"
"/Volumes/External HD/data.csv"

julia> `sort $file`
`sort '/Volumes/External HD/data.csv'`
```

How did the file name get quoted? Julia knows that file is meant to be interpolated as a single argument, so it quotes the word for you. Actually, that is not quite accurate: the value of file is never interpreted by a shell, so there's no need for actual quoting; the quotes are inserted only for presentation to the user. This will even work if you interpolate a value as part of a shell word:

```
julia> path = "/Volumes/External HD"
"/Volumes/External HD"

julia> name = "data"
"data"

julia> ext = "csv"
"csv"

julia> `sort $path/$name.$ext`
`sort '/Volumes/External HD/data.csv'`
```

As you can see, the space in the path variable is appropriately escaped. But what if you *want* to interpolate multiple words? In that case, just use an array (or any other iterable container):

```
julia> files = ["/etc/passwd","/Volumes/External HD/data.csv"]
2-element Array{String,1}:
 "/etc/passwd"
 "/Volumes/External HD/data.csv"

julia> `grep foo $files`
`grep foo /etc/passwd '/Volumes/External HD/data.csv'`
```

If you interpolate an array as part of a shell word, Julia emulates the shell's {a,b,c} argument generation:

```
julia> names = ["foo","bar","baz"]
3-element Array{String,1}:
 "foo"
 "bar"
 "baz"

julia> `grep xylophone $names.txt`
`grep xylophone foo.txt bar.txt baz.txt`
```

Moreover, if you interpolate multiple arrays into the same word, the shell's Cartesian product generation behavior is emulated:

```
julia> names = ["foo","bar","baz"]
3-element Array{String,1}:
 "foo"
 "bar"
 "baz"

julia> exts = ["aux","log"]
2-element Array{String,1}:
 "aux"
 "log"

julia> `rm -f $names.$exts`
`rm -f foo.aux foo.log bar.aux bar.log baz.aux baz.log`
```

Since you can interpolate literal arrays, you can use this generative functionality without needing to create temporary array objects first:

```
julia> `rm -rf $["foo","bar","baz","qux"].$["aux","log","pdf"]`
`rm -rf foo.aux foo.log foo.pdf bar.aux bar.log bar.pdf baz.aux baz.log baz.pdf qux.aux qux.log
↪    qux.pdf`
```

29.2   Quoting

Inevitably, one wants to write commands that aren't quite so simple, and it becomes necessary to use quotes. Here's a simple example of a Perl one-liner at a shell prompt:

```
sh$ perl -le '$|=1; for (0..3) { print }'
0
1
2
3
```

The Perl expression needs to be in single quotes for two reasons: so that spaces don't break the expression into multiple shell words, and so that uses of Perl variables like $| (yes, that's the name of a variable in Perl), don't cause interpolation. In other instances, you may want to use double quotes so that interpolation *does* occur:

```
sh$ first="A"
sh$ second="B"
sh$ perl -le '$|=1; print for @ARGV' "1: $first" "2: $second"
1: A
2: B
```

In general, the Julia backtick syntax is carefully designed so that you can just cut-and-paste shell commands as is into backticks and they will work: the escaping, quoting, and interpolation behaviors are the same as the shell's. The only difference is that the interpolation is integrated and aware of Julia's notion of what is a single string value, and what is a container for multiple values. Let's try the above two examples in Julia:

```
julia> A = `perl -le '$|=1; for (0..3) { print }'`
`perl -le '$|=1; for (0..3) { print }'`

julia> run(A)
0
1
2
3

julia> first = "A"; second = "B";

julia> B = `perl -le 'print for @ARGV' "1: $first" "2: $second"`
`perl -le 'print for @ARGV' '1: A' '2: B'`

julia> run(B)
1: A
2: B
```

The results are identical, and Julia's interpolation behavior mimics the shell's with some improvements due to the fact that Julia supports first-class iterable objects while most shells use strings split on spaces for this, which introduces ambiguities. When trying to port shell commands to Julia, try cut and pasting first. Since Julia shows commands to you before running them, you can easily and safely just examine its interpretation without doing any damage.

29.3   Pipelines

Shell metacharacters, such as |, &, and >, need to be quoted (or escaped) inside of Julia's backticks:

```
julia> run(`echo hello '|' sort`)
hello | sort

julia> run(`echo hello \| sort`)
hello | sort
```

This expression invokes the echo command with three words as arguments: hello, |, and sort. The result is that a single line is printed: hello | sort. How, then, does one construct a pipeline? Instead of using '|' inside of backticks, one uses pipeline():

```
julia> run(pipeline(`echo hello`, `sort`))
hello
```

This pipes the output of the echo command to the sort command. Of course, this isn't terribly interesting since there's only one line to sort, but we can certainly do much more interesting things:

```
julia> run(pipeline(`cut -d: -f3 /etc/passwd`, `sort -n`, `tail -n5`))
210
211
212
213
214
```

This prints the highest five user IDs on a UNIX system. The cut, sort and tail commands are all spawned as immediate children of the current julia process, with no intervening shell process. Julia itself does the work to setup pipes and connect file descriptors that is normally done by the shell. Since Julia does this itself, it retains better control and can do some things that shells cannot.

Julia can run multiple commands in parallel:

```
julia> run(`echo hello` & `echo world`)
world
hello
```

The order of the output here is non-deterministic because the two echo processes are started nearly simultaneously, and race to make the first write to the STDOUT descriptor they share with each other and the julia parent process. Julia lets you pipe the output from both of these processes to another program:

```
julia> run(pipeline(`echo world` & `echo hello`, `sort`))
hello
world
```

In terms of UNIX plumbing, what's happening here is that a single UNIX pipe object is created and written to by both echo processes, and the other end of the pipe is read from by the sort command.

IO redirection can be accomplished by passing keyword arguments stdin, stdout, and stderr to the pipeline function:

```
pipeline(`do_work`, stdout=pipeline(`sort`, "out.txt"), stderr="errs.txt")
```

Avoiding Deadlock in Pipelines

When reading and writing to both ends of a pipeline from a single process, it is important to avoid forcing the kernel to buffer all of the data.

For example, when reading all of the output from a command, call readstring(out), not wait(process), since the former will actively consume all of the data written by the process, whereas the latter will attempt to store the data in the kernel's buffers while waiting for a reader to be connected.

Another common solution is to separate the reader and writer of the pipeline into separate Tasks:

```
writer = @async writeall(process, "data")
reader = @async do_compute(readstring(process))
wait(process)
fetch(reader)
```

Complex Example

The combination of a high-level programming language, a first-class command abstraction, and automatic setup of pipes between processes is a powerful one. To give some sense of the complex pipelines that can be created easily, here are some more sophisticated examples, with apologies for the excessive use of Perl one-liners:

```julia
julia> prefixer(prefix, sleep) = `perl -nle '$|=1; print "'$prefix' ", $_; sleep '$sleep';'`;

julia> run(pipeline(`perl -le '$|=1; for(0..9){ print; sleep 1 }'`, prefixer("A",2) &
 ↪  prefixer("B",2)))
A 0
B 1
A 2
B 3
A 4
B 5
A 6
B 7
A 8
B 9
```

This is a classic example of a single producer feeding two concurrent consumers: one `perl` process generates lines with the numbers 0 through 9 on them, while two parallel processes consume that output, one prefixing lines with the letter "A", the other with the letter "B". Which consumer gets the first line is non-deterministic, but once that race has been won, the lines are consumed alternately by one process and then the other. (Setting $|=1 in Perl causes each print statement to flush the STDOUT handle, which is necessary for this example to work. Otherwise all the output is buffered and printed to the pipe at once, to be read by just one consumer process.)

Here is an even more complex multi-stage producer-consumer example:

```julia
julia> run(pipeline(`perl -le '$|=1; for(0..9){ print; sleep 1 }'`,
          prefixer("X",3) & prefixer("Y",3) & prefixer("Z",3),
          prefixer("A",2) & prefixer("B",2)))
A X 0
B Y 1
A Z 2
B X 3
A Y 4
B Z 5
A X 6
B Y 7
A Z 8
B X 9
```

This example is similar to the previous one, except there are two stages of consumers, and the stages have different latency so they use a different number of parallel workers, to maintain saturated throughput.

We strongly encourage you to try all these examples to see how they work.

Chapter 30

Calling C and Fortran Code

Though most code can be written in Julia, there are many high-quality, mature libraries for numerical computing already written in C and Fortran. To allow easy use of this existing code, Julia makes it simple and efficient to call C and Fortran functions. Julia has a "no boilerplate" philosophy: functions can be called directly from Julia without any "glue" code, code generation, or compilation – even from the interactive prompt. This is accomplished just by making an appropriate call with `ccall` syntax, which looks like an ordinary function call.

The code to be called must be available as a shared library. Most C and Fortran libraries ship compiled as shared libraries already, but if you are compiling the code yourself using GCC (or Clang), you will need to use the `-shared` and `-fPIC` options. The machine instructions generated by Julia's JIT are the same as a native C call would be, so the resulting overhead is the same as calling a library function from C code. (Non-library function calls in both C and Julia can be inlined and thus may have even less overhead than calls to shared library functions. When both libraries and executables are generated by LLVM, it is possible to perform whole-program optimizations that can even optimize across this boundary, but Julia does not yet support that. In the future, however, it may do so, yielding even greater performance gains.)

Shared libraries and functions are referenced by a tuple of the form `(:function, "library")` or `("function", "library")` where `function` is the C-exported function name. `library` refers to the shared library name: shared libraries available in the (platform-specific) load path will be resolved by name, and if necessary a direct path may be specified.

A function name may be used alone in place of the tuple (just `:function` or `"function"`). In this case the name is resolved within the current process. This form can be used to call C library functions, functions in the Julia runtime, or functions in an application linked to Julia.

By default, Fortran compilers generate mangled names (for example, converting function names to lowercase or upper-case, often appending an underscore), and so to call a Fortran function via `ccall` you must pass the mangled identifier corresponding to the rule followed by your Fortran compiler. Also, when calling a Fortran function, all inputs must be passed by reference.

Finally, you can use `ccall` to actually generate a call to the library function. Arguments to `ccall` are as follows:

1. A `(:function, "library")` pair, which must be written as a literal constant,

   OR

   a function pointer (for example, from `dlsym`).

2. Return type (see below for mapping the declared C type to Julia)

   - This argument will be evaluated at compile-time, when the containing method is defined.

3. A tuple of input types. The input types must be written as a literal tuple, not a tuple-valued variable or expression.

    –   This argument will be evaluated at compile-time, when the containing method is defined.

4. The following arguments, if any, are the actual argument values passed to the function.

As a complete but simple example, the following calls the `clock` function from the standard C library:

```julia
julia> t = ccall((:clock, "libc"), Int32, ())
2292761

julia> t
2292761

julia> typeof(ans)
Int32
```

`clock` takes no arguments and returns an `Int32`. One common gotcha is that a 1-tuple must be written with a trailing comma. For example, to call the `getenv` function to get a pointer to the value of an environment variable, one makes a call like this:

```julia
julia> path = ccall((:getenv, "libc"), Cstring, (Cstring,), "SHELL")
Cstring(@0x00007fff5fbffc45)

julia> unsafe_string(path)
"/bin/bash"
```

Note that the argument type tuple must be written as `(Cstring, )`, rather than `(Cstring)`. This is because `(Cstring)` is just the expression `Cstring` surrounded by parentheses, rather than a 1-tuple containing `Cstring`:

```julia
julia> (Cstring)
Cstring

julia> (Cstring,)
(Cstring,)
```

In practice, especially when providing reusable functionality, one generally wraps `ccall` uses in Julia functions that set up arguments and then check for errors in whatever manner the C or Fortran function indicates them, propagating to the Julia caller as exceptions. This is especially important since C and Fortran APIs are notoriously inconsistent about how they indicate error conditions. For example, the `getenv` C library function is wrapped in the following Julia function, which is a simplified version of the actual definition from `env.jl`:

```julia
function getenv(var::AbstractString)
    val = ccall((:getenv, "libc"),
                Cstring, (Cstring,), var)
    if val == C_NULL
        error("getenv: undefined variable: ", var)
    end
    unsafe_string(val)
end
```

The C `getenv` function indicates an error by returning NULL, but other standard C functions indicate errors in various different ways, including by returning -1, 0, 1 and other special values. This wrapper throws an exception clearly indicating the problem if the caller tries to get a non-existent environment variable:

```julia
julia> getenv("SHELL")
"/bin/bash"

julia> getenv("FOOBAR")
getenv: undefined variable: FOOBAR
```

Here is a slightly more complex example that discovers the local machine's hostname:

```julia
function gethostname()
    hostname = Vector{UInt8}(128)
    ccall((:gethostname, "libc"), Int32,
          (Ptr{UInt8}, Csize_t),
          hostname, sizeof(hostname))
    hostname[end] = 0; # ensure null-termination
    return unsafe_string(pointer(hostname))
end
```

This example first allocates an array of bytes, then calls the C library function `gethostname` to fill the array in with the hostname, takes a pointer to the hostname buffer, and converts the pointer to a Julia string, assuming that it is a NUL-terminated C string. It is common for C libraries to use this pattern of requiring the caller to allocate memory to be passed to the callee and filled in. Allocation of memory from Julia like this is generally accomplished by creating an uninitialized array and passing a pointer to its data to the C function. This is why we don't use the `Cstring` type here: as the array is uninitialized, it could contain NUL bytes. Converting to a `Cstring` as part of the `ccall` checks for contained NUL bytes and could therefore throw a conversion error.

30.1 Creating C-Compatible Julia Function Pointers

It is possible to pass Julia functions to native C functions that accept function pointer arguments. For example, to match C prototypes of the form:

```
typedef returntype (*functiontype)(argumenttype,...)
```

The function `cfunction()` generates the C-compatible function pointer for a call to a Julia library function. Arguments to `cfunction()` are as follows:

1. A Julia Function

2. Return type

3. A tuple of input types

A classic example is the standard C library `qsort` function, declared as:

```
void qsort(void *base, size_t nmemb, size_t size,
           int(*compare)(const void *a, const void *b));
```

The `base` argument is a pointer to an array of length nmemb, with elements of `size` bytes each. `compare` is a callback function which takes pointers to two elements a and b and returns an integer less/greater than zero if a should appear before/after b (or zero if any order is permitted). Now, suppose that we have a 1d array A of values in Julia that we want to sort using the `qsort` function (rather than Julia's built-in `sort` function). Before we worry about calling `qsort` and passing arguments, we need to write a comparison function that works for some arbitrary type T:

```
julia> function mycompare(a::T, b::T) where T
           return convert(Cint, a < b ? -1 : a > b ? +1 : 0)::Cint
       end
mycompare (generic function with 1 method)
```

Notice that we have to be careful about the return type: qsort expects a function returning a C int, so we must be sure to return Cint via a call to convert and a typeassert.

In order to pass this function to C, we obtain its address using the function cfunction:

```
julia> const mycompare_c = cfunction(mycompare, Cint, (Ref{Cdouble}, Ref{Cdouble}));
```

cfunction() accepts three arguments: the Julia function (mycompare), the return type (Cint), and a tuple of the argument types, in this case to sort an array of Cdouble (Float64) elements.

The final call to qsort looks like this:

```
julia> A = [1.3, -2.7, 4.4, 3.1]
4-element Array{Float64,1}:
  1.3
 -2.7
  4.4
  3.1

julia> ccall(:qsort, Void, (Ptr{Cdouble}, Csize_t, Csize_t, Ptr{Void}),
             A, length(A), sizeof(eltype(A)), mycompare_c)

julia> A
4-element Array{Float64,1}:
 -2.7
  1.3
  3.1
  4.4
```

As can be seen, A is changed to the sorted array [-2.7, 1.3, 3.1, 4.4]. Note that Julia knows how to convert an array into a Ptr{Cdouble}, how to compute the size of a type in bytes (identical to C's sizeof operator), and so on. For fun, try inserting a println("mycompare($a,$b)") line into mycompare, which will allow you to see the comparisons that qsort is performing (and to verify that it is really calling the Julia function that you passed to it).

30.2   Mapping C Types to Julia

It is critical to exactly match the declared C type with its declaration in Julia. Inconsistencies can cause code that works correctly on one system to fail or produce indeterminate results on a different system.

Note that no C header files are used anywhere in the process of calling C functions: you are responsible for making sure that your Julia types and call signatures accurately reflect those in the C header file. (The Clang package can be used to auto-generate Julia code from a C header file.)

Auto-conversion:

Julia automatically inserts calls to the Base.cconvert() function to convert each argument to the specified type. For example, the following call:

```
ccall((:foo, "libfoo"), Void, (Int32, Float64), x, y)
```

will behave as if the following were written:

```
ccall((:foo, "libfoo"), Void, (Int32, Float64),
      Base.unsafe_convert(Int32, Base.cconvert(Int32, x)),
      Base.unsafe_convert(Float64, Base.cconvert(Float64, y)))
```

Base.cconvert() normally just calls convert(), but can be defined to return an arbitrary new object more appropriate for passing to C. For example, this is used to convert an Array of objects (e.g. strings) to an array of pointers.

Base.unsafe_convert() handles conversion to Ptr types. It is considered unsafe because converting an object to a native pointer can hide the object from the garbage collector, causing it to be freed prematurely.

Type Correspondences:

First, a review of some relevant Julia type terminology:

Syntax / Keyword	Example	Description
mutable struct	String	"Leaf Type" :: A group of related data that includes a type-tag, is managed by the Julia GC, and is defined by object-identity. The type parameters of a leaf type must be fully defined (no TypeVars are allowed) in order for the instance to be constructed.
abstract type	Any, AbstractArray{T, N}, Complex{T}	"Super Type" :: A super-type (not a leaf-type) that cannot be instantiated, but can be used to describe a group of types.
T{A}	Vector{Int}	"Type Parameter" :: A specialization of a type (typically used for dispatch or storage optimization).
		"TypeVar" :: The T in the type parameter declaration is referred to as a TypeVar (short for type variable).
primitive type	Int, Float64	"Primitive Type" :: A type with no fields, but a size. It is stored and defined by-value.
struct	Pair{Int, Int}	"Struct" :: A type with all fields defined to be constant. It is defined by-value, and may be stored with a type-tag.
	Complex128 (isbits)	"Is-Bits" :: A primitive type, or a struct type where all fields are other isbits types. It is defined by-value, and is stored without a type-tag.
struct ...; end	nothing	"Singleton" :: a Leaf Type or Struct with no fields.
(...) or tuple(...)	(1, 2, 3)	"Tuple" :: an immutable data-structure similar to an anonymous struct type, or a constant array. Represented as either an array or a struct.

Bits Types:

There are several special types to be aware of, as no other type can be defined to behave the same:

- Float32

  Exactly corresponds to the float type in C (or REAL*4 in Fortran).

- Float64

  Exactly corresponds to the double type in C (or REAL*8 in Fortran).

- `Complex64`

  Exactly corresponds to the `complex float` type in C (or `COMPLEX*8` in Fortran).

- `Complex128`

  Exactly corresponds to the `complex double` type in C (or `COMPLEX*16` in Fortran).

- `Signed`

  Exactly corresponds to the `signed` type annotation in C (or any `INTEGER` type in Fortran). Any Julia type that is not a subtype of `Signed` is assumed to be unsigned.

- `Ref{T}`

  Behaves like a `Ptr{T}` that can manage its memory via the Julia GC.

- `Array{T,N}`

  When an array is passed to C as a `Ptr{T}` argument, it is not reinterpret-cast: Julia requires that the element type of the array matches `T`, and the address of the first element is passed.

  Therefore, if an `Array` contains data in the wrong format, it will have to be explicitly converted using a call such as `trunc(Int32, a)`.

  To pass an array `A` as a pointer of a different type *without* converting the data beforehand (for example, to pass a `Float64` array to a function that operates on uninterpreted bytes), you can declare the argument as `Ptr{Void}`.

  If an array of eltype `Ptr{T}` is passed as a `Ptr{Ptr{T}}` argument, `Base.cconvert()` will attempt to first make a null-terminated copy of the array with each element replaced by its `Base.cconvert()` version. This allows, for example, passing an `argv` pointer array of type `Vector{String}` to an argument of type `Ptr{Ptr{Cchar}}`.

On all systems we currently support, basic C/C++ value types may be translated to Julia types as follows. Every C type also has a corresponding Julia type with the same name, prefixed by C. This can help for writing portable code (and remembering that an `int` in C is not the same as an `Int` in Julia).

System Independent:

The `Cstring` type is essentially a synonym for `Ptr{UInt8}`, except the conversion to `Cstring` throws an error if the Julia string contains any embedded NUL characters (which would cause the string to be silently truncated if the C routine treats NUL as the terminator). If you are passing a `char*` to a C routine that does not assume NUL termination (e.g. because you pass an explicit string length), or if you know for certain that your Julia string does not contain NUL and want to skip the check, you can use `Ptr{UInt8}` as the argument type. `Cstring` can also be used as the `ccall` return type, but in that case it obviously does not introduce any extra checks and is only meant to improve readability of the call.

System-dependent:

> **Note**
>
> When calling a Fortran function, all inputs must be passed by reference, so all type correspondences above should contain an additional `Ptr{..}` or `Ref{..}` wrapper around their type specification.

> **Warning**
>
> For string arguments (`char*`) the Julia type should be `Cstring` (if NUL- terminated data is expected) or either `Ptr{Cchar}` or `Ptr{UInt8}` otherwise (these two pointer types have the same effect), as described above, not `String`. Similarly, for array arguments (`T[]` or `T*`), the Julia type should again be `Ptr{T}`, not `Vector{T}`.

> **Warning**
>
> Julia's `Char` type is 32 bits, which is not the same as the wide character type (`wchar_t` or `wint_t`) on all platforms.

> **Warning**
>
> A return type of `Union{}` means the function will not return i.e. C++11 `[[noreturn]]` or C11 `_Nore-turn` (e.g. `jl_throw` or `longjmp`). Do not use this for functions that return no value (`void`) but do return, use `Void` instead.

> **Note**
>
> For `wchar_t*` arguments, the Julia type should be `Cwstring` (if the C routine expects a NUL-terminated string) or `Ptr{Cwchar_t}` otherwise. Note also that UTF-8 string data in Julia is internally NUL-terminated, so it can be passed to C functions expecting NUL-terminated data without making a copy (but using the `Cwstring` type will cause an error to be thrown if the string itself contains NUL characters).

> **Note**
>
> C functions that take an argument of the type `char**` can be called by using a `Ptr{Ptr{UInt8}}` type within Julia. For example, C functions of the form:
>
> ```
> int main(int argc, char **argv);
> ```
>
> can be called via the following Julia code:
>
> ```julia
> argv = [ "a.out", "arg1", "arg2" ]
> ccall(:main, Int32, (Int32, Ptr{Ptr{UInt8}}), length(argv), argv)
> ```

> **Note**
>
> A C function declared to return `Void` will return the value `nothing` in Julia.

Struct Type correspondences

Composite types, aka `struct` in C or TYPE in Fortran90 (or STRUCTURE / RECORD in some variants of F77), can be mirrored in Julia by creating a `struct` definition with the same field layout.

When used recursively, `isbits` types are stored inline. All other types are stored as a pointer to the data. When mirroring a struct used by-value inside another struct in C, it is imperative that you do not attempt to manually copy the fields over, as this will not preserve the correct field alignment. Instead, declare an `isbits` struct type and use that instead. Unnamed structs are not possible in the translation to Julia.

Packed structs and union declarations are not supported by Julia.

You can get a near approximation of a `union` if you know, a priori, the field that will have the greatest size (potentially including padding). When translating your fields to Julia, declare the Julia field to be only of that type.

Arrays of parameters can be expressed with `NTuple`:

```
in C:
struct B {
    int A[3];
};
b_a_2 = B.A[2];

in Julia:
```

```
struct B
    A::NTuple{3, CInt}
end
b_a_2 = B.A[3]  # note the difference in indexing (1-based in Julia, 0-based in C)
```

Arrays of unknown size (C99-compliant variable length structs specified by [] or [0]) are not directly supported. Often the best way to deal with these is to deal with the byte offsets directly. For example, if a C library declared a proper string type and returned a pointer to it:

```
struct String {
    int strlen;
    char data[];
};
```

In Julia, we can access the parts independently to make a copy of that string:

```
str = from_c::Ptr{Void}
len = unsafe_load(Ptr{Cint}(str))
unsafe_string(str + Core.sizeof(Cint), len)
```

Type Parameters

The type arguments to `ccall` are evaluated statically, when the method containing the ccall is defined. They therefore must take the form of a literal tuple, not a variable, and cannot reference local variables.

This may sound like a strange restriction, but remember that since C is not a dynamic language like Julia, its functions can only accept argument types with a statically-known, fixed signature.

However, while the type layout must be known statically to compute the `ccall` ABI, the static parameters of the function are considered to be part of this static environment. The static parameters of the function may be used as type parameters in the `ccall` signature, as long as they don't affect the layout of the type. For example, `f(x::T) where {T} = ccall(:valid, Ptr{T}, (Ptr{T},), x)` is valid, since `Ptr` is always a word-size primitive type. But, `g(x::T) where {T} = ccall(:notvalid, T, (T,), x)` is not valid, since the type layout of `T` is not known statically.

SIMD Values

Note: This feature is currently implemented on 64-bit x86 and AArch64 platforms only.

If a C/C++ routine has an argument or return value that is a native SIMD type, the corresponding Julia type is a homogeneous tuple of `VecElement` that naturally maps to the SIMD type. Specifically:

- The tuple must be the same size as the SIMD type. For example, a tuple representing an `__m128` on x86 must have a size of 16 bytes.

- The element type of the tuple must be an instance of `VecElement{T}` where T is a primitive type that is 1, 2, 4 or 8 bytes.

For instance, consider this C routine that uses AVX intrinsics:

```
#include <immintrin.h>

__m256 dist( __m256 a, __m256 b ) {
    return _mm256_sqrt_ps(_mm256_add_ps(_mm256_mul_ps(a, a),
                                        _mm256_mul_ps(b, b)));
}
```

The following Julia code calls `dist` using `ccall`:

```julia
const m256 = NTuple{8, VecElement{Float32}}

a = m256(ntuple(i -> VecElement(sin(Float32(i))), 8))
b = m256(ntuple(i -> VecElement(cos(Float32(i))), 8))

function call_dist(a::m256, b::m256)
    ccall((:dist, "libdist"), m256, (m256, m256), a, b)
end

println(call_dist(a,b))
```

The host machine must have the requisite SIMD registers. For example, the code above will not work on hosts without AVX support.

Memory Ownership

malloc/free

Memory allocation and deallocation of such objects must be handled by calls to the appropriate cleanup routines in the libraries being used, just like in any C program. Do not try to free an object received from a C library with `Libc.free` in Julia, as this may result in the `free` function being called via the wrong `libc` library and cause Julia to crash. The reverse (passing an object allocated in Julia to be freed by an external library) is equally invalid.

When to use T, Ptr{T} and Ref{T}

In Julia code wrapping calls to external C routines, ordinary (non-pointer) data should be declared to be of type T inside the `ccall`, as they are passed by value. For C code accepting pointers, `Ref{T}` should generally be used for the types of input arguments, allowing the use of pointers to memory managed by either Julia or C through the implicit call to `Base.cconvert()`. In contrast, pointers returned by the C function called should be declared to be of output type `Ptr{T}`, reflecting that the memory pointed to is managed by C only. Pointers contained in C structs should be represented as fields of type `Ptr{T}` within the corresponding Julia struct types designed to mimic the internal structure of corresponding C structs.

In Julia code wrapping calls to external Fortran routines, all input arguments should be declared as of type `Ref{T}`, as Fortran passes all variables by reference. The return type should either be `Void` for Fortran subroutines, or a T for Fortran functions returning the type T.

30.3 Mapping C Functions to Julia

`ccall/cfunction` argument translation guide

For translating a C argument list to Julia:

- T, where T is one of the primitive types: `char`, `int`, `long`, `short`, `float`, `double`, `complex`, `enum` or any of their `typedef` equivalents

    - T, where T is an equivalent Julia Bits Type (per the table above)

    - if T is an enum, the argument type should be equivalent to `Cint` or `Cuint`

    - argument value will be copied (passed by value)

- `struct T` (including typedef to a struct)

   – T, where T is a Julia leaf type

   – argument value will be copied (passed by value)

- `void*`

   – depends on how this parameter is used, first translate this to the intended pointer type, then determine
      the Julia equivalent using the remaining rules in this list

   – this argument may be declared as `Ptr{Void}`, if it really is just an unknown pointer

- `jl_value_t*`

   – `Any`

   – argument value must be a valid Julia object

   – currently unsupported by `cfunction()`

- `jl_value_t**`

   – `Ref{Any}`

   – argument value must be a valid Julia object (or C_NULL)

   – currently unsupported by `cfunction()`

- `T*`

   – `Ref{T}`, where T is the Julia type corresponding to T

   – argument value will be copied if it is an `isbits` type otherwise, the value must be a valid Julia object

- `(T*)(...)` (e.g. a pointer to a function)

   – `Ptr{Void}` (you may need to use `cfunction()` explicitly to create this pointer)

- `...` (e.g. a vararg)

   – `T...`, where T is the Julia type

- `va_arg`

   – not supported

`ccall/cfunction` return type translation guide

For translating a C return type to Julia:

- `void`

   – `Void` (this will return the singleton instance `nothing::Void`)

- T, where T is one of the primitive types: `char`, `int`, `long`, `short`, `float`, `double`, `complex`, `enum` or any of
   their `typedef` equivalents

   – T, where T is an equivalent Julia Bits Type (per the table above)

   – if T is an enum, the argument type should be equivalent to `Cint` or `Cuint`

   – argument value will be copied (returned by-value)

- `struct T` (including typedef to a struct)

    - T, where T is a Julia Leaf Type
    - argument value will be copied (returned by-value)

- `void*`

    - depends on how this parameter is used, first translate this to the intended pointer type, then determine the Julia equivalent using the remaining rules in this list
    - this argument may be declared as `Ptr{Void}`, if it really is just an unknown pointer

- `jl_value_t*`

    - Any
    - argument value must be a valid Julia object

- `jl_value_t**`

    - `Ref{Any}`
    - argument value must be a valid Julia object (or C_NULL)

- `T*`

    - If the memory is already owned by Julia, or is an `isbits` type, and is known to be non-null:
        * `Ref{T}`, where T is the Julia type corresponding to T
        * a return type of `Ref{Any}` is invalid, it should either be Any (corresponding to `jl_value_t*`) or `Ptr{Any}` (corresponding to `Ptr{Any}`)
        * C **MUST NOT** modify the memory returned via `Ref{T}` if T is an `isbits` type
    - If the memory is owned by C:
        * `Ptr{T}`, where T is the Julia type corresponding to T

- `(T*)(...)` (e.g. a pointer to a function)

    - `Ptr{Void}` (you may need to use `cfunction()` explicitly to create this pointer)

Passing Pointers for Modifying Inputs

Because C doesn't support multiple return values, often C functions will take pointers to data that the function will modify. To accomplish this within a `ccall`, you need to first encapsulate the value inside an `Ref{T}` of the appropriate type. When you pass this Ref object as an argument, Julia will automatically pass a C pointer to the encapsulated data:

```julia
width = Ref{Cint}(0)
range = Ref{Cfloat}(0)
ccall(:foo, Void, (Ref{Cint}, Ref{Cfloat}), width, range)
```

Upon return, the contents of `width` and `range` can be retrieved (if they were changed by foo) by `width[]` and `range[]`; that is, they act like zero-dimensional arrays.

Special Reference Syntax for ccall (deprecated):

The & syntax is deprecated, use the Ref{T} argument type instead.

A prefix & is used on an argument to ccall to indicate that a pointer to a scalar argument should be passed instead of the scalar value itself (required for all Fortran function arguments, as noted above). The following example computes a dot product using a BLAS function.

```julia
function compute_dot(DX::Vector{Float64}, DY::Vector{Float64})
    @assert length(DX) == length(DY)
    n = length(DX)
    incx = incy = 1
    product = ccall((:ddot_, "libLAPACK"),
                    Float64,
                    (Ptr{Int32}, Ptr{Float64}, Ptr{Int32}, Ptr{Float64}, Ptr{Int32}),
                    &n, DX, &incx, DY, &incy)
    return product
end
```

The meaning of prefix & is not quite the same as in C. In particular, any changes to the referenced variables will not be visible in Julia unless the type is mutable (declared via type). However, even for immutable structs it will not cause any harm for called functions to attempt such modifications (that is, writing through the passed pointers). Moreover, & may be used with any expression, such as &0 or &f(x).

When a scalar value is passed with & as an argument of type Ptr{T}, the value will first be converted to type T.

30.4 Some Examples of C Wrappers

Here is a simple example of a C wrapper that returns a Ptr type:

```julia
mutable struct gsl_permutation
end

# The corresponding C signature is
#     gsl_permutation * gsl_permutation_alloc (size_t n);
function permutation_alloc(n::Integer)
    output_ptr = ccall(
        (:gsl_permutation_alloc, :libgsl), # name of C function and library
        Ptr{gsl_permutation},              # output type
        (Csize_t,),                        # tuple of input types
        n                                  # name of Julia variable to pass in
    )
    if output_ptr == C_NULL # Could not allocate memory
        throw(OutOfMemoryError())
    end
    return output_ptr
end
```

The GNU Scientific Library (here assumed to be accessible through :libgsl) defines an opaque pointer, gsl_permutation *, as the return type of the C function gsl_permutation_alloc(). As user code never has to look inside the gsl_permutation struct, the corresponding Julia wrapper simply needs a new type declaration, gsl_permutation, that has no internal fields and whose sole purpose is to be placed in the type parameter of a Ptr type. The return type of the ccall is declared as Ptr{gsl_permutation}, since the memory allocated and pointed to by output_ptr is controlled by C (and not Julia).

The input n is passed by value, and so the function's input signature is simply declared as (`Csize_t,`) without any `Ref` or `Ptr` necessary. (If the wrapper was calling a Fortran function instead, the corresponding function input signature should instead be (`Ref{Csize_t},`), since Fortran variables are passed by reference.) Furthermore, n can be any type that is convertable to a `Csize_t` integer; the `ccall` implicitly calls `Base.cconvert(Csize_t, n)`.

Here is a second example wrapping the corresponding destructor:

```julia
# The corresponding C signature is
#     void gsl_permutation_free (gsl_permutation * p);
function permutation_free(p::Ref{gsl_permutation})
    ccall(
        (:gsl_permutation_free, :libgsl), # name of C function and library
        Void,                             # output type
        (Ref{gsl_permutation},),          # tuple of input types
        p                                 # name of Julia variable to pass in
    )
end
```

Here, the input p is declared to be of type `Ref{gsl_permutation}`, meaning that the memory that p points to may be managed by Julia or by C. A pointer to memory allocated by C should be of type `Ptr{gsl_permutation}`, but it is convertable using `Base.cconvert()` and therefore can be used in the same (covariant) context of the input argument to a `ccall`. A pointer to memory allocated by Julia must be of type `Ref{gsl_permutation}`, to ensure that the memory address pointed to is valid and that Julia's garbage collector manages the chunk of memory pointed to correctly. Therefore, the `Ref{gsl_permutation}` declaration allows pointers managed by C or Julia to be used.

If the C wrapper never expects the user to pass pointers to memory managed by Julia, then using `p::Ptr{gsl_permutation}` for the method signature of the wrapper and similarly in the `ccall` is also acceptable.

Here is a third example passing Julia arrays:

```julia
# The corresponding C signature is
#    int gsl_sf_bessel_Jn_array (int nmin, int nmax, double x,
#                                double result_array[])
function sf_bessel_Jn_array(nmin::Integer, nmax::Integer, x::Real)
    if nmax < nmin
        throw(DomainError())
    end
    result_array = Vector{Cdouble}(nmax - nmin + 1)
    errorcode = ccall(
        (:gsl_sf_bessel_Jn_array, :libgsl), # name of C function and library
        Cint,                               # output type
        (Cint, Cint, Cdouble, Ref{Cdouble}),# tuple of input types
        nmin, nmax, x, result_array         # names of Julia variables to pass in
    )
    if errorcode != 0
        error("GSL error code $errorcode")
    end
    return result_array
end
```

The C function wrapped returns an integer error code; the results of the actual evaluation of the Bessel J function populate the Julia array `result_array`. This variable can only be used with corresponding input type declaration `Ref{Cdouble}`, since its memory is allocated and managed by Julia, not C. The implicit call to `Base.cconvert(Ref{Cdouble}, result_array)` unpacks the Julia pointer to a Julia array data structure into a form understandable by C.

Note that for this code to work correctly, `result_array` must be declared to be of type `Ref{Cdouble}` and not `Ptr{Cdouble}`. The memory is managed by Julia and the `Ref` signature alerts Julia's garbage collector to keep managing the memory for `result_array` while the `ccall` executes. If `Ptr{Cdouble}` were used instead, the `ccall` may still work, but Julia's garbage collector would not be aware that the memory declared for `result_array` is being used by the external C function. As a result, the code may produce a memory leak if `result_array` never gets freed by the garbage collector, or if the garbage collector prematurely frees `result_array`, the C function may end up throwing an invalid memory access exception.

30.5   Garbage Collection Safety

When passing data to a `ccall`, it is best to avoid using the `pointer()` function. Instead define a convert method and pass the variables directly to the `ccall`. `ccall` automatically arranges that all of its arguments will be preserved from garbage collection until the call returns. If a C API will store a reference to memory allocated by Julia, after the `ccall` returns, you must arrange that the object remains visible to the garbage collector. The suggested way to handle this is to make a global variable of type `Array{Ref,1}` to hold these values, until the C library notifies you that it is finished with them.

Whenever you have created a pointer to Julia data, you must ensure the original data exists until you are done with using the pointer. Many methods in Julia such as `unsafe_load()` and `String()` make copies of data instead of taking ownership of the buffer, so that it is safe to free (or alter) the original data without affecting Julia. A notable exception is `unsafe_wrap()` which, for performance reasons, shares (or can be told to take ownership of) the underlying buffer.

The garbage collector does not guarantee any order of finalization. That is, if a contained a reference to b and both a and b are due for garbage collection, there is no guarantee that b would be finalized after a. If proper finalization of a depends on b being valid, it must be handled in other ways.

30.6   Non-constant Function Specifications

A `(name, library)` function specification must be a constant expression. However, it is possible to use computed values as function names by staging through `eval` as follows:

```
@eval ccall(($(string("a", "b")), "lib"), ...
```

This expression constructs a name using `string`, then substitutes this name into a new `ccall` expression, which is then evaluated. Keep in mind that `eval` only operates at the top level, so within this expression local variables will not be available (unless their values are substituted with $). For this reason, `eval` is typically only used to form top-level definitions, for example when wrapping libraries that contain many similar functions.

If your usage is more dynamic, use indirect calls as described in the next section.

30.7   Indirect Calls

The first argument to `ccall` can also be an expression evaluated at run time. In this case, the expression must evaluate to a `Ptr`, which will be used as the address of the native function to call. This behavior occurs when the first `ccall` argument contains references to non-constants, such as local variables, function arguments, or non-constant globals.

For example, you might look up the function via `dlsym`, then cache it in a global variable for that session. For example:

```
macro dlsym(func, lib)
    z, zlocal = gensym(string(func)), gensym()
    eval(current_module(), :(global $z = C_NULL))
    z = esc(z)
    quote
        let $zlocal::Ptr{Void} = $z::Ptr{Void}
```

```
            if $zlocal == C_NULL
                $zlocal = dlsym($(esc(lib))::Ptr{Void}, $(esc(func)))
                global $z = $zlocal
            end
            $zlocal
        end
    end
end

mylibvar = Libdl.dlopen("mylib")
ccall(@dlsym("myfunc", mylibvar), Void, ())
```

30.8 Calling Convention

The second argument to `ccall` can optionally be a calling convention specifier (immediately preceding return type). Without any specifier, the platform-default C calling convention is used. Other supported conventions are: `stdcall`, `cdecl`, `fastcall`, and `thiscall`. For example (from `base/libc.jl`) we see the same `gethostname` `ccall` as above, but with the correct signature for Windows:

```
hn = Vector{UInt8}(256)
err = ccall(:gethostname, stdcall, Int32, (Ptr{UInt8}, UInt32), hn, length(hn))
```

For more information, please see the LLVM Language Reference.

There is one additional special calling convention `llvmcall`, which allows inserting calls to LLVM intrinsics directly. This can be especially useful when targeting unusual platforms such as GPGPUs. For example, for CUDA, we need to be able to read the thread index:

```
ccall("llvm.nvvm.read.ptx.sreg.tid.x", llvmcall, Int32, ())
```

As with any `ccall`, it is essential to get the argument signature exactly correct. Also, note that there is no compatibility layer that ensures the intrinsic makes sense and works on the current target, unlike the equivalent Julia functions exposed by `Core.Intrinsics`.

30.9 Accessing Global Variables

Global variables exported by native libraries can be accessed by name using the `cglobal()` function. The arguments to `cglobal()` are a symbol specification identical to that used by `ccall`, and a type describing the value stored in the variable:

```
julia> cglobal((:errno, :libc), Int32)
Ptr{Int32} @0x00007f418d0816b8
```

The result is a pointer giving the address of the value. The value can be manipulated through this pointer using `unsafe_load()` and `unsafe_store!()`.

30.10   Accessing Data through a Pointer

The following methods are described as "unsafe" because a bad pointer or type declaration can cause Julia to terminate abruptly.

Given a Ptr{T}, the contents of type T can generally be copied from the referenced memory into a Julia object using unsafe_load(ptr, [index]). The index argument is optional (default is 1), and follows the Julia-convention of 1-based indexing. This function is intentionally similar to the behavior of getindex() and setindex!() (e.g. [] access syntax).

The return value will be a new object initialized to contain a copy of the contents of the referenced memory. The referenced memory can safely be freed or released.

If T is Any, then the memory is assumed to contain a reference to a Julia object (a jl_value_t*), the result will be a reference to this object, and the object will not be copied. You must be careful in this case to ensure that the object was always visible to the garbage collector (pointers do not count, but the new reference does) to ensure the memory is not prematurely freed. Note that if the object was not originally allocated by Julia, the new object will never be finalized by Julia's garbage collector. If the Ptr itself is actually a jl_value_t*, it can be converted back to a Julia object reference by unsafe_pointer_to_objref(ptr). (Julia values v can be converted to jl_value_t* pointers, as Ptr{Void}, by calling pointer_from_objref(v).)

The reverse operation (writing data to a Ptr{T}), can be performed using unsafe_store!(ptr, value, [index]). Currently, this is only supported for primitive types or other pointer-free (isbits) immutable struct types.

Any operation that throws an error is probably currently unimplemented and should be posted as a bug so that it can be resolved.

If the pointer of interest is a plain-data array (primitive type or immutable struct), the function unsafe_wrap(Array, ptr,dims,[own]) may be more useful. The final parameter should be true if Julia should "take ownership" of the underlying buffer and call free(ptr) when the returned Array object is finalized. If the own parameter is omitted or false, the caller must ensure the buffer remains in existence until all access is complete.

Arithmetic on the Ptr type in Julia (e.g. using +) does not behave the same as C's pointer arithmetic. Adding an integer to a Ptr in Julia always moves the pointer by some number of *bytes*, not elements. This way, the address values obtained from pointer arithmetic do not depend on the element types of pointers.

30.11   Thread-safety

Some C libraries execute their callbacks from a different thread, and since Julia isn't thread-safe you'll need to take some extra precautions. In particular, you'll need to set up a two-layered system: the C callback should only *schedule* (via Julia's event loop) the execution of your "real" callback. To do this, create a AsyncCondition object and wait on it:

```
cond = Base.AsyncCondition()
wait(cond)
```

The callback you pass to C should only execute a ccall to :uv_async_send, passing cond.handle as the argument, taking care to avoid any allocations or other interactions with the Julia runtime.

Note that events may be coalesced, so multiple calls to uv_async_send may result in a single wakeup notification to the condition.

30.12   More About Callbacks

For more details on how to pass callbacks to C libraries, see this blog post.

30.13 C++

For direct C++ interfacing, see the Cxx package. For tools to create C++ bindings, see the CxxWrap package.

C name	Fortran name	Standard Julia Alias	Julia Base Type
`unsigned char`	`CHARACTER`	`Cuchar`	`UInt8`
`bool` (only in C++)		`Cuchar`	`UInt8`
`short`	`INTEGER*2,` `LOGICAL*2`	`Cshort`	`Int16`
`unsigned short`		`Cushort`	`UInt16`
`int`, `BOOL` (C, typical)	`INTEGER*4,` `LOGICAL*4`	`Cint`	`Int32`
`unsigned int`		`Cuint`	`UInt32`
`long long`	`INTEGER*8,` `LOGICAL*8`	`Clong-` `long`	`Int64`
`unsigned long long`		`Culong-` `long`	`UInt64`
`intmax_t`		`Cint-` `max_t`	`Int64`
`uintmax_t`		`Cuint-` `max_t`	`UInt64`
`float`	`REAL*4i`	`Cfloat`	`Float32`
`double`	`REAL*8`	`Cdouble`	`Float64`
`complex float`	`COMPLEX*8`	`Com-` `plex64`	`Complex{Float32}`
`complex double`	`COM-` `PLEX*16`	`Com-` `plex128`	`Complex{Float64}`
`ptrdiff_t`		`Cp-` `trdiff_t`	`Int`
`ssize_t`		`Cs-` `size_t`	`Int`
`size_t`		`Csize_t`	`UInt`
`void`			`Void`
`void` and `[[noreturn]]` or `_Noreturn`			`Union{}`
`void*`			`Ptr{Void}`
`T*` (where T represents an appropriately defined type)			`Ref{T}`
`char*` (or `char[]`, e.g. a string)	`CHARAC-` `TER*N`		`Cstring` if NUL-terminated, or `Ptr{UInt8}` if not
`char**` (or `*char[]`)			`Ptr{Ptr{UInt8}}`
`jl_value_t*` (any Julia Type)			`Any`
`jl_value_t**` (a reference to a Julia Type)			`Ref{Any}`
`va_arg`			Not supported
`...` (variadic function specification)			`T...` (where T is one of the above types, variadic functions of different argument types are not supported)

C name	Standard Julia Alias	Julia Base Type
`char`	`Cchar`	`Int8` (x86, x86_64), `UInt8` (powerpc, arm)
`long`	`Clong`	`Int` (UNIX), `Int32` (Windows)
`unsigned long`	`Culong`	`UInt` (UNIX), `UInt32` (Windows)
`wchar_t`	`Cwchar_t`	`Int32` (UNIX), `UInt16` (Windows)

Chapter 31

Handling Operating System Variation

When dealing with platform libraries, it is often necessary to provide special cases for various platforms. The variable `Sys.KERNEL` can be used to write these special cases. There are several functions intended to make this easier: `is_unix`, `is_linux`, `is_apple`, `is_bsd`, and `is_windows`. These may be used as follows:

```
if is_windows()
    some_complicated_thing(a)
end
```

Note that `is_linux` and `is_apple` are mutually exclusive subsets of `is_unix`. Additionally, there is a macro `@static` which makes it possible to use these functions to conditionally hide invalid code, as demonstrated in the following examples.

Simple blocks:

```
ccall( (@static is_windows() ? :_fopen : :fopen), ...)
```

Complex blocks:

```
@static if is_linux()
    some_complicated_thing(a)
else
    some_different_thing(a)
end
```

When chaining conditionals (including if/elseif/end), the `@static` must be repeated for each level (parentheses optional, but recommended for readability):

```
@static is_windows() ? :a : (@static is_apple() ? :b : :c)
```

Chapter 32

Environment Variables

Julia may be configured with a number of environment variables, either in the usual way of the operating system, or in a portable way from within Julia. Suppose you want to set the environment variable JULIA_EDITOR to vim, then either type ENV["JULIA_EDITOR"] = "vim" for instance in the REPL to make this change on a case by case basis, or add the same to the user configuration file .juliarc.jl in the user's home directory to have a permanent effect. The current value of the same environment variable is determined by evaluating ENV["JULIA_EDITOR"].

The environment variables that Julia uses generally start with JULIA. If Base.versioninfo is called with verbose equal to true, then the output will list defined environment variables relevant for Julia, including those for which JULIA appears in the name.

32.1   File locations

JULIA_HOME

The absolute path of the directory containing the Julia executable, which sets the global variable Base.JULIA_HOME. If $JULIA_HOME is not set, then Julia determines the value Base.JULIA_HOME at run-time.

The executable itself is one of

```
$JULIA_HOME/julia
$JULIA_HOME/julia-debug
```

by default.

The global variable Base.DATAROOTDIR determines a relative path from Base.JULIA_HOME to the data directory associated with Julia. Then the path

```
$JULIA_HOME/$DATAROOTDIR/julia/base
```

determines the directory in which Julia initially searches for source files (via Base.find_source_file()).

Likewise, the global variable Base.SYSCONFDIR determines a relative path to the configuration file directory. Then Julia searches for a juliarc.jl file at

```
$JULIA_HOME/$SYSCONFDIR/julia/juliarc.jl
$JULIA_HOME/../etc/julia/juliarc.jl
```

by default (via Base.load_juliarc()).

For example, a Linux installation with a Julia executable located at /bin/julia, a DATAROOTDIR of ../share, and a SYSCONFDIR of ../etc will have JULIA_HOME set to /bin, a source-file search path of

```
/share/julia/base
```

and a global configuration search path of

```
/etc/julia/juliarc.jl
```

JULIA_LOAD_PATH

A separated list of absolute paths that are to be appended to the variable LOAD_PATH. (In Unix-like systems, the path separator is :; in Windows systems, the path separator is ;.) The LOAD_PATH variable is where Base.require and Base.load_in_path() look for code; it defaults to the absolute paths

```
$JULIA_HOME/../local/share/julia/site/v$(VERSION.major).$(VERSION.minor)
$JULIA_HOME/../share/julia/site/v$(VERSION.major).$(VERSION.minor)
```

so that, e.g., version 0.6 of Julia on a Linux system with a Julia executable at /bin/julia will have a default LOAD_PATH of

```
/local/share/julia/site/v0.6
/share/julia/site/v0.6
```

JULIA_PKGDIR

The path of the parent directory Pkg.Dir._pkgroot() for the version-specific Julia package repositories. If the path is relative, then it is taken with respect to the working directory. If $JULIA_PKGDIR is not set, then Pkg.Dir._pkgroot() defaults to

```
$HOME/.julia
```

Then the repository location Pkg.dir for a given Julia version is

```
$JULIA_PKGDIR/v$(VERSION.major).$(VERSION.minor)
```

For example, for a Linux user whose home directory is /home/alice, the directory containing the package repositories would by default be

```
/home/alice/.julia
```

and the package repository for version 0.6 of Julia would be

```
/home/alice/.julia/v0.6
```

JULIA_HISTORY

The absolute path Base.REPL.find_hist_file() of the REPL's history file. If $JULIA_HISTORY is not set, then Base.REPL.find_hist_file() defaults to

```
$HOME/.julia_history
```

JULIA_PKGRESOLVE_ACCURACY

A positive Int that determines how much time the max-sum subroutine MaxSum.maxsum() of the package dependency resolver Base.Pkg.resolve will devote to attempting satisfying constraints before giving up: this value is by default 1, and larger values correspond to larger amounts of time.

Suppose the value of $JULIA_PKGRESOLVE_ACCURACY is n. Then

- the number of pre-decimation iterations is 20*n,

- the number of iterations between decimation steps is 10*n, and

- at decimation steps, at most one in every 20*n packages is decimated.

32.2   External applications

JULIA_SHELL

The absolute path of the shell with which Julia should execute external commands (via `Base.repl_cmd()`). Defaults to the environment variable $SHELL, and falls back to `/bin/sh` if $SHELL is unset.

> **Note**
>
> On Windows, this environment variable is ignored, and external commands are executed directly.

JULIA_EDITOR

The editor returned by `Base.editor()` and used in, e.g., `Base.edit`, referring to the command of the preferred editor, for instance `vim`.

$JULIA_EDITOR takes precedence over $VISUAL, which in turn takes precedence over $EDITOR. If none of these environment variables is set, then the editor is taken to be open on Windows and OS X, or `/etc/alternatives/editor` if it exists, or `emacs` otherwise.

> **Note**
>
> $JULIA_EDITOR is *not* used in the determination of the editor for `Base.Pkg.edit`: this function checks $VISUAL and $EDITOR alone.

32.3   Parallelization

JULIA_CPU_CORES

Overrides the global variable `Base.Sys.CPU_CORES`, the number of logical CPU cores available.

JULIA_WORKER_TIMEOUT

A `Float64` that sets the value of `Base.worker_timeout()` (default: `60.0`). This function gives the number of seconds a worker process will wait for a master process to establish a connection before dying.

JULIA_NUM_THREADS

An unsigned 64-bit integer (`uint64_t`) that sets the maximum number of threads available to Julia. If $JULIA_NUM_THREADS exceeds the number of available physical CPU cores, then the number of threads is set to the number of cores. If $JULIA_NUM_THREADS is not positive or is not set, or if the number of CPU cores cannot be determined through system calls, then the number of threads is set to 1.

JULIA_THREAD_SLEEP_THRESHOLD

If set to a string that starts with the case-insensitive substring `"infinite"`, then spinning threads never sleep. Otherwise, $JULIA_THREAD_SLEEP_THRESHOLD is interpreted as an unsigned 64-bit integer (`uint64_t`) and gives, in nanoseconds, the amount of time after which spinning threads should sleep.

JULIA_EXCLUSIVE

If set to anything besides 0, then Julia's thread policy is consistent with running on a dedicated machine: the master thread is on proc 0, and threads are affinitized. Otherwise, Julia lets the operating system handle thread policy.

32.4   REPL formatting

Environment variables that determine how REPL output should be formatted at the terminal. Generally, these variables should be set to ANSI terminal escape sequences. Julia provides a high-level interface with much of the same functionality: see the section on Interacting With Julia.

JULIA_ERROR_COLOR

The formatting `Base.error_color()` (default: light red, `"\033[91m"`) that errors should have at the terminal.

JULIA_WARN_COLOR

The formatting `Base.warn_color()` (default: yellow, `"\033[93m"`) that warnings should have at the terminal.

JULIA_INFO_COLOR

The formatting `Base.info_color()` (default: cyan, `"\033[36m"`) that info should have at the terminal.

JULIA_INPUT_COLOR

The formatting `Base.input_color()` (default: normal, `"\033[0m"`) that input should have at the terminal.

JULIA_ANSWER_COLOR

The formatting `Base.answer_color()` (default: normal, `"\033[0m"`) that output should have at the terminal.

JULIA_STACKFRAME_LINEINFO_COLOR

The formatting `Base.stackframe_lineinfo_color()` (default: bold, `"\033[1m"`) that line info should have during a stack trace at the terminal.

JULIA_STACKFRAME_FUNCTION_COLOR

The formatting `Base.stackframe_function_color()` (default: bold, `"\033[1m"`) that function calls should have during a stack trace at the terminal.

32.5   Debugging and profiling

JULIA_GC_ALLOC_POOL, JULIA_GC_ALLOC_OTHER, JULIA_GC_ALLOC_PRINT

If set, these environment variables take strings that optionally start with the character `'r'`, followed by a string interpolation of a colon-separated list of three signed 64-bit integers (`int64_t`). This triple of integers `a:b:c` represents the arithmetic sequence a, a + b, a + 2*b, ... c.

- If it's the nth time that `jl_gc_pool_alloc()` has been called, and n belongs to the arithmetic sequence represented by $JULIA_GC_ALLOC_POOL, then garbage collection is forced.

- If it's the nth time that `maybe_collect()` has been called, and n belongs to the arithmetic sequence represented by $JULIA_GC_ALLOC_OTHER, then garbage collection is forced.

- If it's the nth time that `jl_gc_collect()` has been called, and n belongs to the arithmetic sequence represented by $JULIA_GC_ALLOC_PRINT, then counts for the number of calls to `jl_gc_pool_alloc()` and `maybe_collect()` are printed.

If the value of the environment variable begins with the character `'r'`, then the interval between garbage collection events is randomized.

> **Note**
>
> These environment variables only have an effect if Julia was compiled with garbage-collection debugging (that is, if `WITH_GC_DEBUG_ENV` is set to 1 in the build configuration).

JULIA_GC_NO_GENERATIONAL

If set to anything besides 0, then the Julia garbage collector never performs "quick sweeps" of memory.

> **Note**
>
> This environment variable only has an effect if Julia was compiled with garbage-collection debugging (that is, if `WITH_GC_DEBUG_ENV` is set to 1 in the build configuration).

JULIA_GC_WAIT_FOR_DEBUGGER

If set to anything besides 0, then the Julia garbage collector will wait for a debugger to attach instead of aborting whenever there's a critical error.

> **Note**
>
> This environment variable only has an effect if Julia was compiled with garbage-collection debugging (that is, if `WITH_GC_DEBUG_ENV` is set to 1 in the build configuration).

ENABLE_JITPROFILING

If set to anything besides 0, then the compiler will create and register an event listener for just-in-time (JIT) profiling.

> **Note**
>
> This environment variable only has an effect if Julia was compiled with JIT profiling support, using either

- Intel's VTune™ Amplifier (`USE_INTEL_JITEVENTS` set to 1 in the build configuration), or
- OProfile (`USE_OPROFILE_JITEVENTS` set to 1 in the build configuration).

JULIA_LLVM_ARGS

Arguments to be passed to the LLVM backend.

> **Note**
>
> This environment variable has an effect only if Julia was compiled with `JL_DEBUG_BUILD` set — in particular, the `julia-debug` executable is always compiled with this build variable.

JULIA_DEBUG_LOADING

If set, then Julia prints detailed information about the cache in the loading process of `Base.require`.

Chapter 33

Embedding Julia

As we have seen in Calling C and Fortran Code, Julia has a simple and efficient way to call functions written in C. But there are situations where the opposite is needed: calling Julia function from C code. This can be used to integrate Julia code into a larger C/C++ project, without the need to rewrite everything in C/C++. Julia has a C API to make this possible. As almost all programming languages have some way to call C functions, the Julia C API can also be used to build further language bridges (e.g. calling Julia from Python or C#).

33.1   High-Level Embedding

We start with a simple C program that initializes Julia and calls some Julia code:

```
#include <julia.h>

int main(int argc, char *argv[])
{
    /* required: setup the Julia context */
    jl_init();

    /* run Julia commands */
    jl_eval_string("print(sqrt(2.0))");

    /* strongly recommended: notify Julia that the
        program is about to terminate. this allows
        Julia time to cleanup pending write requests
        and run all finalizers
    */
    jl_atexit_hook(0);
    return 0;
}
```

In order to build this program you have to put the path to the Julia header into the include path and link against `libjulia`. For instance, when Julia is installed to $JULIA_DIR, one can compile the above test program `test.c` with gcc using:

```
gcc -o test -fPIC -I$JULIA_DIR/include/julia -L$JULIA_DIR/lib test.c -ljulia $JULIA_DIR/lib/julia
    /libstdc++.so.6
```

Then if the environment variable JULIA_HOME is set to $JULIA_DIR/bin, the output `test` program can be executed.

Alternatively, look at the `embedding.c` program in the Julia source tree in the `examples/` folder. The file `ui/repl.c` program is another simple example of how to set `jl_options` options while linking against `libjulia`.

The first thing that has to be done before calling any other Julia C function is to initialize Julia. This is done by calling `jl_init`, which tries to automatically determine Julia's install location. If you need to specify a custom location, or specify which system image to load, use `jl_init_with_image` instead.

The second statement in the test program evaluates a Julia statement using a call to `jl_eval_string`.

Before the program terminates, it is strongly recommended to call `jl_atexit_hook`. The above example program calls this before returning from `main`.

> **Note**
>
> Currently, dynamically linking with the `libjulia` shared library requires passing the RTLD_GLOBAL option. In Python, this looks like:
>
> ```
> >>> julia=CDLL('./libjulia.dylib',RTLD_GLOBAL)
> >>> julia.jl_init.argtypes = []
> >>> julia.jl_init()
> 250593296
> ```

> **Note**
>
> If the julia program needs to access symbols from the main executable, it may be necessary to add `-Wl,--export-dynamic` linker flag at compile time on Linux in addition to the ones generated by `julia-config.jl` described below. This is not necessary when compiling a shared library.

Using julia-config to automatically determine build parameters

The script `julia-config.jl` was created to aid in determining what build parameters are required by a program that uses embedded Julia. This script uses the build parameters and system configuration of the particular Julia distribution it is invoked by to export the necessary compiler flags for an embedding program to interact with that distribution. This script is located in the Julia shared data directory.

Example

```
#include <julia.h>

int main(int argc, char *argv[])
{
    jl_init();
    (void)jl_eval_string("println(sqrt(2.0))");
    jl_atexit_hook(0);
    return 0;
}
```

On the command line

A simple use of this script is from the command line. Assuming that `julia-config.jl` is located in `/usr/local/julia/share/julia`, it can be invoked on the command line directly and takes any combination of 3 flags:

```
/usr/local/julia/share/julia/julia-config.jl
Usage: julia-config [--cflags|--ldflags|--ldlibs]
```

If the above example source is saved in the file `embed_example.c`, then the following command will compile it into a running program on Linux and Windows (MSYS2 environment), or if on OS/X, then substitute `clang` for `gcc`.:

```
/usr/local/julia/share/julia/julia-config.jl --cflags --ldflags --ldlibs | xargs gcc
    embed_example.c
```

Use in Makefiles

But in general, embedding projects will be more complicated than the above, and so the following allows general makefile support as well – assuming GNU make because of the use of the **shell** macro expansions. Additionally, though many times `julia-config.jl` may be found in the directory `/usr/local`, this is not necessarily the case, but Julia can be used to locate `julia-config.jl` too, and the makefile can be used to take advantage of that. The above example is extended to use a Makefile:

```
JL_SHARE = $(shell julia -e 'print(joinpath(JULIA_HOME,Base.DATAROOTDIR,"julia"))')
CFLAGS   += $(shell $(JL_SHARE)/julia-config.jl --cflags)
CXXFLAGS += $(shell $(JL_SHARE)/julia-config.jl --cflags)
LDFLAGS  += $(shell $(JL_SHARE)/julia-config.jl --ldflags)
LDLIBS   += $(shell $(JL_SHARE)/julia-config.jl --ldlibs)

all: embed_example
```

Now the build command is simply `make`.

33.2 Converting Types

Real applications will not just need to execute expressions, but also return their values to the host program. `jl_eval_string` returns a `jl_value_t*`, which is a pointer to a heap-allocated Julia object. Storing simple data types like Float64 in this way is called `boxing`, and extracting the stored primitive data is called `unboxing`. Our improved sample program that calculates the square root of 2 in Julia and reads back the result in C looks as follows:

```
jl_value_t *ret = jl_eval_string("sqrt(2.0)");

if (jl_typeis(ret, jl_float64_type)) {
    double ret_unboxed = jl_unbox_float64(ret);
    printf("sqrt(2.0) in C: %e \n", ret_unboxed);
}
else {
    printf("ERROR: unexpected return type from sqrt(::Float64)\n");
}
```

In order to check whether `ret` is of a specific Julia type, we can use the `jl_isa`, `jl_typeis`, or `jl_is_...` functions. By typing `typeof(sqrt(2.0))` into the Julia shell we can see that the return type is Float64 (double in C). To convert the boxed Julia value into a C double the `jl_unbox_float64` function is used in the above code snippet.

Corresponding `jl_box_...` functions are used to convert the other way:

```
jl_value_t *a = jl_box_float64(3.0);
jl_value_t *b = jl_box_float32(3.0f);
jl_value_t *c = jl_box_int32(3);
```

As we will see next, boxing is required to call Julia functions with specific arguments.

33.3 Calling Julia Functions

While `jl_eval_string` allows C to obtain the result of a Julia expression, it does not allow passing arguments computed in C to Julia. For this you will need to invoke Julia functions directly, using `jl_call`:

```
jl_function_t *func = jl_get_function(jl_base_module, "sqrt");
jl_value_t *argument = jl_box_float64(2.0);
jl_value_t *ret = jl_call1(func, argument);
```

In the first step, a handle to the Julia function `sqrt` is retrieved by calling `jl_get_function`. The first argument passed to `jl_get_function` is a pointer to the Base module in which `sqrt` is defined. Then, the double value is boxed using `jl_box_float64`. Finally, in the last step, the function is called using `jl_call1`. `jl_call0`, `jl_call2`, and `jl_call3` functions also exist, to conveniently handle different numbers of arguments. To pass more arguments, use `jl_call`:

```
jl_value_t *jl_call(jl_function_t *f, jl_value_t **args, int32_t nargs)
```

Its second argument `args` is an array of `jl_value_t*` arguments and `nargs` is the number of arguments.

33.4   Memory Management

As we have seen, Julia objects are represented in C as pointers. This raises the question of who is responsible for freeing these objects.

Typically, Julia objects are freed by a garbage collector (GC), but the GC does not automatically know that we are holding a reference to a Julia value from C. This means the GC can free objects out from under you, rendering pointers invalid.

The GC can only run when Julia objects are allocated. Calls like `jl_box_float64` perform allocation, and allocation might also happen at any point in running Julia code. However, it is generally safe to use pointers in between `jl_...` calls. But in order to make sure that values can survive `jl_...` calls, we have to tell Julia that we hold a reference to a Julia value. This can be done using the JL_GC_PUSH macros:

```
jl_value_t *ret = jl_eval_string("sqrt(2.0)");
JL_GC_PUSH1(&ret);
// Do something with ret
JL_GC_POP();
```

The JL_GC_POP call releases the references established by the previous JL_GC_PUSH. Note that JL_GC_PUSH is working on the stack, so it must be exactly paired with a JL_GC_POP before the stack frame is destroyed.

Several Julia values can be pushed at once using the JL_GC_PUSH2 , JL_GC_PUSH3 , and JL_GC_PUSH4 macros. To push an array of Julia values one can use the JL_GC_PUSHARGS macro, which can be used as follows:

```
jl_value_t **args;
JL_GC_PUSHARGS(args, 2); // args can now hold 2 `jl_value_t*` objects
args[0] = some_value;
args[1] = some_other_value;
// Do something with args (e.g. call jl_... functions)
JL_GC_POP();
```

The garbage collector also operates under the assumption that it is aware of every old-generation object pointing to a young-generation one. Any time a pointer is updated breaking that assumption, it must be signaled to the collector with the `jl_gc_wb` (write barrier) function like so:

```
jl_value_t *parent = some_old_value, *child = some_young_value;
((some_specific_type*)parent)->field = child;
jl_gc_wb(parent, child);
```

It is in general impossible to predict which values will be old at runtime, so the write barrier must be inserted after all explicit stores. One notable exception is if the `parent` object was just allocated and garbage collection was not run since then. Remember that most `jl_...` functions can sometimes invoke garbage collection.

The write barrier is also necessary for arrays of pointers when updating their data directly. For example:

```
jl_array_t *some_array = ...; // e.g. a Vector{Any}
void **data = (void**)jl_array_data(some_array);
jl_value_t *some_value = ...;
data[0] = some_value;
jl_gc_wb(some_array, some_value);
```

Manipulating the Garbage Collector

There are some functions to control the GC. In normal use cases, these should not be necessary.

Function	Description
`jl_gc_collect()`	Force a GC run
`jl_gc_enable(0)`	Disable the GC, return previous state as int
`jl_gc_enable(1)`	Enable the GC, return previous state as int
`jl_gc_is_enabled()`	Return current state as int

33.5 Working with Arrays

Julia and C can share array data without copying. The next example will show how this works.

Julia arrays are represented in C by the datatype `jl_array_t*`. Basically, `jl_array_t` is a struct that contains:

- Information about the datatype

- A pointer to the data block

- Information about the sizes of the array

To keep things simple, we start with a 1D array. Creating an array containing Float64 elements of length 10 is done by:

```
jl_value_t* array_type = jl_apply_array_type(jl_float64_type, 1);
jl_array_t* x          = jl_alloc_array_1d(array_type, 10);
```

Alternatively, if you have already allocated the array you can generate a thin wrapper around its data:

```
double *existingArray = (double*)malloc(sizeof(double)*10);
jl_array_t *x = jl_ptr_to_array_1d(array_type, existingArray, 10, 0);
```

The last argument is a boolean indicating whether Julia should take ownership of the data. If this argument is non-zero, the GC will call `free` on the data pointer when the array is no longer referenced.

In order to access the data of x, we can use `jl_array_data`:

```
double *xData = (double*)jl_array_data(x);
```

Now we can fill the array:

```
for(size_t i=0; i<jl_array_len(x); i++)
    xData[i] = i;
```

Now let us call a Julia function that performs an in-place operation on x:

```
jl_function_t *func = jl_get_function(jl_base_module, "reverse!");
jl_call1(func, (jl_value_t*)x);
```

By printing the array, one can verify that the elements of x are now reversed.

Accessing Returned Arrays

If a Julia function returns an array, the return value of `jl_eval_string` and `jl_call` can be cast to a `jl_array_t*`:

```
jl_function_t *func  = jl_get_function(jl_base_module, "reverse");
jl_array_t *y = (jl_array_t*)jl_call1(func, (jl_value_t*)x);
```

Now the content of y can be accessed as before using `jl_array_data`. As always, be sure to keep a reference to the array while it is in use.

Multidimensional Arrays

Julia's multidimensional arrays are stored in memory in column-major order. Here is some code that creates a 2D array and accesses its properties:

```
// Create 2D array of float64 type
jl_value_t *array_type = jl_apply_array_type(jl_float64_type, 2);
jl_array_t *x  = jl_alloc_array_2d(array_type, 10, 5);

// Get array pointer
double *p = (double*)jl_array_data(x);
// Get number of dimensions
int ndims = jl_array_ndims(x);
// Get the size of the i-th dim
size_t size0 = jl_array_dim(x,0);
size_t size1 = jl_array_dim(x,1);

// Fill array with data
for(size_t i=0; i<size1; i++)
    for(size_t j=0; j<size0; j++)
        p[j + size0*i] = i + j;
```

Notice that while Julia arrays use 1-based indexing, the C API uses 0-based indexing (for example in calling `jl_array_dim`) in order to read as idiomatic C code.

33.6   Exceptions

Julia code can throw exceptions. For example, consider:

```
jl_eval_string("this_function_does_not_exist()");
```

This call will appear to do nothing. However, it is possible to check whether an exception was thrown:

```
if (jl_exception_occurred())
    printf("%s \n", jl_typeof_str(jl_exception_occurred()));
```

If you are using the Julia C API from a language that supports exceptions (e.g. Python, C#, C++), it makes sense to wrap each call into `libjulia` with a function that checks whether an exception was thrown, and then rethrows the exception in the host language.

Throwing Julia Exceptions

When writing Julia callable functions, it might be necessary to validate arguments and throw exceptions to indicate errors. A typical type check looks like:

```
if (!jl_typeis(val, jl_float64_type)) {
    jl_type_error(function_name, (jl_value_t*)jl_float64_type, val);
}
```

General exceptions can be raised using the functions:

```
void jl_error(const char *str);
void jl_errorf(const char *fmt, ...);
```

`jl_error` takes a C string, and `jl_errorf` is called like `printf`:

```
jl_errorf("argument x = %d is too large", x);
```

where in this example x is assumed to be an integer.

Chapter 34

Packages

Julia has a built-in package manager for installing add-on functionality written in Julia. It can also install external libraries using your operating system's standard system for doing so, or by compiling from source. The list of registered Julia packages can be found at http://pkg.julialang.org. All package manager commands are found in the Pkg module, included in Julia's Base install.

First we'll go over the mechanics of the Pkg family of commands and then we'll provide some guidance on how to get your package registered. Be sure to read the section below on package naming conventions, tagging versions and the importance of a REQUIRE file for when you're ready to add your code to the curated METADATA repository.

34.1 Package Status

The `Pkg.status()` function prints out a summary of the state of packages you have installed. Initially, you'll have no packages installed:

```
julia> Pkg.status()
INFO: Initializing package repository /Users/stefan/.julia/v0.6
INFO: Cloning METADATA from git://github.com/JuliaLang/METADATA.jl
No packages installed.
```

Your package directory is automatically initialized the first time you run a Pkg command that expects it to exist – which includes `Pkg.status()`. Here's an example non-trivial set of required and additional packages:

```
julia> Pkg.status()
Required packages:
 - Distributions            0.2.8
 - SHA                      0.3.2
Additional packages:
 - NumericExtensions        0.2.17
 - Stats                    0.2.6
```

These packages are all on registered versions, managed by Pkg. Packages can be in more complicated states, indicated by annotations to the right of the installed package version; we will explain these states and annotations as we encounter them. For programmatic usage, `Pkg.installed()` returns a dictionary, mapping installed package names to the version of that package which is installed:

```
julia> Pkg.installed()
Dict{String,VersionNumber} with 4 entries:
```

```
"Distributions"    => v"0.2.8"
"Stats"            => v"0.2.6"
"SHA"              => v"0.3.2"
"NumericExtensions" => v"0.2.17"
```

34.2   Adding and Removing Packages

Julia's package manager is a little unusual in that it is declarative rather than imperative. This means that you tell it what you want and it figures out what versions to install (or remove) to satisfy those requirements optimally – and minimally. So rather than installing a package, you just add it to the list of requirements and then "resolve" what needs to be installed. In particular, this means that if some package had been installed because it was needed by a previous version of something you wanted, and a newer version doesn't have that requirement anymore, updating will actually remove that package.

Your package requirements are in the file ~/.julia/v0.6/REQUIRE. You can edit this file by hand and then call Pkg.resolve() to install, upgrade or remove packages to optimally satisfy the requirements, or you can do Pkg.edit(), which will open REQUIRE in your editor (configured via the EDITOR or VISUAL environment variables), and then automatically call Pkg.resolve() afterwards if necessary. If you only want to add or remove the requirement for a single package, you can also use the non-interactive Pkg.add() and Pkg.rm() commands, which add or remove a single requirement to REQUIRE and then call Pkg.resolve().

You can add a package to the list of requirements with the Pkg.add() function, and the package and all the packages that it depends on will be installed:

```
julia> Pkg.status()
No packages installed.

julia> Pkg.add("Distributions")
INFO: Cloning cache of Distributions from git://github.com/JuliaStats/Distributions.jl.git
INFO: Cloning cache of NumericExtensions from git://github.com/lindahua/NumericExtensions.jl.git
INFO: Cloning cache of Stats from git://github.com/JuliaStats/Stats.jl.git
INFO: Installing Distributions v0.2.7
INFO: Installing NumericExtensions v0.2.17
INFO: Installing Stats v0.2.6
INFO: REQUIRE updated.

julia> Pkg.status()
Required packages:
 - Distributions                0.2.7
Additional packages:
 - NumericExtensions            0.2.17
 - Stats                        0.2.6
```

What this is doing is first adding Distributions to your ~/.julia/v0.6/REQUIRE file:

```
$ cat ~/.julia/v0.6/REQUIRE
Distributions
```

It then runs Pkg.resolve() using these new requirements, which leads to the conclusion that the Distributions package should be installed since it is required but not installed. As stated before, you can accomplish the same thing by editing your ~/.julia/v0.6/REQUIRE file by hand and then running Pkg.resolve() yourself:

```
$ echo SHA >> ~/.julia/v0.6/REQUIRE
```

```
julia> Pkg.resolve()
INFO: Cloning cache of SHA from git://github.com/staticfloat/SHA.jl.git
INFO: Installing SHA v0.3.2

julia> Pkg.status()
Required packages:
 - Distributions             0.2.7
 - SHA                       0.3.2
Additional packages:
 - NumericExtensions         0.2.17
 - Stats                     0.2.6
```

This is functionally equivalent to calling `Pkg.add("SHA")`, except that `Pkg.add()` doesn't change REQUIRE until *after* installation has completed, so if there are problems, REQUIRE will be left as it was before calling `Pkg.add()`. The format of the REQUIRE file is described in Requirements Specification; it allows, among other things, requiring specific ranges of versions of packages.

When you decide that you don't want to have a package around any more, you can use `Pkg.rm()` to remove the requirement for it from the REQUIRE file:

```
julia> Pkg.rm("Distributions")
INFO: Removing Distributions v0.2.7
INFO: Removing Stats v0.2.6
INFO: Removing NumericExtensions v0.2.17
INFO: REQUIRE updated.

julia> Pkg.status()
Required packages:
 - SHA                       0.3.2

julia> Pkg.rm("SHA")
INFO: Removing SHA v0.3.2
INFO: REQUIRE updated.

julia> Pkg.status()
No packages installed.
```

Once again, this is equivalent to editing the REQUIRE file to remove the line with each package name on it then running `Pkg.resolve()` to update the set of installed packages to match. While `Pkg.add()` and `Pkg.rm()` are convenient for adding and removing requirements for a single package, when you want to add or remove multiple packages, you can call `Pkg.edit()` to manually change the contents of REQUIRE and then update your packages accordingly. `Pkg.edit()` does not roll back the contents of REQUIRE if `Pkg.resolve()` fails – rather, you have to run `Pkg.edit()` again to fix the files contents yourself.

Because the package manager uses libgit2 internally to manage the package git repositories, users may run into protocol issues (if behind a firewall, for example), when running `Pkg.add()`. By default, all GitHub-hosted packages wil be accessed via 'https'; this default can be modified by calling `Pkg.setprotocol!()`. The following command can be run from the command line in order to tell git to use 'https' instead of the 'git' protocol when cloning all repositories, wherever they are hosted:

```
git config --global url."https://".insteadOf git://
```

However, this change will be system-wide and thus the use of `Pkg.setprotocol!()` is preferable.

Note

The package manager functions also accept the `.jl` suffix on package names, though the suffix is stripped internally. For example:

```
Pkg.add("Distributions.jl")
Pkg.rm("Distributions.jl")
```

34.3   Offline Installation of Packages

For machines with no Internet connection, packages may be installed by copying the package root directory (given by `Pkg.dir()`) from a machine with the same operating system and environment.

`Pkg.add()` does the following within the package root directory:

1. Adds the name of the package to REQUIRE.

2. Downloads the package to `.cache`, then copies the package to the package root directory.

3. Recursively performs step 2 against all the packages listed in the package's REQUIRE file.

4. Runs `Pkg.build()`

Warning

Copying installed packages from a different machine is brittle for packages requiring binary external dependencies. Such packages may break due to differences in operating system versions, build environments, and/or absolute path dependencies.

34.4   Installing Unregistered Packages

Julia packages are simply git repositories, clonable via any of the protocols that git supports, and containing Julia code that follows certain layout conventions. Official Julia packages are registered in the METADATA.jl repository, available at a well-known location [1]. The `Pkg.add()` and `Pkg.rm()` commands in the previous section interact with registered packages, but the package manager can install and work with unregistered packages too. To install an unregistered package, use `Pkg.clone(url)`, where `url` is a git URL from which the package can be cloned:

```
julia> Pkg.clone("git://example.com/path/to/Package.jl.git")
INFO: Cloning Package from git://example.com/path/to/Package.jl.git
Cloning into 'Package'...
remote: Counting objects: 22, done.
remote: Compressing objects: 100% (10/10), done.
remote: Total 22 (delta 8), reused 22 (delta 8)
Receiving objects: 100% (22/22), 2.64 KiB, done.
Resolving deltas: 100% (8/8), done.
```

By convention, Julia repository names end with `.jl` (the additional `.git` indicates a "bare" git repository), which keeps them from colliding with repositories for other languages, and also makes Julia packages easy to find in search engines. When packages are installed in your `.julia/v0.6` directory, however, the extension is redundant so we leave it off.

If unregistered packages contain a REQUIRE file at the top of their source tree, that file will be used to determine which registered packages the unregistered package depends on, and they will automatically be installed. Unregistered

packages participate in the same version resolution logic as registered packages, so installed package versions will be adjusted as necessary to satisfy the requirements of both registered and unregistered packages.

34.5 Updating Packages

When package developers publish new registered versions of packages that you're using, you will, of course, want the new shiny versions. To get the latest and greatest versions of all your packages, just do `Pkg.update()`:

```julia
julia> Pkg.update()
INFO: Updating METADATA...
INFO: Computing changes...
INFO: Upgrading Distributions: v0.2.8 => v0.2.10
INFO: Upgrading Stats: v0.2.7 => v0.2.8
```

The first step of updating packages is to pull new changes to `~/.julia/v0.6/METADATA` and see if any new registered package versions have been published. After this, `Pkg.update()` attempts to update packages that are checked out on a branch and not dirty (i.e. no changes have been made to files tracked by git) by pulling changes from the package's upstream repository. Upstream changes will only be applied if no merging or rebasing is necessary – i.e. if the branch can be "fast-forwarded". If the branch cannot be fast-forwarded, it is assumed that you're working on it and will update the repository yourself.

Finally, the update process recomputes an optimal set of package versions to have installed to satisfy your top-level requirements and the requirements of "fixed" packages. A package is considered fixed if it is one of the following:

1. **Unregistered:** the package is not in `METADATA` – you installed it with `Pkg.clone()`.

2. **Checked out:** the package repo is on a development branch.

3. **Dirty:** changes have been made to files in the repo.

If any of these are the case, the package manager cannot freely change the installed version of the package, so its requirements must be satisfied by whatever other package versions it picks. The combination of top-level requirements in `~/.julia/v0.6/REQUIRE` and the requirement of fixed packages are used to determine what should be installed.

You can also update only a subset of the installed packages, by providing arguments to the `Pkg.update` function. In that case, only the packages provided as arguments and their dependencies will be updated:

```julia
julia> Pkg.update("Example")
INFO: Updating METADATA...
INFO: Computing changes...
INFO: Upgrading Example: v0.4.0 => 0.4.1
```

This partial update process still computes the new set of package versions according to top-level requirements and "fixed" packages, but it additionally considers all other packages except those explicitly provided, and their dependencies, as fixed.

---

[1]The official set of packages is at https://github.com/JuliaLang/METADATA.jl, but individuals and organizations can easily use a different metadata repository. This allows control which packages are available for automatic installation. One can allow only audited and approved package versions, and make private packages or forks available. See Custom METADATA Repository for details.

34.6   Checkout, Pin and Free

You may want to use the `master` version of a package rather than one of its registered versions. There might be fixes or functionality on master that you need that aren't yet published in any registered versions, or you may be a developer of the package and need to make changes on `master` or some other development branch. In such cases, you can do `Pkg.checkout(pkg)` to checkout the `master` branch of pkg or `Pkg.checkout(pkg,branch)` to checkout some other branch:

```
julia> Pkg.add("Distributions")
INFO: Installing Distributions v0.2.9
INFO: Installing NumericExtensions v0.2.17
INFO: Installing Stats v0.2.7
INFO: REQUIRE updated.

julia> Pkg.status()
Required packages:
 - Distributions                  0.2.9
Additional packages:
 - NumericExtensions              0.2.17
 - Stats                          0.2.7

julia> Pkg.checkout("Distributions")
INFO: Checking out Distributions master...
INFO: No packages to install, update or remove.

julia> Pkg.status()
Required packages:
 - Distributions                  0.2.9+            master
Additional packages:
 - NumericExtensions              0.2.17
 - Stats                          0.2.7
```

Immediately after installing `Distributions` with `Pkg.add()` it is on the current most recent registered version – `0.2.9` at the time of writing this. Then after running `Pkg.checkout("Distributions")`, you can see from the output of `Pkg.status()` that `Distributions` is on an unregistered version greater than `0.2.9`, indicated by the "pseudo-version" number `0.2.9+`.

When you checkout an unregistered version of a package, the copy of the REQUIRE file in the package repo takes precedence over any requirements registered in METADATA, so it is important that developers keep this file accurate and up-to-date, reflecting the actual requirements of the current version of the package. If the REQUIRE file in the package repo is incorrect or missing, dependencies may be removed when the package is checked out. This file is also used to populate newly published versions of the package if you use the API that Pkg provides for this (described below).

When you decide that you no longer want to have a package checked out on a branch, you can "free" it back to the control of the package manager with `Pkg.free(pkg)`:

```
julia> Pkg.free("Distributions")
INFO: Freeing Distributions...
INFO: No packages to install, update or remove.

julia> Pkg.status()
Required packages:
 - Distributions                  0.2.9
Additional packages:
```

```
 - NumericExtensions          0.2.17
 - Stats                      0.2.7
```

After this, since the package is on a registered version and not on a branch, its version will be updated as new registered versions of the package are published.

If you want to pin a package at a specific version so that calling `Pkg.update()` won't change the version the package is on, you can use the `Pkg.pin()` function:

```
julia> Pkg.pin("Stats")
INFO: Creating Stats branch pinned.47c198b1.tmp

julia> Pkg.status()
Required packages:
 - Distributions              0.2.9
Additional packages:
 - NumericExtensions          0.2.17
 - Stats                      0.2.7              pinned.47c198b1.tmp
```

After this, the `Stats` package will remain pinned at version `0.2.7` – or more specifically, at commit `47c198b1`, but since versions are permanently associated a given git hash, this is the same thing. `Pkg.pin()` works by creating a throw-away branch for the commit you want to pin the package at and then checking that branch out. By default, it pins a package at the current commit, but you can choose a different version by passing a second argument:

```
julia> Pkg.pin("Stats",v"0.2.5")
INFO: Creating Stats branch pinned.1fd0983b.tmp
INFO: No packages to install, update or remove.

julia> Pkg.status()
Required packages:
 - Distributions              0.2.9
Additional packages:
 - NumericExtensions          0.2.17
 - Stats                      0.2.5              pinned.1fd0983b.tmp
```

Now the `Stats` package is pinned at commit `1fd0983b`, which corresponds to version `0.2.5`. When you decide to "unpin" a package and let the package manager update it again, you can use `Pkg.free()` like you would to move off of any branch:

```
julia> Pkg.free("Stats")
INFO: Freeing Stats...
INFO: No packages to install, update or remove.

julia> Pkg.status()
Required packages:
 - Distributions              0.2.9
Additional packages:
 - NumericExtensions          0.2.17
 - Stats                      0.2.7
```

After this, the `Stats` package is managed by the package manager again, and future calls to `Pkg.update()` will upgrade it to newer versions when they are published. The throw-away `pinned.1fd0983b.tmp` branch remains in

your local `Stats` repo, but since git branches are extremely lightweight, this doesn't really matter; if you feel like cleaning them up, you can go into the repo and delete those branches [2].

34.7   Custom METADATA Repository

By default, Julia assumes you will be using the official METADATA.jl repository for downloading and installing packages. You can also provide a different metadata repository location. A common approach is to keep your `metadata-v2` branch up to date with the Julia official branch and add another branch with your custom packages. You can initialize your local metadata repository using that custom location and branch and then periodically rebase your custom branch with the official `metadata-v2` branch. In order to use a custom repository and branch, issue the following command:

```julia
julia> Pkg.init("https://me.example.com/METADATA.jl.git", "branch")
```

The branch argument is optional and defaults to `metadata-v2`. Once initialized, a file named `META_BRANCH` in your `~/.julia/vX.Y/` path will track the branch that your METADATA repository was initialized with. If you want to change branches, you will need to either modify the `META_BRANCH` file directly (be careful!) or remove the `vX.Y` directory and re-initialize your METADATA repository using the `Pkg.init` command.

---

[2]Packages that aren't on branches will also be marked as dirty if you make changes in the repo, but that's a less common thing to do.

Chapter 35

Package Development

Julia's package manager is designed so that when you have a package installed, you are already in a position to look at its source code and full development history. You are also able to make changes to packages, commit them using git, and easily contribute fixes and enhancements upstream. Similarly, the system is designed so that if you want to create a new package, the simplest way to do so is within the infrastructure provided by the package manager.

35.1   Initial Setup

Since packages are git repositories, before doing any package development you should setup the following standard global git configuration settings:

```
$ git config --global user.name "FULL NAME"
$ git config --global user.email "EMAIL"
```

where FULL NAME is your actual full name (spaces are allowed between the double quotes) and EMAIL is your actual email address. Although it isn't necessary to use GitHub to create or publish Julia packages, most Julia packages as of writing this are hosted on GitHub and the package manager knows how to format origin URLs correctly and otherwise work with the service smoothly. We recommend that you create a free account on GitHub and then do:

```
$ git config --global github.user "USERNAME"
```

where USERNAME is your actual GitHub user name. Once you do this, the package manager knows your GitHub user name and can configure things accordingly. You should also upload your public SSH key to GitHub and set up an SSH agent on your development machine so that you can push changes with minimal hassle. In the future, we will make this system extensible and support other common git hosting options like BitBucket and allow developers to choose their favorite. Since the package development functions has been moved to the PkgDev package, you need to run Pkg.add("PkgDev"); import PkgDev to access the functions starting with PkgDev. in the document below.

35.2   Making changes to an existing package

Documentation changes

If you want to improve the online documentation of a package, the easiest approach (at least for small changes) is to use GitHub's online editing functionality. First, navigate to the repository's GitHub "home page," find the file (e.g., README.md) within the repository's folder structure, and click on it. You'll see the contents displayed, along with a small "pencil" icon in the upper right hand corner. Clicking that icon opens the file in edit mode. Make your changes, write a brief summary describing the changes you want to make (this is your *commit message*), and then hit "Propose file change." Your changes will be submitted for consideration by the package owner(s) and collaborators.

For larger documentation changes–and especially ones that you expect to have to update in response to feedback–you might find it easier to use the procedure for code changes described below.

Code changes

Executive summary

Here we assume you've already set up git on your local machine and have a GitHub account (see above). Let's imagine you're fixing a bug in the Images package:

```
Pkg.checkout("Images")            # check out the master branch
<here, make sure your bug is still a bug and hasn't been fixed already>
cd(Pkg.dir("Images"))
;git checkout -b myfixes          # create a branch for your changes
<edit code>                       # be sure to add a test for your bug
Pkg.test("Images")                # make sure everything works now
;git commit -a -m "Fix foo by calling bar"  # write a descriptive message
using PkgDev
PkgDev.submit("Images")
```

The last line will present you with a link to submit a pull request to incorporate your changes.

Detailed description

If you want to fix a bug or add new functionality, you want to be able to test your changes before you submit them for consideration. You also need to have an easy way to update your proposal in response to the package owner's feedback. Consequently, in this case the strategy is to work locally on your own machine; once you are satisfied with your changes, you submit them for consideration. This process is called a *pull request* because you are asking to "pull" your changes into the project's main repository. Because the online repository can't see the code on your private machine, you first *push* your changes to a publicly-visible location, your own online *fork* of the package (hosted on your own personal GitHub account).

Let's assume you already have the Foo package installed. In the description below, anything starting with `Pkg.` or `PkgDev.` is meant to be typed at the Julia prompt; anything starting with `git` is meant to be typed in julia's shell mode (or using the shell that comes with your operating system). Within Julia, you can combine these two modes:

```
julia> cd(Pkg.dir("Foo"))          # go to Foo's folder

shell> git command arguments...    # command will apply to Foo
```

Now suppose you're ready to make some changes to Foo. While there are several possible approaches, here is one that is widely used:

- From the Julia prompt, type `Pkg.checkout("Foo")`. This ensures you're running the latest code (the `master` branch), rather than just whatever "official release" version you have installed. (If you're planning to fix a bug, at this point it's a good idea to check again whether the bug has already been fixed by someone else. If it has, you can request that a new official release be tagged so that the fix gets distributed to the rest of the community.) If you receive an error `Foo is dirty, bailing`, see Dirty packages below.

- Create a branch for your changes: navigate to the package folder (the one that Julia reports from `Pkg.dir("Foo")`) and (in shell mode) create a new branch using `git checkout -b <newbranch>`, where <newbranch> might be some descriptive name (e.g., `fixbar`). By creating a branch, you ensure that you can easily go back and forth between your new work and the current `master` branch (see https://git-scm.com/book/en/v2/Git-Branching-Branches-in-a-Nutshell).

  If you forget to do this step until after you've already made some changes, don't worry: see more detail about branching below.

- Make your changes. Whether it's fixing a bug or adding new functionality, in most cases your change should include updates to both the `src/` and `test/` folders. If you're fixing a bug, add your minimal example demonstrating the bug (on the current code) to the test suite; by contributing a test for the bug, you ensure that the bug won't accidentally reappear at some later time due to other changes. If you're adding new functionality, creating tests demonstrates to the package owner that you've made sure your code works as intended.

- Run the package's tests and make sure they pass. There are several ways to run the tests:

  - From Julia, run `Pkg.test("Foo")`: this will run your tests in a separate (new) `julia` process.

  - From Julia, `include("runtests.jl")` from the package's `test/` folder (it's possible the file has a different name, look for one that runs all the tests): this allows you to run the tests repeatedly in the same session without reloading all the package code; for packages that take a while to load, this can be much faster. With this approach, you do have to do some extra work to make changes in the package code.

  - From the shell, run `julia ../test/runtests.jl` from within the package's `src/` folder.

- Commit your changes: see https://git-scm.com/book/en/v2/Git-Basics-Recording-Changes-to-the-Repository.

- Submit your changes: From the Julia prompt, type `PkgDev.submit("Foo")`. This will push your changes to your GitHub fork, creating it if it doesn't already exist. (If you encounter an error, make sure you've set up your SSH keys.) Julia will then give you a hyperlink; open that link, edit the message, and then click "submit." At that point, the package owner will be notified of your changes and may initiate discussion. (If you are comfortable with git, you can also do these steps manually from the shell.)

- The package owner may suggest additional improvements. To respond to those suggestions, you can easily update the pull request (this only works for changes that have not already been merged; for merged pull requests, make new changes by starting a new branch):

  - If you've changed branches in the meantime, make sure you go back to the same branch with `git checkout fixbar` (from shell mode) or `Pkg.checkout("Foo", "fixbar")` (from the Julia prompt).

  - As above, make your changes, run the tests, and commit your changes.

  - From the shell, type `git push`. This will add your new commit(s) to the same pull request; you should see them appear automatically on the page holding the discussion of your pull request.

  One potential type of change the owner may request is that you squash your commits. See Squashing below.

Dirty packages

If you can't change branches because the package manager complains that your package is dirty, it means you have some changes that have not been committed. From the shell, use `git diff` to see what these changes are; you can either discard them (`git checkout changedfile.jl`) or commit them before switching branches. If you can't easily resolve the problems manually, as a last resort you can delete the entire `"Foo"` folder and reinstall a fresh copy with `Pkg.add("Foo")`. Naturally, this deletes any changes you've made.

Making a branch *post hoc*

Especially for newcomers to git, one often forgets to create a new branch until after some changes have already been made. If you haven't yet staged or committed your changes, you can create a new branch with `git checkout -b <newbranch>` just as usual–git will kindly show you that some files have been modified and create the new branch for you. *Your changes have not yet been committed to this new branch*, so the normal work rules still apply.

However, if you've already made a commit to `master` but wish to go back to the official `master` (called `origin/master`), use the following procedure:

- Create a new branch. This branch will hold your changes.

- Make sure everything is committed to this branch.

- `git checkout master`. If this fails, *do not* proceed further until you have resolved the problems, or you may lose your changes.

- *Reset* master (your current branch) back to an earlier state with `git reset --hard origin/master` (see https://git-scm.com/blog/2011/07/11/reset.html).

This requires a bit more familiarity with git, so it's much better to get in the habit of creating a branch at the outset.

Squashing and rebasing

Depending on the tastes of the package owner (s)he may ask you to "squash" your commits. This is especially likely if your change is quite simple but your commit history looks like this:

```
WIP: add new 1-line whizbang function (currently breaks package)
Finish whizbang function
Fix typo in variable name
Oops, don't forget to supply default argument
Split into two 1-line functions
Rats, forgot to export the second function
...
```

This gets into the territory of more advanced git usage, and you're encouraged to do some reading (https://git-scm.com/-book/en/v2/Git-Branching-Rebasing). However, a brief summary of the procedure is as follows:

- To protect yourself from error, start from your `fixbar` branch and create a new branch with `git checkout -b fixbar_backup`. Since you started from `fixbar`, this will be a copy. Now go back to the one you intend to modify with `git checkout fixbar`.

- From the shell, type `git rebase -i origin/master`.

- To combine commits, change `pick` to `squash` (for additional options, consult other sources). Save the file and close the editor window.

- Edit the combined commit message.

If the rebase goes badly, you can go back to the beginning to try again like this:

```
git checkout fixbar
git reset --hard fixbar_backup
```

Now let's assume you've rebased successfully. Since your `fixbar` repository has now diverged from the one in your GitHub fork, you're going to have to do a *force push*:

- To make it easy to refer to your GitHub fork, create a "handle" for it with `git remote add myfork https://github.com/my-account/Foo.jl.git`, where the URL comes from the "clone URL" on your GitHub fork's page.

- Force-push to your fork with `git push myfork +fixbar`. The + indicates that this should replace the `fixbar` branch found at `myfork`.

35.3 Creating a new Package

REQUIRE speaks for itself

You should have a REQUIRE file in your package repository, with a bare minimum directive of what Julia version you expect your users to be running for the package to work. Putting a floor on what Julia version your package supports is done by simply adding `julia 0.x` in this file. While this line is partly informational, it also has the consequence of whether `Pkg.update()` will update code found in `.julia` version directories. It will not update code found in version directories beneath the floor of what's specified in your REQUIRE.

As the development version `0.y` matures, you may find yourself using it more frequently, and wanting your package to support it. Be warned, the development branch of Julia is the land of breakage, and you can expect things to break. When you go about fixing whatever broke your package in the development `0.y` branch, you will likely find that you just broke your package on the stable version.

There is a mechanism found in the Compat package that will enable you to support both the stable version and breaking changes found in the development version. Should you decide to use this solution, you will need to add `Compat` to your REQUIRE file. In this case, you will still have `julia 0.x` in your REQUIRE. The x is the floor version of what your package supports.

You might also have no interest in supporting the development version of Julia. Just as you can add a floor to the version you expect your users to be on, you can set an upper bound. In this case, you would put `julia 0.x 0.y-` in your REQUIRE file. The – at the end of the version number means pre-release versions of that specific version from the very first commit. By setting it as the ceiling, you mean the code supports everything up to but not including the ceiling version.

Another scenario is that you are writing the bulk of the code for your package with Julia `0.y` and do not want to support the current stable version of Julia. If you choose to do this, simply add `julia 0.y-` to your REQUIRE. Just remember to change the `julia 0.y-` to `julia 0.y` in your REQUIRE file once `0.y` is officially released. If you don't edit the dash cruft you are suggesting that you support both the development and stable versions of the same version number! That would be madness. See the Requirements Specification for the full format of REQUIRE.

Lastly, in many cases you may need extra packages for testing. Additional packages which are only required for tests should be specified in the `test/REQUIRE` file. This REQUIRE file has the same specification as the standard REQUIRE file.

Guidelines for naming a package

Package names should be sensible to most Julia users, *even to those who are not domain experts*. When you submit your package to METADATA, you can expect a little back and forth about the package name with collaborators, especially if it's ambiguous or can be confused with something other than what it is. During this bike-shedding, it's not uncommon to get a range of *different* name suggestions. These are only suggestions though, with the intent being to keep a tidy namespace in the curated METADATA repository. Since this repository belongs to the entire community, there will likely be a few collaborators who care about your package name. Here are some guidelines to follow in naming your package:

1. Avoid jargon. In particular, avoid acronyms unless there is minimal possibility of confusion.

    – It's ok to say USA if you're talking about the USA.

    – It's not ok to say PMA, even if you're talking about positive mental attitude.

2. Avoid using `Julia` in your package name.

    – It is usually clear from context and to your users that the package is a Julia package.

- Having Julia in the name can imply that the package is connected to, or endorsed by, contributors to the Julia language itself.

3. Packages that provide most of their functionality in association with a new type should have pluralized names.

   - `DataFrames` provides the `DataFrame` type.
   - `BloomFilters` provides the `BloomFilter` type.
   - In contrast, `JuliaParser` provides no new type, but instead new functionality in the `JuliaParser.parse()` function.

4. Err on the side of clarity, even if clarity seems long-winded to you.

   - `RandomMatrices` is a less ambiguous name than `RndMat` or `RMT`, even though the latter are shorter.

5. A less systematic name may suit a package that implements one of several possible approaches to its domain.

   - Julia does not have a single comprehensive plotting package. Instead, `Gadfly`, `PyPlot`, `Winston` and other packages each implement a unique approach based on a particular design philosophy.
   - In contrast, `SortingAlgorithms` provides a consistent interface to use many well-established sorting algorithms.

6. Packages that wrap external libraries or programs should be named after those libraries or programs.

   - `CPLEX.jl` wraps the `CPLEX` library, which can be identified easily in a web search.
   - `MATLAB.jl` provides an interface to call the MATLAB engine from within Julia.

Generating the package

Suppose you want to create a new Julia package called `FooBar`. To get started, do `PkgDev.generate(pkg,license)` where `pkg` is the new package name and `license` is the name of a license that the package generator knows about:

```
julia> PkgDev.generate("FooBar","MIT")
INFO: Initializing FooBar repo: /Users/stefan/.julia/v0.6/FooBar
INFO: Origin: git://github.com/StefanKarpinski/FooBar.jl.git
INFO: Generating LICENSE.md
INFO: Generating README.md
INFO: Generating src/FooBar.jl
INFO: Generating test/runtests.jl
INFO: Generating REQUIRE
INFO: Generating .travis.yml
INFO: Generating appveyor.yml
INFO: Generating .gitignore
INFO: Committing FooBar generated files
```

This creates the directory `~/.julia/v0.6/FooBar`, initializes it as a git repository, generates a bunch of files that all packages should have, and commits them to the repository:

```
$ cd ~/.julia/v0.6/FooBar && git show --stat

commit 84b8e266dae6de30ab9703150b3bf771ec7b6285
Author: Stefan Karpinski <stefan@karpinski.org>
Date:   Wed Oct 16 17:57:58 2013 -0400
```

```
    FooBar.jl generated files.

        license: MIT
        authors: Stefan Karpinski
        years:   2013
        user:    StefanKarpinski

    Julia Version 0.3.0-prerelease+3217 [5fcfb13*]

.gitignore       |  2 ++
.travis.yml      | 13 +++++++++++++
LICENSE.md       | 22 +++++++++++++++++++++++
README.md        |  3 +++
REQUIRE          |  1 +
appveyor.yml     | 34 ++++++++++++++++++++++++++++++++++
src/FooBar.jl    |  5 +++++
test/runtests.jl |  5 +++++
8 files changed, 85 insertions(+)
```

At the moment, the package manager knows about the MIT "Expat" License, indicated by `"MIT"`, the Simplified BSD License, indicated by `"BSD"`, and version 2.0 of the Apache Software License, indicated by `"ASL"`. If you want to use a different license, you can ask us to add it to the package generator, or just pick one of these three and then modify the `~/.julia/v0.6/PACKAGE/LICENSE.md` file after it has been generated.

If you created a GitHub account and configured git to know about it, `PkgDev.generate()` will set an appropriate origin URL for you. It will also automatically generate a `.travis.yml` file for using the Travis automated testing service, and an `appveyor.yml` file for using AppVeyor. You will have to enable testing on the Travis and AppVeyor websites for your package repository, but once you've done that, it will already have working tests. Of course, all the default testing does is verify that `using FooBar` in Julia works.

Loading Static Non-Julia Files

If your package code needs to load static files which are not Julia code, e.g. an external library or data files, and are located within the package directory, use the `@__DIR__` macro to determine the directory of the current source file. For example if `FooBar/src/FooBar.jl` needs to load `FooBar/data/foo.csv`, use the following code:

```
datapath = joinpath(@__DIR__, "..", "data")
foo = readcsv(joinpath(datapath, "foo.csv"))
```

Making Your Package Available

Once you've made some commits and you're happy with how `FooBar` is working, you may want to get some other people to try it out. First you'll need to create the remote repository and push your code to it; we don't yet automatically do this for you, but we will in the future and it's not too hard to figure out [3]. Once you've done this, letting people try out your code is as simple as sending them the URL of the published repo – in this case:

```
git://github.com/StefanKarpinski/FooBar.jl.git
```

For your package, it will be your GitHub user name and the name of your package, but you get the idea. People you send this URL to can use `Pkg.clone()` to install the package and try it out:

```
julia> Pkg.clone("git://github.com/StefanKarpinski/FooBar.jl.git")
INFO: Cloning FooBar from git@github.com:StefanKarpinski/FooBar.jl.git
```

Tagging and Publishing Your Package

> **Tip**
>
> If you are hosting your package on GitHub, you can use the attobot integration to handle package regis-
> tration, tagging and publishing.

Once you've decided that `FooBar` is ready to be registered as an official package, you can add it to your local copy of
METADATA using `PkgDev.register()`:

```julia
julia> PkgDev.register("FooBar")
INFO: Registering FooBar at git://github.com/StefanKarpinski/FooBar.jl.git
INFO: Committing METADATA for FooBar
```

This creates a commit in the `~/.julia/v0.6/METADATA` repo:

```
$ cd ~/.julia/v0.6/METADATA && git show

commit 9f71f4becb05cadacb983c54a72eed744e5c019d
Author: Stefan Karpinski <stefan@karpinski.org>
Date:   Wed Oct 16 18:46:02 2013 -0400

    Register FooBar

diff --git a/FooBar/url b/FooBar/url
new file mode 100644
index 0000000..30e525e
--- /dev/null
+++ b/FooBar/url
@@ -0,0 +1 @@
+git://github.com/StefanKarpinski/FooBar.jl.git
```

This commit is only locally visible, however. To make it visible to the Julia community, you need to merge your local
METADATA upstream into the official repo. The `PkgDev.publish()` command will fork the METADATA repository on
GitHub, push your changes to your fork, and open a pull request:

```julia
julia> PkgDev.publish()
INFO: Validating METADATA
INFO: No new package versions to publish
INFO: Submitting METADATA changes
INFO: Forking JuliaLang/METADATA.jl to StefanKarpinski
INFO: Pushing changes as branch pull-request/ef45f54b
INFO: To create a pull-request open:

  https://github.com/StefanKarpinski/METADATA.jl/compare/pull-request/ef45f54b
```

> **Tip**
>
> If `PkgDev.publish()` fails with error:
>
> `ERROR: key not found: "token"`

---

[3]Installing and using GitHub's "hub" tool is highly recommended. It allows you to do things like run hub `create` in the package repo and have
it automatically created via GitHub's API.

then you may have encountered an issue from using the GitHub API on multiple systems. The solution is to delete the "Julia Package Manager" personal access token from your Github account and try again.

Other failures may require you to circumvent PkgDev.publish() by creating a pull request on GitHub. See: Publishing METADATA manually below.

Once the package URL for FooBar is registered in the official METADATA repo, people know where to clone the package from, but there still aren't any registered versions available. You can tag and register it with the PkgDev.tag() command:

```
julia> PkgDev.tag("FooBar")
INFO: Tagging FooBar v0.0.1
INFO: Committing METADATA for FooBar
```

This tags v0.0.1 in the FooBar repo:

```
$ cd ~/.julia/v0.6/FooBar && git tag
v0.0.1
```

It also creates a new version entry in your local METADATA repo for FooBar:

```
$ cd ~/.julia/v0.6/FooBar && git show
commit de77ee4dc0689b12c5e8b574aef7f70e8b311b0e
Author: Stefan Karpinski <stefan@karpinski.org>
Date:   Wed Oct 16 23:06:18 2013 -0400

    Tag FooBar v0.0.1

diff --git a/FooBar/versions/0.0.1/sha1 b/FooBar/versions/0.0.1/sha1
new file mode 100644
index 0000000..c1cb1c1
--- /dev/null
+++ b/FooBar/versions/0.0.1/sha1
@@ -0,0 +1 @@
+84b8e266dae6de30ab9703150b3bf771ec7b6285
```

The PkgDev.tag() command takes an optional second argument that is either an explicit version number object like v"0.0.1" or one of the symbols :patch, :minor or :major. These increment the patch, minor or major version number of your package intelligently.

Adding a tagged version of your package will expedite the official registration into METADATA.jl by collaborators. It is strongly recommended that you complete this process, regardless if your package is completely ready for an official release.

As a general rule, packages should be tagged 0.0.1 first. Since Julia itself hasn't achieved 1.0 status, it's best to be conservative in your package's tagged versions.

As with PkgDev.register(), these changes to METADATA aren't available to anyone else until they've been included upstream. Again, use the PkgDev.publish() command, which first makes sure that individual package repos have been tagged, pushes them if they haven't already been, and then opens a pull request to METADATA:

```
julia> PkgDev.publish()
INFO: Validating METADATA
INFO: Pushing FooBar permanent tags: v0.0.1
INFO: Submitting METADATA changes
INFO: Forking JuliaLang/METADATA.jl to StefanKarpinski
```

```
INFO: Pushing changes as branch pull-request/3ef4f5c4
INFO: To create a pull-request open:

  https://github.com/StefanKarpinski/METADATA.jl/compare/pull-request/3ef4f5c4
```

Publishing METADATA manually

If `PkgDev.publish()` fails you can follow these instructions to manually publish your package.

By "forking" the main METADATA repository, you can create a personal copy (of METADATA.jl) under your GitHub account. Once that copy exists, you can push your local changes to your copy (just like any other GitHub project).

1. go to https://github.com/login?return_to=https%3A%2F%2Fgithub.com%2FJuliaLang%2FMETADATA.jl%2Ffork

and create your own fork.

2. add your fork as a remote repository for the METADATA repository on your local computer (in

the terminal where USERNAME is your github username):

```
cd ~/.julia/v0.6/METADATA
git remote add USERNAME https://github.com/USERNAME/METADATA.jl.git
```

1. push your changes to your fork:

   ```
   git push USERNAME metadata-v2
   ```

2. If all of that works, then go back to the GitHub page for your fork, and click the "pull request"

link.

35.4   Fixing Package Requirements

If you need to fix the registered requirements of an already-published package version, you can do so just by editing the metadata for that version, which will still have the same commit hash – the hash associated with a version is permanent:

```
$ cd ~/.julia/v0.6/METADATA/FooBar/versions/0.0.1 && cat requires
julia 0.3-
$ vi requires
```

Since the commit hash stays the same, the contents of the `REQUIRE` file that will be checked out in the repo will **not** match the requirements in METADATA after such a change; this is unavoidable. When you fix the requirements in METADATA for a previous version of a package, however, you should also fix the `REQUIRE` file in the current version of the package.

35.5 Requirements Specification

The `~/.julia/v0.6/REQUIRE` file, the `REQUIRE` file inside packages, and the `METADATA` package `requires` files use a simple line-based format to express the ranges of package versions which need to be installed. Package `REQUIRE` and `METADATA requires` files should also include the range of versions of `julia` the package is expected to work with. Additionally, packages can include a `test/REQUIRE` file to specify additional packages which are only required for testing.

Here's how these files are parsed and interpreted.

- Everything after a # mark is stripped from each line as a comment.

- If nothing but whitespace is left, the line is ignored.

- If there are non-whitespace characters remaining, the line is a requirement and the is split on whitespace into words.

The simplest possible requirement is just the name of a package name on a line by itself:

```
Distributions
```

This requirement is satisfied by any version of the `Distributions` package. The package name can be followed by zero or more version numbers in ascending order, indicating acceptable intervals of versions of that package. One version opens an interval, while the next closes it, and the next opens a new interval, and so on; if an odd number of version numbers are given, then arbitrarily large versions will satisfy; if an even number of version numbers are given, the last one is an upper limit on acceptable version numbers. For example, the line:

```
Distributions 0.1
```

is satisfied by any version of `Distributions` greater than or equal to `0.1.0`. Suffixing a version with – allows any pre-release versions as well. For example:

```
Distributions 0.1-
```

is satisfied by pre-release versions such as `0.1-dev` or `0.1-rc1`, or by any version greater than or equal to `0.1.0`.

This requirement entry:

```
Distributions 0.1 0.2.5
```

is satisfied by versions from `0.1.0` up to, but not including `0.2.5`. If you want to indicate that any `0.1.x` version will do, you will want to write:

```
Distributions 0.1 0.2-
```

If you want to start accepting versions after `0.2.7`, you can write:

```
Distributions 0.1 0.2- 0.2.7
```

If a requirement line has leading words that begin with @, it is a system-dependent requirement. If your system matches these system conditionals, the requirement is included, if not, the requirement is ignored. For example:

```
@osx Homebrew
```

will require the `Homebrew` package only on systems where the operating system is OS X. The system conditions that are currently supported are (hierarchically):

- @unix

    - @linux
    - @bsd

        * @osx

- @windows

The @unix condition is satisfied on all UNIX systems, including Linux and BSD. Negated system conditionals are also supported by adding a ! after the leading @. Examples:

```
@!windows
@unix @!osx
```

The first condition applies to any system but Windows and the second condition applies to any UNIX system besides OS X.

Runtime checks for the current version of Julia can be made using the built-in VERSION variable, which is of type VersionNumber. Such code is occasionally necessary to keep track of new or deprecated functionality between various releases of Julia. Examples of runtime checks:

```
VERSION < v"0.3-" #exclude all pre-release versions of 0.3

v"0.2-" <= VERSION < v"0.3-" #get all 0.2 versions, including pre-releases, up to the above

v"0.2" <= VERSION < v"0.3-" #To get only stable 0.2 versions (Note v"0.2" == v"0.2.0")

VERSION >= v"0.2.1" #get at least version 0.2.1
```

See the section on version number literals for a more complete description.

Chapter 36

Profiling

The `Profile` module provides tools to help developers improve the performance of their code. When used, it takes measurements on running code, and produces output that helps you understand how much time is spent on individual line(s). The most common usage is to identify "bottlenecks" as targets for optimization.

`Profile` implements what is known as a "sampling" or statistical profiler. It works by periodically taking a backtrace during the execution of any task. Each backtrace captures the currently-running function and line number, plus the complete chain of function calls that led to this line, and hence is a "snapshot" of the current state of execution.

If much of your run time is spent executing a particular line of code, this line will show up frequently in the set of all backtraces. In other words, the "cost" of a given line–or really, the cost of the sequence of function calls up to and including this line–is proportional to how often it appears in the set of all backtraces.

A sampling profiler does not provide complete line-by-line coverage, because the backtraces occur at intervals (by default, 1 ms on Unix systems and 10 ms on Windows, although the actual scheduling is subject to operating system load). Moreover, as discussed further below, because samples are collected at a sparse subset of all execution points, the data collected by a sampling profiler is subject to statistical noise.

Despite these limitations, sampling profilers have substantial strengths:

- You do not have to make any modifications to your code to take timing measurements (in contrast to the alternative instrumenting profiler).

- It can profile into Julia's core code and even (optionally) into C and Fortran libraries.

- By running "infrequently" there is very little performance overhead; while profiling, your code can run at nearly native speed.

For these reasons, it's recommended that you try using the built-in sampling profiler before considering any alternatives.

36.1 Basic usage

Let's work with a simple test case:

```julia
julia> function myfunc()
           A = rand(200, 200, 400)
           maximum(A)
       end
```

It's a good idea to first run the code you intend to profile at least once (unless you want to profile Julia's JIT-compiler):

```
julia> myfunc() # run once to force compilation
```

Now we're ready to profile this function:

```
julia> @profile myfunc()
```

To see the profiling results, there is a graphical browser available, but here we'll use the text-based display that comes with the standard library:

```
julia> Profile.print()
80 ./event.jl:73; (::Base.REPL.##1#2{Base.REPL.REPLBackend})()
 80 ./REPL.jl:97; macro expansion
  80 ./REPL.jl:66; eval_user_input(::Any, ::Base.REPL.REPLBackend)
   80 ./boot.jl:235; eval(::Module, ::Any)
    80 ./<missing>:?; anonymous
     80 ./profile.jl:23; macro expansion
      52 ./REPL[1]:2; myfunc()
        38 ./random.jl:431; rand!(::MersenneTwister, ::Array{Float64,3}, ::Int64, ::Type{B...
         38 ./dSFMT.jl:84; dsfmt_fill_array_close_open!(::Base.dSFMT.DSFMT_state, ::Ptr{F...
        14 ./random.jl:278; rand
         14 ./random.jl:277; rand
          14 ./random.jl:366; rand
           14 ./random.jl:369; rand
      28 ./REPL[1]:3; myfunc()
       28 ./reduce.jl:270; _mapreduce(::Base.#identity, ::Base.#scalarmax, ::IndexLinear,...
        3  ./reduce.jl:426; mapreduce_impl(::Base.#identity, ::Base.#scalarmax, ::Array{F...
        25 ./reduce.jl:428; mapreduce_impl(::Base.#identity, ::Base.#scalarmax, ::Array{F...
```

Each line of this display represents a particular spot (line number) in the code. Indentation is used to indicate the nested sequence of function calls, with more-indented lines being deeper in the sequence of calls. In each line, the first "field" is the number of backtraces (samples) taken *at this line or in any functions executed by this line*. The second field is the file name and line number and the third field is the function name. Note that the specific line numbers may change as Julia's code changes; if you want to follow along, it's best to run this example yourself.

In this example, we can see that the top level function called is in the file event.jl. This is the function that runs the REPL when you launch Julia. If you examine line 97 of REPL.jl, you'll see this is where the function eval_user_input() is called. This is the function that evaluates what you type at the REPL, and since we're working interactively these functions were invoked when we entered @profile myfunc(). The next line reflects actions taken in the @profile macro.

The first line shows that 80 backtraces were taken at line 73 of event.jl, but it's not that this line was "expensive" on its own: the third line reveals that all 80 of these backtraces were actually triggered inside its call to eval_user_input, and so on. To find out which operations are actually taking the time, we need to look deeper in the call chain.

The first "important" line in this output is this one:

```
52 ./REPL[1]:2; myfunc()
```

REPL refers to the fact that we defined myfunc in the REPL, rather than putting it in a file; if we had used a file, this would show the file name. The [1] shows that the function myfunc was the first expression evaluated in this REPL session. Line 2 of myfunc() contains the call to rand, and there were 52 (out of 80) backtraces that occurred at this line. Below that, you can see a call to dsfmt_fill_array_close_open! inside dSFMT.jl.

A little further down, you see:

```
28 ./REPL[1]:3; myfunc()
```

Line 3 of `myfunc` contains the call to `maximum`, and there were 28 (out of 80) backtraces taken here. Below that, you can see the specific places in `base/reduce.jl` that carry out the time-consuming operations in the `maximum` function for this type of input data.

Overall, we can tentatively conclude that generating the random numbers is approximately twice as expensive as finding the maximum element. We could increase our confidence in this result by collecting more samples:

```julia
julia> @profile (for i = 1:100; myfunc(); end)

julia> Profile.print()
[....]
 3821 ./REPL[1]:2; myfunc()
  3511 ./random.jl:431; rand!(::MersenneTwister, ::Array{Float64,3}, ::Int64, ::Type...
   3511 ./dSFMT.jl:84; dsfmt_fill_array_close_open!(::Base.dSFMT.DSFMT_state, ::Ptr...
  310  ./random.jl:278; rand
   [....]
 2893 ./REPL[1]:3; myfunc()
  2893 ./reduce.jl:270; _mapreduce(::Base.#identity, ::Base.#scalarmax, ::IndexLinea...
   [....]
```

In general, if you have N samples collected at a line, you can expect an uncertainty on the order of `sqrt(N)` (barring other sources of noise, like how busy the computer is with other tasks). The major exception to this rule is garbage collection, which runs infrequently but tends to be quite expensive. (Since Julia's garbage collector is written in C, such events can be detected using the `C=true` output mode described below, or by using ProfileView.jl.)

This illustrates the default "tree" dump; an alternative is the "flat" dump, which accumulates counts independent of their nesting:

```julia
julia> Profile.print(format=:flat)
 Count File          Line Function
  6714 ./<missing>     -1 anonymous
  6714 ./REPL.jl       66 eval_user_input(::Any, ::Base.REPL.REPLBackend)
  6714 ./REPL.jl       97 macro expansion
  3821 ./REPL[1]        2 myfunc()
  2893 ./REPL[1]        3 myfunc()
  6714 ./REPL[7]        1 macro expansion
  6714 ./boot.jl      235 eval(::Module, ::Any)
  3511 ./dSFMT.jl      84 dsfmt_fill_array_close_open!(::Base.dSFMT.DSFMT_s...
  6714 ./event.jl      73 (::Base.REPL.##1#2{Base.REPL.REPLBackend})()
  6714 ./profile.jl    23 macro expansion
  3511 ./random.jl    431 rand!(::MersenneTwister, ::Array{Float64,3}, ::In...
   310 ./random.jl    277 rand
   310 ./random.jl    278 rand
   310 ./random.jl    366 rand
   310 ./random.jl    369 rand
  2893 ./reduce.jl    270 _mapreduce(::Base.#identity, ::Base.#scalarmax, :...
     5 ./reduce.jl    420 mapreduce_impl(::Base.#identity, ::Base.#scalarma...
   253 ./reduce.jl    426 mapreduce_impl(::Base.#identity, ::Base.#scalarma...
  2592 ./reduce.jl    428 mapreduce_impl(::Base.#identity, ::Base.#scalarma...
    43 ./reduce.jl    429 mapreduce_impl(::Base.#identity, ::Base.#scalarma...
```

If your code has recursion, one potentially-confusing point is that a line in a "child" function can accumulate more counts than there are total backtraces. Consider the following function definitions:

```
dumbsum(n::Integer) = n == 1 ? 1 : 1 + dumbsum(n-1)
dumbsum3() = dumbsum(3)
```

If you were to profile dumbsum3, and a backtrace was taken while it was executing dumbsum(1), the backtrace would look like this:

```
dumbsum3
    dumbsum(3)
        dumbsum(2)
            dumbsum(1)
```

Consequently, this child function gets 3 counts, even though the parent only gets one. The "tree" representation makes this much clearer, and for this reason (among others) is probably the most useful way to view the results.

36.2   Accumulation and clearing

Results from @profile accumulate in a buffer; if you run multiple pieces of code under @profile, then Profile.print() will show you the combined results. This can be very useful, but sometimes you want to start fresh; you can do so with Profile.clear().

36.3   Options for controlling the display of profile results

Profile.print() has more options than we've described so far. Let's see the full declaration:

```
function print(io::IO = STDOUT, data = fetch(); kwargs...)
```

Let's first discuss the two positional arguments, and later the keyword arguments:

- io – Allows you to save the results to a buffer, e.g. a file, but the default is to print to STDOUT (the console).

- data – Contains the data you want to analyze; by default that is obtained from Profile.fetch(), which pulls out the backtraces from a pre-allocated buffer. For example, if you want to profile the profiler, you could say:

  ```
  data = copy(Profile.fetch())
  Profile.clear()
  @profile Profile.print(STDOUT, data) # Prints the previous results
  Profile.print()                      # Prints results from Profile.print()
  ```

The keyword arguments can be any combination of:

- format – Introduced above, determines whether backtraces are printed with (default, :tree) or without (:flat) indentation indicating tree structure.

- C – If true, backtraces from C and Fortran code are shown (normally they are excluded). Try running the introductory example with Profile.print(C = true). This can be extremely helpful in deciding whether it's Julia code or C code that is causing a bottleneck; setting C = true also improves the interpretability of the nesting, at the cost of longer profile dumps.

- combine – Some lines of code contain multiple operations; for example, `s += A[i]` contains both an array reference (`A[i]`) and a sum operation. These correspond to different lines in the generated machine code, and hence there may be two or more different addresses captured during backtraces on this line. `combine = true` lumps them together, and is probably what you typically want, but you can generate an output separately for each unique instruction pointer with `combine = false`.

- maxdepth – Limits frames at a depth higher than `maxdepth` in the `:tree` format.

- sortedby – Controls the order in `:flat` format. `:filefuncline` (default) sorts by the source line, whereas `:count` sorts in order of number of collected samples.

- noisefloor – Limits frames that are below the heuristic noise floor of the sample (only applies to format `:tree`). A suggested value to try for this is 2.0 (the default is 0). This parameter hides samples for which n `<= noisefloor * √N`, where n is the number of samples on this line, and N is the number of samples for the callee.

- mincount – Limits frames with less than `mincount` occurrences.

File/function names are sometimes truncated (with `...`), and indentation is truncated with a `+n` at the beginning, where n is the number of extra spaces that would have been inserted, had there been room. If you want a complete profile of deeply-nested code, often a good idea is to save to a file using a wide `displaysize` in an `IOContext`:

```julia
open("/tmp/prof.txt", "w") do s
    Profile.print(IOContext(s, :displaysize => (24, 500)))
end
```

36.4 Configuration

`@profile` just accumulates backtraces, and the analysis happens when you call `Profile.print()`. For a long-running computation, it's entirely possible that the pre-allocated buffer for storing backtraces will be filled. If that happens, the backtraces stop but your computation continues. As a consequence, you may miss some important profiling data (you will get a warning when that happens).

You can obtain and configure the relevant parameters this way:

```julia
Profile.init() # returns the current settings
Profile.init(n = 10^7, delay = 0.01)
```

n is the total number of instruction pointers you can store, with a default value of `10^6`. If your typical backtrace is 20 instruction pointers, then you can collect 50000 backtraces, which suggests a statistical uncertainty of less than 1%. This may be good enough for most applications.

Consequently, you are more likely to need to modify `delay`, expressed in seconds, which sets the amount of time that Julia gets between snapshots to perform the requested computations. A very long-running job might not need frequent backtraces. The default setting is `delay = 0.001`. Of course, you can decrease the delay as well as increase it; however, the overhead of profiling grows once the delay becomes similar to the amount of time needed to take a backtrace (~30 microseconds on the author's laptop).

Chapter 37

Memory allocation analysis

One of the most common techniques to improve performance is to reduce memory allocation. The total amount of allocation can be measured with `@time` and `@allocated`, and specific lines triggering allocation can often be inferred from profiling via the cost of garbage collection that these lines incur. However, sometimes it is more efficient to directly measure the amount of memory allocated by each line of code.

To measure allocation line-by-line, start Julia with the `--track-allocation=<setting>` command-line option, for which you can choose none (the default, do not measure allocation), `user` (measure memory allocation everywhere except Julia's core code), or `all` (measure memory allocation at each line of Julia code). Allocation gets measured for each line of compiled code. When you quit Julia, the cumulative results are written to text files with `.mem` appended after the file name, residing in the same directory as the source file. Each line lists the total number of bytes allocated. The `Coverage package` contains some elementary analysis tools, for example to sort the lines in order of number of bytes allocated.

In interpreting the results, there are a few important details. Under the `user` setting, the first line of any function directly called from the REPL will exhibit allocation due to events that happen in the REPL code itself. More significantly, JIT-compilation also adds to allocation counts, because much of Julia's compiler is written in Julia (and compilation usually requires memory allocation). The recommended procedure is to force compilation by executing all the commands you want to analyze, then call `Profile.clear_malloc_data()` to reset all allocation counters. Finally, execute the desired commands and quit Julia to trigger the generation of the `.mem` files.

Chapter 38

Stack Traces

The `StackTraces` module provides simple stack traces that are both human readable and easy to use programmatically.

38.1   Viewing a stack trace

The primary function used to obtain a stack trace is `stacktrace()`:

```julia
julia> stacktrace()
4-element Array{StackFrame,1}:
 eval(::Module, ::Any) at boot.jl:236
 eval_user_input(::Any, ::Base.REPL.REPLBackend) at REPL.jl:66
 macro expansion at REPL.jl:97 [inlined]
 (::Base.REPL.##1#2{Base.REPL.REPLBackend})() at event.jl:73
```

Calling `stacktrace()` returns a vector of `StackFrame`s. For ease of use, the alias `StackTrace` can be used in place of `Vector{StackFrame}`. (Examples with `[...]` indicate that output may vary depending on how the code is run.)

```julia
julia> example() = stacktrace()
example (generic function with 1 method)

julia> example()
5-element Array{StackFrame,1}:
 example() at REPL[1]:1
 eval(::Module, ::Any) at boot.jl:236
[...]

julia> @noinline child() = stacktrace()
child (generic function with 1 method)

julia> @noinline parent() = child()
parent (generic function with 1 method)

julia> grandparent() = parent()
grandparent (generic function with 1 method)

julia> grandparent()
7-element Array{StackFrame,1}:
 child() at REPL[3]:1
```

```
 parent() at REPL[4]:1
 grandparent() at REPL[5]:1
[...]
```

Note that when calling `stacktrace()` you'll typically see a frame with `eval(...)` at `boot.jl`. When calling `stacktrace()` from the REPL you'll also have a few extra frames in the stack from `REPL.jl`, usually looking something like this:

```
julia> example() = stacktrace()
example (generic function with 1 method)

julia> example()
5-element Array{StackFrame,1}:
 example() at REPL[1]:1
 eval(::Module, ::Any) at boot.jl:236
 eval_user_input(::Any, ::Base.REPL.REPLBackend) at REPL.jl:66
 macro expansion at REPL.jl:97 [inlined]
 (::Base.REPL.##1#2{Base.REPL.REPLBackend})() at event.jl:73
```

38.2   Extracting useful information

Each `StackFrame` contains the function name, file name, line number, lambda info, a flag indicating whether the frame has been inlined, a flag indicating whether it is a C function (by default C functions do not appear in the stack trace), and an integer representation of the pointer returned by `backtrace()`:

```
julia> top_frame = stacktrace()[1]
eval(::Module, ::Any) at boot.jl:236

julia> top_frame.func
:eval

julia> top_frame.file
Symbol("./boot.jl")

julia> top_frame.line
236

julia> top_frame.linfo
Nullable{Core.MethodInstance}(MethodInstance for eval(::Module, ::Any))

julia> top_frame.inlined
false

julia> top_frame.from_c
false


julia> top_frame.pointer
0x00007f390d152a59
```

This makes stack trace information available programmatically for logging, error handling, and more.

38.3   Error handling

While having easy access to information about the current state of the callstack can be helpful in many places, the most obvious application is in error handling and debugging.

```julia
julia> @noinline bad_function() = undeclared_variable
bad_function (generic function with 1 method)

julia> @noinline example() = try
           bad_function()
       catch
           stacktrace()
       end
example (generic function with 1 method)

julia> example()
5-element Array{StackFrame,1}:
 example() at REPL[2]:4
 eval(::Module, ::Any) at boot.jl:236
 [...]
```

You may notice that in the example above the first stack frame points points at line 4, where `stacktrace()` is called, rather than line 2, where *bad_function* is called, and `bad_function`'s frame is missing entirely. This is understandable, given that `stacktrace()` is called from the context of the *catch*. While in this example it's fairly easy to find the actual source of the error, in complex cases tracking down the source of the error becomes nontrivial.

This can be remedied by calling `catch_stacktrace()` instead of `stacktrace()`. Instead of returning callstack information for the current context, `catch_stacktrace()` returns stack information for the context of the most recent exception:

```julia
julia> @noinline bad_function() = undeclared_variable
bad_function (generic function with 1 method)

julia> @noinline example() = try
           bad_function()
       catch
           catch_stacktrace()
       end
example (generic function with 1 method)

julia> example()
6-element Array{StackFrame,1}:
 bad_function() at REPL[1]:1
 example() at REPL[2]:2
 [...]
```

Notice that the stack trace now indicates the appropriate line number and the missing frame.

```julia
julia> @noinline child() = error("Whoops!")
child (generic function with 1 method)

julia> @noinline parent() = child()
parent (generic function with 1 method)
```

```julia
julia> @noinline function grandparent()
           try
               parent()
           catch err
               println("ERROR: ", err.msg)
               catch_stacktrace()
           end
       end
grandparent (generic function with 1 method)

julia> grandparent()
ERROR: Whoops!
7-element Array{StackFrame,1}:
 child() at REPL[1]:1
 parent() at REPL[2]:1
 grandparent() at REPL[3]:3
[...]
```

38.4   Comparison with `backtrace()`

A call to `backtrace()` returns a vector of `Ptr{Void}`, which may then be passed into `stacktrace()` for translation:

```julia
julia> trace = backtrace()
21-element Array{Ptr{Void},1}:
 Ptr{Void} @0x00007f10049d5b2f
 Ptr{Void} @0x00007f0ffeb4d29c
 Ptr{Void} @0x00007f0ffeb4d2a9
 Ptr{Void} @0x00007f1004993fe7
 Ptr{Void} @0x00007f10049a92be
 Ptr{Void} @0x00007f10049a823a
 Ptr{Void} @0x00007f10049a9fb0
 Ptr{Void} @0x00007f10049aa718
 Ptr{Void} @0x00007f10049c0d5e
 Ptr{Void} @0x00007f10049a3286
 Ptr{Void} @0x00007f0ffe9ba3ba
 Ptr{Void} @0x00007f0ffe9ba3d0
 Ptr{Void} @0x00007f1004993fe7
 Ptr{Void} @0x00007f0ded34583d
 Ptr{Void} @0x00007f0ded345a87
 Ptr{Void} @0x00007f1004993fe7
 Ptr{Void} @0x00007f0ded34308f
 Ptr{Void} @0x00007f0ded343320
 Ptr{Void} @0x00007f1004993fe7
 Ptr{Void} @0x00007f10049aeb67
 Ptr{Void} @0x0000000000000000

julia> stacktrace(trace)
5-element Array{StackFrame,1}:
 backtrace() at error.jl:46
 eval(::Module, ::Any) at boot.jl:236
 eval_user_input(::Any, ::Base.REPL.REPLBackend) at REPL.jl:66
 macro expansion at REPL.jl:97 [inlined]
 (::Base.REPL.##1#2{Base.REPL.REPLBackend})() at event.jl:73
```

Notice that the vector returned by `backtrace()` had 21 pointers, while the vector returned by `stacktrace()` only

has 5. This is because, by default, `stacktrace()` removes any lower-level C functions from the stack. If you want to include stack frames from C calls, you can do it like this:

```
julia> stacktrace(trace, true)
27-element Array{StackFrame,1}:
 jl_backtrace_from_here at stackwalk.c:103
 backtrace() at error.jl:46
 backtrace() at sys.so:?
 jl_call_method_internal at julia_internal.h:248 [inlined]
 jl_apply_generic at gf.c:2215
 do_call at interpreter.c:75
 eval at interpreter.c:215
 eval_body at interpreter.c:519
 jl_interpret_toplevel_thunk at interpreter.c:664
 jl_toplevel_eval_flex at toplevel.c:592
 jl_toplevel_eval_in at builtins.c:614
 eval(::Module, ::Any) at boot.jl:236
 eval(::Module, ::Any) at sys.so:?
 jl_call_method_internal at julia_internal.h:248 [inlined]
 jl_apply_generic at gf.c:2215
 eval_user_input(::Any, ::Base.REPL.REPLBackend) at REPL.jl:66
 ip:0x7f1c707f1846
 jl_call_method_internal at julia_internal.h:248 [inlined]
 jl_apply_generic at gf.c:2215
 macro expansion at REPL.jl:97 [inlined]
 (::Base.REPL.##1#2{Base.REPL.REPLBackend})() at event.jl:73
 ip:0x7f1c707ea1ef
 jl_call_method_internal at julia_internal.h:248 [inlined]
 jl_apply_generic at gf.c:2215
 jl_apply at julia.h:1411 [inlined]
 start_task at task.c:261
 ip:0xffffffffffffffff
```

Individual pointers returned by `backtrace()` can be translated into `StackFrame` s by passing them into `StackTraces.lookup()`:

```
julia> pointer = backtrace()[1];

julia> frame = StackTraces.lookup(pointer)
1-element Array{StackFrame,1}:
 jl_backtrace_from_here at stackwalk.c:103

julia> println("The top frame is from $(frame[1].func)!")
The top frame is from jl_backtrace_from_here!
```

Chapter 39

Performance Tips

In the following sections, we briefly go through a few techniques that can help make your Julia code run as fast as possible.

39.1 Avoid global variables

A global variable might have its value, and therefore its type, change at any point. This makes it difficult for the compiler to optimize code using global variables. Variables should be local, or passed as arguments to functions, whenever possible.

Any code that is performance critical or being benchmarked should be inside a function.

We find that global names are frequently constants, and declaring them as such greatly improves performance:

```julia
const DEFAULT_VAL = 0
```

Uses of non-constant globals can be optimized by annotating their types at the point of use:

```julia
global x
y = f(x::Int + 1)
```

Writing functions is better style. It leads to more reusable code and clarifies what steps are being done, and what their inputs and outputs are.

> **Note**
>
> All code in the REPL is evaluated in global scope, so a variable defined and assigned at toplevel will be a **global** variable.

In the following REPL session:

```julia
julia> x = 1.0
```

is equivalent to:

```julia
julia> global x = 1.0
```

so all the performance issues discussed previously apply.

39.2   Measure performance with `@time` and pay attention to memory allocation

A useful tool for measuring performance is the `@time` macro. The following example illustrates good working style:

```julia
julia> function f(n)
           s = 0
           for i = 1:n
               s += i/2
           end
           s
       end
f (generic function with 1 method)

julia> @time f(1)
  0.012686 seconds (2.09 k allocations: 103.421 KiB)
0.5

julia> @time f(10^6)
  0.021061 seconds (3.00 M allocations: 45.777 MiB, 11.69% gc time)
2.5000025e11
```

On the first call (`@time f(1)`), `f` gets compiled. (If you've not yet used `@time` in this session, it will also compile functions needed for timing.) You should not take the results of this run seriously. For the second run, note that in addition to reporting the time, it also indicated that a large amount of memory was allocated. This is the single biggest advantage of `@time` vs. functions like `tic()` and `toc()`, which only report time.

Unexpected memory allocation is almost always a sign of some problem with your code, usually a problem with type-stability. Consequently, in addition to the allocation itself, it's very likely that the code generated for your function is far from optimal. Take such indications seriously and follow the advice below.

For more serious benchmarking, consider the BenchmarkTools.jl package which evaluates the function multiple times in order to reduce noise.

As a teaser, an improved version of this function allocates no memory (the allocation reported below is due to running the `@time` macro in global scope) and has an order of magnitude faster execution after the first call:

```julia
julia> @time f_improved(1)
  0.007008 seconds (1.32 k allocations: 63.640 KiB)
0.5

julia> @time f_improved(10^6)
  0.002997 seconds (6 allocations: 192 bytes)
2.5000025e11
```

Below you'll learn how to spot the problem with `f` and how to fix it.

In some situations, your function may need to allocate memory as part of its operation, and this can complicate the simple picture above. In such cases, consider using one of the tools below to diagnose problems, or write a version of your function that separates allocation from its algorithmic aspects (see Pre-allocating outputs).

39.3   Tools

Julia and its package ecosystem includes tools that may help you diagnose problems and improve the performance of your code:

- Profiling allows you to measure the performance of your running code and identify lines that serve as bottlenecks. For complex projects, the ProfileView package can help you visualize your profiling results.

- Unexpectedly-large memory allocations–as reported by @time, @allocated, or the profiler (through calls to the garbage-collection routines)–hint that there might be issues with your code. If you don't see another reason for the allocations, suspect a type problem. You can also start Julia with the `--track-allocation=user` option and examine the resulting `*.mem` files to see information about where those allocations occur. See Memory allocation analysis.

- @code_warntype generates a representation of your code that can be helpful in finding expressions that result in type uncertainty. See @code_warntype below.

- The Lint package can also warn you of certain types of programming errors.

39.4   Avoid containers with abstract type parameters

When working with parameterized types, including arrays, it is best to avoid parameterizing with abstract types where possible.

Consider the following:

```julia
a = Real[]    # typeof(a) = Array{Real,1}
if (f = rand()) < .8
    push!(a, f)
end
```

Because a is a an array of abstract type Real, it must be able to hold any Real value. Since Real objects can be of arbitrary size and structure, a must be represented as an array of pointers to individually allocated Real objects. Because f will always be a Float64, we should instead, use:

```julia
a = Float64[] # typeof(a) = Array{Float64,1}
```

which will create a contiguous block of 64-bit floating-point values that can be manipulated efficiently.

See also the discussion under Parametric Types.

39.5   Type declarations

In many languages with optional type declarations, adding declarations is the principal way to make code run faster. This is *not* the case in Julia. In Julia, the compiler generally knows the types of all function arguments, local variables, and expressions. However, there are a few specific instances where declarations are helpful.

Avoid fields with abstract type

Types can be declared without specifying the types of their fields:

```julia
julia> struct MyAmbiguousType
           a
       end
```

This allows a to be of any type. This can often be useful, but it does have a downside: for objects of type MyAmbiguousType, the compiler will not be able to generate high-performance code. The reason is that the compiler uses the types of objects, not their values, to determine how to build code. Unfortunately, very little can be inferred about an object of type MyAmbiguousType:

```
julia> b = MyAmbiguousType("Hello")
MyAmbiguousType("Hello")

julia> c = MyAmbiguousType(17)
MyAmbiguousType(17)

julia> typeof(b)
MyAmbiguousType

julia> typeof(c)
MyAmbiguousType
```

b and c have the same type, yet their underlying representation of data in memory is very different. Even if you stored just numeric values in field a, the fact that the memory representation of a `UInt8` differs from a `Float64` also means that the CPU needs to handle them using two different kinds of instructions. Since the required information is not available in the type, such decisions have to be made at run-time. This slows performance.

You can do better by declaring the type of a. Here, we are focused on the case where a might be any one of several types, in which case the natural solution is to use parameters. For example:

```
julia> mutable struct MyType{T<:AbstractFloat}
           a::T
       end
```

This is a better choice than

```
julia> mutable struct MyStillAmbiguousType
           a::AbstractFloat
       end
```

because the first version specifies the type of a from the type of the wrapper object. For example:

```
julia> m = MyType(3.2)
MyType{Float64}(3.2)

julia> t = MyStillAmbiguousType(3.2)
MyStillAmbiguousType(3.2)

julia> typeof(m)
MyType{Float64}

julia> typeof(t)
MyStillAmbiguousType
```

The type of field a can be readily determined from the type of m, but not from the type of t. Indeed, in t it's possible to change the type of field a:

```
julia> typeof(t.a)
Float64

julia> t.a = 4.5f0
4.5f0

julia> typeof(t.a)
Float32
```

In contrast, once m is constructed, the type of m.a cannot change:

```
julia> m.a = 4.5f0
4.5f0

julia> typeof(m.a)
Float64
```

The fact that the type of `m.a` is known from `m`'s type–coupled with the fact that its type cannot change mid-function–allows the compiler to generate highly-optimized code for objects like `m` but not for objects like `t`.

Of course, all of this is true only if we construct `m` with a concrete type. We can break this by explicitly constructing it with an abstract type:

```
julia> m = MyType{AbstractFloat}(3.2)
MyType{AbstractFloat}(3.2)

julia> typeof(m.a)
Float64

julia> m.a = 4.5f0
4.5f0

julia> typeof(m.a)
Float32
```

For all practical purposes, such objects behave identically to those of `MyStillAmbiguousType`.

It's quite instructive to compare the sheer amount code generated for a simple function

```
func(m::MyType) = m.a+1
```

using

```
code_llvm(func, Tuple{MyType{Float64}})
code_llvm(func, Tuple{MyType{AbstractFloat}})
code_llvm(func, Tuple{MyType})
```

For reasons of length the results are not shown here, but you may wish to try this yourself. Because the type is fully-specified in the first case, the compiler doesn't need to generate any code to resolve the type at run-time. This results in shorter and faster code.

Avoid fields with abstract containers

The same best practices also work for container types:

```
julia> mutable struct MySimpleContainer{A<:AbstractVector}
           a::A
       end

julia> mutable struct MyAmbiguousContainer{T}
           a::AbstractVector{T}
       end
```

For example:

```
julia> c = MySimpleContainer(1:3);

julia> typeof(c)
MySimpleContainer{UnitRange{Int64}}

julia> c = MySimpleContainer([1:3;]);

julia> typeof(c)
MySimpleContainer{Array{Int64,1}}

julia> b = MyAmbiguousContainer(1:3);

julia> typeof(b)
MyAmbiguousContainer{Int64}

julia> b = MyAmbiguousContainer([1:3;]);

julia> typeof(b)
MyAmbiguousContainer{Int64}
```

For MySimpleContainer, the object is fully-specified by its type and parameters, so the compiler can generate optimized functions. In most instances, this will probably suffice.

While the compiler can now do its job perfectly well, there are cases where *you* might wish that your code could do different things depending on the *element type* of a. Usually the best way to achieve this is to wrap your specific operation (here, foo) in a separate function:

```
julia> function sumfoo(c::MySimpleContainer)
           s = 0
           for x in c.a
               s += foo(x)
           end
           s
       end
sumfoo (generic function with 1 method)

julia> foo(x::Integer) = x
foo (generic function with 1 method)

julia> foo(x::AbstractFloat) = round(x)
foo (generic function with 2 methods)
```

This keeps things simple, while allowing the compiler to generate optimized code in all cases.

However, there are cases where you may need to declare different versions of the outer function for different element types of a. You could do it like this:

```
function myfun(c::MySimpleContainer{Vector{T}}) where T<:AbstractFloat
    ...
end
function myfun(c::MySimpleContainer{Vector{T}}) where T<:Integer
    ...
end
```

This works fine for Vector{T}, but we'd also have to write explicit versions for UnitRange{T} or other abstract types. To prevent such tedium, you can use two parameters in the declaration of MyContainer:

```
julia> mutable struct MyContainer{T, A<:AbstractVector}
```

```
            a::A
        end

julia> MyContainer(v::AbstractVector) = MyContainer{eltype(v), typeof(v)}(v)
MyContainer

julia> b = MyContainer(1:5);

julia> typeof(b)
MyContainer{Int64,UnitRange{Int64}}
```

Note the somewhat surprising fact that T doesn't appear in the declaration of field a, a point that we'll return to in a moment. With this approach, one can write functions such as:

```
julia> function myfunc(c::MyContainer{<:Integer, <:AbstractArray})
           return c.a[1]+1
       end
myfunc (generic function with 1 method)

julia> function myfunc(c::MyContainer{<:AbstractFloat})
           return c.a[1]+2
       end
myfunc (generic function with 2 methods)

julia> function myfunc(c::MyContainer{T,Vector{T}}) where T<:Integer
           return c.a[1]+3
       end
myfunc (generic function with 3 methods)
```

> **Note**
>
> Because we can only define MyContainer for A<:AbstractArray, and any unspecified parameters are arbitrary, the first function above could have been written more succinctly as function myfunc{T<:Integer}(c::MyContainer{T})

```
julia> myfunc(MyContainer(1:3))
2

julia> myfunc(MyContainer(1.0:3))
3.0

julia> myfunc(MyContainer([1:3;]))
4
```

As you can see, with this approach it's possible to specialize on both the element type T and the array type A.

However, there's one remaining hole: we haven't enforced that A has element type T, so it's perfectly possible to construct an object like this:

```
julia> b = MyContainer{Int64, UnitRange{Float64}}(UnitRange(1.3, 5.0));

julia> typeof(b)
MyContainer{Int64,UnitRange{Float64}}
```

To prevent this, we can add an inner constructor:

```julia
julia> mutable struct MyBetterContainer{T<:Real, A<:AbstractVector}
           a::A
           MyBetterContainer{T,A}(v::AbstractVector{T}) where {T,A} = new(v)
       end

julia> MyBetterContainer(v::AbstractVector) = MyBetterContainer{eltype(v),typeof(v)}(v)
MyBetterContainer

julia> b = MyBetterContainer(UnitRange(1.3, 5.0));

julia> typeof(b)
MyBetterContainer{Float64,UnitRange{Float64}}

julia> b = MyBetterContainer{Int64, UnitRange{Float64}}(UnitRange(1.3, 5.0));
ERROR: MethodError: Cannot `convert` an object of type UnitRange{Float64} to an object of type
    MyBetterContainer{Int64,UnitRange{Float64}}
[...]
```

The inner constructor requires that the element type of A be T.

Annotate values taken from untyped locations

It is often convenient to work with data structures that may contain values of any type (arrays of type `Array{Any}`). But, if you're using one of these structures and happen to know the type of an element, it helps to share this knowledge with the compiler:

```julia
function foo(a::Array{Any,1})
    x = a[1]::Int32
    b = x+1
    ...
end
```

Here, we happened to know that the first element of a would be an `Int32`. Making an annotation like this has the added benefit that it will raise a run-time error if the value is not of the expected type, potentially catching certain bugs earlier.

Declare types of keyword arguments

Keyword arguments can have declared types:

```julia
function with_keyword(x; name::Int = 1)
    ...
end
```

Functions are specialized on the types of keyword arguments, so these declarations will not affect performance of code inside the function. However, they will reduce the overhead of calls to the function that include keyword arguments.

Functions with keyword arguments have near-zero overhead for call sites that pass only positional arguments.

Passing dynamic lists of keyword arguments, as in `f(x; keywords...)`, can be slow and should be avoided in performance-sensitive code.

39.6 Break functions into multiple definitions

Writing a function as many small definitions allows the compiler to directly call the most applicable code, or even inline it.

Here is an example of a "compound function" that should really be written as multiple definitions:

```
function norm(A)
    if isa(A, Vector)
        return sqrt(real(dot(A,A)))
    elseif isa(A, Matrix)
        return maximum(svd(A)[2])
    else
        error("norm: invalid argument")
    end
end
```

This can be written more concisely and efficiently as:

```
norm(x::Vector) = sqrt(real(dot(x,x)))
norm(A::Matrix) = maximum(svd(A)[2])
```

39.7 Write "type-stable" functions

When possible, it helps to ensure that a function always returns a value of the same type. Consider the following definition:

```
pos(x) = x < 0 ? 0 : x
```

Although this seems innocent enough, the problem is that 0 is an integer (of type `Int`) and `x` might be of any type. Thus, depending on the value of `x`, this function might return a value of either of two types. This behavior is allowed, and may be desirable in some cases. But it can easily be fixed as follows:

```
pos(x) = x < 0 ? zero(x) : x
```

There is also a `one()` function, and a more general `oftype(x, y)` function, which returns y converted to the type of x.

39.8 Avoid changing the type of a variable

An analogous "type-stability" problem exists for variables used repeatedly within a function:

```
function foo()
    x = 1
    for i = 1:10
        x = x/bar()
    end
    return x
end
```

Local variable x starts as an integer, and after one loop iteration becomes a floating-point number (the result of `/` operator). This makes it more difficult for the compiler to optimize the body of the loop. There are several possible fixes:

- Initialize x with `x = 1.0`

- Declare the type of x: `x::Float64 = 1`

- Use an explicit conversion: `x = oneunit(T)`

- Initialize with the first loop iteration, to `x = 1/bar()`, then loop `for i = 2:10`

39.9   Separate kernel functions (aka, function barriers)

Many functions follow a pattern of performing some set-up work, and then running many iterations to perform a core computation. Where possible, it is a good idea to put these core computations in separate functions. For example, the following contrived function returns an array of a randomly-chosen type:

```julia
julia> function strange_twos(n)
           a = Vector{rand(Bool) ? Int64 : Float64}(n)
           for i = 1:n
               a[i] = 2
           end
           return a
       end
strange_twos (generic function with 1 method)

julia> strange_twos(3)
3-element Array{Float64,1}:
 2.0
 2.0
 2.0
```

This should be written as:

```julia
julia> function fill_twos!(a)
           for i=1:length(a)
               a[i] = 2
           end
       end
fill_twos! (generic function with 1 method)

julia> function strange_twos(n)
           a = Array{rand(Bool) ? Int64 : Float64}(n)
           fill_twos!(a)
           return a
       end
strange_twos (generic function with 1 method)

julia> strange_twos(3)
3-element Array{Float64,1}:
 2.0
 2.0
 2.0
```

Julia's compiler specializes code for argument types at function boundaries, so in the original implementation it does not know the type of a during the loop (since it is chosen randomly). Therefore the second version is generally faster since the inner loop can be recompiled as part of `fill_twos!` for different types of a.

The second form is also often better style and can lead to more code reuse.

This pattern is used in several places in the standard library. For example, see `hvcat_fill` in `abstractarray.jl`, or the `fill!` function, which we could have used instead of writing our own `fill_twos!`.

Functions like `strange_twos` occur when dealing with data of uncertain type, for example data loaded from an input file that might contain either integers, floats, strings, or something else.

39.10   Types with values-as-parameters

Let's say you want to create an N-dimensional array that has size 3 along each axis. Such arrays can be created like this:

```julia
julia> A = fill(5.0, (3, 3))
3×3 Array{Float64,2}:
 5.0  5.0  5.0
 5.0  5.0  5.0
 5.0  5.0  5.0
```

This approach works very well: the compiler can figure out that A is an `Array{Float64,2}` because it knows the type of the fill value (`5.0::Float64`) and the dimensionality (`(3, 3)::NTuple{2,Int}`). This implies that the compiler can generate very efficient code for any future usage of A in the same function.

But now let's say you want to write a function that creates a 3×3×... array in arbitrary dimensions; you might be tempted to write a function

```julia
julia> function array3(fillval, N)
           fill(fillval, ntuple(d->3, N))
       end
array3 (generic function with 1 method)

julia> array3(5.0, 2)
3×3 Array{Float64,2}:
 5.0  5.0  5.0
 5.0  5.0  5.0
 5.0  5.0  5.0
```

This works, but (as you can verify for yourself using `@code_warntype array3(5.0, 2)`) the problem is that the output type cannot be inferred: the argument N is a *value* of type `Int`, and type-inference does not (and cannot) predict its value in advance. This means that code using the output of this function has to be conservative, checking the type on each access of A; such code will be very slow.

Now, one very good way to solve such problems is by using the function-barrier technique. However, in some cases you might want to eliminate the type-instability altogether. In such cases, one approach is to pass the dimensionality as a parameter, for example through `Val{T}` (see "Value types"):

```julia
julia> function array3(fillval, ::Type{Val{N}}) where N
           fill(fillval, ntuple(d->3, Val{N}))
       end
array3 (generic function with 1 method)
```

```julia
julia> array3(5.0, Val{2})
3×3 Array{Float64,2}:
 5.0  5.0  5.0
 5.0  5.0  5.0
 5.0  5.0  5.0
```

Julia has a specialized version of `ntuple` that accepts a `Val{::Int}` as the second parameter; by passing N as a type-parameter, you make its "value" known to the compiler. Consequently, this version of `array3` allows the compiler to predict the return type.

However, making use of such techniques can be surprisingly subtle. For example, it would be of no help if you called `array3` from a function like this:

```julia
function call_array3(fillval, n)
    A = array3(fillval, Val{n})
end
```

Here, you've created the same problem all over again: the compiler can't guess the type of n, so it doesn't know the type of `Val{n}`. Attempting to use Val, but doing so incorrectly, can easily make performance *worse* in many situations. (Only in situations where you're effectively combining `Val` with the function-barrier trick, to make the kernel function more efficient, should code like the above be used.)

An example of correct usage of `Val` would be:

```julia
function filter3(A::AbstractArray{T,N}) where {T,N}
    kernel = array3(1, Val{N})
    filter(A, kernel)
end
```

In this example, N is passed as a parameter, so its "value" is known to the compiler. Essentially, `Val{T}` works only when T is either hard-coded (`Val{3}`) or already specified in the type-domain.

39.11   The dangers of abusing multiple dispatch (aka, more on types with values-as-parameters)

Once one learns to appreciate multiple dispatch, there's an understandable tendency to go crazy and try to use it for everything. For example, you might imagine using it to store information, e.g.

```julia
struct Car{Make,Model}
    year::Int
    ...more fields...
end
```

and then dispatch on objects like `Car{:Honda, :Accord}(year, args...)`.

This might be worthwhile when the following are true:

- You require CPU-intensive processing on each `Car`, and it becomes vastly more efficient if you know the `Make` and `Model` at compile time.

- You have homogenous lists of the same type of `Car` to process, so that you can store them all in an `Array{Car{:Honda, :Accord},N}`.

When the latter holds, a function processing such a homogenous array can be productively specialized: Julia knows the type of each element in advance (all objects in the container have the same concrete type), so Julia can "look up" the correct method calls when the function is being compiled (obviating the need to check at run-time) and thereby emit efficient code for processing the whole list.

When these do not hold, then it's likely that you'll get no benefit; worse, the resulting "combinatorial explosion of types" will be counterproductive. If `items[i+1]` has a different type than `item[i]`, Julia has to look up the type at run-time, search for the appropriate method in method tables, decide (via type intersection) which one matches, determine whether it has been JIT-compiled yet (and do so if not), and then make the call. In essence, you're asking the full type- system and JIT-compilation machinery to basically execute the equivalent of a switch statement or dictionary lookup in your own code.

Some run-time benchmarks comparing (1) type dispatch, (2) dictionary lookup, and (3) a "switch" statement can be found on the mailing list.

Perhaps even worse than the run-time impact is the compile-time impact: Julia will compile specialized functions for each different `Car{Make, Model}`; if you have hundreds or thousands of such types, then every function that accepts such an object as a parameter (from a custom `get_year` function you might write yourself, to the generic `push!` function in the standard library) will have hundreds or thousands of variants compiled for it. Each of these increases the size of the cache of compiled code, the length of internal lists of methods, etc. Excess enthusiasm for values-as-parameters can easily waste enormous resources.

39.12 Access arrays in memory order, along columns

Multidimensional arrays in Julia are stored in column-major order. This means that arrays are stacked one column at a time. This can be verified using the `vec` function or the syntax `[:]` as shown below (notice that the array is ordered `[1 3 2 4]`, not `[1 2 3 4]`):

```julia
julia> x = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> x[:]
4-element Array{Int64,1}:
 1
 3
 2
 4
```

This convention for ordering arrays is common in many languages like Fortran, Matlab, and R (to name a few). The alternative to column-major ordering is row-major ordering, which is the convention adopted by C and Python (numpy) among other languages. Remembering the ordering of arrays can have significant performance effects when looping over arrays. A rule of thumb to keep in mind is that with column-major arrays, the first index changes most rapidly. Essentially this means that looping will be faster if the inner-most loop index is the first to appear in a slice expression.

Consider the following contrived example. Imagine we wanted to write a function that accepts a `Vector` and returns a square `Matrix` with either the rows or the columns filled with copies of the input vector. Assume that it is not important whether rows or columns are filled with these copies (perhaps the rest of the code can be easily adapted accordingly). We could conceivably do this in at least four ways (in addition to the recommended call to the built-in `repmat()`):

```julia
function copy_cols(x::Vector{T}) where T
    n = size(x, 1)
```

```julia
    out = Array{T}(n, n)
    for i = 1:n
        out[:, i] = x
    end
    out
end

function copy_rows(x::Vector{T}) where T
    n = size(x, 1)
    out = Array{T}(n, n)
    for i = 1:n
        out[i, :] = x
    end
    out
end

function copy_col_row(x::Vector{T}) where T
    n = size(x, 1)
    out = Array{T}(n, n)
    for col = 1:n, row = 1:n
        out[row, col] = x[row]
    end
    out
end

function copy_row_col(x::Vector{T}) where T
    n = size(x, 1)
    out = Array{T}(n, n)
    for row = 1:n, col = 1:n
        out[row, col] = x[col]
    end
    out
end
```

Now we will time each of these functions using the same random `10000` by `1` input vector:

```julia
julia> x = randn(10000);

julia> fmt(f) = println(rpad(string(f)*": ", 14, ' '), @elapsed f(x))

julia> map(fmt, Any[copy_cols, copy_rows, copy_col_row, copy_row_col]);
copy_cols:    0.331706323
copy_rows:    1.799009911
copy_col_row: 0.415630047
copy_row_col: 1.721531501
```

Notice that `copy_cols` is much faster than `copy_rows`. This is expected because `copy_cols` respects the column-based memory layout of the `Matrix` and fills it one column at a time. Additionally, `copy_col_row` is much faster than `copy_row_col` because it follows our rule of thumb that the first element to appear in a slice expression should be coupled with the inner-most loop.

39.13   Pre-allocating outputs

If your function returns an `Array` or some other complex type, it may have to allocate memory. Unfortunately, of-tentimes allocation and its converse, garbage collection, are substantial bottlenecks.

Sometimes you can circumvent the need to allocate memory on each function call by preallocating the output. As a trivial example, compare

```julia
function xinc(x)
    return [x, x+1, x+2]
end

function loopinc()
    y = 0
    for i = 1:10^7
        ret = xinc(i)
        y += ret[2]
    end
    y
end
```

with

```julia
function xinc!(ret::AbstractVector{T}, x::T) where T
    ret[1] = x
    ret[2] = x+1
    ret[3] = x+2
    nothing
end

function loopinc_prealloc()
    ret = Array{Int}(3)
    y = 0
    for i = 1:10^7
        xinc!(ret, i)
        y += ret[2]
    end
    y
end
```

Timing results:

```julia
julia> @time loopinc()
  0.529894 seconds (40.00 M allocations: 1.490 GiB, 12.14% gc time)
50000015000000

julia> @time loopinc_prealloc()
  0.030850 seconds (6 allocations: 288 bytes)
50000015000000
```

Preallocation has other advantages, for example by allowing the caller to control the "output" type from an algorithm. In the example above, we could have passed a SubArray rather than an Array, had we so desired.

Taken to its extreme, pre-allocation can make your code uglier, so performance measurements and some judgment may be required. However, for "vectorized" (element-wise) functions, the convenient syntax x .= f.(y) can be used for in-place operations with fused loops and no temporary arrays (see the dot syntax for vectorizing functions).

39.14   More dots: Fuse vectorized operations

Julia has a special dot syntax that converts any scalar function into a "vectorized" function call, and any operator into a "vectorized" operator, with the special property that nested "dot calls" are *fusing*: they are combined at the syntax level into a single loop, without allocating temporary arrays. If you use .= and similar assignment operators, the result can also be stored in-place in a pre-allocated array (see above).

In a linear-algebra context, this means that even though operations like vector + vector and vector * scalar are defined, it can be advantageous to instead use vector .+ vector and vector .* scalar because the resulting loops can be fused with surrounding computations. For example, consider the two functions:

```
f(x) = 3x.^2 + 4x + 7x.^3

fdot(x) = @. 3x^2 + 4x + 7x^3 # equivalent to 3 .* x.^2 .+ 4 .* x .+ 7 .* x.^3
```

Both f and fdot compute the same thing. However, fdot (defined with the help of the @. macro) is significantly faster when applied to an array:

```
julia> x = rand(10^6);

julia> @time f(x);
  0.010986 seconds (18 allocations: 53.406 MiB, 11.45% gc time)

julia> @time fdot(x);
  0.003470 seconds (6 allocations: 7.630 MiB)

julia> @time f.(x);
  0.003297 seconds (30 allocations: 7.631 MiB)
```

That is, fdot(x) is three times faster and allocates 1/7 the memory of f(x), because each * and + operation in f(x) allocates a new temporary array and executes in a separate loop. (Of course, if you just do f.(x) then it is as fast as fdot(x) in this example, but in many contexts it is more convenient to just sprinkle some dots in your expressions rather than defining a separate function for each vectorized operation.)

39.15   Consider using views for slices

In Julia, an array "slice" expression like array[1:5, :] creates a copy of that data (except on the left-hand side of an assignment, where array[1:5, :] = ... assigns in-place to that portion of array). If you are doing many operations on the slice, this can be good for performance because it is more efficient to work with a smaller contiguous copy than it would be to index into the original array. On the other hand, if you are just doing a few simple operations on the slice, the cost of the allocation and copy operations can be substantial.

An alternative is to create a "view" of the array, which is an array object (a SubArray) that actually references the data of the original array in-place, without making a copy. (If you write to a view, it modifies the original array's data as well.) This can be done for individual slices by calling view(), or more simply for a whole expression or block of code by putting @views in front of that expression. For example:

```
julia> fcopy(x) = sum(x[2:end-1])

julia> @views fview(x) = sum(x[2:end-1])

julia> x = rand(10^6);
```

```
julia> @time fcopy(x);
  0.003051 seconds (7 allocations: 7.630 MB)

julia> @time fview(x);
  0.001020 seconds (6 allocations: 224 bytes)
```

Notice both the 3× speedup and the decreased memory allocation of the fview version of the function.

39.16 Avoid string interpolation for I/O

When writing data to a file (or other I/O device), forming extra intermediate strings is a source of overhead. Instead of:

```
println(file, "$a $b")
```

use:

```
println(file, a, " ", b)
```

The first version of the code forms a string, then writes it to the file, while the second version writes values directly to the file. Also notice that in some cases string interpolation can be harder to read. Consider:

```
println(file, "$(f(a))$(f(b))")
```

versus:

```
println(file, f(a), f(b))
```

39.17 Optimize network I/O during parallel execution

When executing a remote function in parallel:

```
responses = Vector{Any}(nworkers())
@sync begin
    for (idx, pid) in enumerate(workers())
        @async responses[idx] = remotecall_fetch(pid, foo, args...)
    end
end
```

is faster than:

```
refs = Vector{Any}(nworkers())
for (idx, pid) in enumerate(workers())
    refs[idx] = @spawnat pid foo(args...)
end
responses = [fetch(r) for r in refs]
```

The former results in a single network round-trip to every worker, while the latter results in two network calls - first by the @spawnat and the second due to the fetch (or even a wait). The fetch/wait is also being executed serially resulting in an overall poorer performance.

39.18   Fix deprecation warnings

A deprecated function internally performs a lookup in order to print a relevant warning only once. This extra lookup can cause a significant slowdown, so all uses of deprecated functions should be modified as suggested by the warnings.

39.19   Tweaks

These are some minor points that might help in tight inner loops.

- Avoid unnecessary arrays. For example, instead of `sum([x,y,z])` use x+y+z.

- Use `abs2(z)` instead of `abs(z)^2` for complex z. In general, try to rewrite code to use `abs2()` instead of `abs()` for complex arguments.

- Use `div(x,y)` for truncating division of integers instead of `trunc(x/y)`, `fld(x,y)` instead of `floor(x/y)`, and `cld(x,y)` instead of `ceil(x/y)`.

39.20   Performance Annotations

Sometimes you can enable better optimization by promising certain program properties.

- Use `@inbounds` to eliminate array bounds checking within expressions. Be certain before doing this. If the subscripts are ever out of bounds, you may suffer crashes or silent corruption.

- Use `@fastmath` to allow floating point optimizations that are correct for real numbers, but lead to differences for IEEE numbers. Be careful when doing this, as this may change numerical results. This corresponds to the `-ffast-math` option of clang.

- Write `@simd` in front of `for` loops that are amenable to vectorization. **This feature is experimental** and could change or disappear in future versions of Julia.

Note: While `@simd` needs to be placed directly in front of a loop, both `@inbounds` and `@fastmath` can be applied to several statements at once, e.g. using `begin` ... `end`, or even to a whole function.

Here is an example with both `@inbounds` and `@simd` markup:

```julia
function inner(x, y)
    s = zero(eltype(x))
    for i=1:length(x)
        @inbounds s += x[i]*y[i]
    end
    s
end

function innersimd(x, y)
    s = zero(eltype(x))
    @simd for i=1:length(x)
        @inbounds s += x[i]*y[i]
    end
    s
end

function timeit(n, reps)
```

```julia
    x = rand(Float32,n)
    y = rand(Float32,n)
    s = zero(Float64)
    time = @elapsed for j in 1:reps
        s+=inner(x,y)
    end
    println("GFlop/sec        = ",2.0*n*reps/time*1E-9)
    time = @elapsed for j in 1:reps
        s+=innersimd(x,y)
    end
    println("GFlop/sec (SIMD) = ",2.0*n*reps/time*1E-9)
end

timeit(1000,1000)
```

On a computer with a 2.4GHz Intel Core i5 processor, this produces:

```
GFlop/sec        = 1.9467069505224963
GFlop/sec (SIMD) = 17.578554163920018
```

(GFlop/sec measures the performance, and larger numbers are better.) The range for a `@simd for` loop should be a one-dimensional range. A variable used for accumulating, such as s in the example, is called a *reduction variable*. By using `@simd`, you are asserting several properties of the loop:

- It is safe to execute iterations in arbitrary or overlapping order, with special consideration for reduction variables.

- Floating-point operations on reduction variables can be reordered, possibly causing different results than without `@simd`.

- No iteration ever waits on another iteration to make forward progress.

A loop containing `break`, `continue`, or `@goto` will cause a compile-time error.

Using `@simd` merely gives the compiler license to vectorize. Whether it actually does so depends on the compiler. To actually benefit from the current implementation, your loop should have the following additional properties:

- The loop must be an innermost loop.

- The loop body must be straight-line code. This is why `@inbounds` is currently needed for all array accesses. The compiler can sometimes turn short `&&`, `||`, and `?:` expressions into straight-line code, if it is safe to evaluate all operands unconditionally. Consider using `ifelse()` instead of `?:` in the loop if it is safe to do so.

- Accesses must have a stride pattern and cannot be "gathers" (random-index reads) or "scatters" (random-index writes).

- The stride should be unit stride.

- In some simple cases, for example with 2-3 arrays accessed in a loop, the LLVM auto-vectorization may kick in automatically, leading to no further speedup with `@simd`.

Here is an example with all three kinds of markup. This program first calculates the finite difference of a one-dimensional array, and then evaluates the L2-norm of the result:

```julia
function init!(u)
    n = length(u)
    dx = 1.0 / (n-1)
    @fastmath @inbounds @simd for i in 1:n
        u[i] = sin(2pi*dx*i)
    end
end

function deriv!(u, du)
    n = length(u)
    dx = 1.0 / (n-1)
    @fastmath @inbounds du[1] = (u[2] - u[1]) / dx
    @fastmath @inbounds @simd for i in 2:n-1
        du[i] = (u[i+1] - u[i-1]) / (2*dx)
    end
    @fastmath @inbounds du[n] = (u[n] - u[n-1]) / dx
end

function norm(u)
    n = length(u)
    T = eltype(u)
    s = zero(T)
    @fastmath @inbounds @simd for i in 1:n
        s += u[i]^2
    end
    @fastmath @inbounds return sqrt(s/n)
end

function main()
    n = 2000
    u = Array{Float64}(n)
    init!(u)
    du = similar(u)

    deriv!(u, du)
    nu = norm(du)

    @time for i in 1:10^6
        deriv!(u, du)
        nu = norm(du)
    end

    println(nu)
end

main()
```

On a computer with a 2.7 GHz Intel Core i7 processor, this produces:

```
$ julia wave.jl;
elapsed time: 1.207814709 seconds (0 bytes allocated)

$ julia --math-mode=ieee wave.jl;
elapsed time: 4.487083643 seconds (0 bytes allocated)
```

Here, the option `--math-mode=ieee` disables the `@fastmath` macro, so that we can compare results.

In this case, the speedup due to @fastmath is a factor of about 3.7. This is unusually large – in general, the speedup will be smaller. (In this particular example, the working set of the benchmark is small enough to fit into the L1 cache of the processor, so that memory access latency does not play a role, and computing time is dominated by CPU usage. In many real world programs this is not the case.) Also, in this case this optimization does not change the result – in general, the result will be slightly different. In some cases, especially for numerically unstable algorithms, the result can be very different.

The annotation @fastmath re-arranges floating point expressions, e.g. changing the order of evaluation, or assuming that certain special cases (inf, nan) cannot occur. In this case (and on this particular computer), the main difference is that the expression 1 / (2*dx) in the function deriv is hoisted out of the loop (i.e. calculated outside the loop), as if one had written idx = 1 / (2*dx). In the loop, the expression ... / (2*dx) then becomes ... * idx, which is much faster to evaluate. Of course, both the actual optimization that is applied by the compiler as well as the resulting speedup depend very much on the hardware. You can examine the change in generated code by using Julia's code_native() function.

39.21 Treat Subnormal Numbers as Zeros

Subnormal numbers, formerly called denormal numbers, are useful in many contexts, but incur a performance penalty on some hardware. A call set_zero_subnormals(true) grants permission for floating-point operations to treat subnormal inputs or outputs as zeros, which may improve performance on some hardware. A call set_zero_subnormals(false) enforces strict IEEE behavior for subnormal numbers.

Below is an example where subnormals noticeably impact performance on some hardware:

```julia
function timestep(b::Vector{T}, a::Vector{T}, Δt::T) where T
    @assert length(a)==length(b)
    n = length(b)
    b[1] = 1                            # Boundary condition
    for i=2:n-1
        b[i] = a[i] + (a[i-1] - T(2)*a[i] + a[i+1]) * Δt
    end
    b[n] = 0                            # Boundary condition
end

function heatflow(a::Vector{T}, nstep::Integer) where T
    b = similar(a)
    for t=1:div(nstep,2)                # Assume nstep is even
        timestep(b,a,T(0.1))
        timestep(a,b,T(0.1))
    end
end

heatflow(zeros(Float32,10),2)          # Force compilation
for trial=1:6
    a = zeros(Float32,1000)
    set_zero_subnormals(iseven(trial))  # Odd trials use strict IEEE arithmetic
    @time heatflow(a,1000)
end
```

This example generates many subnormal numbers because the values in a become an exponentially decreasing curve, which slowly flattens out over time.

Treating subnormals as zeros should be used with caution, because doing so breaks some identities, such as x-y == 0 implies x == y:

```julia
julia> x = 3f-38; y = 2f-38;

julia> set_zero_subnormals(true); (x - y, x == y)
(0.0f0, false)

julia> set_zero_subnormals(false); (x - y, x == y)
(1.0000001f-38, false)
```

In some applications, an alternative to zeroing subnormal numbers is to inject a tiny bit of noise. For example, instead of initializing a with zeros, initialize it with:

```julia
a = rand(Float32,1000) * 1.f-9
```

39.22   @code_warntype

The macro `@code_warntype` (or its function variant `code_warntype()`) can sometimes be helpful in diagnosing type-related problems. Here's an example:

```julia
pos(x) = x < 0 ? 0 : x

function f(x)
    y = pos(x)
    sin(y*x+1)
end

julia> @code_warntype f(3.2)
Variables:
  #self#::#f
  x::Float64
  y::UNION{FLOAT64,INT64}
  fy::Float64
  #temp#@_5::UNION{FLOAT64,INT64}
  #temp#@_6::Core.MethodInstance
  #temp#@_7::Float64

Body:
  begin
      $(Expr(:inbounds, false))
      # meta: location REPL[1] pos 1
      # meta: location float.jl < 487
      fy::Float64 = (Core.typeassert)((Base.sitofp)(Float64,0)::Float64,Float64)::Float64
      # meta: pop location
      unless
↪   (Base.or_int)((Base.lt_float)(x::Float64,fy::Float64)::Bool,(Base.and_int)((Base.and_int)((Base.eq_float)(x::Float64,f
↪   goto 9
      #temp#@_5::UNION{FLOAT64,INT64} = 0
      goto 11
      9:
      #temp#@_5::UNION{FLOAT64,INT64} = x::Float64
      11:
      # meta: pop location
      $(Expr(:inbounds, :pop))
      y::UNION{FLOAT64,INT64} = #temp#@_5::UNION{FLOAT64,INT64} # line 3:
      unless (y::UNION{FLOAT64,INT64} isa Int64)::ANY goto 19
```

```
    #temp#@_6::Core.MethodInstance = MethodInstance for *(::Int64, ::Float64)
    goto 28
    19:
    unless (y::UNION{FLOAT64,INT64} isa Float64)::ANY goto 23
    #temp#@_6::Core.MethodInstance = MethodInstance for *(::Float64, ::Float64)
    goto 28
    23:
    goto 25
    25:
    #temp#@_7::Float64 = (y::UNION{FLOAT64,INT64} * x::Float64)::Float64
    goto 30
    28:
    #temp#@_7::Float64 = $(Expr(:invoke, :(#temp#@_6), :(Main.*), :(y), :(x)))
    30:
    return $(Expr(:invoke, MethodInstance for sin(::Float64), :(Main.sin),
↪   :((Base.add_float)(#temp#@_7,(Base.sitofp)(Float64,1)::Float64)::Float64)))
  end::Float64
```

Interpreting the output of @code_warntype, like that of its cousins @code_lowered, @code_typed, @code_llvm, and @code_native, takes a little practice. Your code is being presented in form that has been partially digested on its way to generating compiled machine code. Most of the expressions are annotated by a type, indicated by the ::T (where T might be Float64, for example). The most important characteristic of @code_warntype is that non-concrete types are displayed in red; in the above example, such output is shown in all-caps.

The top part of the output summarizes the type information for the different variables internal to the function. You can see that y, one of the variables you created, is a Union{Int64,Float64}, due to the type-instability of pos. There is another variable, _var4, which you can see also has the same type.

The next lines represent the body of f. The lines starting with a number followed by a colon (1:, 2:) are labels, and represent targets for jumps (via goto) in your code. Looking at the body, you can see that pos has been *inlined* into f–everything before 2: comes from code defined in pos.

Starting at 2:, the variable y is defined, and again annotated as a Union type. Next, we see that the compiler created the temporary variable _var1 to hold the result of y*x. Because a Float64 times *either* an Int64 or Float64 yields a Float64, all type-instability ends here. The net result is that f(x::Float64) will not be type-unstable in its output, even if some of the intermediate computations are type-unstable.

How you use this information is up to you. Obviously, it would be far and away best to fix pos to be type-stable: if you did so, all of the variables in f would be concrete, and its performance would be optimal. However, there are circumstances where this kind of *ephemeral* type instability might not matter too much: for example, if pos is never used in isolation, the fact that f's output is type-stable (for Float64 inputs) will shield later code from the propagating effects of type instability. This is particularly relevant in cases where fixing the type instability is difficult or impossible: for example, currently it's not possible to infer the return type of an anonymous function. In such cases, the tips above (e.g., adding type annotations and/or breaking up functions) are your best tools to contain the "damage" from type instability.

The following examples may help you interpret expressions marked as containing non-leaf types:

- Function body ending in end::Union{T1,T2})

    - Interpretation: function with unstable return type
    - Suggestion: make the return value type-stable, even if you have to annotate it

- f(x::T)::Union{T1,T2}

    - Interpretation: call to a type-unstable function

     – Suggestion: fix the function, or if necessary annotate the return value

- `(top(arrayref))(A::Array{Any,1},1)::Any`

     – Interpretation: accessing elements of poorly-typed arrays

     – Suggestion: use arrays with better-defined types, or if necessary annotate the type of individual element accesses

- `(top(getfield))(A::ArrayContainer{Float64},:data)::Array{Float64,N}`

     – Interpretation: getting a field that is of non-leaf type. In this case, `ArrayContainer` had a field `data::Array{T}`. But `Array` needs the dimension `N`, too, to be a concrete type.

     – Suggestion: use concrete types like `Array{T,3}` or `Array{T,N}`, where `N` is now a parameter of `ArrayContainer`

Chapter 40

Workflow Tips

Here are some tips for working with Julia efficiently.

40.1 REPL-based workflow

As already elaborated in Interacting With Julia, Julia's REPL provides rich functionality that facilitates an efficient interactive workflow. Here are some tips that might further enhance your experience at the command line.

A basic editor/REPL workflow

The most basic Julia workflows involve using a text editor in conjunction with the `julia` command line. A common pattern includes the following elements:

- **Put code under development in a temporary module.** Create a file, say `Tmp.jl`, and include within it

  ```
  module Tmp

  <your definitions here>

  end
  ```

- **Put your test code in another file.** Create another file, say `tst.jl`, which begins with

  ```
  import Tmp
  ```

  and includes tests for the contents of Tmp. The value of using `import` versus `using` is that you can call `reload("Tmp")` instead of having to restart the REPL when your definitions change. Of course, the cost is the need to prepend `Tmp.` to uses of names defined in your module. (You can lower that cost by keeping your module name short.)

  Alternatively, you can wrap the contents of your test file in a module, as

  ```
  module Tst
      using Tmp

      <scratch work>

  end
  ```

  The advantage is that you can now do `using  Tmp` in your test code and can therefore avoid prepending `Tmp.` everywhere. The disadvantage is that code can no longer be selectively copied to the REPL without some tweaking.

- **Lather. Rinse. Repeat.** Explore ideas at the `julia` command prompt. Save good ideas in `tst.jl`. Occasionally restart the REPL, issuing

  ```
  reload("Tmp")
  include("tst.jl")
  ```

Simplify initialization

To simplify restarting the REPL, put project-specific initialization code in a file, say `_init.jl`, which you can run on startup by issuing the command:

```
julia -L _init.jl
```

If you further add the following to your `.juliarc.jl` file

```
isfile("_init.jl") && include(joinpath(pwd(), "_init.jl"))
```

then calling `julia` from that directory will run the initialization code without the additional command line argument.

40.2   Browser-based workflow

It is also possible to interact with a Julia REPL in the browser via IJulia. See the package home for details.

Chapter 41

Style Guide

The following sections explain a few aspects of idiomatic Julia coding style. None of these rules are absolute; they are only suggestions to help familiarize you with the language and to help you choose among alternative designs.

41.1   Write functions, not just scripts

Writing code as a series of steps at the top level is a quick way to get started solving a problem, but you should try to divide a program into functions as soon as possible. Functions are more reusable and testable, and clarify what steps are being done and what their inputs and outputs are. Furthermore, code inside functions tends to run much faster than top level code, due to how Julia's compiler works.

It is also worth emphasizing that functions should take arguments, instead of operating directly on global variables (aside from constants like `pi`).

41.2   Avoid writing overly-specific types

Code should be as generic as possible. Instead of writing:

```
convert(Complex{Float64}, x)
```

it's better to use available generic functions:

```
complex(float(x))
```

The second version will convert x to an appropriate type, instead of always the same type.

This style point is especially relevant to function arguments. For example, don't declare an argument to be of type Int or `Int32` if it really could be any integer, expressed with the abstract type `Integer`. In fact, in many cases you can omit the argument type altogether, unless it is needed to disambiguate from other method definitions, since a `MethodError` will be thrown anyway if a type is passed that does not support any of the requisite operations. (This is known as duck typing.)

For example, consider the following definitions of a function addone that returns one plus its argument:

```
addone(x::Int) = x + 1              # works only for Int
addone(x::Integer) = x + oneunit(x) # any integer type
addone(x::Number) = x + oneunit(x)  # any numeric type
addone(x) = x + oneunit(x)          # any type supporting + and oneunit
```

The last definition of addone handles any type supporting `oneunit` (which returns 1 in the same type as x, which avoids unwanted type promotion) and the `+` function with those arguments. The key thing to realize is that there is *no performance penalty* to defining *only* the general addone(x) = x + oneunit(x), because Julia will automatically compile specialized versions as needed. For example, the first time you call addone(12), Julia will automatically compile a specialized addone function for x::Int arguments, with the call to oneunit replaced by its inlined value 1. Therefore, the first three definitions of addone above are completely redundant with the fourth definition.

41.3   Handle excess argument diversity in the caller

Instead of:

```julia
function foo(x, y)
    x = Int(x); y = Int(y)
    ...
end
foo(x, y)
```

use:

```julia
function foo(x::Int, y::Int)
    ...
end
foo(Int(x), Int(y))
```

This is better style because foo does not really accept numbers of all types; it really needs Int s.

One issue here is that if a function inherently requires integers, it might be better to force the caller to decide how non-integers should be converted (e.g. floor or ceiling). Another issue is that declaring more specific types leaves more "space" for future method definitions.

41.4   Append ! to names of functions that modify their arguments

Instead of:

```julia
function double(a::AbstractArray{<:Number})
    for i = 1:endof(a)
        a[i] *= 2
    end
    return a
end
```

use:

```julia
function double!(a::AbstractArray{<:Number})
    for i = 1:endof(a)
        a[i] *= 2
    end
    return a
end
```

The Julia standard library uses this convention throughout and contains examples of functions with both copying and modifying forms (e.g., `sort()` and `sort!()`), and others which are just modifying (e.g., `push!()`, `pop!()`, `splice!()`). It is typical for such functions to also return the modified array for convenience.

41.5 Avoid strange type `Unions`

Types such as `Union{Function,AbstractString}` are often a sign that some design could be cleaner.

41.6 Avoid type Unions in fields

When creating a type such as:

```
mutable struct MyType
    ...
    x::Union{Void,T}
end
```

ask whether the option for x to be `nothing` (of type `Void`) is really necessary. Here are some alternatives to consider:

- Find a safe default value to initialize x with

- Introduce another type that lacks x

- If there are many fields like x, store them in a dictionary

- Determine whether there is a simple rule for when x is `nothing`. For example, often the field will start as `nothing` but get initialized at some well-defined point. In that case, consider leaving it undefined at first.

- If x really needs to hold no value at some times, define it as `::Nullable{T}` instead, as this guarantees type-stability in the code accessing this field (see Nullable types).

41.7 Avoid elaborate container types

It is usually not much help to construct arrays like the following:

```
a = Array{Union{Int,AbstractString,Tuple,Array}}(n)
```

In this case `Array{Any}(n)` is better. It is also more helpful to the compiler to annotate specific uses (e.g. `a[i]::Int`) than to try to pack many alternatives into one type.

41.8 Use naming conventions consistent with Julia's `base/`

- modules and type names use capitalization and camel case: `module SparseArrays`, `struct UnitRange`.

- functions are lowercase (`maximum()`, `convert()`) and, when readable, with multiple words squashed together (`isequal()`, `haskey()`). When necessary, use underscores as word separators. Underscores are also used to indicate a combination of concepts (`remotecall_fetch()` as a more efficient implementation of `fetch(re-motecall(...)))` or as modifiers (`sum_kbn()`).

- conciseness is valued, but avoid abbreviation (`indexin()` rather than `indxin()`) as it becomes difficult to remember whether and how particular words are abbreviated.

If a function name requires multiple words, consider whether it might represent more than one concept and might be better split into pieces.

41.9   Don't overuse try-catch

It is better to avoid errors than to rely on catching them.

41.10   Don't parenthesize conditions

Julia doesn't require parens around conditions in `if` and `while`. Write:

```julia
if a == b
```

instead of:

```julia
if (a == b)
```

41.11   Don't overuse . . .

Splicing function arguments can be addictive. Instead of `[a..., b...]`, use simply `[a; b]`, which already concatenates arrays. `collect(a)` is better than `[a...]`, but since a is already iterable it is often even better to leave it alone, and not convert it to an array.

41.12   Don't use unnecessary static parameters

A function signature:

```julia
foo(x::T) where {T<:Real} = ...
```

should be written as:

```julia
foo(x::Real) = ...
```

instead, especially if T is not used in the function body. Even if T is used, it can be replaced with `typeof(x)` if convenient. There is no performance difference. Note that this is not a general caution against static parameters, just against uses where they are not needed.

Note also that container types, specifically may need type parameters in function calls. See the FAQ Avoid fields with abstract containers for more information.

41.13   Avoid confusion about whether something is an instance or a type

Sets of definitions like the following are confusing:

```julia
foo(::Type{MyType}) = ...
foo(::MyType) = foo(MyType)
```

Decide whether the concept in question will be written as `MyType` or `MyType()`, and stick to it.

The preferred style is to use instances by default, and only add methods involving `Type{MyType}` later if they become necessary to solve some problem.

If a type is effectively an enumeration, it should be defined as a single (ideally immutable struct or primitive) type, with the enumeration values being instances of it. Constructors and conversions can check whether values are valid. This design is preferred over making the enumeration an abstract type, with the "values" as subtypes.

41.14 Don't overuse macros

Be aware of when a macro could really be a function instead.

Calling `eval()` inside a macro is a particularly dangerous warning sign; it means the macro will only work when called at the top level. If such a macro is written as a function instead, it will naturally have access to the run-time values it needs.

41.15 Don't expose unsafe operations at the interface level

If you have a type that uses a native pointer:

```julia
mutable struct NativeType
    p::Ptr{UInt8}
    ...
end
```

don't write definitions like the following:

```julia
getindex(x::NativeType, i) = unsafe_load(x.p, i)
```

The problem is that users of this type can write `x[i]` without realizing that the operation is unsafe, and then be susceptible to memory bugs.

Such a function should either check the operation to ensure it is safe, or have `unsafe` somewhere in its name to alert callers.

41.16 Don't overload methods of base container types

It is possible to write definitions like the following:

```julia
show(io::IO, v::Vector{MyType}) = ...
```

This would provide custom showing of vectors with a specific new element type. While tempting, this should be avoided. The trouble is that users will expect a well-known type like `Vector()` to behave in a certain way, and overly customizing its behavior can make it harder to work with.

41.17 Avoid type piracy

"Type piracy" refers to the practice of extending or redefining methods in Base or other packages on types that you have not defined. In some cases, you can get away with type piracy with little ill effect. In extreme cases, however, you can even crash Julia (e.g. if your method extension or redefinition causes invalid input to be passed to a `ccall`). Type piracy can complicate reasoning about code, and may introduce incompatibilities that are hard to predict and diagnose.

As an example, suppose you wanted to define multiplication on symbols in a module:

```julia
module A
import Base.*
*(x::Symbol, y::Symbol) = Symbol(x,y)
end
```

The problem is that now any other module that uses `Base.*` will also see this definition. Since `Symbol` is defined in Base and is used by other modules, this can change the behavior of unrelated code unexpectedly. There are several alternatives here, including using a different function name, or wrapping the `Symbol`s in another type that you define.

Sometimes, coupled packages may engage in type piracy to separate features from definitions, especially when the packages were designed by collaborating authors, and when the definitions are reusable. For example, one package might provide some types useful for working with colors; another package could define methods for those types that enable conversions between color spaces. Another example might be a package that acts as a thin wrapper for some C code, which another package might then pirate to implement a higher-level, Julia-friendly API.

41.18   Be careful with type equality

You generally want to use `isa()` and `<:` (`issubtype()`) for testing types, not ==. Checking types for exact equality typically only makes sense when comparing to a known concrete type (e.g. `T == Float64`), or if you *really*, *really* know what you're doing.

41.19   Do not write `x->f(x)`

Since higher-order functions are often called with anonymous functions, it is easy to conclude that this is desirable or even necessary. But any function can be passed directly, without being "wrapped" in an anonymous function. Instead of writing `map(x->f(x), a)`, write `map(f, a)`.

41.20   Avoid using floats for numeric literals in generic code when possible

If you write generic code which handles numbers, and which can be expected to run with many different numeric type arguments, try using literals of a numeric type that will affect the arguments as little as possible through promotion.

For example,

```julia
julia> f(x) = 2.0 * x
f (generic function with 1 method)

julia> f(1//2)
1.0

julia> f(1/2)
1.0

julia> f(1)
2.0
```

while

```julia
julia> g(x) = 2 * x
g (generic function with 1 method)

julia> g(1//2)
1//1

julia> g(1/2)
1.0

julia> g(1)
2
```

As you can see, the second version, where we used an `Int` literal, preserved the type of the input argument, while the first didn't. This is because e.g. `promote_type(Int, Float64) == Float64`, and promotion happens with the multiplication. Similarly, `Rational` literals are less type disruptive than `Float64` literals, but more disruptive than `Int`s:

```
julia> h(x) = 2//1 * x
h (generic function with 1 method)

julia> h(1//2)
1//1

julia> h(1/2)
1.0

julia> h(1)
2//1
```

Thus, use `Int` literals when possible, with `Rational{Int}` for literal non-integer numbers, in order to make it easier to use your code.

Chapter 42

Frequently Asked Questions

42.1 Sessions and the REPL

How do I delete an object in memory?

Julia does not have an analog of MATLAB's `clear` function; once a name is defined in a Julia session (technically, in module `Main`), it is always present.

If memory usage is your concern, you can always replace objects with ones that consume less memory. For example, if A is a gigabyte-sized array that you no longer need, you can free the memory with `A = 0`. The memory will be released the next time the garbage collector runs; you can force this to happen with `gc()`.

How can I modify the declaration of a type in my session?

Perhaps you've defined a type and then realize you need to add a new field. If you try this at the REPL, you get the error:

```
ERROR: invalid redefinition of constant MyType
```

Types in module `Main` cannot be redefined.

While this can be inconvenient when you are developing new code, there's an excellent workaround. Modules can be replaced by redefining them, and so if you wrap all your new code inside a module you can redefine types and constants. You can't import the type names into `Main` and then expect to be able to redefine them there, but you can use the module name to resolve the scope. In other words, while developing you might use a workflow something like this:

```julia
include("mynewcode.jl")              # this defines a module MyModule
obj1 = MyModule.ObjConstructor(a, b)
obj2 = MyModule.somefunction(obj1)
# Got an error. Change something in "mynewcode.jl"
include("mynewcode.jl")              # reload the module
obj1 = MyModule.ObjConstructor(a, b) # old objects are no longer valid, must reconstruct
obj2 = MyModule.somefunction(obj1)   # this time it worked!
obj3 = MyModule.someotherfunction(obj2, c)
...
```

42.2 Functions

I passed an argument x to a function, modified it inside that function, but on the outside,

the variable x is still unchanged. Why?

Suppose you call a function like this:

```julia
julia> x = 10
10

julia> function change_value!(y)
           y = 17
       end
change_value! (generic function with 1 method)

julia> change_value!(x)
17

julia> x # x is unchanged!
10
```

In Julia, the binding of a variable x cannot be changed by passing x as an argument to a function. When calling change_value!(x) in the above example, y is a newly created variable, bound initially to the value of x, i.e. 10; then y is rebound to the constant 17, while the variable x of the outer scope is left untouched.

But here is a thing you should pay attention to: suppose x is bound to an object of type Array (or any other *mutable* type). From within the function, you cannot "unbind" x from this Array, but you can change its content. For example:

```julia
julia> x = [1,2,3]
3-element Array{Int64,1}:
 1
 2
 3

julia> function change_array!(A)
           A[1] = 5
       end
change_array! (generic function with 1 method)

julia> change_array!(x)
5

julia> x
3-element Array{Int64,1}:
 5
 2
 3
```

Here we created a function change_array!(), that assigns 5 to the first element of the passed array (bound to x at the call site, and bound to A within the function). Notice that, after the function call, x is still bound to the same array, but the content of that array changed: the variables A and x were distinct bindings refering to the same mutable Array object.

Can I use using or import inside a function?

No, you are not allowed to have a using or import statement inside a function. If you want to import a module but only use its symbols inside a specific function or set of functions, you have two options:

1. Use import:

```
import Foo
function bar(...)
    # ... refer to Foo symbols via Foo.baz ...
end
```

This loads the module Foo and defines a variable Foo that refers to the module, but does not import any of the other symbols from the module into the current namespace. You refer to the Foo symbols by their qualified names Foo.bar etc.

2. Wrap your function in a module:

```
module Bar
export bar
using Foo
function bar(...)
    # ... refer to Foo.baz as simply baz ....
end
end
using Bar
```

This imports all the symbols from Foo, but only inside the module Bar.

What does the ... operator do?

The two uses of the ... operator: slurping and splatting

Many newcomers to Julia find the use of ... operator confusing. Part of what makes the ... operator confusing is that it means two different things depending on context.

... combines many arguments into one argument in function definitions

In the context of function definitions, the ... operator is used to combine many different arguments into a single argument. This use of ... for combining many different arguments into a single argument is called slurping:

```
julia> function printargs(args...)
           @printf("%s\n", typeof(args))
           for (i, arg) in enumerate(args)
               @printf("Arg %d = %s\n", i, arg)
           end
       end
printargs (generic function with 1 method)

julia> printargs(1, 2, 3)
Tuple{Int64,Int64,Int64}
Arg 1 = 1
Arg 2 = 2
Arg 3 = 3
```

If Julia were a language that made more liberal use of ASCII characters, the slurping operator might have been written as <-... instead of ....

. . . splits one argument into many different arguments in function calls

In contrast to the use of the `. . .` operator to denote slurping many different arguments into one argument when defining a function, the `. . .` operator is also used to cause a single function argument to be split apart into many different arguments when used in the context of a function call. This use of `. . .` is called splatting:

```julia
julia> function threeargs(a, b, c)
           @printf("a = %s::%s\n", a, typeof(a))
           @printf("b = %s::%s\n", b, typeof(b))
           @printf("c = %s::%s\n", c, typeof(c))
       end
threeargs (generic function with 1 method)

julia> vec = [1, 2, 3]
3-element Array{Int64,1}:
 1
 2
 3

julia> threeargs(vec...)
a = 1::Int64
b = 2::Int64
c = 3::Int64
```

If Julia were a language that made more liberal use of ASCII characters, the splatting operator might have been written as `. . .->` instead of `. . . .`.

42.3   Types, type declarations, and constructors

What does "type-stable" mean?

It means that the type of the output is predictable from the types of the inputs. In particular, it means that the type of the output cannot vary depending on the *values* of the inputs. The following code is *not* type-stable:

```julia
julia> function unstable(flag::Bool)
           if flag
               return 1
           else
               return 1.0
           end
       end
unstable (generic function with 1 method)
```

It returns either an Int or a `Float64` depending on the value of its argument. Since Julia can't predict the return type of this function at compile-time, any computation that uses it will have to guard against both types possibly occurring, making generation of fast machine code difficult.

Why does Julia give a `DomainError` for certain seemingly-sensible operations?

Certain operations make mathematical sense but result in errors:

```julia
julia> sqrt(-2.0)
ERROR: DomainError:
```

```
sqrt will only return a complex result if called with a complex argument. Try sqrt(complex(x)).
Stacktrace:
 [1] sqrt(::Float64) at ./math.jl:425

julia> 2^-5
ERROR: DomainError:
Cannot raise an integer x to a negative power -n.
Make x a float by adding a zero decimal (e.g. 2.0^-n instead of 2^-n), or write 1/x^n,
↪  float(x)^-n, or (x//1)^-n.
Stacktrace:
 [1] power_by_squaring(::Int64, ::Int64) at ./intfuncs.jl:170
 [2] literal_pow(::Base.#^, ::Int64, ::Type{Val{-5}}) at ./intfuncs.jl:205
```

This behavior is an inconvenient consequence of the requirement for type-stability. In the case of `sqrt()`, most users want `sqrt(2.0)` to give a real number, and would be unhappy if it produced the complex number 1.4142135623730951 + 0.0im. One could write the `sqrt()` function to switch to a complex-valued output only when passed a negative number (which is what `sqrt()` does in some other languages), but then the result would not be type-stable and the `sqrt()` function would have poor performance.

In these and other cases, you can get the result you want by choosing an *input type* that conveys your willingness to accept an *output type* in which the result can be represented:

```
julia> sqrt(-2.0+0im)
0.0 + 1.4142135623730951im

julia> 2.0^-5
0.03125
```

Why does Julia use native machine integer arithmetic?

Julia uses machine arithmetic for integer computations. This means that the range of Int values is bounded and wraps around at either end so that adding, subtracting and multiplying integers can overflow or underflow, leading to some results that can be unsettling at first:

```
julia> typemax(Int)
9223372036854775807

julia> ans+1
-9223372036854775808

julia> -ans
-9223372036854775808

julia> 2*ans
0
```

Clearly, this is far from the way mathematical integers behave, and you might think it less than ideal for a high-level programming language to expose this to the user. For numerical work where efficiency and transparency are at a premium, however, the alternatives are worse.

One alternative to consider would be to check each integer operation for overflow and promote results to bigger integer types such as `Int128` or `BigInt` in the case of overflow. Unfortunately, this introduces major overhead on every integer operation (think incrementing a loop counter) – it requires emitting code to perform run-time overflow checks after arithmetic instructions and branches to handle potential overflows. Worse still, this would cause every

computation involving integers to be type-unstable. As we mentioned above, type-stability is crucial for effective generation of efficient code. If you can't count on the results of integer operations being integers, it's impossible to generate fast, simple code the way C and Fortran compilers do.

A variation on this approach, which avoids the appearance of type instability is to merge the `Int` and `BigInt` types into a single hybrid integer type, that internally changes representation when a result no longer fits into the size of a machine integer. While this superficially avoids type-instability at the level of Julia code, it just sweeps the problem under the rug by foisting all of the same difficulties onto the C code implementing this hybrid integer type. This approach *can* be made to work and can even be made quite fast in many cases, but has several drawbacks. One problem is that the in-memory representation of integers and arrays of integers no longer match the natural representation used by C, Fortran and other languages with native machine integers. Thus, to interoperate with those languages, we would ultimately need to introduce native integer types anyway. Any unbounded representation of integers cannot have a fixed number of bits, and thus cannot be stored inline in an array with fixed-size slots – large integer values will always require separate heap-allocated storage. And of course, no matter how clever a hybrid integer implementation one uses, there are always performance traps – situations where performance degrades unexpectedly. Complex representation, lack of interoperability with C and Fortran, the inability to represent integer arrays without additional heap storage, and unpredictable performance characteristics make even the cleverest hybrid integer implementations a poor choice for high-performance numerical work.

An alternative to using hybrid integers or promoting to BigInts is to use saturating integer arithmetic, where adding to the largest integer value leaves it unchanged and likewise for subtracting from the smallest integer value. This is precisely what Matlab™ does:

```
>> int64(9223372036854775807)

ans =

  9223372036854775807

>> int64(9223372036854775807) + 1

ans =

  9223372036854775807

>> int64(-9223372036854775808)

ans =

 -9223372036854775808

>> int64(-9223372036854775808) - 1

ans =

 -9223372036854775808
```

At first blush, this seems reasonable enough since 9223372036854775807 is much closer to 9223372036854775808 than -9223372036854775808 is and integers are still represented with a fixed size in a natural way that is compatible with C and Fortran. Saturated integer arithmetic, however, is deeply problematic. The first and most obvious issue is that this is not the way machine integer arithmetic works, so implementing saturated operations requires emitting instructions after each machine integer operation to check for underflow or overflow and replace the result with `typemin(Int)` or `typemax(Int)` as appropriate. This alone expands each integer operation from a single, fast instruction into half a dozen instructions, probably including branches. Ouch. But it gets worse – saturating integer arithmetic isn't associative. Consider this Matlab computation:

```
>> n = int64(2)^62
4611686018427387904

>> n + (n - 1)
9223372036854775807

>> (n + n) - 1
9223372036854775806
```

This makes it hard to write many basic integer algorithms since a lot of common techniques depend on the fact that machine addition with overflow *is* associative. Consider finding the midpoint between integer values `lo` and `hi` in Julia using the expression `(lo + hi) >>> 1`:

```
julia> n = 2^62
4611686018427387904

julia> (n + 2n) >>> 1
6917529027641081856
```

See? No problem. That's the correct midpoint between 2^62 and 2^63, despite the fact that n + 2n is -4611686018427387904. Now try it in Matlab:

```
>> (n + 2*n)/2

ans =

   4611686018427387904
```

Oops. Adding a `>>>` operator to Matlab wouldn't help, because saturation that occurs when adding n and 2n has already destroyed the information necessary to compute the correct midpoint.

Not only is lack of associativity unfortunate for programmers who cannot rely it for techniques like this, but it also defeats almost anything compilers might want to do to optimize integer arithmetic. For example, since Julia integers use normal machine integer arithmetic, LLVM is free to aggressively optimize simple little functions like `f(k) = 5k-1`. The machine code for this function is just this:

```
julia> code_native(f, Tuple{Int})
  .text
Filename: none
  pushq %rbp
  movq  %rsp, %rbp
Source line: 1
  leaq  -1(%rdi,%rdi,4), %rax
  popq  %rbp
  retq
  nopl  (%rax,%rax)
```

The actual body of the function is a single `leaq` instruction, which computes the integer multiply and add at once. This is even more beneficial when `f` gets inlined into another function:

```
julia> function g(k, n)
           for i = 1:n
               k = f(k)
           end
```

```
            return k
        end
g (generic function with 1 methods)

julia> code_native(g, Tuple{Int,Int})
  .text
Filename: none
  pushq %rbp
  movq  %rsp, %rbp
Source line: 2
  testq %rsi, %rsi
  jle L26
  nopl  (%rax)
Source line: 3
L16:
  leaq  -1(%rdi,%rdi,4), %rdi
Source line: 2
  decq  %rsi
  jne L16
Source line: 5
L26:
  movq  %rdi, %rax
  popq  %rbp
  retq
  nop
```

Since the call to f gets inlined, the loop body ends up being just a single `leaq` instruction. Next, consider what happens if we make the number of loop iterations fixed:

```
julia> function g(k)
           for i = 1:10
               k = f(k)
           end
           return k
       end
g (generic function with 2 methods)

julia> code_native(g,(Int,))
  .text
Filename: none
  pushq %rbp
  movq  %rsp, %rbp
Source line: 3
  imulq $9765625, %rdi, %rax     # imm = 0x9502F9
  addq  $-2441406, %rax          # imm = 0xFFDABF42
Source line: 5
  popq  %rbp
  retq
  nopw  %cs:(%rax,%rax)
```

Because the compiler knows that integer addition and multiplication are associative and that multiplication distributes over addition – neither of which is true of saturating arithmetic – it can optimize the entire loop down to just a multiply and an add. Saturated arithmetic completely defeats this kind of optimization since associativity and distributivity can fail at each loop iteration, causing different outcomes depending on which iteration the failure occurs in. The compiler can unroll the loop, but it cannot algebraically reduce multiple operations into fewer equivalent operations.

The most reasonable alternative to having integer arithmetic silently overflow is to do checked arithmetic everywhere, raising errors when adds, subtracts, and multiplies overflow, producing values that are not value-correct. In this blog post, Dan Luu analyzes this and finds that rather than the trivial cost that this approach should in theory have, it ends up having a substantial cost due to compilers (LLVM and GCC) not gracefully optimizing around the added overflow checks. If this improves in the future, we could consider defaulting to checked integer arithmetic in Julia, but for now, we have to live with the possibility of overflow.

What are the possible causes of an `UndefVarError` during remote execution?

As the error states, an immediate cause of an `UndefVarError` on a remote node is that a binding by that name does not exist. Let us explore some of the possible causes.

```julia
julia> module Foo
           foo() = remotecall_fetch(x->x, 2, "Hello")
       end

julia> Foo.foo()
ERROR: On worker 2:
UndefVarError: Foo not defined
[...]
```

The closure `x->x` carries a reference to `Foo`, and since `Foo` is unavailable on node 2, an `UndefVarError` is thrown.

Globals under modules other than `Main` are not serialized by value to the remote node. Only a reference is sent. Functions which create global bindings (except under `Main`) may cause an `UndefVarError` to be thrown later.

```julia
julia> @everywhere module Foo
           function foo()
               global gvar = "Hello"
               remotecall_fetch(()->gvar, 2)
           end
       end

julia> Foo.foo()
ERROR: On worker 2:
UndefVarError: gvar not defined
[...]
```

In the above example, `@everywhere module Foo` defined Foo on all nodes. However the call to `Foo.foo()` created a new global binding `gvar` on the local node, but this was not found on node 2 resulting in an `UndefVarError` error.

Note that this does not apply to globals created under module `Main`. Globals under module `Main` are serialized and new bindings created under `Main` on the remote node.

```julia
julia> gvar_self = "Node1"
"Node1"

julia> remotecall_fetch(()->gvar_self, 2)
"Node1"

julia> remotecall_fetch(whos, 2)
        From worker 2:                                  Base  41762 KB    Module
        From worker 2:                                  Core  27337 KB    Module
        From worker 2:                                   Foo   2477 bytes  Module
```

```
    From worker 2:                                    Main   46191 KB     Module
    From worker 2:                          gvar_self    13 bytes  String
```

This does not apply to `function` or `type` declarations. However, anonymous functions bound to global variables are serialized as can be seen below.

```
julia> bar() = 1
bar (generic function with 1 method)

julia> remotecall_fetch(bar, 2)
ERROR: On worker 2:
UndefVarError: #bar not defined
[...]

julia> anon_bar  = ()->1
(::#21) (generic function with 1 method)

julia> remotecall_fetch(anon_bar, 2)
1
```

42.4   Packages and Modules

What is the difference between "using" and "importall"?

There is only one difference, and on the surface (syntax-wise) it may seem very minor. The difference between `using` and `importall` is that with `using` you need to say `function Foo.bar(..` to extend module Foo's function bar with a new method, but with `importall` or `import Foo.bar`, you only need to say `function bar(...` and it automatically extends module Foo's function bar.

If you use `importall`, then `function Foo.bar(...` and `function bar(...` become equivalent. If you use `using`, then they are different.

The reason this is important enough to have been given separate syntax is that you don't want to accidentally extend a function that you didn't know existed, because that could easily cause a bug. This is most likely to happen with a method that takes a common type like a string or integer, because both you and the other module could define a method to handle such a common type. If you use `importall`, then you'll replace the other module's implementation of `bar(s::AbstractString)` with your new implementation, which could easily do something completely different (and break all/many future usages of the other functions in module Foo that depend on calling bar).

42.5   Nothingness and missing values

How does "null" or "nothingness" work in Julia?

Unlike many languages (for example, C and Java), Julia does not have a "null" value. When a reference (variable, object field, or array element) is uninitialized, accessing it will immediately throw an error. This situation can be detected using the `isdefined` function.

Some functions are used only for their side effects, and do not need to return a value. In these cases, the convention is to return the value `nothing`, which is just a singleton object of type `Void`. This is an ordinary type with no fields; there is nothing special about it except for this convention, and that the REPL does not print anything for it. Some language constructs that would not otherwise have a value also yield `nothing`, for example `if false; end`.

For situations where a value exists only sometimes (for example, missing statistical data), it is best to use the `Nullable{T}` type, which allows specifying the type of a missing value.

The empty tuple (()) is another form of nothingness. But, it should not really be thought of as nothing but rather a tuple of zero values.

In code written for Julia prior to version 0.4 you may occasionally see None, which is quite different. It is the empty (or "bottom") type, a type with no values and no subtypes (except itself). This is now written as Union{} (an empty union type). You will generally not need to use this type.

42.6 Memory

Why does x += y allocate memory when x and y are arrays?

In Julia, x += y gets replaced during parsing by x = x + y. For arrays, this has the consequence that, rather than storing the result in the same location in memory as x, it allocates a new array to store the result.

While this behavior might surprise some, the choice is deliberate. The main reason is the presence of immutable objects within Julia, which cannot change their value once created. Indeed, a number is an immutable object; the statements x = 5; x += 1 do not modify the meaning of 5, they modify the value bound to x. For an immutable, the only way to change the value is to reassign it.

To amplify a bit further, consider the following function:

```julia
function power_by_squaring(x, n::Int)
    ispow2(n) || error("This implementation only works for powers of 2")
    while n >= 2
        x *= x
        n >>= 1
    end
    x
end
```

After a call like x = 5; y = power_by_squaring(x, 4), you would get the expected result: x == 5 && y == 625. However, now suppose that *=, when used with matrices, instead mutated the left hand side. There would be two problems:

- For general square matrices, A = A*B cannot be implemented without temporary storage: A[1,1] gets computed and stored on the left hand side before you're done using it on the right hand side.

- Suppose you were willing to allocate a temporary for the computation (which would eliminate most of the point of making *= work in-place); if you took advantage of the mutability of x, then this function would behave differently for mutable vs. immutable inputs. In particular, for immutable x, after the call you'd have (in general) y != x, but for mutable x you'd have y == x.

Because supporting generic programming is deemed more important than potential performance optimizations that can be achieved by other means (e.g., using explicit loops), operators like += and *= work by rebinding new values.

42.7 Asynchronous IO and concurrent synchronous writes

Why do concurrent writes to the same stream result in inter-mixed output?

While the streaming I/O API is synchronous, the underlying implementation is fully asynchronous.

Consider the printed output from the following:

```julia
julia> @sync for i in 1:3
           @async write(STDOUT, string(i), " Foo ", " Bar ")
       end
123 Foo  Foo  Foo  Bar  Bar  Bar
```

This is happening because, while the `write` call is synchronous, the writing of each argument yields to other tasks while waiting for that part of the I/O to complete.

`print` and `println` "lock" the stream during a call. Consequently changing `write` to `println` in the above example results in:

```julia
julia> @sync for i in 1:3
           @async println(STDOUT, string(i), " Foo ", " Bar ")
       end
1 Foo  Bar
2 Foo  Bar
3 Foo  Bar
```

You can lock your writes with a `ReentrantLock` like this:

```julia
julia> l = ReentrantLock()
ReentrantLock(Nullable{Task}(), Condition(Any[]), 0)

julia> @sync for i in 1:3
           @async begin
               lock(l)
               try
                   write(STDOUT, string(i), " Foo ", " Bar ")
               finally
                   unlock(l)
               end
           end
       end
1 Foo  Bar 2 Foo  Bar 3 Foo  Bar
```

42.8   Julia Releases

Do I want to use a release, beta, or nightly version of Julia?

You may prefer the release version of Julia if you are looking for a stable code base. Releases generally occur every 6 months, giving you a stable platform for writing code.

You may prefer the beta version of Julia if you don't mind being slightly behind the latest bugfixes and changes, but find the slightly faster rate of changes more appealing. Additionally, these binaries are tested before they are published to ensure they are fully functional.

You may prefer the nightly version of Julia if you want to take advantage of the latest updates to the language, and don't mind if the version available today occasionally doesn't actually work.

Finally, you may also consider building Julia from source for yourself. This option is mainly for those individuals who are comfortable at the command line, or interested in learning. If this describes you, you may also be interested in reading our guidelines for contributing.

Links to each of these download types can be found on the download page at https://julialang.org/downloads/. Note that not all versions of Julia are available for all platforms.

When are deprecated functions removed?

Deprecated functions are removed after the subsequent release. For example, functions marked as deprecated in the 0.1 release will not be available starting with the 0.2 release.

Chapter 43

Noteworthy Differences from other Languages

43.1 Noteworthy differences from MATLAB

Although MATLAB users may find Julia's syntax familiar, Julia is not a MATLAB clone. There are major syntactic and functional differences. The following are some noteworthy differences that may trip up Julia users accustomed to MATLAB:

- Julia arrays are indexed with square brackets, A[i,j].

- Julia arrays are assigned by reference. After A=B, changing elements of B will modify A as well.

- Julia values are passed and assigned by reference. If a function modifies an array, the changes will be visible in the caller.

- Julia does not automatically grow arrays in an assignment statement. Whereas in MATLAB a(4) = 3.2 can create the array a = [0 0 0 3.2] and a(5) = 7 can grow it into a = [0 0 0 3.2 7], the corresponding Julia statement a[5] = 7 throws an error if the length of a is less than 5 or if this statement is the first use of the identifier a. Julia has push!() and append!(), which grow Vectors much more efficiently than MATLAB's a(end+1) = val.

- The imaginary unit sqrt(-1) is represented in Julia as im, not i or j as in MATLAB.

- In Julia, literal numbers without a decimal point (such as 42) create integers instead of floating point numbers. Arbitrarily large integer literals are supported. As a result, some operations such as 2^-1 will throw a domain error as the result is not an integer (see the FAQ entry on domain errors for details).

- In Julia, multiple values are returned and assigned as tuples, e.g. (a, b) = (1, 2) or a, b = 1, 2. MAT-LAB's nargout, which is often used in MATLAB to do optional work based on the number of returned values, does not exist in Julia. Instead, users can use optional and keyword arguments to achieve similar capabilities.

- Julia has true one-dimensional arrays. Column vectors are of size N, not Nx1. For example, rand(N) makes a 1-dimensional array.

- In Julia, [x,y,z] will always construct a 3-element array containing x, y and z.

  - To concatenate in the first ("vertical") dimension use either vcat(x,y,z) or separate with semicolons ([x; y; z]).
  - To concatenate in the second ("horizontal") dimension use either hcat(x,y,z) or separate with spaces ([x y z]).
  - To construct block matrices (concatenating in the first two dimensions), use either hvcat() or combine spaces and semicolons ([a b; c d]).

- In Julia, a:b and a:b:c construct Range objects. To construct a full vector like in MATLAB, use `collect(a:b)`. Generally, there is no need to call `collect` though. Range will act like a normal array in most cases but is more efficient because it lazily computes its values. This pattern of creating specialized objects instead of full arrays is used frequently, and is also seen in functions such as `linspace`, or with iterators such as `enumerate`, and `zip`. The special objects can mostly be used as if they were normal arrays.

- Functions in Julia return values from their last expression or the `return` keyword instead of listing the names of variables to return in the function definition (see The return Keyword for details).

- A Julia script may contain any number of functions, and all definitions will be externally visible when the file is loaded. Function definitions can be loaded from files outside the current working directory.

- In Julia, reductions such as `sum()`, `prod()`, and `max()` are performed over every element of an array when called with a single argument, as in `sum(A)`, even if A has more than one dimension.

- In Julia, functions such as `sort()` that operate column-wise by default (`sort(A)` is equivalent to `sort(A,1)`) do not have special behavior for 1xN arrays; the argument is returned unmodified since it still performs `sort(A,1)`. To sort a 1xN matrix like a vector, use `sort(A,2)`.

- In Julia, if A is a 2-dimensional array, `fft(A)` computes a 2D FFT. In particular, it is not equivalent to `fft(A,1)`, which computes a 1D FFT acting column-wise.

- In Julia, parentheses must be used to call a function with zero arguments, like in `tic()` and `toc()`.

- Julia discourages the used of semicolons to end statements. The results of statements are not automatically printed (except at the interactive prompt), and lines of code do not need to end with semicolons. `println()` or `@printf()` can be used to print specific output.

- In Julia, if A and B are arrays, logical comparison operations like A == B do not return an array of booleans. Instead, use A .== B, and similarly for the other boolean operators like <, > and =.

- In Julia, the operators &, |, and (xor) perform the bitwise operations equivalent to and, or, and xor respectively in MATLAB, and have precedence similar to Python's bitwise operators (unlike C). They can operate on scalars or element-wise across arrays and can be used to combine logical arrays, but note the difference in order of operations: parentheses may be required (e.g., to select elements of A equal to 1 or 2 use (A .== 1) | (A .== 2)).

- In Julia, the elements of a collection can be passed as arguments to a function using the splat operator ..., as in `xs=[1,2]; f(xs...)`.

- Julia's `svd()` returns singular values as a vector instead of as a dense diagonal matrix.

- In Julia, ... is not used to continue lines of code. Instead, incomplete expressions automatically continue onto the next line.

- In both Julia and MATLAB, the variable ans is set to the value of the last expression issued in an interactive session. In Julia, unlike MATLAB, ans is not set when Julia code is run in non-interactive mode.

- Julia's types do not support dynamically adding fields at runtime, unlike MATLAB's classes. Instead, use a `Dict`.

- In Julia each module has its own global scope/namespace, whereas in MATLAB there is just one global scope.

- In MATLAB, an idiomatic way to remove unwanted values is to use logical indexing, like in the expression x(x>3) or in the statement x(x>3) = [] to modify x in-place. In contrast, Julia provides the higher order functions `filter()` and `filter!()`, allowing users to write `filter(z->z>3, x)` and `filter!(z->z>3, x)` as alternatives to the corresponding transliterations `x[x.>3]` and `x = x[x.>3]`. Using `filter!()` reduces the use of temporary arrays.

- The analogue of extracting (or "dereferencing") all elements of a cell array, e.g. in `vertcat(A{:})` in MATLAB, is written using the splat operator in Julia, e.g. as `vcat(A...)`.

43.2 Noteworthy differences from R

One of Julia's goals is to provide an effective language for data analysis and statistical programming. For users coming to Julia from R, these are some noteworthy differences:

- Julia's single quotes enclose characters, not strings.

- Julia can create substrings by indexing into strings. In R, strings must be converted into character vectors before creating substrings.

- In Julia, like Python but unlike R, strings can be created with triple quotes `""" ... """`. This syntax is convenient for constructing strings that contain line breaks.

- In Julia, varargs are specified using the splat operator `...`, which always follows the name of a specific variable, unlike R, for which `...` can occur in isolation.

- In Julia, modulus is `mod(a, b)`, not `a %% b`. `%` in Julia is the remainder operator.

- In Julia, not all data structures support logical indexing. Furthermore, logical indexing in Julia is supported only with vectors of length equal to the object being indexed. For example:

    - In R, `c(1, 2, 3, 4)[c(TRUE, FALSE)]` is equivalent to `c(1, 3)`.
    - In R, `c(1, 2, 3, 4)[c(TRUE, FALSE, TRUE, FALSE)]` is equivalent to `c(1, 3)`.
    - In Julia, `[1, 2, 3, 4][[true, false]]` throws a `BoundsError`.
    - In Julia, `[1, 2, 3, 4][[true, false, true, false]]` produces `[1, 3]`.

- Like many languages, Julia does not always allow operations on vectors of different lengths, unlike R where the vectors only need to share a common index range. For example, `c(1, 2, 3, 4) + c(1, 2)` is valid R but the equivalent `[1, 2, 3, 4] + [1, 2]` will throw an error in Julia.

- Julia's `map()` takes the function first, then its arguments, unlike `lapply(<structure>, function, ...)` in R. Similarly Julia's equivalent of `apply(X, MARGIN, FUN, ...)` in R is `mapslices()` where the function is the first argument.

- Multivariate apply in R, e.g. `mapply(choose, 11:13, 1:3)`, can be written as `broadcast(binomial, 11:13, 1:3)` in Julia. Equivalently Julia offers a shorter dot syntax for vectorizing functions `binomial.(11:13, 1:3)`.

- Julia uses end to denote the end of conditional blocks, like `if`, loop blocks, like `while`/ `for`, and functions. In lieu of the one-line `if ( cond ) statement`, Julia allows statements of the form `if cond; statement; end`, `cond && statement` and `!cond || statement`. Assignment statements in the latter two syntaxes must be explicitly wrapped in parentheses, e.g. `cond && (x = value)`.

- In Julia, `<-`, `<<-` and `->` are not assignment operators.

- Julia's `->` creates an anonymous function, like Python.

- Julia constructs vectors using brackets. Julia's `[1, 2, 3]` is the equivalent of R's `c(1, 2, 3)`.

- Julia's `*` operator can perform matrix multiplication, unlike in R. If A and B are matrices, then `A * B` denotes a matrix multiplication in Julia, equivalent to R's `A %*% B`. In R, this same notation would perform an element-wise (Hadamard) product. To get the element-wise multiplication operation, you need to write `A .* B` in Julia.

- Julia performs matrix transposition using the `.'` operator and conjugated transposition using the `'` operator. Julia's `A.'` is therefore equivalent to R's `t(A)`.

- Julia does not require parentheses when writing `if` statements or `for`/`while` loops: use `for i in [1, 2, 3]` instead of `for (i in c(1, 2, 3))` and `if i == 1` instead of `if (i == 1)`.

- Julia does not treat the numbers `0` and `1` as Booleans. You cannot write `if (1)` in Julia, because `if` statements accept only booleans. Instead, you can write `if true`, `if Bool(1)`, or `if 1==1`.

- Julia does not provide `nrow` and `ncol`. Instead, use `size(M, 1)` for `nrow(M)` and `size(M, 2)` for `ncol(M)`.

- Julia is careful to distinguish scalars, vectors and matrices. In R, `1` and `c(1)` are the same. In Julia, they can not be used interchangeably. One potentially confusing result of this is that `x' * y` for vectors `x` and `y` is a 1-element vector, not a scalar. To get a scalar, use `dot(x, y)`.

- Julia's `diag()` and `diagm()` are not like R's.

- Julia cannot assign to the results of function calls on the left hand side of an assignment operation: you cannot write `diag(M) = ones(n)`.

- Julia discourages populating the main namespace with functions. Most statistical functionality for Julia is found in packages under the JuliaStats organization. For example:

    - Functions pertaining to probability distributions are provided by the Distributions package.
    - The DataFrames package provides data frames.
    - Generalized linear models are provided by the GLM package.

- Julia provides tuples and real hash tables, but not R-style lists. When returning multiple items, you should typically use a tuple: instead of `list(a = 1, b = 2)`, use `(1, 2)`.

- Julia encourages users to write their own types, which are easier to use than S3 or S4 objects in R. Julia's multiple dispatch system means that `table(x::TypeA)` and `table(x::TypeB)` act like R's `table.TypeA(x)` and `table.TypeB(x)`.

- In Julia, values are passed and assigned by reference. If a function modifies an array, the changes will be visible in the caller. This is very different from R and allows new functions to operate on large data structures much more efficiently.

- In Julia, vectors and matrices are concatenated using `hcat()`, `vcat()` and `hvcat()`, not c, rbind and cbind like in R.

- In Julia, a range like `a:b` is not shorthand for a vector like in R, but is a specialized Range that is used for iteration without high memory overhead. To convert a range into a vector, use `collect(a:b)`.

- Julia's `max()` and `min()` are the equivalent of `pmax` and `pmin` respectively in R, but both arguments need to have the same dimensions. While `maximum()` and `minimum()` replace max and min in R, there are important differences.

- Julia's `sum()`, `prod()`, `maximum()`, and `minimum()` are different from their counterparts in R. They all accept one or two arguments. The first argument is an iterable collection such as an array. If there is a second argument, then this argument indicates the dimensions, over which the operation is carried out. For instance, let `A=[[1 2],[3 4]]` in Julia and `B=rbind(c(1,2),c(3,4))` be the same matrix in R. Then `sum(A)` gives the same result as `sum(B)`, but `sum(A, 1)` is a row vector containing the sum over each column and `sum(A, 2)` is a column vector containing the sum over each row. This contrasts to the behavior of R, where `sum(B,1)=11` and `sum(B,2)=12`. If the second argument is a vector, then it specifies all the dimensions over which the sum is performed, e.g., `sum(A,[1,2])=10`. It should be noted that there is no error checking regarding the second argument.

- Julia has several functions that can mutate their arguments. For example, it has both `sort()` and `sort!()`.

- In R, performance requires vectorization. In Julia, almost the opposite is true: the best performing code is often achieved by using devectorized loops.

- Julia is eagerly evaluated and does not support R-style lazy evaluation. For most users, this means that there are very few unquoted expressions or column names.

- Julia does not support the `NULL` type.

- Julia lacks the equivalent of R's `assign` or `get`.

- In Julia, `return` does not require parentheses.

- In R, an idiomatic way to remove unwanted values is to use logical indexing, like in the expression `x[x>3]` or in the statement `x = x[x>3]` to modify x in-place. In contrast, Julia provides the higher order functions `filter()` and `filter!()`, allowing users to write `filter(z->z>3, x)` and `filter!(z->z>3, x)` as alternatives to the corresponding transliterations `x[x.>3]` and `x = x[x.>3]`. Using `filter!()` reduces the use of temporary arrays.

43.3 Noteworthy differences from Python

- Julia requires `end` to end a block. Unlike Python, Julia has no `pass` keyword.

- In Julia, indexing of arrays, strings, etc. is 1-based not 0-based.

- Julia's slice indexing includes the last element, unlike in Python. `a[2:3]` in Julia is `a[1:3]` in Python.

- Julia does not support negative indexes. In particular, the last element of a list or array is indexed with end in Julia, not -1 as in Python.

- Julia's `for`, `if`, `while`, etc. blocks are terminated by the `end` keyword. Indentation level is not significant as it is in Python.

- Julia has no line continuation syntax: if, at the end of a line, the input so far is a complete expression, it is considered done; otherwise the input continues. One way to force an expression to continue is to wrap it in parentheses.

- Julia arrays are column major (Fortran ordered) whereas NumPy arrays are row major (C-ordered) by default. To get optimal performance when looping over arrays, the order of the loops should be reversed in Julia relative to NumPy (see relevant section of Performance Tips).

- Julia's updating operators (e.g. +=, -=, ...) are *not in-place* whereas NumPy's are. This means `A = ones(4); B = A; B += 3` doesn't change values in A, it rather rebinds the name B to the result of the right-hand side `B = B + 3`, which is a new array. Use `B[:] += 3`, explicit loops, or `InplaceOps.jl`.

- Julia evaluates default values of function arguments every time the method is invoked, unlike in Python where the default values are evaluated only once when the function is defined. For example, the function `f(x=rand()) = x` returns a new random number every time it is invoked without argument. On the other hand, the function `g(x=[1,2]) = push!(x,3)` returns `[1,2,3]` every time it is called as `g()`.

- In Julia % is the remainder operator, whereas in Python it is the modulus.

43.4   Noteworthy differences from C/C++

- Julia arrays are indexed with square brackets, and can have more than one dimension `A[i,j]`. This syntax is not just syntactic sugar for a reference to a pointer or address as in C/C++. See the Julia documentation for the syntax for array construction (it has changed between versions).

- In Julia, indexing of arrays, strings, etc. is 1-based not 0-based.

- Julia arrays are assigned by reference. After A=B, changing elements of B will modify A as well. Updating operators like += do not operate in-place, they are equivalent to A = A + B which rebinds the left-hand side to the result of the right-hand side expression.

- Julia arrays are column major (Fortran ordered) whereas C/C++ arrays are row major ordered by default. To get optimal performance when looping over arrays, the order of the loops should be reversed in Julia relative to C/C++ (see relevant section of Performance Tips).

- Julia values are passed and assigned by reference. If a function modifies an array, the changes will be visible in the caller.

- In Julia, whitespace is significant, unlike C/C++, so care must be taken when adding/removing whitespace from a Julia program.

- In Julia, literal numbers without a decimal point (such as 42) create signed integers, of type `Int`, but literals too large to fit in the machine word size will automatically be promoted to a larger size type, such as `Int64` (if `Int` is `Int32`), `Int128`, or the arbitrarily large `BigInt` type. There are no numeric literal suffixes, such as L, LL, U, UL, ULL to indicate unsigned and/or signed vs. unsigned. Decimal literals are always signed, and hexadecimal literals (which start with 0x like C/C++), are unsigned. Hexadecimal literals also, unlike C/C++/Java and unlike decimal literals in Julia, have a type based on the *length* of the literal, including leading 0s. For example, 0x0 and 0x00 have type `UInt8`, 0x000 and 0x0000 have type `UInt16`, then literals with 5 to 8 hex digits have type UInt32, 9 to 16 hex digits type UInt64 and 17 to 32 hex digits type UInt128. This needs to be taken into account when defining hexadecimal masks, for example ~0xf == 0xf0 is very different from ~0x000f == 0xfff0. 64 bit Float64 and 32 bit `Float32` bit literals are expressed as 1.0 and 1.0f0 respectively. Floating point literals are rounded (and not promoted to the BigFloat type) if they can not be exactly represented. Floating point literals are closer in behavior to C/C++. Octal (prefixed with 0o) and binary (prefixed with 0b) literals are also treated as unsigned.

- String literals can be delimited with either " or """, """ delimited literals can contain " characters without quoting it like "\"" String literals can have values of other variables or expressions interpolated into them, indicated by $variablename or $(expression), which evaluates the variable name or the expression in the context of the function.

- // indicates a `Rational` number, and not a single-line comment (which is # in Julia)

- #= indicates the start of a multiline comment, and =# ends it.

- Functions in Julia return values from their last expression(s) or the `return` keyword. Multiple values can be returned from functions and assigned as tuples, e.g. (a, b) = myfunction() or a, b = myfunction(), instead of having to pass pointers to values as one would have to do in C/C++ (i.e. a = myfunction(&b).

- Julia does not require the use of semicolons to end statements. The results of expressions are not automatically printed (except at the interactive prompt, i.e. the REPL), and lines of code do not need to end with semicolons. println() or @printf() can be used to print specific output. In the REPL, ; can be used to suppress output. ; also has a different meaning within [ ], something to watch out for. ; can be used to separate expressions on a single line, but are not strictly necessary in many cases, and are more an aid to readability.

- In Julia, the operator (`xor`) performs the bitwise XOR operation, i.e. `^` in C/C++. Also, the bitwise operators do not have the same precedence as C/++, so parenthesis may be required.

- Julia's `^` is exponentiation (pow), not bitwise XOR as in C/C++ (use , or `xor`, in Julia)

- Julia has two right-shift operators, `>>` and `>>>`. `>>>` performs an arithmetic shift, `>>` always performs a logical shift, unlike C/C++, where the meaning of `>>` depends on the type of the value being shifted.

- Julia's `->` creates an anonymous function, it does not access a member via a pointer.

- Julia does not require parentheses when writing `if` statements or `for`/`while` loops: use `for i in [1, 2, 3]` instead of `for (int i=1; i <= 3; i++)` and `if i == 1` instead of `if (i == 1)`.

- Julia does not treat the numbers `0` and `1` as Booleans. You cannot write `if (1)` in Julia, because `if` statements accept only booleans. Instead, you can write `if true`, `if Bool(1)`, or `if 1==1`.

- Julia uses end to denote the end of conditional blocks, like `if`, loop blocks, like `while`/ `for`, and functions. In lieu of the one-line `if ( cond ) statement`, Julia allows statements of the form `if cond; statement; end`, `cond && statement` and `!cond || statement`. Assignment statements in the latter two syntaxes must be explicitly wrapped in parentheses, e.g. `cond && (x = value)`, because of the operator precedence.

- Julia has no line continuation syntax: if, at the end of a line, the input so far is a complete expression, it is considered done; otherwise the input continues. One way to force an expression to continue is to wrap it in parentheses.

- Julia macros operate on parsed expressions, rather than the text of the program, which allows them to perform sophisticated transformations of Julia code. Macro names start with the @ character, and have both a function-like syntax, `@mymacro(arg1, arg2, arg3)`, and a statement-like syntax, `@mymacro arg1 arg2 arg3`. The forms are interchangable; the function-like form is particularly useful if the macro appears within another expression, and is often clearest. The statement-like form is often used to annotate blocks, as in the parallel `for` construct: `@parallel for i in 1:n; #= body =#; end`. Where the end of the macro construct may be unclear, use the function-like form.

- Julia now has an enumeration type, expressed using the macro `@enum(name, value1, value2, ...)` For example: `@enum(Fruit, banana=1, apple, pear)`

- By convention, functions that modify their arguments have a `!` at the end of the name, for example `push!`.

- In C++, by default, you have static dispatch, i.e. you need to annotate a function as virtual, in order to have dynamic dispatch. On the other hand, in Julia every method is "virtual" (although it's more general than that since methods are dispatched on every argument type, not only `this`, using the most-specific-declaration rule).

Chapter 44

Unicode Input

The following table lists Unicode characters that can be entered via tab completion of LaTeX-like abbreviations in the Julia REPL (and in various other editing environments). You can also get information on how to type a symbol by entering it in the REPL help, i.e. by typing ? and then entering the symbol in the REPL (e.g., by copy-paste from somewhere you saw the symbol).

> **Warning**
>
> This table may appear to contain missing characters in the second column, or even show characters that are inconsistent with the characters as they are rendered in the Julia REPL. In these cases, users are strongly advised to check their choice of fonts in their browser and REPL environment, as there are known issues with glyphs in many fonts.

Part IV

Standard Library

Chapter 45

Essentials

45.1 Introduction

The Julia standard library contains a range of functions and macros appropriate for performing scientific and numerical computing, but is also as broad as those of many general purpose programming languages. Additional functionality is available from a growing collection of available packages. Functions are grouped by topic below.

Some general notes:

- Except for functions in built-in modules (Pkg, `Collections`, `Test` and `Profile`), all functions documented here are directly available for use in programs.

- To use module functions, use `import Module` to import the module, and `Module.fn(x)` to use the functions.

- Alternatively, `using Module` will import all exported `Module` functions into the current namespace.

- By convention, function names ending with an exclamation point (!) modify their arguments. Some functions have both modifying (e.g., `sort!`) and non-modifying (`sort`) versions.

45.2 Getting Around

Base.exit – Function.

```
exit([code])
```

Quit (or control-D at the prompt). The default exit code is zero, indicating that the processes completed successfully.

source

Base.quit – Function.

```
quit()
```

Quit the program indicating that the processes completed successfully. This function calls `exit(0)` (see exit).

source

Base.atexit – Function.

```
atexit(f)
```

Register a zero-argument function `f()` to be called at process exit. `atexit()` hooks are called in last in first out (LIFO) order and run before object finalizers.

source

`Base.atreplinit` – Function.

```
atreplinit(f)
```

Register a one-argument function to be called before the REPL interface is initialized in interactive sessions; this is useful to customize the interface. The argument of `f` is the REPL object. This function should be called from within the `.juliarc.jl` initialization file.

source

`Base.isinteractive` – Function.

```
isinteractive() -> Bool
```

Determine whether Julia is running an interactive session.

source

`Base.whos` – Function.

```
whos(io::IO=STDOUT, m::Module=current_module(), pattern::Regex=r"")
```

Print information about exported global variables in a module, optionally restricted to those matching `pattern`.

The memory consumption estimate is an approximate lower bound on the size of the internal structure of the object.

source

`Base.summarysize` – Function.

```
Base.summarysize(obj; exclude=Union{...}, chargeall=Union{...}) -> Int
```

Compute the amount of memory used by all unique objects reachable from the argument.

Keyword Arguments

- `exclude`: specifies the types of objects to exclude from the traversal.
- `chargeall`: specifies the types of objects to always charge the size of all of their fields, even if those fields would normally be excluded.

source

`Base.edit` – Method.

```
edit(path::AbstractString, line::Integer=0)
```

Edit a file or directory optionally providing a line number to edit the file at. Returns to the `julia` prompt when you quit the editor. The editor can be changed by setting JULIA_EDITOR, VISUAL or EDITOR as an environment variable.

source

`Base.edit` – Method.

```
edit(function, [types])
```

Edit the definition of a function, optionally specifying a tuple of types to indicate which method to edit. The editor can be changed by setting JULIA_EDITOR, VISUAL or EDITOR as an environment variable.

source

Base.@edit – Macro.

```
@edit
```

Evaluates the arguments to the function or macro call, determines their types, and calls the edit function on the resulting expression.

source

Base.less – Method.

```
less(file::AbstractString, [line::Integer])
```

Show a file using the default pager, optionally providing a starting line number. Returns to the julia prompt when you quit the pager.

source

Base.less – Method.

```
less(function, [types])
```

Show the definition of a function using the default pager, optionally specifying a tuple of types to indicate which method to see.

source

Base.@less – Macro.

```
@less
```

Evaluates the arguments to the function or macro call, determines their types, and calls the less function on the resulting expression.

source

Base.clipboard – Method.

```
clipboard(x)
```

Send a printed form of x to the operating system clipboard ("copy").

source

Base.clipboard – Method.

```
clipboard() -> AbstractString
```

Return a string with the contents of the operating system clipboard ("paste").

source

Base.reload – Function.

```
reload(name::AbstractString)
```

Force reloading of a package, even if it has been loaded before. This is intended for use during package development as code is modified.

source

Base.require – Function.

```
require(module::Symbol)
```

This function is part of the implementation of using / import, if a module is not already defined in Main. It can also be called directly to force reloading a module, regardless of whether it has been loaded before (for example, when interactively developing libraries).

Loads a source file, in the context of the Main module, on every active node, searching standard locations for files. require is considered a top-level operation, so it sets the current include path but does not use it to search for files (see help for include). This function is typically used to load library code, and is implicitly called by using to load packages.

When searching for files, require first looks for package code under Pkg.dir(), then tries paths in the global array LOAD_PATH. require is case-sensitive on all platforms, including those with case-insensitive filesystems like macOS and Windows.

source

Base.compilecache – Function.

```
Base.compilecache(module::String)
```

Creates a precompiled cache file for a module and all of its dependencies. This can be used to reduce package load times. Cache files are stored in LOAD_CACHE_PATH[1], which defaults to ~/.julia/lib/VERSION. See Module initialization and precompilation for important notes.

source

Base.__precompile__ – Function.

```
__precompile__(isprecompilable::Bool=true)
```

Specify whether the file calling this function is precompilable. If isprecompilable is true, then __precompile__ throws an exception when the file is loaded by using/import/require *unless* the file is being precompiled, and in a module file it causes the module to be automatically precompiled when it is imported. Typically, __precompile__() should occur before the module declaration in the file, or better yet VERSION >= v"0.4" && __precompile__() in order to be backward-compatible with Julia 0.3.

If a module or file is *not* safely precompilable, it should call __precompile__(false) in order to throw an error if Julia attempts to precompile it.

__precompile__() should *not* be used in a module unless all of its dependencies are also using __precompile__(). Failure to do so can result in a runtime error when loading the module.

source

Base.include – Function.

```
include(path::AbstractString)
```

Evaluate the contents of the input source file in the current context. Returns the result of the last evaluated expression of the input file. During including, a task-local include path is set to the directory containing the file. Nested calls to `include` will search relative to that path. All paths refer to files on node 1 when running in parallel, and files will be fetched from node 1. This function is typically used to load source interactively, or to combine files in packages that are broken into multiple source files.

source

`Base.include_string` – Function.

```
include_string(code::AbstractString, filename::AbstractString="string")
```

Like `include`, except reads code from the given string rather than from a file. Since there is no file path involved, no path processing or fetching from node 1 is done.

source

`Base.include_dependency` – Function.

```
include_dependency(path::AbstractString)
```

In a module, declare that the file specified by `path` (relative or absolute) is a dependency for precompilation; that is, the module will need to be recompiled if this file changes.

This is only needed if your module depends on a file that is not used via `include`. It has no effect outside of compilation.

source

`Base.Docs.apropos` – Function.

```
apropos(string)
```

Search through all documentation for a string, ignoring case.

source

`Base.which` – Method.

```
which(f, types)
```

Returns the method of `f` (a `Method` object) that would be called for arguments of the given `types`.

If `types` is an abstract type, then the method that would be called by `invoke` is returned.

source

`Base.which` – Method.

```
which(symbol)
```

Return the module in which the binding for the variable referenced by `symbol` was created.

source

`Base.@which` – Macro.

```
@which
```

Applied to a function or macro call, it evaluates the arguments to the specified call, and returns the `Method` object for the method that would be called for those arguments. Applied to a variable, it returns the module in which the variable was bound. It calls out to the `which` function.

   source

`Base.methods` – Function.

```
methods(f, [types])
```

Returns the method table for `f`.

If `types` is specified, returns an array of methods whose types match.

   source

`Base.methodswith` – Function.

```
methodswith(typ[, module or function][, showparents::Bool=false])
```

Return an array of methods with an argument of type `typ`.

The optional second argument restricts the search to a particular module or function (the default is all modules, starting from Main).

If optional `showparents` is `true`, also return arguments with a parent type of `typ`, excluding type Any.

   source

`Base.@show` – Macro.

```
@show
```

Show an expression and result, returning the result.

   source

`Base.versioninfo` – Function.

```
versioninfo(io::IO=STDOUT, verbose::Bool=false)
```

Print information about the version of Julia in use. If the `verbose` argument is `true`, detailed system information is shown as well.

   source

`Base.workspace` – Function.

```
workspace()
```

Replace the top-level module (`Main`) with a new one, providing a clean workspace. The previous `Main` module is made available as `LastMain`. A previously-loaded package can be accessed using a statement such as `using LastMain.Package`.

This function should only be used interactively.

   source

`ans` – Keyword.

```
ans
```

A variable referring to the last computed value, automatically set at the interactive prompt.

   source

45.3   All Objects

Core.:=== – Function.

```
===(x,y) -> Bool
(x,y) -> Bool
```

Determine whether x and y are identical, in the sense that no program could distinguish them. Compares mutable objects by address in memory, and compares immutable objects (such as numbers) by contents at the bit level. This function is sometimes called egal.

```
julia> a = [1, 2]; b = [1, 2];

julia> a == b
true

julia> a === b
false

julia> a === a
true
```

source

Core.isa – Function.

```
isa(x, type) -> Bool
```

Determine whether x is of the given type. Can also be used as an infix operator, e.g. x isa type.

source

Base.isequal – Method.

```
isequal(x, y)
```

Similar to ==, except treats all floating-point NaN values as equal to each other, and treats -0.0 as unequal to 0.0. The default implementation of isequal calls ==, so if you have a type that doesn't have these floating-point subtleties then you probably only need to define ==.

isequal is the comparison function used by hash tables (Dict). isequal(x,y) must imply that hash(x) == hash(y).

This typically means that if you define your own == function then you must define a corresponding hash (and vice versa). Collections typically implement isequal by calling isequal recursively on all contents.

Scalar types generally do not need to implement isequal separate from ==, unless they represent floating-point numbers amenable to a more efficient implementation than that provided as a generic fallback (based on isnan, signbit, and ==).

```
julia> isequal([1., NaN], [1., NaN])
true

julia> [1., NaN] == [1., NaN]
false

julia> 0.0 == -0.0
```

```
true

julia> isequal(0.0, -0.0)
false
```

source

`Base.isequal` – Method.

```
isequal(x, y)
```

Similar to ==, except treats all floating-point NaN values as equal to each other, and treats -0.0 as unequal to 0.0. The default implementation of `isequal` calls ==, so if you have a type that doesn't have these floating-point subtleties then you probably only need to define ==.

`isequal` is the comparison function used by hash tables (`Dict`). `isequal(x,y)` must imply that `hash(x) == hash(y)`.

This typically means that if you define your own == function then you must define a corresponding `hash` (and vice versa). Collections typically implement `isequal` by calling `isequal` recursively on all contents.

Scalar types generally do not need to implement `isequal` separate from ==, unless they represent floating-point numbers amenable to a more efficient implementation than that provided as a generic fallback (based on `isnan`, `signbit`, and ==).

```
julia> isequal([1., NaN], [1., NaN])
true

julia> [1., NaN] == [1., NaN]
false

julia> 0.0 == -0.0
true

julia> isequal(0.0, -0.0)
false
```

source

```
isequal(x::Nullable, y::Nullable)
```

If neither x nor y is null, compare them according to their values (i.e. `isequal(get(x), get(y))`). Else, return `true` if both arguments are null, and `false` if one is null but not the other: nulls are considered equal.

source

`Base.isless` – Function.

```
isless(x, y)
```

Test whether x is less than y, according to a canonical total order. Values that are normally unordered, such as NaN, are ordered in an arbitrary but consistent fashion. This is the default comparison used by `sort`. Non-numeric types with a canonical total order should implement this function. Numeric types only need to implement it if they have special values such as NaN.

source

`Base.isless` – Method.

```
isless(x::Nullable, y::Nullable)
```

If neither x nor y is null, compare them according to their values (i.e. `isless(get(x), get(y))`). Else, return `true` if only y is null, and `false` otherwise: nulls are always considered greater than non-nulls, but not greater than another null.

source

Base.ifelse – Function.

```
ifelse(condition::Bool, x, y)
```

Return x if `condition` is `true`, otherwise return y. This differs from `?` or `if` in that it is an ordinary function, so all the arguments are evaluated first. In some cases, using `ifelse` instead of an `if` statement can eliminate the branch in generated code and provide higher performance in tight loops.

```julia
julia> ifelse(1 > 2, 1, 2)
2
```

source

Base.lexcmp – Function.

```
lexcmp(x, y)
```

Compare x and y lexicographically and return -1, 0, or 1 depending on whether x is less than, equal to, or greater than y, respectively. This function should be defined for lexicographically comparable types, and `lexless` will call `lexcmp` by default.

```julia
julia> lexcmp("abc", "abd")
-1

julia> lexcmp("abc", "abc")
0
```

source

Base.lexless – Function.

```
lexless(x, y)
```

Determine whether x is lexicographically less than y.

```julia
julia> lexless("abc", "abd")
true
```

source

Core.typeof – Function.

```
typeof(x)
```

Get the concrete type of x.

source

Core.tuple – Function.

```
tuple(xs...)
```

Construct a tuple of the given objects.

```
julia> tuple(1, 'a', pi)
(1, 'a', π = 3.1415926535897...)
```

source

Base.ntuple – Function.

```
ntuple(f::Function, n::Integer)
```

Create a tuple of length n, computing each element as f(i), where i is the index of the element.

```
julia> ntuple(i -> 2*i, 4)
(2, 4, 6, 8)
```

source

Base.object_id – Function.

```
object_id(x)
```

Get a hash value for x based on object identity. object_id(x)==object_id(y) if x === y.

source

Base.hash – Function.

```
hash(x[, h::UInt])
```

Compute an integer hash code such that isequal(x,y) implies hash(x)==hash(y). The optional second argument h is a hash code to be mixed with the result.

New types should implement the 2-argument form, typically by calling the 2-argument hash method recursively in order to mix hashes of the contents with each other (and with h). Typically, any type that implements hash should also implement its own == (hence isequal) to guarantee the property mentioned above.

source

Base.finalizer – Function.

```
finalizer(x, f)
```

Register a function f(x) to be called when there are no program-accessible references to x. The type of x must be a mutable struct, otherwise the behavior of this function is unpredictable.

source

Base.finalize – Function.

```
finalize(x)
```

Immediately run finalizers registered for object x.

source

Base.copy – Function.

```
copy(x)
```

Create a shallow copy of x: the outer structure is copied, but not all internal values. For example, copying an array produces a new array with identically-same elements as the original.

source

Base.deepcopy – Function.

```
deepcopy(x)
```

Create a deep copy of x: everything is copied recursively, resulting in a fully independent object. For example, deep-copying an array produces a new array whose elements are deep copies of the original elements. Calling deepcopy on an object should generally have the same effect as serializing and then deserializing it.

As a special case, functions can only be actually deep-copied if they are anonymous, otherwise they are just copied. The difference is only relevant in the case of closures, i.e. functions which may contain hidden internal references.

While it isn't normally necessary, user-defined types can override the default deepcopy behavior by defining a specialized version of the function deepcopy_internal(x::T, dict::ObjectIdDict) (which shouldn't otherwise be used), where T is the type to be specialized for, and dict keeps track of objects copied so far within the recursion. Within the definition, deepcopy_internal should be used in place of deepcopy, and the dict variable should be updated as appropriate before returning.

source

Core.isdefined – Function.

```
isdefined([m::Module,] s::Symbol)
isdefined(object, s::Symbol)
isdefined(object, index::Int)
```

Tests whether an assignable location is defined. The arguments can be a module and a symbol or a composite object and field name (as a symbol) or index. With a single symbol argument, tests whether a global variable with that name is defined in current_module().

source

Base.convert – Function.

```
convert(T, x)
```

Convert x to a value of type T.

If T is an Integer type, an InexactError will be raised if x is not representable by T, for example if x is not integer-valued, or is outside the range supported by T.

```
julia> convert(Int, 3.0)
3

julia> convert(Int, 3.5)
ERROR: InexactError()
Stacktrace:
 [1] convert(::Type{Int64}, ::Float64) at ./float.jl:679
```

If T is a AbstractFloat or Rational type, then it will return the closest value to x representable by T.

```julia
julia> x = 1/3
0.3333333333333333

julia> convert(Float32, x)
0.33333334f0

julia> convert(Rational{Int32}, x)
1//3

julia> convert(Rational{Int64}, x)
6004799503160661//18014398509481984
```

If T is a collection type and x a collection, the result of convert(T, x) may alias x.

```julia
julia> x = Int[1,2,3];

julia> y = convert(Vector{Int}, x);

julia> y === x
true
```

Similarly, if T is a composite type and x a related instance, the result of convert(T, x) may alias part or all of x.

```julia
julia> x = speye(5);

julia> typeof(x)
SparseMatrixCSC{Float64,Int64}

julia> y = convert(SparseMatrixCSC{Float64,Int64}, x);

julia> z = convert(SparseMatrixCSC{Float32,Int64}, y);

julia> y === x
true

julia> z === x
false

julia> z.colptr === x.colptr
true
```

source

Base.promote – Function.

```julia
promote(xs...)
```

Convert all arguments to their common promotion type (if any), and return them all (as a tuple).

source

Base.oftype – Function.

```julia
oftype(x, y)
```

Convert y to the type of x (convert(typeof(x), y)).

source

Base.widen – Function.

```
widen(x)
```

If x is a type, return a "larger" type (for numeric types, this will be a type with at least as much range and precision as the argument, and usually more). Otherwise x is converted to `widen(typeof(x))`.

```
julia> widen(Int32)
Int64

julia> widen(1.5f0)
1.5
```

source

Base.identity – Function.

```
identity(x)
```

The identity function. Returns its argument.

```
julia> identity("Well, what did you expect?")
"Well, what did you expect?"
```

source

45.4   Types

Base.supertype – Function.

```
supertype(T::DataType)
```

Return the supertype of DataType T.

```
julia> supertype(Int32)
Signed
```

source

Core.issubtype – Function.

```
issubtype(type1, type2)
```

Return `true` if and only if all values of `type1` are also of `type2`. Can also be written using the `<:` infix operator as `type1 <: type2`.

```
julia> issubtype(Int8, Int32)
false

julia> Int8 <: Integer
true
```

source

Base.:<: – Function.

```
<:(T1, T2)
```

Subtype operator, equivalent to `issubtype(T1, T2)`.

```julia
julia> Float64 <: AbstractFloat
true

julia> Vector{Int} <: AbstractArray
true

julia> Matrix{Float64} <: Matrix{AbstractFloat}
false
```

source

Base.:>: – Function.

```
>:(T1, T2)
```

Supertype operator, equivalent to `issubtype(T2, T1)`.

source

Base.subtypes – Function.

```
subtypes(T::DataType)
```

Return a list of immediate subtypes of DataType T. Note that all currently loaded subtypes are included, including those not visible in the current module.

```julia
julia> subtypes(Integer)
4-element Array{Union{DataType, UnionAll},1}:
 BigInt
 Bool
 Signed
 Unsigned
```

source

Base.typemin – Function.

```
typemin(T)
```

The lowest value representable by the given (real) numeric DataType T.

```julia
julia> typemin(Float16)
-Inf16

julia> typemin(Float32)
-Inf32
```

source

Base.typemax – Function.

```
typemax(T)
```

The highest value representable by the given (real) numeric `DataType`.

source

`Base.realmin` – Function.

```
realmin(T)
```

The smallest in absolute value non-subnormal value representable by the given floating-point DataType T.

source

`Base.realmax` – Function.

```
realmax(T)
```

The highest finite value representable by the given floating-point DataType T.

```
julia> realmax(Float16)
Float16(6.55e4)

julia> realmax(Float32)
3.4028235f38
```

source

`Base.maxintfloat` – Function.

```
maxintfloat(T)
```

The largest integer losslessly representable by the given floating-point DataType T.

source

`Base.sizeof` – Method.

```
sizeof(T)
```

Size, in bytes, of the canonical binary representation of the given DataType T, if any.

```
julia> sizeof(Float32)
4

julia> sizeof(Complex128)
16
```

If T does not have a specific size, an error is thrown.

```
julia> sizeof(Base.LinAlg.LU)
ERROR: argument is an abstract type; size is indeterminate
Stacktrace:
 [1] sizeof(::Type{T} where T) at ./essentials.jl:159
```

source

`Base.eps` – Method.

```
eps(::Type{T}) where T<:AbstractFloat
eps()
```

Returns the *machine epsilon* of the floating point type T (T = Float64 by default). This is defined as the gap between 1 and the next largest value representable by T, and is equivalent to eps(one(T)).

```julia
julia> eps()
2.220446049250313e-16

julia> eps(Float32)
1.1920929f-7

julia> 1.0 + eps()
1.0000000000000002

julia> 1.0 + eps()/2
1.0
```

source

Base.eps – Method.

```julia
eps(x::AbstractFloat)
```

Returns the *unit in last place* (ulp) of x. This is the distance between consecutive representable floating point values at x. In most cases, if the distance on either side of x is different, then the larger of the two is taken, that is

```julia
eps(x) == max(x-prevfloat(x), nextfloat(x)-x)
```

The exceptions to this rule are the smallest and largest finite values (e.g. nextfloat(-Inf) and prevfloat(Inf) for Float64), which round to the smaller of the values.

The rationale for this behavior is that eps bounds the floating point rounding error. Under the default Round-Nearest rounding mode, if y is a real number and x is the nearest floating point number to y, then

$$|y - x| \leq \mathrm{eps}(x)/2.$$

```julia
julia> eps(1.0)
2.220446049250313e-16

julia> eps(prevfloat(2.0))
2.220446049250313e-16

julia> eps(2.0)
4.440892098500626e-16

julia> x = prevfloat(Inf)      # largest finite Float64
1.7976931348623157e308

julia> x + eps(x)/2            # rounds up
Inf

julia> x + prevfloat(eps(x)/2) # rounds down
1.7976931348623157e308
```

source

Base.promote_type – Function.

```
promote_type(type1, type2)
```

Determine a type big enough to hold values of each argument type without loss, whenever possible. In some cases, where no type exists to which both types can be promoted losslessly, some loss is tolerated; for example, `promote_type(Int64, Float64)` returns `Float64` even though strictly, not all `Int64` values can be represented exactly as `Float64` values.

```julia
julia> promote_type(Int64, Float64)
Float64

julia> promote_type(Int32, Int64)
Int64

julia> promote_type(Float32, BigInt)
BigFloat
```

source

`Base.promote_rule` – Function.

```
promote_rule(type1, type2)
```

Specifies what type should be used by `promote` when given values of types `type1` and `type2`. This function should not be called directly, but should have definitions added to it for new types as appropriate.

source

`Core.getfield` – Function.

```
getfield(value, name::Symbol)
```

Extract a named field from a `value` of composite type. The syntax `a.b` calls `getfield(a, :b)`.

```julia
julia> a = 1//2
1//2

julia> getfield(a, :num)
1
```

source

`Core.setfield!` – Function.

```
setfield!(value, name::Symbol, x)
```

Assign x to a named field in `value` of composite type. The syntax `a.b = c` calls `setfield!(a, :b, c)`.

source

`Base.fieldoffset` – Function.

```
fieldoffset(type, i)
```

The byte offset of field i of a type relative to the data start. For example, we could use it in the following manner to summarize information about a struct:

```julia
julia> structinfo(T) = [(fieldoffset(T,i), fieldname(T,i), fieldtype(T,i)) for i =
↪   1:nfields(T)];

julia> structinfo(Base.Filesystem.StatStruct)
12-element Array{Tuple{UInt64,Symbol,DataType},1}:
 (0x0000000000000000, :device, UInt64)
 (0x0000000000000008, :inode, UInt64)
 (0x0000000000000010, :mode, UInt64)
 (0x0000000000000018, :nlink, Int64)
 (0x0000000000000020, :uid, UInt64)
 (0x0000000000000028, :gid, UInt64)
 (0x0000000000000030, :rdev, UInt64)
 (0x0000000000000038, :size, Int64)
 (0x0000000000000040, :blksize, Int64)
 (0x0000000000000048, :blocks, Int64)
 (0x0000000000000050, :mtime, Float64)
 (0x0000000000000058, :ctime, Float64)
```

source

Core.fieldtype – Function.

```julia
fieldtype(T, name::Symbol | index::Int)
```

Determine the declared type of a field (specified by name or index) in a composite DataType T.

```julia
julia> struct Foo
           x::Int64
           y::String
       end

julia> fieldtype(Foo, :x)
Int64

julia> fieldtype(Foo, 2)
String
```

source

Base.isimmutable – Function.

```julia
isimmutable(v)
```

Return `true` iff value v is immutable. See Mutable Composite Types for a discussion of immutability. Note that this function works on values, so if you give it a type, it will tell you that a value of `DataType` is mutable.

```julia
julia> isimmutable(1)
true

julia> isimmutable([1,2])
false
```

source

Base.isbits – Function.

```julia
isbits(T)
```

Return `true` if T is a "plain data" type, meaning it is immutable and contains no references to other values. Typical examples are numeric types such as `UInt8`, `Float64`, and `Complex{Float64}`.

```julia
julia> isbits(Complex{Float64})
true

julia> isbits(Complex)
false
```

source

`Base.isleaftype` – Function.

```
isleaftype(T)
```

Determine whether T's only subtypes are itself and `Union{}`. This means T is a concrete type that can have instances.

```julia
julia> isleaftype(Complex)
false

julia> isleaftype(Complex{Float32})
true

julia> isleaftype(Vector{Complex})
true

julia> isleaftype(Vector{Complex{Float32}})
true
```

source

`Base.typejoin` – Function.

```
typejoin(T, S)
```

Compute a type that contains both T and S.

source

`Base.typeintersect` – Function.

```
typeintersect(T, S)
```

Compute a type that contains the intersection of T and S. Usually this will be the smallest such type or one close to it.

source

`Base.Val` – Type.

```
Val{c}
```

Create a "value type" out of c, which must be an `isbits` value. The intent of this construct is to be able to dispatch on constants, e.g., `f(Val{false})` allows you to dispatch directly (at compile-time) to an implementation `f(::Type{Val{false}})`, without having to test the boolean value at runtime.

source

Base.Enums.@enum – Macro.

```
@enum EnumName[::BaseType] EnumValue1[=x] EnumValue2[=y]
```

Create an Enum{BaseType} subtype with name EnumName and enum member values of EnumValue1 and Enum-Value2 with optional assigned values of x and y, respectively. EnumName can be used just like other types and enum member values as regular values, such as

```
julia> @enum Fruit apple=1 orange=2 kiwi=3

julia> f(x::Fruit) = "I'm a Fruit with value: $(Int(x))"
f (generic function with 1 method)

julia> f(apple)
"I'm a Fruit with value: 1"
```

BaseType, which defaults to Int32, must be a primitive subtype of Integer. Member values can be converted between the enum type and BaseType. read and write perform these conversions automatically.

source

Base.instances – Function.

```
instances(T::Type)
```

Return a collection of all instances of the given type, if applicable. Mostly used for enumerated types (see @enum).

```
julia> @enum Colors Red Blue Green

julia> instances(Colors)
(Red::Colors = 0, Blue::Colors = 1, Green::Colors = 2)
```

source

45.5   Generic Functions

Core.Function – Type.

```
Function
```

Abstract type of all functions.

```
julia> isa(+, Function)
true

julia> typeof(sin)
Base.#sin

julia> ans <: Function
true
```

source

Base.method_exists – Function.

```
method_exists(f, Tuple type, world=typemax(UInt)) -> Bool
```

Determine whether the given generic function has a method matching the given `Tuple` of argument types with the upper bound of world age given by `world`.

```julia
julia> method_exists(length, Tuple{Array})
true
```

source

Core.applicable – Function.

```julia
applicable(f, args...) -> Bool
```

Determine whether the given generic function has a method applicable to the given arguments.

```julia
julia> function f(x, y)
           x + y
       end;

julia> applicable(f, 1)
false

julia> applicable(f, 1, 2)
true
```

source

Core.invoke – Function.

```julia
invoke(f, types <: Tuple, args...)
```

Invoke a method for the given generic function matching the specified types, on the specified arguments. The arguments must be compatible with the specified types. This allows invoking a method other than the most specific matching method, which is useful when the behavior of a more general definition is explicitly needed (often as part of the implementation of a more specific method of the same function).

source

Base.invokelatest – Function.

```julia
invokelatest(f, args...)
```

Calls `f(args...)`, but guarantees that the most recent method of f will be executed. This is useful in specialized circumstances, e.g. long-running event loops or callback functions that may call obsolete versions of a function f. (The drawback is that `invokelatest` is somewhat slower than calling f directly, and the type of the result cannot be inferred by the compiler.)

source

Base.:|> – Function.

```julia
|>(x, f)
```

Applies a function to the preceding argument. This allows for easy function chaining.

```julia
julia> [1:5;] |> x->x.^2 |> sum |> inv
0.01818181818181818
```

Base.: – Function.

```
f  g
```

Compose functions: i.e. (f g)(args...) means f(g(args...)). The symbol can be entered in the Julia
REPL (and most editors, appropriately configured) by typing \circ<tab>. Example:

```
julia> map(uppercasehex, 250:255)
6-element Array{String,1}:
 "FA"
 "FB"
 "FC"
 "FD"
 "FE"
 "FF"
```

45.6   Syntax

Core.eval – Function.

```
eval([m::Module], expr::Expr)
```

Evaluate an expression in the given module and return the result. Every Module (except those defined with
baremodule) has its own 1-argument definition of eval, which evaluates expressions in that module.

Base.@eval – Macro.

```
@eval [mod,] ex
```

Evaluate an expression with values interpolated into it using eval. If two arguments are provided, the first is the
module to evaluate in.

Base.evalfile – Function.

```
evalfile(path::AbstractString, args::Vector{String}=String[])
```

Load the file using include, evaluate all expressions, and return the value of the last one.

Base.esc – Function.

```
esc(e::ANY)
```

Only valid in the context of an Expr returned from a macro. Prevents the macro hygiene pass from turning em-
bedded variables into gensym variables. See the Macros section of the Metaprogramming chapter of the manual
for more details and examples.

Base.@inbounds – Macro.

```
@inbounds(blk)
```

Eliminates array bounds checking within expressions.

In the example below the bound check of array A is skipped to improve performance.

```julia
function sum(A::AbstractArray)
    r = zero(eltype(A))
    for i = 1:length(A)
        @inbounds r += A[i]
    end
    return r
end
```

> **Warning**
>
> Using @inbounds may return incorrect results/crashes/corruption for out-of-bounds indices. The user is responsible for checking it manually.

source

Base.@inline – Macro.

```
@inline
```

Give a hint to the compiler that this function is worth inlining.

Small functions typically do not need the @inline annotation, as the compiler does it automatically. By using @inline on bigger functions, an extra nudge can be given to the compiler to inline it. This is shown in the following example:

```julia
@inline function bigfunction(x)
    #=
        Function Definition
    =#
end
```

source

Base.@noinline – Macro.

```
@noinline
```

Prevents the compiler from inlining a function.

Small functions are typically inlined automatically. By using @noinline on small functions, auto-inlining can be prevented. This is shown in the following example:

```julia
@noinline function smallfunction(x)
    #=
        Function Definition
    =#
end
```

source

`Base.gensym` – Function.

```
gensym([tag])
```

Generates a symbol which will not conflict with other variable names.

source

`Base.@gensym` – Macro.

```
@gensym
```

Generates a gensym symbol for a variable. For example, `@gensym  x  y` is transformed into `x = gensym("x");`
`y = gensym("y")`.

source

`Base.@polly` – Macro.

```
@polly
```

Tells the compiler to apply the polyhedral optimizer Polly to a function.

source

`Base.parse` – Method.

```
parse(str, start; greedy=true, raise=true)
```

Parse the expression string and return an expression (which could later be passed to eval for execution). `start`
is the index of the first character to start parsing. If `greedy` is `true` (default), `parse` will try to consume as
much input as it can; otherwise, it will stop as soon as it has parsed a valid expression. Incomplete but otherwise
syntactically valid expressions will return `Expr(:incomplete,  "(error message)")`. If `raise` is `true` (de-
fault), syntax errors other than incomplete expressions will raise an error. If `raise` is `false`, `parse` will return an
expression that will raise an error upon evaluation.

source

`Base.parse` – Method.

```
parse(str; raise=true)
```

Parse the expression string greedily, returning a single expression. An error is thrown if there are additional
characters after the first expression. If `raise` is `true` (default), syntax errors will raise an error; otherwise, `parse`
will return an expression that will raise an error upon evaluation.

source

45.7   Nullables

`Base.Nullable` – Type.

```
Nullable(x, hasvalue::Bool=true)
```

Wrap value x in an object of type `Nullable`, which indicates whether a value is present. `Nullable(x)` yields
a non-empty wrapper and `Nullable{T}()` yields an empty instance of a wrapper that might contain a value of
type T.

`Nullable(x, false)` yields `Nullable{typeof(x)}()` with x stored in the result's `value` field.

Examples

```julia
julia> Nullable(1)
Nullable{Int64}(1)

julia> Nullable{Int64}()
Nullable{Int64}()

julia> Nullable(1, false)
Nullable{Int64}()

julia> dump(Nullable(1, false))
Nullable{Int64}
  hasvalue: Bool false
  value: Int64 1
```

source

`Base.get` – Method.

```julia
get(x::Nullable[, y])
```

Attempt to access the value of x. Returns the value if it is present; otherwise, returns y if provided, or throws a `NullException` if not.

source

`Base.isnull` – Function.

```julia
isnull(x)
```

Return whether or not x is null for `Nullable` x; return `false` for all other x.

Examples

```julia
julia> x = Nullable(1, false)
Nullable{Int64}()

julia> isnull(x)
true

julia> x = Nullable(1, true)
Nullable{Int64}(1)

julia> isnull(x)
false

julia> x = 1
1

julia> isnull(x)
false
```

source

`Base.unsafe_get` – Function.

```julia
unsafe_get(x)
```

Return the value of x for `Nullable` x; return x for all other x.

This method does not check whether or not x is null before attempting to access the value of x for x::Nullable (hence "unsafe").

```julia
julia> x = Nullable(1)
Nullable{Int64}(1)

julia> unsafe_get(x)
1

julia> x = Nullable{String}()
Nullable{String}()

julia> unsafe_get(x)
ERROR: UndefRefError: access to undefined reference
Stacktrace:
 [1] unsafe_get(::Nullable{String}) at ./nullable.jl:125

julia> x = 1
1

julia> unsafe_get(x)
1
```

source

45.8   System

`Base.run` – Function.

```
run(command, args...)
```

Run a command object, constructed with backticks. Throws an error if anything goes wrong, including the process exiting with a non-zero status.

source

`Base.spawn` – Function.

```
spawn(command)
```

Run a command object asynchronously, returning the resulting `Process` object.

source

`Base.DevNull` – Constant.

```
DevNull
```

Used in a stream redirect to discard all data written to it. Essentially equivalent to /dev/null on Unix or NUL on Windows. Usage:

```
run(pipeline(`cat test.txt`, DevNull))
```

source

`Base.success` – Function.

```
success(command)
```

Run a command object, constructed with backticks, and tell whether it was successful (exited with a code of 0). An exception is raised if the process cannot be started.

source

`Base.process_running` – Function.

```
process_running(p::Process)
```

Determine whether a process is currently running.

source

`Base.process_exited` – Function.

```
process_exited(p::Process)
```

Determine whether a process has exited.

source

`Base.kill` – Method.

```
kill(p::Process, signum=SIGTERM)
```

Send a signal to a process. The default is to terminate the process.

source

`Base.Sys.set_process_title` – Function.

```
Sys.set_process_title(title::AbstractString)
```

Set the process title. No-op on some operating systems.

source

`Base.Sys.get_process_title` – Function.

```
Sys.get_process_title()
```

Get the process title. On some systems, will always return an empty string.

source

`Base.readandwrite` – Function.

```
readandwrite(command)
```

Starts running a command asynchronously, and returns a tuple (stdout,stdin,process) of the output stream and input stream of the process, and the process object itself.

source

`Base.ignorestatus` – Function.

```
ignorestatus(command)
```

Mark a command object so that running it will not throw an error if the result code is non-zero.

> source

Base.detach – Function.

```
detach(command)
```

Mark a command object so that it will be run in a new process group, allowing it to outlive the julia process, and not have Ctrl-C interrupts passed to it.

> source

Base.Cmd – Type.

```
Cmd(cmd::Cmd; ignorestatus, detach, windows_verbatim, windows_hide, env, dir)
```

Construct a new Cmd object, representing an external program and arguments, from cmd, while changing the settings of the optional keyword arguments:

- ignorestatus::Bool: If true (defaults to false), then the Cmd will not throw an error if the return code is nonzero.
- detach::Bool: If true (defaults to false), then the Cmd will be run in a new process group, allowing it to outlive the julia process and not have Ctrl-C passed to it.
- windows_verbatim::Bool: If true (defaults to false), then on Windows the Cmd will send a command-line string to the process with no quoting or escaping of arguments, even arguments containing spaces. (On Windows, arguments are sent to a program as a single "command-line" string, and programs are responsible for parsing it into arguments. By default, empty arguments and arguments with spaces or tabs are quoted with double quotes " in the command line, and \ or " are preceded by backslashes. windows_verbatim=true is useful for launching programs that parse their command line in nonstandard ways.) Has no effect on non-Windows systems.
- windows_hide::Bool: If true (defaults to false), then on Windows no new console window is displayed when the Cmd is executed. This has no effect if a console is already open or on non-Windows systems.
- env: Set environment variables to use when running the Cmd. env is either a dictionary mapping strings to strings, an array of strings of the form "var=val", an array or tuple of "var"=>val pairs, or nothing. In order to modify (rather than replace) the existing environment, create env by copy(ENV) and then set env["var"]=val as desired.
- dir::AbstractString: Specify a working directory for the command (instead of the current directory).

For any keywords that are not specified, the current settings from cmd are used. Normally, to create a Cmd object in the first place, one uses backticks, e.g.

```
Cmd(`echo "Hello world"`, ignorestatus=true, detach=false)
```

> source

Base.setenv – Function.

```
setenv(command::Cmd, env; dir="")
```

Set environment variables to use when running the given command. env is either a dictionary mapping strings to strings, an array of strings of the form "var=val", or zero or more "var"=>val pair arguments. In order to modify (rather than replace) the existing environment, create env by copy(ENV) and then setting env["var"]=val as desired, or use withenv.

The dir keyword argument can be used to specify a working directory for the command.

> source

Base.withenv – Function.

```
withenv(f::Function, kv::Pair...)
```

Execute `f()` in an environment that is temporarily modified (not replaced as in `setenv`) by zero or more `"var"=>val` arguments `kv`. `withenv` is generally used via the `withenv(kv...) do ... end` syntax. A value of `nothing` can be used to temporarily unset an environment variable (if it is set). When `withenv` returns, the original environment has been restored.

source

Base.pipeline – Method.

```
pipeline(from, to, ...)
```

Create a pipeline from a data source to a destination. The source and destination can be commands, I/O streams, strings, or results of other `pipeline` calls. At least one argument must be a command. Strings refer to filenames. When called with more than two arguments, they are chained together from left to right. For example `pipeline(a,b,c)` is equivalent to `pipeline(pipeline(a,b),c)`. This provides a more concise way to specify multi-stage pipelines.

Examples:

```
run(pipeline(`ls`, `grep xyz`))
run(pipeline(`ls`, "out.txt"))
run(pipeline("out.txt", `grep xyz`))
```

source

Base.pipeline – Method.

```
pipeline(command; stdin, stdout, stderr, append=false)
```

Redirect I/O to or from the given command. Keyword arguments specify which of the command's streams should be redirected. `append` controls whether file output appends to the file. This is a more general version of the 2-argument `pipeline` function. `pipeline(from, to)` is equivalent to `pipeline(from, stdout=to)` when `from` is a command, and to `pipeline(to, stdin=from)` when `from` is another kind of data source.

Examples:

```
run(pipeline(`dothings`, stdout="out.txt", stderr="errs.txt"))
run(pipeline(`update`, stdout="log.txt", append=true))
```

source

Base.Libc.gethostname – Function.

```
gethostname() -> AbstractString
```

Get the local machine's host name.

source

Base.getipaddr – Function.

```
getipaddr() -> IPAddr
```

Get the IP address of the local machine.

source

Base.Libc.getpid – Function.

```
getpid() -> Int32
```

Get Julia's process ID.

source

Base.Libc.time – Method.

```
time()
```

Get the system time in seconds since the epoch, with fairly high (typically, microsecond) resolution.

source

Base.time_ns – Function.

```
time_ns()
```

Get the time in nanoseconds. The time corresponding to 0 is undefined, and wraps every 5.8 years.

source

Base.tic – Function.

```
tic()
```

Set a timer to be read by the next call to toc or toq. The macro call @time expr can also be used to time evaluation.

```julia
julia> tic()
0x0000c45bc7abac95

julia> sleep(0.3)

julia> toc()
elapsed time: 0.302745944 seconds
0.302745944
```

source

Base.toc – Function.

```
toc()
```

Print and return the time elapsed since the last tic. The macro call @time expr can also be used to time evaluation.

```julia
julia> tic()
0x0000c45bc7abac95

julia> sleep(0.3)

julia> toc()
elapsed time: 0.302745944 seconds
0.302745944
```

source

`Base.toq` – Function.

```
toq()
```

Return, but do not print, the time elapsed since the last `tic`. The macro calls `@timed expr` and `@elapsed expr` also return evaluation time.

```julia
julia> tic()
0x0000c46477a9675d

julia> sleep(0.3)

julia> toq()
0.302251004
```

source

`Base.@time` – Macro.

```
@time
```

A macro to execute an expression, printing the time it took to execute, the number of allocations, and the total number of bytes its execution caused to be allocated, before returning the value of the expression.

See also `@timev`, `@timed`, `@elapsed`, and `@allocated`.

```julia
julia> @time rand(10^6);
  0.001525 seconds (7 allocations: 7.630 MiB)

julia> @time begin
           sleep(0.3)
           1+1
       end
  0.301395 seconds (8 allocations: 336 bytes)
```

source

`Base.@timev` – Macro.

```
@timev
```

This is a verbose version of the `@time` macro. It first prints the same information as `@time`, then any non-zero memory allocation counters, and then returns the value of the expression.

See also `@time`, `@timed`, `@elapsed`, and `@allocated`.

```julia
julia> @timev rand(10^6);
  0.001006 seconds (7 allocations: 7.630 MiB)
elapsed time (ns): 1005567
bytes allocated:   8000256
pool allocs:       6
malloc() calls:    1
```

source

Base.@timed – Macro.

> @timed

A macro to execute an expression, and return the value of the expression, elapsed time, total bytes allocated, garbage collection time, and an object with various memory allocation counters.

See also @time, @timev, @elapsed, and @allocated.

```
julia> val, t, bytes, gctime, memallocs = @timed rand(10^6);

julia> t
0.006634834

julia> bytes
8000256

julia> gctime
0.0055765

julia> fieldnames(typeof(memallocs))
9-element Array{Symbol,1}:
 :allocd
 :malloc
 :realloc
 :poolalloc
 :bigalloc
 :freecall
 :total_time
 :pause
 :full_sweep

julia> memallocs.total_time
5576500
```

source

Base.@elapsed – Macro.

> @elapsed

A macro to evaluate an expression, discarding the resulting value, instead returning the number of seconds it took to execute as a floating-point number.

See also @time, @timev, @timed, and @allocated.

```
julia> @elapsed sleep(0.3)
0.301391426
```

source

Base.@allocated – Macro.

> @allocated

A macro to evaluate an expression, discarding the resulting value, instead returning the total number of bytes allocated during evaluation of the expression. Note: the expression is evaluated inside a local function, instead of

the current context, in order to eliminate the effects of compilation, however, there still may be some allocations due to JIT compilation. This also makes the results inconsistent with the `@time` macros, which do not try to adjust for the effects of compilation.

See also `@time`, `@timev`, `@timed`, and `@elapsed`.

```julia
julia> @allocated rand(10^6)
8000080
```

source

`Base.EnvHash` – Type.

```
EnvHash() -> EnvHash
```

A singleton of this type provides a hash table interface to environment variables.

source

`Base.ENV` – Constant.

```
ENV
```

Reference to the singleton EnvHash, providing a dictionary interface to system environment variables.

source

`Base.is_unix` – Function.

```
is_unix([os])
```

Predicate for testing if the OS provides a Unix-like interface. See documentation in Handling Operating System Variation.

source

`Base.is_apple` – Function.

```
is_apple([os])
```

Predicate for testing if the OS is a derivative of Apple Macintosh OS X or Darwin. See documentation in Handling Operating System Variation.

source

`Base.is_linux` – Function.

```
is_linux([os])
```

Predicate for testing if the OS is a derivative of Linux. See documentation in Handling Operating System Variation.

source

`Base.is_bsd` – Function.

```
is_bsd([os])
```

Predicate for testing if the OS is a derivative of BSD. See documentation in Handling Operating System Variation.

source

Base.is_windows – Function.

```
is_windows([os])
```

Predicate for testing if the OS is a derivative of Microsoft Windows NT. See documentation in Handling Operating System Variation.

source

Base.Sys.windows_version – Function.

```
Sys.windows_version()
```

Returns the version number for the Windows NT Kernel as a (major, minor) pair, or (0, 0) if this is not running on Windows.

source

Base.@static – Macro.

```
@static
```

Partially evaluates an expression at parse time.

For example, `@static is_windows() ? foo : bar` will evaluate `is_windows()` and insert either `foo` or `bar` into the expression. This is useful in cases where a construct would be invalid on other platforms, such as a `ccall` to a non-existent function. `@static if is_apple() foo end` and `@static foo <&&,||> bar` are also valid syntax.

source

45.9 Errors

Base.error – Function.

```
error(message::AbstractString)
```

Raise an `ErrorException` with the given message.

source

Core.throw – Function.

```
throw(e)
```

Throw an object as an exception.

source

Base.rethrow – Function.

```
rethrow([e])
```

Throw an object without changing the current exception backtrace. The default argument is the current exception (if called within a `catch` block).

source

Base.backtrace – Function.

```
backtrace()
```

Get a backtrace object for the current program point.

source

Base.catch_backtrace – Function.

```
catch_backtrace()
```

Get the backtrace of the current exception, for use within `catch` blocks.

source

Base.assert – Function.

```
assert(cond)
```

Throw an AssertionError if cond is `false`. Also available as the macro `@assert expr`.

source

Base.@assert – Macro.

```
@assert cond [text]
```

Throw an AssertionError if cond is `false`. Preferred syntax for writing assertions. Message `text` is optionally displayed upon assertion failure.

source

Base.ArgumentError – Type.

```
ArgumentError(msg)
```

The parameters to a function call do not match a valid signature. Argument `msg` is a descriptive error string.

source

Base.AssertionError – Type.

```
AssertionError([msg])
```

The asserted condition did not evaluate to `true`. Optional argument `msg` is a descriptive error string.

source

Core.BoundsError – Type.

```
BoundsError([a],[i])
```

An indexing operation into an array, a, tried to access an out-of-bounds element, `i`.

source

Base.DimensionMismatch – Type.

```
DimensionMismatch([msg])
```

The objects called do not have matching dimensionality. Optional argument `msg` is a descriptive error string.

source

Core.DivideError – Type.

> DivideError()

Integer division was attempted with a denominator value of 0.

> source

Core.DomainError – Type.

> DomainError()

The arguments to a function or constructor are outside the valid domain.

> source

Base.EOFError – Type.

> EOFError()

No more data was available to read from a file or stream.

> source

Core.ErrorException – Type.

> ErrorException(msg)

Generic error type. The error message, in the `.msg` field, may provide more specific details.

> source

Core.InexactError – Type.

> InexactError()

Type conversion cannot be done exactly.

> source

Core.InterruptException – Type.

> InterruptException()

The process was stopped by a terminal interrupt (CTRL+C).

> source

Base.KeyError – Type.

> KeyError(key)

An indexing operation into an `Associative` (`Dict`) or `Set` like object tried to access or delete a non-existent element.

> source

Base.LoadError – Type.

> LoadError(file::AbstractString, line::Int, error)

An error occurred while `includeing`, `requireing`, or `using` a file. The error specifics should be available in the `.error` field.

source

Base.MethodError – Type.

> MethodError(f, args)

A method with the required type signature does not exist in the given generic function. Alternatively, there is no unique most-specific method.

source

Base.NullException – Type.

> NullException()

An attempted access to a Nullable with no defined value.

source

Core.OutOfMemoryError – Type.

> OutOfMemoryError()

An operation allocated too much memory for either the system or the garbage collector to handle properly.

source

Core.ReadOnlyMemoryError – Type.

> ReadOnlyMemoryError()

An operation tried to write to memory that is read-only.

source

Core.OverflowError – Type.

> OverflowError()

The result of an expression is too large for the specified type and will cause a wraparound.

source

Base.ParseError – Type.

> ParseError(msg)

The expression passed to the `parse` function could not be interpreted as a valid Julia expression.

source

Base.Distributed.ProcessExitedException – Type.

> ProcessExitedException()

After a client Julia process has exited, further attempts to reference the dead child will throw this exception.

source

`Core.StackOverflowError` – Type.

> StackOverflowError()

The function call grew beyond the size of the call stack. This usually happens when a call recurses infinitely.

> source

`Base.SystemError` – Type.

> SystemError(prefix::AbstractString, [errno::Int32])

A system call failed with an error code (in the `errno` global variable).

> source

`Core.TypeError` – Type.

> TypeError(func::Symbol, context::AbstractString, expected::Type, got)

A type assertion failure, or calling an intrinsic function with an incorrect argument type.

> source

`Core.UndefRefError` – Type.

> UndefRefError()

The item or field is not defined for the given object.

> source

`Core.UndefVarError` – Type.

> UndefVarError(var::Symbol)

A symbol in the current scope is not defined.

> source

`Base.InitError` – Type.

> InitError(mod::Symbol, error)

An error occurred when running a module's `__init__` function. The actual error thrown is available in the `.error` field.

> source

`Base.retry` – Function.

> retry(f::Function; delays=ExponentialBackOff(), check=nothing) -> Function

Returns an anonymous function that calls function `f`. If an exception arises, `f` is repeatedly called again, each time `check` returns `true`, after waiting the number of seconds specified in `delays`. `check` should input `delays`'s current state and the `Exception`.

Examples

```
retry(f, delays=fill(5.0, 3))
retry(f, delays=rand(5:10, 2))
retry(f, delays=Base.ExponentialBackOff(n=3, first_delay=5, max_delay=1000))
retry(http_get, check=(s,e)->e.status == "503")(url)
retry(read, check=(s,e)->isa(e, UVError))(io, 128; all=false)
```

source

Base.ExponentialBackOff – Type.

```
ExponentialBackOff(; n=1, first_delay=0.05, max_delay=10.0, factor=5.0, jitter=0.1)
```

A Float64 iterator of length n whose elements exponentially increase at a rate in the interval factor * (1 ± jitter). The first element is first_delay and all elements are clamped to max_delay.

source

45.10 Events

Base.Timer – Method.

```
Timer(callback::Function, delay, repeat=0)
```

Create a timer to call the given callback function. The callback is passed one argument, the timer object itself. The callback will be invoked after the specified initial delay, and then repeating with the given repeat interval. If repeat is 0, the timer is only triggered once. Times are in seconds. A timer is stopped and has its resources freed by calling close on it.

source

Base.Timer – Type.

```
Timer(delay, repeat=0)
```

Create a timer that wakes up tasks waiting for it (by calling wait on the timer object) at a specified interval. Times are in seconds. Waiting tasks are woken with an error when the timer is closed (by close. Use isopen to check whether a timer is still active.

source

Base.AsyncCondition – Type.

```
AsyncCondition()
```

Create a async condition that wakes up tasks waiting for it (by calling wait on the object) when notified from C by a call to uv_async_send. Waiting tasks are woken with an error when the object is closed (by close. Use isopen to check whether it is still active.

source

Base.AsyncCondition – Method.

```
AsyncCondition(callback::Function)
```

Create a async condition that calls the given callback function. The callback is passed one argument, the async condition object itself.

source

45.11   Reflection

Base.module_name – Function.

```
module_name(m::Module) -> Symbol
```

Get the name of a `Module` as a `Symbol`.

```
julia> module_name(Base.LinAlg)
:LinAlg
```

source

Base.module_parent – Function.

```
module_parent(m::Module) -> Module
```

Get a module's enclosing `Module`. `Main` is its own parent, as is `LastMain` after `workspace()`.

```
julia> module_parent(Main)
Main

julia> module_parent(Base.LinAlg.BLAS)
Base.LinAlg
```

source

Base.current_module – Function.

```
current_module() -> Module
```

Get the *dynamically* current `Module`, which is the `Module` code is currently being read from. In general, this is not the same as the module containing the call to this function.

source

Base.fullname – Function.

```
fullname(m::Module)
```

Get the fully-qualified name of a module as a tuple of symbols. For example,

```
julia> fullname(Base.Pkg)
(:Base, :Pkg)

julia> fullname(Main)
()
```

source

Base.names – Function.

```
names(x::Module, all::Bool=false, imported::Bool=false)
```

Get an array of the names exported by a `Module`, excluding deprecated names. If `all` is true, then the list also includes non-exported names defined in the module, deprecated names, and compiler-generated names. If `imported` is true, then names explicitly imported from other modules are also included.

As a special case, all names defined in `Main` are considered "exported", since it is not idiomatic to explicitly export names from `Main`.

source

`Core.nfields` – Function.

```
nfields(x::DataType) -> Int
```

Get the number of fields of a `DataType`.

source

`Base.fieldnames` – Function.

```
fieldnames(x::DataType)
```

Get an array of the fields of a `DataType`.

```julia
julia> fieldnames(Hermitian)
2-element Array{Symbol,1}:
 :data
 :uplo
```

source

`Base.fieldname` – Function.

```
fieldname(x::DataType, i::Integer)
```

Get the name of field `i` of a `DataType`.

```julia
julia> fieldname(SparseMatrixCSC,1)
 :m

julia> fieldname(SparseMatrixCSC,5)
 :nzval
```

source

`Base.datatype_module` – Function.

```
Base.datatype_module(t::DataType) -> Module
```

Determine the module containing the definition of a `DataType`.

source

`Base.datatype_name` – Function.

```
Base.datatype_name(t) -> Symbol
```

Get the name of a (potentially UnionAll-wrapped) `DataType` (without its parent module) as a symbol.

source

Base.isconst – Function.

> isconst([m::Module], s::Symbol) -> Bool

Determine whether a global is declared const in a given Module. The default Module argument is current_module().

source

Base.function_name – Function.

> Base.function_name(f::Function) -> Symbol

Get the name of a generic Function as a symbol, or :anonymous.

source

Base.function_module – Method.

> Base.function_module(f::Function) -> Module

Determine the module containing the (first) definition of a generic function.

source

Base.function_module – Method.

> Base.function_module(f::Function, types) -> Module

Determine the module containing a given definition of a generic function.

source

Base.functionloc – Method.

> functionloc(f::Function, types)

Returns a tuple (filename,line) giving the location of a generic Function definition.

source

Base.functionloc – Method.

> functionloc(m::Method)

Returns a tuple (filename,line) giving the location of a Method definition.

source

Base.@functionloc – Macro.

> @functionloc

Applied to a function or macro call, it evaluates the arguments to the specified call, and returns a tuple (filename,line) giving the location for the method that would be called for those arguments. It calls out to the functionloc function.

source

45.12   Internals

Base.gc – Function.

```
gc()
```

Perform garbage collection. This should not generally be used.

source

Base.gc_enable – Function.

```
gc_enable(on::Bool)
```

Control whether garbage collection is enabled using a boolean argument (`true` for enabled, `false` for disabled). Returns previous GC state. Disabling garbage collection should be used only with extreme caution, as it can cause memory use to grow without bound.

source

Base.macroexpand – Function.

```
macroexpand(x)
```

Takes the expression x and returns an equivalent expression with all macros removed (expanded).

source

Base.@macroexpand – Macro.

```
@macroexpand
```

Return equivalent expression with all macros removed (expanded).

There is a subtle difference between @macroexpand and macroexpand in that expansion takes place in different contexts. This is best seen in the following example:

```julia
julia> module M
           macro m()
               1
           end
           function f()
               (@macroexpand(@m), macroexpand(:(@m)))
           end
       end
M

julia> macro m()
           2
       end
@m (macro with 1 method)

julia> M.f()
(1, 2)
```

With @macroexpand the expression expands where @macroexpand appears in the code (module M in the example). With macroexpand the expression expands in the current module where the code was finally called (REPL in the example). Note that when calling macroexpand or @macroexpand directly from the REPL, both of these contexts coincide, hence there is no difference.

source

`Base.expand` – Function.

```
expand(x)
```

Takes the expression x and returns an equivalent expression in lowered form. See also `code_lowered`.

source

`Base.code_lowered` – Function.

```
code_lowered(f, types)
```

Returns an array of lowered ASTs for the methods matching the given generic function and type signature.

source

`Base.@code_lowered` – Macro.

```
@code_lowered
```

Evaluates the arguments to the function or macro call, determines their types, and calls `code_lowered` on the resulting expression.

source

`Base.code_typed` – Function.

```
code_typed(f, types; optimize=true)
```

Returns an array of lowered and type-inferred ASTs for the methods matching the given generic function and type signature. The keyword argument `optimize` controls whether additional optimizations, such as inlining, are also applied.

source

`Base.@code_typed` – Macro.

```
@code_typed
```

Evaluates the arguments to the function or macro call, determines their types, and calls `code_typed` on the resulting expression.

source

`Base.code_warntype` – Function.

```
code_warntype([io::IO], f, types)
```

Prints lowered and type-inferred ASTs for the methods matching the given generic function and type signature to `io` which defaults to STDOUT. The ASTs are annotated in such a way as to cause "non-leaf" types to be emphasized (if color is available, displayed in red). This serves as a warning of potential type instability. Not all non-leaf types are particularly problematic for performance, so the results need to be used judiciously. See `@code_warntype` for more information.

source

`Base.@code_warntype` – Macro.

```
@code_warntype
```

Evaluates the arguments to the function or macro call, determines their types, and calls `code_warntype` on the resulting expression.

source

`Base.code_llvm` – Function.

```
code_llvm([io], f, types)
```

Prints the LLVM bitcodes generated for running the method matching the given generic function and type signature to `io` which defaults to STDOUT.

All metadata and dbg.* calls are removed from the printed bitcode. Use code_llvm_raw for the full IR.

source

`Base.@code_llvm` – Macro.

```
@code_llvm
```

Evaluates the arguments to the function or macro call, determines their types, and calls `code_llvm` on the resulting expression.

source

`Base.code_native` – Function.

```
code_native([io], f, types, [syntax])
```

Prints the native assembly instructions generated for running the method matching the given generic function and type signature to `io` which defaults to STDOUT. Switch assembly syntax using `syntax` symbol parameter set to `:att` for AT&T syntax or `:intel` for Intel syntax. Output is AT&T syntax by default.

source

`Base.@code_native` – Macro.

```
@code_native
```

Evaluates the arguments to the function or macro call, determines their types, and calls `code_native` on the resulting expression.

source

`Base.precompile` – Function.

```
precompile(f,args::Tuple{Vararg{Any}})
```

Compile the given function `f` for the argument tuple (of types) `args`, but do not execute it.

source

Chapter 46

Collections and Data Structures

46.1 Iteration

Sequential iteration is implemented by the methods `start()`, `done()`, and `next()`. The general `for` loop:

```
for i = I   # or  "for i in I"
    # body
end
```

is translated into:

```
state = start(I)
while !done(I, state)
    (i, state) = next(I, state)
    # body
end
```

The `state` object may be anything, and should be chosen appropriately for each iterable type. See the manual section on the iteration interface for more details about defining a custom iterable type.

`Base.start` – Function.

```
start(iter) -> state
```

Get initial iteration state for an iterable object.

```
julia> start(1:5)
1

julia> start([1;2;3])
1

julia> start([4;2;3])
1
```

source

`Base.done` – Function.

```
done(iter, state) -> Bool
```

Test whether we are done iterating.

```julia
julia> done(1:5, 3)
false

julia> done(1:5, 5)
false

julia> done(1:5, 6)
true
```

source

`Base.next` – Function.

```
next(iter, state) -> item, state
```

For a given iterable object and iteration state, return the current item and the next iteration state.

```julia
julia> next(1:5, 3)
(3, 4)

julia> next(1:5, 5)
(5, 6)
```

source

`Base.iteratorsize` – Function.

```
iteratorsize(itertype::Type) -> IteratorSize
```

Given the type of an iterator, returns one of the following values:

- `SizeUnknown()` if the length (number of elements) cannot be determined in advance.
- `HasLength()` if there is a fixed, finite length.
- `HasShape()` if there is a known length plus a notion of multidimensional shape (as for an array). In this case the `size` function is valid for the iterator.
- `IsInfinite()` if the iterator yields values forever.

The default value (for iterators that do not define this function) is `HasLength()`. This means that most iterators are assumed to implement `length`.

This trait is generally used to select between algorithms that pre-allocate space for their result, and algorithms that resize their result incrementally.

```julia
julia> Base.iteratorsize(1:5)
Base.HasShape()

julia> Base.iteratorsize((2,3))
Base.HasLength()
```

source

`Base.iteratoreltype` – Function.

```
iteratoreltype(itertype::Type) -> IteratorEltype
```

Given the type of an iterator, returns one of the following values:

- `EltypeUnknown()` if the type of elements yielded by the iterator is not known in advance.
- `HasEltype()` if the element type is known, and `eltype` would return a meaningful value.

`HasEltype()` is the default, since iterators are assumed to implement `eltype`.

This trait is generally used to select between algorithms that pre-allocate a specific type of result, and algorithms that pick a result type based on the types of yielded values.

```
julia> Base.iteratoreltype(1:5)
Base.HasEltype()
```

source

Fully implemented by:

- `Range`
- `UnitRange`
- `Tuple`
- `Number`
- `AbstractArray`
- `IntSet`
- `ObjectIdDict`
- `Dict`
- `WeakKeyDict`
- `EachLine`
- `AbstractString`
- `Set`

46.2   General Collections

`Base.isempty` – Function.

```
isempty(collection) -> Bool
```

Determine whether a collection is empty (has no elements).

```
julia> isempty([])
true

julia> isempty([1 2 3])
false
```

source

`Base.empty!` – Function.

```
empty!(collection) -> collection
```

Remove all elements from a `collection`.

```
julia> A = Dict("a" => 1, "b" => 2)
Dict{String,Int64} with 2 entries:
  "b" => 2
  "a" => 1

julia> empty!(A);

julia> A
Dict{String,Int64} with 0 entries
```

source

`Base.length` – Method.

```
length(collection) -> Integer
```

For ordered, indexable collections, returns the maximum index i for which `getindex(collection, i)` is valid. For unordered collections, returns the number of elements.

```
julia> length(1:5)
5

julia> length([1; 2; 3; 4])
4
```

source

`Base.endof` – Function.

```
endof(collection) -> Integer
```

Returns the last index of the collection.

```
julia> endof([1,2,4])
3
```

source

Fully implemented by:

- `Range`

- `UnitRange`

- `Tuple`

- `Number`

- `AbstractArray`

- IntSet

- ObjectIdDict

- Dict

- WeakKeyDict

- AbstractString

- Set

46.3   Iterable Collections

Base.in – Function.

```
in(item, collection) -> Bool
(item,collection) -> Bool
(collection,item) -> Bool
(item,collection) -> Bool
(collection,item) -> Bool
```

Determine whether an item is in the given collection, in the sense that it is == to one of the values generated by iterating over the collection. Some collections need a slightly different definition; for example Sets check whether the item isequal to one of the elements. Dicts look for (key, value) pairs, and the key is compared using isequal. To test for the presence of a key in a dictionary, use haskey or k in keys(dict).

```
julia> a = 1:3:20
1:3:19

julia> 4 in a
true

julia> 5 in a
false
```

source

Base.eltype – Function.

```
eltype(type)
```

Determine the type of the elements generated by iterating a collection of the given type. For associative collection types, this will be a Pair{KeyType,ValType}. The definition eltype(x) = eltype(typeof(x)) is provided for convenience so that instances can be passed instead of types. However the form that accepts a type argument should be defined for new types.

```
julia> eltype(ones(Float32,2,2))
Float32

julia> eltype(ones(Int8,2,2))
Int8
```

source

Base.indexin – Function.

```
indexin(a, b)
```

Returns a vector containing the highest index in b for each value in a that is a member of b . The output vector contains 0 wherever a is not a member of b.

```julia
julia> a = ['a', 'b', 'c', 'b', 'd', 'a'];

julia> b = ['a','b','c'];

julia> indexin(a,b)
6-element Array{Int64,1}:
 1
 2
 3
 2
 0
 1

julia> indexin(b,a)
3-element Array{Int64,1}:
 6
 4
 3
```

source

Base.findin – Function.

```
findin(a, b)
```

Returns the indices of elements in collection a that appear in collection b.

```julia
julia> a = collect(1:3:15)
5-element Array{Int64,1}:
  1
  4
  7
 10
 13

julia> b = collect(2:4:10)
3-element Array{Int64,1}:
  2
  6
 10

julia> findin(a,b) # 10 is the only common element
1-element Array{Int64,1}:
 4
```

source

Base.unique – Function.

```
unique(itr)
```

Returns an array containing one value from itr for each unique value, as determined by isequal.

```
julia> unique([1; 2; 2; 6])
3-element Array{Int64,1}:
 1
 2
 6
```

source

```
unique(f, itr)
```

Returns an array containing one value from `itr` for each unique value produced by `f` applied to elements of `itr`.

```
julia> unique(isodd, [1; 2; 2; 6])
2-element Array{Int64,1}:
 1
 2
```

source

```
unique(itr[, dim])
```

Returns an array containing only the unique elements of the iterable `itr`, in the order that the first of each set of equivalent elements originally appears. If `dim` is specified, returns unique regions of the array `itr` along `dim`.

```
julia> A = map(isodd, reshape(collect(1:8), (2,2,2)))
2×2×2 Array{Bool,3}:
[:, :, 1] =
  true   true
 false  false

[:, :, 2] =
  true   true
 false  false

julia> unique(A)
2-element Array{Bool,1}:
  true
 false

julia> unique(A, 2)
2×1×2 Array{Bool,3}:
[:, :, 1] =
  true
 false

[:, :, 2] =
  true
 false

julia> unique(A, 3)
2×2×1 Array{Bool,3}:
[:, :, 1] =
  true   true
 false  false
```

source

`Base.allunique` – Function.

```
allunique(itr) -> Bool
```

Return `true` if all values from `itr` are distinct when compared with `isequal`.

```
julia> a = [1; 2; 3]
3-element Array{Int64,1}:
 1
 2
 3

julia> allunique([a, a])
false
```

source

`Base.reduce` – Method.

```
reduce(op, v0, itr)
```

Reduce the given collection `itr` with the given binary operator `op`. `v0` must be a neutral element for `op` that will be returned for empty collections. It is unspecified whether `v0` is used for non-empty collections.

Reductions for certain commonly-used operators have special implementations which should be used instead: `maximum(itr)`, `minimum(itr)`, `sum(itr)`, `prod(itr)`, `any(itr)`, `all(itr)`.

The associativity of the reduction is implementation dependent. This means that you can't use non-associative operations like – because it is undefined whether `reduce(-,[1,2,3])` should be evaluated as `(1-2)-3` or `1-(2-3)`. Use `foldl` or `foldr` instead for guaranteed left or right associativity.

Some operations accumulate error, and parallelism will also be easier if the reduction can be executed in groups. Future versions of Julia might change the algorithm. Note that the elements are not reordered if you use an ordered collection.

Examples

```
julia> reduce(*, 1, [2; 3; 4])
24
```

source

`Base.reduce` – Method.

```
reduce(op, itr)
```

Like `reduce(op, v0, itr)`. This cannot be used with empty collections, except for some special cases (e.g. when op is one of `+`, `*`, `max`, `min`, `&`, `|`) when Julia can determine the neutral element of op.

```
julia> reduce(*, [2; 3; 4])
24
```

source

`Base.foldl` – Method.

```
foldl(op, v0, itr)
```

Like `reduce`, but with guaranteed left associativity. `v0` will be used exactly once.

```
julia> foldl(-, 1, 2:5)
-13
```

source

`Base.foldl` – Method.

```
foldl(op, itr)
```

Like `foldl(op, v0, itr)`, but using the first element of `itr` as `v0`. In general, this cannot be used with empty collections (see `reduce(op, itr)`).

```
julia> foldl(-, 2:5)
-10
```

source

`Base.foldr` – Method.

```
foldr(op, v0, itr)
```

Like `reduce`, but with guaranteed right associativity. `v0` will be used exactly once.

```
julia> foldr(-, 1, 2:5)
-1
```

source

`Base.foldr` – Method.

```
foldr(op, itr)
```

Like `foldr(op, v0, itr)`, but using the last element of `itr` as `v0`. In general, this cannot be used with empty collections (see `reduce(op, itr)`).

```
julia> foldr(-, 2:5)
-2
```

source

`Base.maximum` – Method.

```
maximum(itr)
```

Returns the largest element in a collection.

```
julia> maximum(-20.5:10)
9.5

julia> maximum([1,2,3])
3
```

source

`Base.maximum` – Method.

```
maximum(A, dims)
```

Compute the maximum value of an array over the given dimensions. See also the `max(a,b)` function to take the maximum of two or more arguments, which can be applied elementwise to arrays via `max.(a,b)`.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> maximum(A, 1)
1×2 Array{Int64,2}:
 3  4

julia> maximum(A, 2)
2×1 Array{Int64,2}:
 2
 4
```

source

`Base.maximum!` – Function.

```
maximum!(r, A)
```

Compute the maximum value of `A` over the singleton dimensions of `r`, and write results to `r`.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> maximum!([1; 1], A)
2-element Array{Int64,1}:
 2
 4

julia> maximum!([1 1], A)
1×2 Array{Int64,2}:
 3  4
```

source

`Base.minimum` – Method.

```
minimum(itr)
```

Returns the smallest element in a collection.

```
julia> minimum(-20.5:10)
-20.5

julia> minimum([1,2,3])
1
```

source

Base.minimum – Method.

```
minimum(A, dims)
```

Compute the minimum value of an array over the given dimensions. See also the `min(a,b)` function to take the minimum of two or more arguments, which can be applied elementwise to arrays via `min.(a,b)`.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> minimum(A, 1)
1×2 Array{Int64,2}:
 1  2

julia> minimum(A, 2)
2×1 Array{Int64,2}:
 1
 3
```

source

Base.minimum! – Function.

```
minimum!(r, A)
```

Compute the minimum value of A over the singleton dimensions of r, and write results to r.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> minimum!([1; 1], A)
2-element Array{Int64,1}:
 1
 3

julia> minimum!([1 1], A)
1×2 Array{Int64,2}:
 1  2
```

source

Base.extrema – Method.

```
extrema(itr) -> Tuple
```

Compute both the minimum and maximum element in a single pass, and return them as a 2-tuple.

```
julia> extrema(2:10)
(2, 10)

julia> extrema([9,pi,4.5])
(3.141592653589793, 9.0)
```

source

Base.extrema – Method.

```
extrema(A, dims) -> Array{Tuple}
```

Compute the minimum and maximum elements of an array over the given dimensions.

Example

```julia
julia> A = reshape(collect(1:2:16), (2,2,2))
2×2×2 Array{Int64,3}:
[:, :, 1] =
 1  5
 3  7

[:, :, 2] =
  9  13
 11  15

julia> extrema(A, (1,2))
1×1×2 Array{Tuple{Int64,Int64},3}:
[:, :, 1] =
 (1, 7)

[:, :, 2] =
 (9, 15)
```

source

Base.indmax – Function.

```
indmax(itr) -> Integer
```

Returns the index of the maximum element in a collection. If there are multiple maximal elements, then the first one will be returned. NaN values are ignored, unless all elements are NaN.

The collection must not be empty.

```julia
julia> indmax([8,0.1,-9,pi])
1

julia> indmax([1,7,7,6])
2

julia> indmax([1,7,7,NaN])
2
```

source

Base.indmin – Function.

```
indmin(itr) -> Integer
```

Returns the index of the minimum element in a collection. If there are multiple minimal elements, then the first one will be returned. NaN values are ignored, unless all elements are NaN.

The collection must not be empty.

```julia
julia> indmin([8,0.1,-9,pi])
3

julia> indmin([7,1,1,6])
2

julia> indmin([7,1,1,NaN])
2
```

source

`Base.findmax` – Method.

```
findmax(itr) -> (x, index)
```

Returns the maximum element of the collection `itr` and its index. If there are multiple maximal elements, then the first one will be returned. NaN values are ignored, unless all elements are NaN.

The collection must not be empty.

```julia
julia> findmax([8,0.1,-9,pi])
(8.0, 1)

julia> findmax([1,7,7,6])
(7, 2)

julia> findmax([1,7,7,NaN])
(7.0, 2)
```

source

`Base.findmax` – Method.

```
findmax(A, region) -> (maxval, index)
```

For an array input, returns the value and index of the maximum over the given region.

```julia
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> findmax(A,1)
([3 4], [2 4])

julia> findmax(A,2)
([2; 4], [3; 4])
```

source

`Base.findmin` – Method.

```
findmin(itr) -> (x, index)
```

Returns the minimum element of the collection `itr` and its index. If there are multiple minimal elements, then the first one will be returned. NaN values are ignored, unless all elements are NaN.

The collection must not be empty.

```
julia> findmin([8,0.1,-9,pi])
(-9.0, 3)

julia> findmin([7,1,1,6])
(1, 2)

julia> findmin([7,1,1,NaN])
(1.0, 2)
```

source

Base.findmin – Method.

```
findmin(A, region) -> (minval, index)
```

For an array input, returns the value and index of the minimum over the given region.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> findmin(A, 1)
([1 2], [1 3])

julia> findmin(A, 2)
([1; 3], [1; 2])
```

source

Base.findmax! – Function.

```
findmax!(rval, rind, A, [init=true]) -> (maxval, index)
```

Find the maximum of A and the corresponding linear index along singleton dimensions of rval and rind, and store the results in rval and rind.

source

Base.findmin! – Function.

```
findmin!(rval, rind, A, [init=true]) -> (minval, index)
```

Find the minimum of A and the corresponding linear index along singleton dimensions of rval and rind, and store the results in rval and rind.

source

Base.sum – Function.

```
sum(f, itr)
```

Sum the results of calling function f on each element of itr.

```
julia> sum(abs2, [2; 3; 4])
29
```

source

```
sum(itr)
```

Returns the sum of all elements in a collection.

```
julia> sum(1:20)
210
```

source

```
sum(A, dims)
```

Sum elements of an array over the given dimensions.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> sum(A, 1)
1×2 Array{Int64,2}:
 4  6

julia> sum(A, 2)
2×1 Array{Int64,2}:
 3
 7
```

source

Base.sum! – Function.

```
sum!(r, A)
```

Sum elements of A over the singleton dimensions of r, and write results to r.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> sum!([1; 1], A)
2-element Array{Int64,1}:
 3
 7

julia> sum!([1 1], A)
1×2 Array{Int64,2}:
 4  6
```

source

Base.prod – Function.

```
prod(f, itr)
```

Returns the product of f applied to each element of itr.

```
julia> prod(abs2, [2; 3; 4])
576
```

source

```
prod(itr)
```

Returns the product of all elements of a collection.

```
julia> prod(1:20)
2432902008176640000
```

source

```
prod(A, dims)
```

Multiply elements of an array over the given dimensions.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> prod(A, 1)
1×2 Array{Int64,2}:
 3  8

julia> prod(A, 2)
2×1 Array{Int64,2}:
  2
 12
```

source

Base.prod! – Function.

```
prod!(r, A)
```

Multiply elements of A over the singleton dimensions of r, and write results to r.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> prod!([1; 1], A)
2-element Array{Int64,1}:
  2
 12

julia> prod!([1 1], A)
1×2 Array{Int64,2}:
 3  8
```

source

Base.any – Method.

```
any(itr) -> Bool
```

Test whether any elements of a boolean collection are `true`, returning `true` as soon as the first `true` value in `itr` is encountered (short-circuiting).

```julia
julia> a = [true, false, false, true]
4-element Array{Bool,1}:
  true
 false
 false
  true

julia> any(a)
true

julia> any((println(i); v) for (i, v) in enumerate(a))
1
true
```

source

`Base.any` – Method.

```
any(A, dims)
```

Test whether any values along the given dimensions of an array are `true`.

```julia
julia> A = [true false; true false]
2×2 Array{Bool,2}:
 true  false
 true  false

julia> any(A, 1)
1×2 Array{Bool,2}:
 true  false

julia> any(A, 2)
2×1 Array{Bool,2}:
 true
 true
```

source

`Base.any!` – Function.

```
any!(r, A)
```

Test whether any values in A along the singleton dimensions of r are `true`, and write results to r.

```julia
julia> A = [true false; true false]
2×2 Array{Bool,2}:
 true  false
 true  false

julia> any!([1; 1], A)
2-element Array{Int64,1}:
```

```
  1
  1

julia> any!([1 1], A)
1×2 Array{Int64,2}:
  1  0
```

source

Base.all – Method.

```
all(itr) -> Bool
```

Test whether all elements of a boolean collection are `true`, returning `false` as soon as the first `false` value in `itr` is encountered (short-circuiting).

```
julia> a = [true,false,false,true]
4-element Array{Bool,1}:
  true
 false
 false
  true

julia> all(a)
false

julia> all((println(i); v) for (i, v) in enumerate(a))
1
2
false
```

source

Base.all – Method.

```
all(A, dims)
```

Test whether all values along the given dimensions of an array are `true`.

```
julia> A = [true false; true true]
2×2 Array{Bool,2}:
 true  false
 true   true

julia> all(A, 1)
1×2 Array{Bool,2}:
 true  false

julia> all(A, 2)
2×1 Array{Bool,2}:
 false
  true
```

source

Base.all! – Function.

```
all!(r, A)
```

Test whether all values in A along the singleton dimensions of r are true, and write results to r.

```julia
julia> A = [true false; true false]
2×2 Array{Bool,2}:
 true  false
 true  false

julia> all!([1; 1], A)
2-element Array{Int64,1}:
 0
 0

julia> all!([1 1], A)
1×2 Array{Int64,2}:
 1  0
```

source

Base.count – Function.

```
count(p, itr) -> Integer
count(itr) -> Integer
```

Count the number of elements in itr for which predicate p returns true. If p is omitted, counts the number of true elements in itr (which should be a collection of boolean values).

```julia
julia> count(i->(4<=i<=6), [2,3,4,5,6])
3

julia> count([true, false, true, true])
3
```

source

Base.any – Method.

```
any(p, itr) -> Bool
```

Determine whether predicate p returns true for any elements of itr, returning true as soon as the first item in itr for which p returns true is encountered (short-circuiting).

```julia
julia> any(i->(4<=i<=6), [3,5,7])
true

julia> any(i -> (println(i); i > 3), 1:10)
1
2
3
4
true
```

source

Base.all – Method.

```
all(p, itr) -> Bool
```

Determine whether predicate p returns `true` for all elements of `itr`, returning `false` as soon as the first item in `itr` for which p returns `false` is encountered (short-circuiting).

```julia
julia> all(i->(4<=i<=6), [4,5,6])
true

julia> all(i -> (println(i); i < 3), 1:10)
1
2
3
false
```

source

Base.`foreach` – Function.

```
foreach(f, c...) -> Void
```

Call function f on each element of iterable c. For multiple iterable arguments, f is called elementwise. `foreach` should be used instead of `map` when the results of f are not needed, for example in `foreach(println, array)`.

```julia
julia> a = 1:3:7;

julia> foreach(x -> println(x^2), a)
1
16
49
```

source

Base.`map` – Function.

```
map(f, c...) -> collection
```

Transform collection c by applying f to each element. For multiple collection arguments, apply f elementwise.

```julia
julia> map(x -> x * 2, [1, 2, 3])
3-element Array{Int64,1}:
 2
 4
 6

julia> map(+, [1, 2, 3], [10, 20, 30])
3-element Array{Int64,1}:
 11
 22
 33
```

source

```
map(f, x::Nullable)
```

Return `f` applied to the value of `x` if it has one, as a `Nullable`. If `x` is null, then return a null value of type `Nullable{S}`. S is guaranteed to be either `Union{}` or a concrete type. Whichever of these is chosen is an implementation detail, but typically the choice that maximizes performance would be used. If `x` has a value, then the return type is guaranteed to be of type `Nullable{typeof(f(x))}`.

source

`Base.map!` – Function.

```
map!(function, destination, collection...)
```

Like map, but stores the result in `destination` rather than a new collection. `destination` must be at least as large as the first collection.

```
julia> x = zeros(3);

julia> map!(x -> x * 2, x, [1, 2, 3]);

julia> x
3-element Array{Float64,1}:
 2.0
 4.0
 6.0
```

source

`Base.mapreduce` – Method.

```
mapreduce(f, op, v0, itr)
```

Apply function `f` to each element in `itr`, and then reduce the result using the binary function `op`. `v0` must be a neutral element for `op` that will be returned for empty collections. It is unspecified whether `v0` is used for non-empty collections.

mapreduce is functionally equivalent to calling `reduce(op, v0, map(f, itr))`, but will in general execute faster since no intermediate collection needs to be created. See documentation for reduce and map.

```
julia> mapreduce(x->x^2, +, [1:3;]) # == 1 + 4 + 9
14
```

The associativity of the reduction is implementation-dependent. Additionally, some implementations may reuse the return value of `f` for elements that appear multiple times in `itr`. Use mapfoldl or mapfoldr instead for guaranteed left or right associativity and invocation of `f` for every value.

source

`Base.mapreduce` – Method.

```
mapreduce(f, op, itr)
```

Like `mapreduce(f, op, v0, itr)`. In general, this cannot be used with empty collections (see `reduce(op, itr)`).

source

`Base.mapfoldl` – Method.

```
mapfoldl(f, op, v0, itr)
```

Like `mapreduce`, but with guaranteed left associativity, as in `foldl`. `v0` will be used exactly once.

> source

`Base.mapfoldl` – Method.

```
mapfoldl(f, op, itr)
```

Like `mapfoldl(f, op, v0, itr)`, but using the first element of `itr` as `v0`. In general, this cannot be used with empty collections (see `reduce(op, itr)`).

> source

`Base.mapfoldr` – Method.

```
mapfoldr(f, op, v0, itr)
```

Like `mapreduce`, but with guaranteed right associativity, as in `foldr`. `v0` will be used exactly once.

> source

`Base.mapfoldr` – Method.

```
mapfoldr(f, op, itr)
```

Like `mapfoldr(f, op, v0, itr)`, but using the first element of `itr` as `v0`. In general, this cannot be used with empty collections (see `reduce(op, itr)`).

> source

`Base.first` – Function.

```
first(coll)
```

Get the first element of an iterable collection. Returns the start point of a Range even if it is empty.

```
julia> first(2:2:10)
2

julia> first([1; 2; 3; 4])
1
```

> source

`Base.last` – Function.

```
last(coll)
```

Get the last element of an ordered collection, if it can be computed in O(1) time. This is accomplished by calling `endof` to get the last index. Returns the end point of a Range even if it is empty.

```
julia> last(1:2:10)
9

julia> last([1; 2; 3; 4])
4
```

> source

Base.step – Function.

```
step(r)
```

Get the step size of a Range object.

```
julia> step(1:10)
1

julia> step(1:2:10)
2

julia> step(2.5:0.3:10.9)
0.3

julia> step(linspace(2.5,10.9,85))
0.1
```

source

Base.collect – Method.

```
collect(collection)
```

Return an `Array` of all items in a collection or iterator. For associative collections, returns `Pair{KeyType, ValType}`. If the argument is array-like or is an iterator with the `HasShape()` trait, the result will have the same shape and number of dimensions as the argument.

```
julia> collect(1:2:13)
7-element Array{Int64,1}:
  1
  3
  5
  7
  9
 11
 13
```

source

Base.collect – Method.

```
collect(element_type, collection)
```

Return an `Array` with the given element type of all items in a collection or iterable. The result has the same shape and number of dimensions as `collection`.

```
julia> collect(Float64, 1:2:5)
3-element Array{Float64,1}:
 1.0
 3.0
 5.0
```

source

Base.issubset – Method.

```
issubset(a, b)
(a,b) -> Bool
(a,b) -> Bool
(a,b) -> Bool
```

Determine whether every element of a is also in b, using in.

source

Base.filter – Function.

```
filter(function, collection)
```

Return a copy of `collection`, removing elements for which `function` is `false`. For associative collections, the function is passed two arguments (key and value).

```
julia> a = 1:10
1:10

julia> filter(isodd, a)
5-element Array{Int64,1}:
 1
 3
 5
 7
 9
```

source

```
filter(p, x::Nullable)
```

Return null if either x is null or p(get(x)) is false, and x otherwise.

source

Base.filter! – Function.

```
filter!(function, collection)
```

Update `collection`, removing elements for which `function` is `false`. For associative collections, the function is passed two arguments (key and value).

```
julia> filter!(isodd, collect(1:10))
5-element Array{Int64,1}:
 1
 3
 5
 7
 9
```

source

46.4   Indexable Collections

Base.getindex – Method.

```
getindex(collection, key...)
```

Retrieve the value(s) stored at the given key or index within a collection. The syntax `a[i,j,...]` is converted by the compiler to `getindex(a, i, j, ...)`.

```
julia> A = Dict("a" => 1, "b" => 2)
Dict{String,Int64} with 2 entries:
  "b" => 2
  "a" => 1

julia> getindex(A, "a")
1
```

source

`Base.setindex!` – Method.

```
setindex!(collection, value, key...)
```

Store the given value at the given key or index within a collection. The syntax `a[i,j,...] = x` is converted by the compiler to `(setindex!(a, x, i, j, ...); x)`.

source

Fully implemented by:

- `Array`

- `BitArray`

- `AbstractArray`

- `SubArray`

- `ObjectIdDict`

- `Dict`

- `WeakKeyDict`

- `AbstractString`

Partially implemented by:

- `Range`

- `UnitRange`

- `Tuple`

46.5   Associative Collections

Dict is the standard associative collection. Its implementation uses hash() as the hashing function for the key, and isequal() to determine equality. Define these two functions for custom types to override how they are stored in a hash table.

ObjectIdDict is a special hash table where the keys are always object identities.

WeakKeyDict is a hash table implementation where the keys are weak references to objects, and thus may be garbage collected even when referenced in a hash table.

Dicts can be created by passing pair objects constructed with =>() to a Dict constructor: Dict("A"=>1, "B"=>2). This call will attempt to infer type information from the keys and values (i.e. this example creates a Dict{String, Int64}). To explicitly specify types use the syntax Dict{KeyType,ValueType}(...). For example, Dict{String,Int32}("A"=>1, "B"=>2).

Associative collections may also be created with generators. For example, Dict(i => f(i) for i = 1:10).

Given a dictionary D, the syntax D[x] returns the value of key x (if it exists) or throws an error, and D[x] = y stores the key-value pair x => y in D (replacing any existing value for the key x). Multiple arguments to D[...] are converted to tuples; for example, the syntax D[x,y] is equivalent to D[(x,y)], i.e. it refers to the value keyed by the tuple (x,y).

Base.Dict – Type.

```
Dict([itr])
```

Dict{K,V}() constructs a hash table with keys of type K and values of type V.

Given a single iterable argument, constructs a Dict whose key-value pairs are taken from 2-tuples (key,value) generated by the argument.

```
julia> Dict([("A", 1), ("B", 2)])
Dict{String,Int64} with 2 entries:
  "B" => 2
  "A" => 1
```

Alternatively, a sequence of pair arguments may be passed.

```
julia> Dict("A"=>1, "B"=>2)
Dict{String,Int64} with 2 entries:
  "B" => 2
  "A" => 1
```

source

Base.ObjectIdDict – Type.

```
ObjectIdDict([itr])
```

ObjectIdDict() constructs a hash table where the keys are (always) object identities. Unlike Dict it is not parameterized on its key and value type and thus its eltype is always Pair{Any,Any}.

See Dict for further help.

source

Base.WeakKeyDict – Type.

```
WeakKeyDict([itr])
```

`WeakKeyDict()` constructs a hash table where the keys are weak references to objects, and thus may be garbage collected even when referenced in a hash table.

See `Dict` for further help.

source

`Base.haskey` – Function.

```
haskey(collection, key) -> Bool
```

Determine whether a collection has a mapping for a given key.

```julia
julia> a = Dict('a'=>2, 'b'=>3)
Dict{Char,Int64} with 2 entries:
  'b' => 3
  'a' => 2

julia> haskey(a,'a')
true

julia> haskey(a,'c')
false
```

source

`Base.get` – Method.

```
get(collection, key, default)
```

Return the value stored for the given key, or the given default value if no mapping for the key is present.

source

`Base.get` – Function.

```
get(f::Function, collection, key)
```

Return the value stored for the given key, or if no mapping for the key is present, return `f()`. Use `get!` to also store the default value in the dictionary.

This is intended to be called using do block syntax

```
get(dict, key) do
    # default value calculated here
    time()
end
```

source

`Base.get!` – Method.

```
get!(collection, key, default)
```

Return the value stored for the given key, or if no mapping for the key is present, store `key => default`, and return `default`.

```julia
julia> d = Dict("a"=>1, "b"=>2, "c"=>3);

julia> get!(d, "a", 5)
1

julia> get!(d, "d", 4)
4

julia> d
Dict{String,Int64} with 4 entries:
  "c" => 3
  "b" => 2
  "a" => 1
  "d" => 4
```

source

Base.get! – Method.

```julia
get!(f::Function, collection, key)
```

Return the value stored for the given key, or if no mapping for the key is present, store `key => f()`, and return `f()`.

This is intended to be called using do block syntax:

```julia
get!(dict, key) do
    # default value calculated here
    time()
end
```

source

Base.getkey – Function.

```julia
getkey(collection, key, default)
```

Return the key matching argument `key` if one exists in `collection`, otherwise return `default`.

```julia
julia> a = Dict('a'=>2, 'b'=>3)
Dict{Char,Int64} with 2 entries:
  'b' => 3
  'a' => 2

julia> getkey(a,'a',1)
'a': ASCII/Unicode U+0061 (category Ll: Letter, lowercase)

julia> getkey(a,'d','a')
'a': ASCII/Unicode U+0061 (category Ll: Letter, lowercase)
```

source

Base.delete! – Function.

```julia
delete!(collection, key)
```

Delete the mapping for the given key in a collection, and return the collection.

source

Base.pop! – Method.

```
pop!(collection, key[, default])
```

Delete and return the mapping for `key` if it exists in `collection`, otherwise return `default`, or throw an error if `default` is not specified.

```
julia> d = Dict("a"=>1, "b"=>2, "c"=>3);

julia> pop!(d, "a")
1

julia> pop!(d, "d")
ERROR: KeyError: key "d" not found
Stacktrace:
 [1] pop!(::Dict{String,Int64}, ::String) at ./dict.jl:539

julia> pop!(d, "e", 4)
4
```

source

Base.keys – Function.

```
keys(a::Associative)
```

Return an iterator over all keys in a collection. `collect(keys(a))` returns an array of keys. Since the keys are stored internally in a hash table, the order in which they are returned may vary. But `keys(a)` and `values(a)` both iterate a and return the elements in the same order.

```
julia> a = Dict('a'=>2, 'b'=>3)
Dict{Char,Int64} with 2 entries:
  'b' => 3
  'a' => 2

julia> collect(keys(a))
2-element Array{Char,1}:
 'b'
 'a'
```

source

Base.values – Function.

```
values(a::Associative)
```

Return an iterator over all values in a collection. `collect(values(a))` returns an array of values. Since the values are stored internally in a hash table, the order in which they are returned may vary. But `keys(a)` and `values(a)` both iterate a and return the elements in the same order.

```
julia> a = Dict('a'=>2, 'b'=>3)
Dict{Char,Int64} with 2 entries:
  'b' => 3
  'a' => 2

julia> collect(values(a))
```

```
2-element Array{Int64,1}:
 3
 2
```

source

Base.merge – Function.

```
merge(d::Associative, others::Associative...)
```

Construct a merged collection from the given collections. If necessary, the types of the resulting collection will be promoted to accommodate the types of the merged collections. If the same key is present in another collection, the value for that key will be the value it has in the last collection listed.

```julia
julia> a = Dict("foo" => 0.0, "bar" => 42.0)
Dict{String,Float64} with 2 entries:
  "bar" => 42.0
  "foo" => 0.0

julia> b = Dict("baz" => 17, "bar" => 4711)
Dict{String,Int64} with 2 entries:
  "bar" => 4711
  "baz" => 17

julia> merge(a, b)
Dict{String,Float64} with 3 entries:
  "bar" => 4711.0
  "baz" => 17.0
  "foo" => 0.0

julia> merge(b, a)
Dict{String,Float64} with 3 entries:
  "bar" => 42.0
  "baz" => 17.0
  "foo" => 0.0
```

source

```
merge(combine, d::Associative, others::Associative...)
```

Construct a merged collection from the given collections. If necessary, the types of the resulting collection will be promoted to accommodate the types of the merged collections. Values with the same key will be combined using the combiner function.

```julia
julia> a = Dict("foo" => 0.0, "bar" => 42.0)
Dict{String,Float64} with 2 entries:
  "bar" => 42.0
  "foo" => 0.0

julia> b = Dict("baz" => 17, "bar" => 4711)
Dict{String,Int64} with 2 entries:
  "bar" => 4711
  "baz" => 17

julia> merge(+, a, b)
Dict{String,Float64} with 3 entries:
```

```
  "bar" => 4753.0
  "baz" => 17.0
  "foo" => 0.0
```

source

Base.merge! – Function.

Merge changes into current head

source

Internal implementation of merge. Returns `true` if merge was successful, otherwise `false`

source

```
merge!(repo::GitRepo; kwargs...) -> Bool
```

Perform a git merge on the repository `repo`, merging commits with diverging history into the current branch. Returns `true` if the merge succeeded, `false` if not.

The keyword arguments are:

- `committish::AbstractString=""`: Merge the named commit(s) in `committish`.
- `branch::AbstractString=""`: Merge the branch `branch` and all its commits since it diverged from the current branch.
- `fastforward::Bool=false`: If `fastforward` is `true`, only merge if the merge is a fast-forward (the current branch head is an ancestor of the commits to be merged), otherwise refuse to merge and return `false`. This is equivalent to the git CLI option `--ff-only`.
- `merge_opts::MergeOptions=MergeOptions()`: `merge_opts` specifies options for the merge, such as merge strategy in case of conflicts.
- `checkout_opts::CheckoutOptions=CheckoutOptions()`: `checkout_opts` specifies options for the checkout step.

Equivalent to `git merge [--ff-only] [<committish> | <branch>]`.

> **Note**
> If you specify a `branch`, this must be done in reference format, since the string will be turned into a `GitReference`. For example, if you wanted to merge branch `branch_a`, you would call `merge!(repo, branch="refs/heads/branch_a")`.

source

```
merge!(d::Associative, others::Associative...)
```

Update collection with pairs from the other collections. See also merge.

```
julia> d1 = Dict(1 => 2, 3 => 4);

julia> d2 = Dict(1 => 4, 4 => 5);

julia> merge!(d1, d2);

julia> d1
Dict{Int64,Int64} with 3 entries:
  4 => 5
  3 => 4
  1 => 4
```

source

```
merge!(combine, d::Associative, others::Associative...)
```

Update collection with pairs from the other collections. Values with the same key will be combined using the combiner function.

```
julia> d1 = Dict(1 => 2, 3 => 4);

julia> d2 = Dict(1 => 4, 4 => 5);

julia> merge!(+, d1, d2);

julia> d1
Dict{Int64,Int64} with 3 entries:
  4 => 5
  3 => 4
  1 => 6

julia> merge!(-, d1, d1);

julia> d1
Dict{Int64,Int64} with 3 entries:
  4 => 0
  3 => 0
  1 => 0
```

source

`Base.sizehint!` – Function.

```
sizehint!(s, n)
```

Suggest that collection s reserve capacity for at least n elements. This can improve performance.

source

`Base.keytype` – Function.

```
keytype(type)
```

Get the key type of an associative collection type. Behaves similarly to `eltype`.

```
julia> keytype(Dict(Int32(1) => "foo"))
Int32
```

source

`Base.valtype` – Function.

```
valtype(type)
```

Get the value type of an associative collection type. Behaves similarly to `eltype`.

```
julia> valtype(Dict(Int32(1) => "foo"))
String
```

source

Fully implemented by:

- `ObjectIdDict`

- `Dict`

- `WeakKeyDict`

Partially implemented by:

- `IntSet`

- `Set`

- `EnvHash`

- `Array`

- `BitArray`

46.6   Set-Like Collections

`Base.Set` – Type.

```
Set([itr])
```

Construct a `Set` of the values generated by the given iterable object, or an empty set. Should be used instead of `IntSet` for sparse integer sets, or for sets of arbitrary objects.

source

`Base.IntSet` – Type.

```
IntSet([itr])
```

Construct a sorted set of positive `Int`s generated by the given iterable object, or an empty set. Implemented as a bit string, and therefore designed for dense integer sets. Only `Int`s greater than 0 can be stored. If the set will be sparse (for example holding a few very large integers), use `Set` instead.

source

`Base.union` – Function.

```
union(s1,s2...)
(s1,s2...)
```

Construct the union of two or more sets. Maintains order with arrays.

source

`Base.union!` – Function.

```
union!(s, iterable)
```

Union each element of `iterable` into set s in-place.

source

Base.intersect – Function.

```
intersect(s1,s2...)
(s1,s2)
```

Construct the intersection of two or more sets. Maintains order and multiplicity of the first argument for arrays and ranges.

source

Base.setdiff – Function.

```
setdiff(a, b)
```

Construct the set of elements in a but not b. Maintains order with arrays. Note that both arguments must be collections, and both will be iterated over. In particular, `setdiff(set,element)` where `element` is a potential member of `set`, will not work in general.

```
julia> setdiff([1,2,3],[3,4,5])
2-element Array{Int64,1}:
 1
 2
```

source

Base.setdiff! – Function.

```
setdiff!(s, iterable)
```

Remove each element of `iterable` from set `s` in-place.

source

Base.symdiff – Function.

```
symdiff(a, b, rest...)
```

Construct the symmetric difference of elements in the passed in sets or arrays. Maintains order with arrays.

```
julia> symdiff([1,2,3],[3,4,5],[4,5,6])
3-element Array{Int64,1}:
 1
 2
 6
```

source

Base.symdiff! – Method.

```
symdiff!(s, n)
```

The set `s` is destructively modified to toggle the inclusion of integer `n`.

source

Base.symdiff! – Method.

```
symdiff!(s, itr)
```

For each element in `itr`, destructively toggle its inclusion in set `s`.

source

`Base.symdiff!` – Method.

```
symdiff!(s, itr)
```

For each element in `itr`, destructively toggle its inclusion in set `s`.

source

`Base.intersect!` – Function.

```
intersect!(s1::IntSet, s2::IntSet)
```

Intersects sets `s1` and `s2` and overwrites the set `s1` with the result. If needed, `s1` will be expanded to the size of `s2`.

source

`Base.issubset` – Function.

```
issubset(A, S) -> Bool
(A,S) -> Bool
```

Return `true` if `A` is a subset of or equal to `S`.

source

Fully implemented by:

- IntSet
- Set

Partially implemented by:

- Array

46.7 Dequeues

`Base.push!` – Function.

```
push!(collection, items...) -> collection
```

Insert one or more `items` at the end of `collection`.

```
julia> push!([1, 2, 3], 4, 5, 6)
6-element Array{Int64,1}:
 1
 2
 3
 4
 5
 6
```

Use append! to add all the elements of another collection to collection. The result of the preceding example is equivalent to append!([1, 2, 3], [4, 5, 6]).

source

Base.pop! – Method.

```
pop!(collection) -> item
```

Remove the last item in collection and return it.

```
julia> A=[1, 2, 3, 4, 5, 6]
6-element Array{Int64,1}:
 1
 2
 3
 4
 5
 6

julia> pop!(A)
6

julia> A
5-element Array{Int64,1}:
 1
 2
 3
 4
 5
```

source

Base.unshift! – Function.

```
unshift!(collection, items...) -> collection
```

Insert one or more items at the beginning of collection.

```
julia> unshift!([1, 2, 3, 4], 5, 6)
6-element Array{Int64,1}:
 5
 6
 1
 2
 3
 4
```

source

Base.shift! – Function.

```
shift!(collection) -> item
```

Remove the first item from collection.

```
julia> A = [1, 2, 3, 4, 5, 6]
6-element Array{Int64,1}:
 1
 2
 3
 4
 5
 6

julia> shift!(A)
1

julia> A
5-element Array{Int64,1}:
 2
 3
 4
 5
 6
```

source

Base.insert! – Function.

```
insert!(a::Vector, index::Integer, item)
```

Insert an `item` into a at the given `index`. `index` is the index of `item` in the resulting a.

```
julia> insert!([6, 5, 4, 2, 1], 4, 3)
6-element Array{Int64,1}:
 6
 5
 4
 3
 2
 1
```

source

Base.deleteat! – Function.

```
deleteat!(a::Vector, i::Integer)
```

Remove the item at the given `i` and return the modified a. Subsequent items are shifted to fill the resulting gap.

```
julia> deleteat!([6, 5, 4, 3, 2, 1], 2)
5-element Array{Int64,1}:
 6
 4
 3
 2
 1
```

source

```
deleteat!(a::Vector, inds)
```

Remove the items at the indices given by inds, and return the modified a. Subsequent items are shifted to fill the resulting gap.

inds can be either an iterator or a collection of sorted and unique integer indices, or a boolean vector of the same length as a with true indicating entries to delete.

```julia
julia> deleteat!([6, 5, 4, 3, 2, 1], 1:2:5)
3-element Array{Int64,1}:
 5
 3
 1

julia> deleteat!([6, 5, 4, 3, 2, 1], [true, false, true, false, true, false])
3-element Array{Int64,1}:
 5
 3
 1

julia> deleteat!([6, 5, 4, 3, 2, 1], (2, 2))
ERROR: ArgumentError: indices must be unique and sorted
Stacktrace:
 [1] _deleteat!(::Array{Int64,1}, ::Tuple{Int64,Int64}) at ./array.jl:885
 [2] deleteat!(::Array{Int64,1}, ::Tuple{Int64,Int64}) at ./array.jl:872
```

source

Base.splice! – Function.

```
splice!(a::Vector, index::Integer, [replacement]) -> item
```

Remove the item at the given index, and return the removed item. Subsequent items are shifted left to fill the resulting gap. If specified, replacement values from an ordered collection will be spliced in place of the removed item.

```
julia> A = [6, 5, 4, 3, 2, 1]; splice!(A, 5)
2

julia> A
5-element Array{Int64,1}:
 6
 5
 4
 3
 1

julia> splice!(A, 5, -1)
1

julia> A
5-element Array{Int64,1}:
  6
  5
  4
  3
 -1

julia> splice!(A, 1, [-1, -2, -3])
```

```
6

julia> A
7-element Array{Int64,1}:
 -1
 -2
 -3
  5
  4
  3
 -1
```

To insert `replacement` before an index n without removing any items, use `splice!(collection, n:n-1, replacement)`.

```
splice!(a::Vector, range, [replacement]) -> items
```

Remove items in the specified index range, and return a collection containing the removed items. Subsequent items are shifted left to fill the resulting gap. If specified, replacement values from an ordered collection will be spliced in place of the removed items.

To insert `replacement` before an index n without removing any items, use `splice!(collection, n:n-1, replacement)`.

```
julia> splice!(A, 4:3, 2)
0-element Array{Int64,1}

julia> A
8-element Array{Int64,1}:
 -1
 -2
 -3
  2
  5
  4
  3
 -1
```

`Base.resize!` – Function.

```
resize!(a::Vector, n::Integer) -> Vector
```

Resize a to contain n elements. If n is smaller than the current collection length, the first n elements will be retained. If n is larger, the new elements are not guaranteed to be initialized.

```
julia> resize!([6, 5, 4, 3, 2, 1], 3)
3-element Array{Int64,1}:
 6
 5
 4

julia> resize!([6, 5, 4, 3, 2, 1], 8)
8-element Array{Int64,1}:
```

```
6
5
4
3
2
1
0
0
```

source

Base.append! – Function.

```
append!(collection, collection2) -> collection.
```

Add the elements of `collection2` to the end of `collection`.

```
julia> append!([1],[2,3])
3-element Array{Int64,1}:
 1
 2
 3
```

```
julia> append!([1, 2, 3], [4, 5, 6])
6-element Array{Int64,1}:
 1
 2
 3
 4
 5
 6
```

Use `push!` to add individual items to `collection` which are not already themselves in another collection. The result is of the preceding example is equivalent to `push!([1, 2, 3], 4, 5, 6)`.

source

Base.prepend! – Function.

```
prepend!(a::Vector, items) -> collection
```

Insert the elements of `items` to the beginning of `a`.

```
julia> prepend!([3],[1,2])
3-element Array{Int64,1}:
 1
 2
 3
```

source

Fully implemented by:

- `Vector` (a.k.a. 1-dimensional `Array`)

- `BitVector` (a.k.a. 1-dimensional `BitArray`)

Chapter 47

Mathematics

47.1 Mathematical Operators

`Base.:-` – Method.

```
-(x)
```

Unary minus operator.

`source`

`Base.:+` – Function.

```
+(x, y...)
```

Addition operator. `x+y+z+...` calls this function with all arguments, i.e. `+(x, y, z, ...)`.

`source`

`Base.:-` – Method.

```
-(x, y)
```

Subtraction operator.

`source`

`Base.:*` – Method.

```
*(x, y...)
```

Multiplication operator. `x*y*z*...` calls this function with all arguments, i.e. `*(x, y, z, ...)`.

`source`

`Base.:/` – Function.

```
/(x, y)
```

Right division operator: multiplication of x by the inverse of y on the right. Gives floating-point results for integer arguments.

`source`

`Base.:\` – Method.

```
| \(x, y)
```

Left division operator: multiplication of y by the inverse of x on the left. Gives floating-point results for integer arguments.

```julia
julia> 3 \ 6
2.0

julia> inv(3) * 6
2.0

julia> A = [1 2; 3 4]; x = [5, 6];

julia> A \ x
2-element Array{Float64,1}:
 -4.0
  4.5

julia> inv(A) * x
2-element Array{Float64,1}:
 -4.0
  4.5
```

source

Base.:^ – Method.

```
| ^(x, y)
```

Exponentiation operator. If x is a matrix, computes matrix exponentiation.

If y is an Int literal (e.g. 2 in x^2 or -3 in x^-3), the Julia code x^y is transformed by the compiler to `Base.literal_pow(^, x, Val{y})`, to enable compile-time specialization on the value of the exponent. (As a default fallback we have `Base.literal_pow(^, x, Val{y}) = ^(x,y)`, where usually `^ == Base.^` unless `^` has been defined in the calling namespace.)

```julia
julia> 3^5
243

julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> A^3
2×2 Array{Int64,2}:
 37   54
 81  118
```

source

Base.fma – Function.

```
| fma(x, y, z)
```

Computes x*y+z without rounding the intermediate result x*y. On some systems this is significantly more expensive than x*y+z. fma is used to improve accuracy in certain algorithms. See muladd.

source

Base.muladd – Function.

```
muladd(x, y, z)
```

Combined multiply-add, computes x*y+z in an efficient manner. This may on some systems be equivalent to x*y+z, or to fma(x,y,z). muladd is used to improve performance. See fma.

Example

```
julia> muladd(3, 2, 1)
7

julia> 3 * 2 + 1
7
```

source

Base.div – Function.

```
div(x, y)
÷(x, y)
```

The quotient from Euclidean division. Computes x/y, truncated to an integer.

```
julia> 9 ÷ 4
2

julia> -5 ÷ 3
-1
```

source

Base.fld – Function.

```
fld(x, y)
```

Largest integer less than or equal to x/y.

```
julia> fld(7.3,5.5)
1.0
```

source

Base.cld – Function.

```
cld(x, y)
```

Smallest integer larger than or equal to x/y.

```
julia> cld(5.5,2.2)
3.0
```

Base.mod – Function.

```
mod(x, y)
rem(x, y, RoundDown)
```

The reduction of x modulo y, or equivalently, the remainder of x after floored division by y, i.e.

```
x - y*fld(x,y)
```

if computed without intermediate rounding.

The result will have the same sign as y, and magnitude less than `abs(y)` (with some exceptions, see note below).

> **Note**
>
> When used with floating point values, the exact result may not be representable by the type, and so rounding error may occur. In particular, if the exact result is very close to y, then it may be rounded to y.

```
julia> mod(8, 3)
2

julia> mod(9, 3)
0

julia> mod(8.9, 3)
2.9000000000000004

julia> mod(eps(), 3)
2.220446049250313e-16

julia> mod(-eps(), 3)
3.0
```

```
rem(x::Integer, T::Type{<:Integer}) -> T
mod(x::Integer, T::Type{<:Integer}) -> T
%(x::Integer, T::Type{<:Integer}) -> T
```

Find `y::T` such that $x \equiv y \pmod{n}$, where n is the number of integers representable in T, and y is an integer in `[typemin(T),typemax(T)]`. If T can represent any integer (e.g. `T == BigInt`), then this operation corresponds to a conversion to T.

```
julia> 129 % Int8
-127
```

Base.rem – Function.

```
rem(x, y)
%(x, y)
```

Remainder from Euclidean division, returning a value of the same sign as x, and smaller in magnitude than y. This value is always exact.

```
julia> x = 15; y = 4;

julia> x % y
3

julia> x == div(x, y) * y + rem(x, y)
true
```

source

Base.Math.rem2pi – Function.

```
rem2pi(x, r::RoundingMode)
```

Compute the remainder of x after integer division by 2π, with the quotient rounded according to the rounding mode r. In other words, the quantity

```
x - π2*round(xπ/(2),r)
```

without any intermediate rounding. This internally uses a high precision approximation of 2π, and so will give a more accurate result than `rem(x,2π,r)`

- if r == RoundNearest, then the result is in the interval $[-,]$. This will generally be the most accurate result.
- if r == RoundToZero, then the result is in the interval $[0, 2]$ if x is positive,. or $[-2, 0]$ otherwise.
- if r == RoundDown, then the result is in the interval $[0, 2]$.
- if r == RoundUp, then the result is in the interval $[-2, 0]$.

Example

```
julia> rem2pi(7pi/4, RoundNearest)
-0.7853981633974485

julia> rem2pi(7pi/4, RoundDown)
5.497787143782138
```

source

Base.Math.mod2pi – Function.

```
mod2pi(x)
```

Modulus after division by 2π, returning in the range $[0, 2)$.

This function computes a floating point representation of the modulus after division by numerically exact 2π, and is therefore not exactly the same as `mod(x,2π)`, which would compute the modulus of x relative to division by the floating-point number 2π.

Example

```
julia> mod2pi(9*pi/4)
0.7853981633974481
```

source

`Base.divrem` – Function.

```
divrem(x, y)
```

The quotient and remainder from Euclidean division. Equivalent to `(div(x,y), rem(x,y))` or `(x÷y, x%y)`.

```
julia> divrem(3,7)
(0, 3)

julia> divrem(7,3)
(2, 1)
```

source

`Base.fldmod` – Function.

```
fldmod(x, y)
```

The floored quotient and modulus after division. Equivalent to `(fld(x,y), mod(x,y))`.

source

`Base.fld1` – Function.

```
fld1(x, y)
```

Flooring division, returning a value consistent with `mod1(x,y)`

See also: mod1.

```
julia> x = 15; y = 4;

julia> fld1(x, y)
4

julia> x == fld(x, y) * y + mod(x, y)
true

julia> x == (fld1(x, y) - 1) * y + mod1(x, y)
true
```

source

`Base.mod1` – Function.

```
mod1(x, y)
```

Modulus after flooring division, returning a value r such that `mod(r, y) == mod(x, y)` in the range $(0, y]$ for positive y and in the range $[y, 0)$ for negative y.

```
julia> mod1(4, 2)
2

julia> mod1(4, 3)
1
```

source

Base.`fldmod1` – Function.

```
fldmod1(x, y)
```

Return (`fld1(x,y)`, `mod1(x,y)`).

See also: `fld1`, `mod1`.

source

Base.`://` – Function.

```
//(num, den)
```

Divide two integers or rational numbers, giving a `Rational` result.

```julia
julia> 3 // 5
3//5

julia> (3 // 5) // (2 // 1)
3//10
```

source

Base.`rationalize` – Function.

```
rationalize([T<:Integer=Int,] x; tol::Real=eps(x))
```

Approximate floating point number x as a `Rational` number with components of the given integer type. The result will differ from x by no more than `tol`. If T is not provided, it defaults to `Int`.

```julia
julia> rationalize(5.6)
28//5

julia> a = rationalize(BigInt, 10.3)
103//10

julia> typeof(numerator(a))
BigInt
```

source

Base.`numerator` – Function.

```
numerator(x)
```

Numerator of the rational representation of x.

```julia
julia> numerator(2//3)
2

julia> numerator(4)
4
```

source

Base.`denominator` – Function.

```
denominator(x)
```

Denominator of the rational representation of x.

```
julia> denominator(2//3)
3

julia> denominator(4)
1
```

source

Base.:<< – Function.

```
<<(x, n)
```

Left bit shift operator, x << n. For n >= 0, the result is x shifted left by n bits, filling with 0s. This is equivalent to x * 2^n. For n < 0, this is equivalent to x >> -n.

```
julia> Int8(3) << 2
12

julia> bits(Int8(3))
"00000011"

julia> bits(Int8(12))
"00001100"
```

See also >>, >>>.

source

```
<<(B::BitVector, n) -> BitVector
```

Left bit shift operator, B << n. For n >= 0, the result is B with elements shifted n positions backwards, filling with false values. If n < 0, elements are shifted forwards. Equivalent to B >> -n.

Examples

```
julia> B = BitVector([true, false, true, false, false])
5-element BitArray{1}:
  true
 false
  true
 false
 false

julia> B << 1
5-element BitArray{1}:
 false
  true
 false
 false
 false

julia> B << -1
5-element BitArray{1}:
```

```
  false
   true
  false
   true
  false
```

source

`Base.:>>` – Function.

```
>>(x, n)
```

Right bit shift operator, x >> n. For n >= 0, the result is x shifted right by n bits, where n >= 0, filling with 0s if x >= 0, 1s if x < 0, preserving the sign of x. This is equivalent to fld(x, 2^n). For n < 0, this is equivalent to x << -n.

```julia
julia> Int8(13) >> 2
3

julia> bits(Int8(13))
"00001101"

julia> bits(Int8(3))
"00000011"

julia> Int8(-14) >> 2
-4

julia> bits(Int8(-14))
"11110010"

julia> bits(Int8(-4))
"11111100"
```

See also >>>, <<.

source

```
>>(B::BitVector, n) -> BitVector
```

Right bit shift operator, B >> n. For n >= 0, the result is B with elements shifted n positions forward, filling with `false` values. If n < 0, elements are shifted backwards. Equivalent to B << -n.

Example

```julia
julia> B = BitVector([true, false, true, false, false])
5-element BitArray{1}:
   true
  false
   true
  false
  false

julia> B >> 1
5-element BitArray{1}:
  false
   true
```

```
  false
   true
  false

julia> B >> -1
5-element BitArray{1}:
  false
   true
  false
  false
  false
```

source

Base.:>>> – Function.

```
>>>(x, n)
```

Unsigned right bit shift operator, x >>> n. For n >= 0, the result is x shifted right by n bits, where n >= 0, filling with 0s. For n < 0, this is equivalent to x << -n.

For Unsigned integer types, this is equivalent to >>. For Signed integer types, this is equivalent to signed(unsigned(x) >> n).

```
julia> Int8(-14) >>> 2
60

julia> bits(Int8(-14))
"11110010"

julia> bits(Int8(60))
"00111100"
```

BigInts are treated as if having infinite size, so no filling is required and this is equivalent to >>.

See also >>, <<.

source

```
>>>(B::BitVector, n) -> BitVector
```

Unsigned right bitshift operator, B >>> n. Equivalent to B >> n. See >> for details and examples.

source

Base.colon – Function.

```
colon(start, [step], stop)
```

Called by : syntax for constructing ranges.

```
julia> colon(1, 2, 5)
1:2:5
```

source

```
:(start, [step], stop)
```

Range operator. `a:b` constructs a range from a to b with a step size of 1, and `a:s:b` is similar but uses a step size of s. These syntaxes call the function `colon`. The colon is also used in indexing to select whole dimensions.

source

Base.range – Function.

```
range(start, [step], length)
```

Construct a range by length, given a starting value and optional step (defaults to 1).

source

Base.OneTo – Type.

```
Base.OneTo(n)
```

Define an `AbstractUnitRange` that behaves like `1:n`, with the added distinction that the lower limit is guaranteed (by the type system) to be 1.

source

Base.StepRangeLen – Type.

```
StepRangeLen{T,R,S}(ref::R, step::S, len, [offset=1])
```

A range r where `r[i]` produces values of type T, parametrized by a `reference` value, a `step`, and the `length`. By default `ref` is the starting value `r[1]`, but alternatively you can supply it as the value of `r[offset]` for some other index `1 <= offset <= len`. In conjunction with `TwicePrecision` this can be used to implement ranges that are free of roundoff error.

source

Base.:== – Function.

```
==(x, y)
```

Generic equality operator, giving a single `Bool` result. Falls back to ===. Should be implemented for all types with a notion of equality, based on the abstract value that an instance represents. For example, all numeric types are compared by numeric value, ignoring type. Strings are compared as sequences of characters, ignoring encoding.

Follows IEEE semantics for floating-point numbers.

Collections should generally implement == by calling == recursively on all contents.

New numeric types should implement this function for two arguments of the new type, and handle comparison to other types via promotion rules where possible.

source

Base.:!= – Function.

```
!=(x, y)≠
(x,y)
```

Not-equals comparison operator. Always gives the opposite answer as ==. New types should generally not implement this, and rely on the fallback definition `!=(x,y) = !(x==y)` instead.

```julia
julia> 3 != 2
true

julia> "foo" ≠ "foo"
false
```

source

Base.:!== – Function.

```julia
!==(x, y)
(x,y)
```

Equivalent to !(x === y).

```julia
julia> a = [1, 2]; b = [1, 2];

julia> a ≢ b
true

julia> a ≢ a
false
```

source

Base.:< – Function.

```julia
<(x, y)
```

Less-than comparison operator. New numeric types should implement this function for two arguments of the new type. Because of the behavior of floating-point NaN values, < implements a partial order. Types with a canonical partial order should implement <, and types with a canonical total order should implement isless.

```julia
julia> 'a' < 'b'
true

julia> "abc" < "abd"
true

julia> 5 < 3
false
```

source

Base.:<= – Function.

```julia
<=(x, y)≤
(x,y)
```

Less-than-or-equals comparison operator.

```julia
julia> 'a' <= 'b'
true

julia> 7 ≤ 7 ≤ 9
true
```

```
julia> "abc" ≤ "abc"
true

julia> 5 <= 3
false
```

source

Base.:> – Function.

```
>(x, y)
```

Greater-than comparison operator. Generally, new types should implement < instead of this function, and rely on the fallback definition >(x, y) = y < x.

```
julia> 'a' > 'b'
false

julia> 7 > 3 > 1
true

julia> "abc" > "abd"
false

julia> 5 > 3
true
```

source

Base.:>= – Function.

```
>=(x, y)≥
(x,y)
```

Greater-than-or-equals comparison operator.

```
julia> 'a' >= 'b'
false

julia> 7 ≥ 7 ≥ 3
true

julia> "abc" ≥ "abc"
true

julia> 5 >= 3
true
```

source

Base.cmp – Function.

```
cmp(x,y)
```

Return -1, 0, or 1 depending on whether x is less than, equal to, or greater than y, respectively. Uses the total order implemented by `isless`. For floating-point numbers, uses < but throws an error for unordered arguments.

```
julia> cmp(1, 2)
-1

julia> cmp(2, 1)
1

julia> cmp(2+im, 3-im)
ERROR: MethodError: no method matching isless(::Complex{Int64}, ::Complex{Int64})
[...]
```

source

Base.:~ – Function.

```
~(x)
```

Bitwise not.

```
julia> ~4
-5

julia> ~10
-11

julia> ~true
false
```

source

Base.:& – Function.

```
&(x, y)
```

Bitwise and.

```
julia> 4 & 10
0

julia> 4 & 12
4
```

source

Base.:| – Function.

```
|(x, y)
```

Bitwise or.

```
julia> 4 | 10
14

julia> 4 | 1
5
```

source

Base.xor – Function.

```
xor(x, y)
⊻(x, y)
```

Bitwise exclusive or of x and y. The infix operation a ⊻ b is a synonym for `xor(a,b)`, and ⊻ can be typed by tab-completing `\xor` or `\veebar` in the Julia REPL.

```
julia> [true; true; false] .⊻ [true; false; false]
3-element BitArray{1}:
 false
  true
 false
```

source

Base.:! – Function.

```
!(x)
```

Boolean not.

```
julia> !true
false

julia> !false
true

julia> .![true false true]
1×3 BitArray{2}:
 false  true  false
```

source

```
!f::Function
```

Predicate function negation: when the argument of `!` is a function, it returns a function which computes the boolean negation of f. Example:

```
julia> str = "∀ ε > 0, ∃ δ > 0: |x-y| < δ ⟹ |f(x)-f(y)| < ε"
"∀ ε > 0, ∃ δ > 0: |x-y| < δ ⟹ |f(x)-f(y)| < ε"

julia> filter(isalpha, str)
"εδxyδfxfyε"

julia> filter(!isalpha, str)
"∀ > 0, ∃ > 0: |-| < ⟹ |()-()| < "
```

source

&& – Keyword.

```
x && y
```

Short-circuiting boolean AND.

source

|| – Keyword.

```
x || y
```

Short-circuiting boolean OR.

source

47.2   Mathematical Functions

Base.isapprox – Function.

```
isapprox(x, y; rtol::Real=sqrt(eps), atol::Real=0, nans::Bool=false, norm::Function)
```

Inexact equality comparison: `true` if `norm(x-y) <= atol + rtol*max(norm(x), norm(y))`. The default `atol` is zero and the default `rtol` depends on the types of `x` and `y`. The keyword argument `nans` determines whether or not NaN values are considered equal (defaults to false).

For real or complex floating-point values, `rtol` defaults to `sqrt(eps(typeof(real(x-y))))`. This corresponds to requiring equality of about half of the significand digits. For other types, `rtol` defaults to zero.

`x` and `y` may also be arrays of numbers, in which case `norm` defaults to `vecnorm` but may be changed by passing a `norm::Function` keyword argument. (For numbers, `norm` is the same thing as `abs`.) When `x` and `y` are arrays, if `norm(x-y)` is not finite (i.e. `±Inf` or `NaN`), the comparison falls back to checking whether all elements of `x` and `y` are approximately equal component-wise.

The binary operator ≈ is equivalent to `isapprox` with the default arguments, and x   y is equivalent to `!isapprox(x,y)`.

```
julia> 0.1 ≈ (0.1 - 1e-10)
true

julia> isapprox(10, 11; atol = 2)
true

julia> isapprox([10.0^9, 1.0], [10.0^9, 2.0])
true
```

source

Base.sin – Function.

```
sin(x)
```

Compute sine of x, where x is in radians.

source

Base.cos – Function.

```
cos(x)
```

Compute cosine of x, where x is in radians.

source

Base.tan – Function.

```
tan(x)
```

Compute tangent of x, where x is in radians.

source

Base.Math.sind – Function.

```
sind(x)
```

Compute sine of x, where x is in degrees.

source

Base.Math.cosd – Function.

```
cosd(x)
```

Compute cosine of x, where x is in degrees.

source

Base.Math.tand – Function.

```
tand(x)
```

Compute tangent of x, where x is in degrees.

source

Base.Math.sinpi – Function.

```
sinpi(x)
```

Compute $\sin(\pi x)$ more accurately than `sin(pi*x)`, especially for large x.

source

Base.Math.cospi – Function.

```
cospi(x)
```

Compute $\cos(\pi x)$ more accurately than `cos(pi*x)`, especially for large x.

source

Base.sinh – Function.

```
sinh(x)
```

Compute hyperbolic sine of x.

source

Base.cosh – Function.

```
cosh(x)
```

Compute hyperbolic cosine of x.

source

Base.tanh – Function.

```
tanh(x)
```

Compute hyperbolic tangent of x.

source

Base.asin – Function.

```
asin(x)
```

Compute the inverse sine of x, where the output is in radians.

source

Base.acos – Function.

```
acos(x)
```

Compute the inverse cosine of x, where the output is in radians

source

Base.atan – Function.

```
atan(x)
```

Compute the inverse tangent of x, where the output is in radians.

source

Base.Math.atan2 – Function.

```
atan2(y, x)
```

Compute the inverse tangent of y/x, using the signs of both x and y to determine the quadrant of the return value.

source

Base.Math.asind – Function.

```
asind(x)
```

Compute the inverse sine of x, where the output is in degrees.

source

Base.Math.acosd – Function.

```
acosd(x)
```

Compute the inverse cosine of x, where the output is in degrees.

source

Base.Math.atand – Function.

```
atand(x)
```

Compute the inverse tangent of x, where the output is in degrees.

source

Base.Math.sec – Function.

```
sec(x)
```

Compute the secant of x, where x is in radians.

source

Base.Math.csc – Function.

```
csc(x)
```

Compute the cosecant of x, where x is in radians.

source

Base.Math.cot – Function.

```
cot(x)
```

Compute the cotangent of x, where x is in radians.

source

Base.Math.secd – Function.

```
secd(x)
```

Compute the secant of x, where x is in degrees.

source

Base.Math.cscd – Function.

```
cscd(x)
```

Compute the cosecant of x, where x is in degrees.

source

Base.Math.cotd – Function.

```
cotd(x)
```

Compute the cotangent of x, where x is in degrees.

source

Base.Math.asec – Function.

```
asec(x)
```

Compute the inverse secant of x, where the output is in radians.

source

Base.Math.acsc – Function.

```
acsc(x)
```

Compute the inverse cosecant of x, where the output is in radians.

source

Base.Math.acot – Function.

```
acot(x)
```

Compute the inverse cotangent of x, where the output is in radians.

source

Base.Math.asecd – Function.

```
asecd(x)
```

Compute the inverse secant of x, where the output is in degrees.

source

Base.Math.acscd – Function.

```
acscd(x)
```

Compute the inverse cosecant of x, where the output is in degrees.

source

Base.Math.acotd – Function.

```
acotd(x)
```

Compute the inverse cotangent of x, where the output is in degrees.

source

Base.Math.sech – Function.

```
sech(x)
```

Compute the hyperbolic secant of x

source

Base.Math.csch – Function.

```
csch(x)
```

Compute the hyperbolic cosecant of x.

source

Base.Math.coth – Function.

```
coth(x)
```

Compute the hyperbolic cotangent of x.

source

Base.asinh – Function.

> asinh(x)

Compute the inverse hyperbolic sine of x.

> source

Base.acosh – Function.

> acosh(x)

Compute the inverse hyperbolic cosine of x.

> source

Base.atanh – Function.

> atanh(x)

Compute the inverse hyperbolic tangent of x.

> source

Base.Math.asech – Function.

> asech(x)

Compute the inverse hyperbolic secant of x.

> source

Base.Math.acsch – Function.

> acsch(x)

Compute the inverse hyperbolic cosecant of x.

> source

Base.Math.acoth – Function.

> acoth(x)

Compute the inverse hyperbolic cotangent of x.

> source

Base.Math.sinc – Function.

> sinc(x)

Compute $\sin(\pi x)/(\pi x)$ if $x \neq 0$, and 1 if $x = 0$.

> source

Base.Math.cosc – Function.

> cosc(x)

Compute $\cos(\pi x)/x - \sin(\pi x)/(\pi x^2)$ if $x \neq 0$, and 0 if $x = 0$. This is the derivative of sinc(x).

> source

Base.Math.deg2rad – Function.

```
deg2rad(x)
```

Convert x from degrees to radians.

```
julia> deg2rad(90)
1.5707963267948966
```

source

Base.Math.rad2deg – Function.

```
rad2deg(x)
```

Convert x from radians to degrees.

```
julia> rad2deg(pi)
180.0
```

source

Base.Math.hypot – Function.

```
hypot(x, y)
```

Compute the hypotenuse $\sqrt{x^2 + y^2}$ avoiding overflow and underflow.

Examples

```
julia> a = 10^10;

julia> hypot(a, a)
1.4142135623730951e10

julia> √(a^2 + a^2) # a^2 overflows
ERROR: DomainError:
sqrt will only return a complex result if called with a complex argument. Try
↪   sqrt(complex(x)).
Stacktrace:
 [1] sqrt(::Int64) at ./math.jl:434
```

source

```
hypot(x...)
```

Compute the hypotenuse $\sqrt{\sum x_i^2}$ avoiding overflow and underflow.

source

Base.log – Method.

```
log(x)
```

Compute the natural logarithm of x. Throws DomainError for negative Real arguments. Use complex negative arguments to obtain complex results.

There is an experimental variant in the Base.Math.JuliaLibm module, which is typically faster and more accurate.

source

Base.log – Method.

```
log(b,x)
```

Compute the base b logarithm of x. Throws DomainError for negative Real arguments.

```
julia> log(4,8)
1.5

julia> log(4,2)
0.5
```

> **Note**
> If b is a power of 2 or 10, log2 or log10 should be used, as these will typically be faster and more accurate. For example,
>
> ```
> julia> log(100,1000000)
> 2.9999999999999996
>
> julia> log10(1000000)/2
> 3.0
> ```

source

Base.log2 – Function.

```
log2(x)
```

Compute the logarithm of x to base 2. Throws DomainError for negative Real arguments.

Example

```
julia> log2(4)
2.0

julia> log2(10)
3.321928094887362
```

source

Base.log10 – Function.

```
log10(x)
```

Compute the logarithm of x to base 10. Throws DomainError for negative Real arguments.

Example

```
julia> log10(100)
2.0

julia> log10(2)
0.3010299956639812
```

source

Base.log1p – Function.

```
log1p(x)
```

Accurate natural logarithm of 1+x. Throws `DomainError` for `Real` arguments less than -1.

There is an experimental variant in the `Base.Math.JuliaLibm` module, which is typically faster and more accurate.

Examples

```julia
julia> log1p(-0.5)
-0.6931471805599453

julia> log1p(0)
0.0
```

source

`Base.Math.frexp` – Function.

```
frexp(val)
```

Return `(x, exp)` such that x has a magnitude in the interval $[1/2, 1)$ or 0, and `val` is equal to $x \times 2^{exp}$.

source

`Base.exp` – Function.

```
exp(x)
```

Compute the natural base exponential of x, in other words e^x.

source

`Base.exp2` – Function.

```
exp2(x)
```

Compute 2^x.

```julia
julia> exp2(5)
32.0
```

source

`Base.exp10` – Function.

```
exp10(x)
```

Compute 10^x.

```julia
julia> exp10(2)
100.0

julia> exp10(0.2)
1.5848931924611136
```

source

`Base.Math.ldexp` – Function.

```
ldexp(x, n)
```

Compute $x \times 2^n$.

Example

```
julia> ldexp(5., 2)
20.0
```

`Base.Math.modf` – Function.

```
modf(x)
```

Return a tuple (fpart,ipart) of the fractional and integral parts of a number. Both parts have the same sign as the argument.

Example

```
julia> modf(3.5)
(0.5, 3.0)
```

`Base.expm1` – Function.

```
expm1(x)
```

Accurately compute $e^x - 1$.

`Base.round` – Method.

```
round([T,] x, [digits, [base]], [r::RoundingMode])
```

Rounds x to an integer value according to the provided `RoundingMode`, returning a value of the same type as x. When not specifying a rounding mode the global mode will be used (see `rounding`), which by default is round to the nearest integer (`RoundNearest` mode), with ties (fractional values of 0.5) being rounded to the nearest even integer.

```
julia> round(1.7)
2.0

julia> round(1.5)
2.0

julia> round(2.5)
2.0
```

The optional `RoundingMode` argument will change how the number gets rounded.

`round(T, x, [r::RoundingMode])` converts the result to type T, throwing an `InexactError` if the value is not representable.

`round(x, digits)` rounds to the specified number of digits after the decimal place (or before if negative). `round(x, digits, base)` rounds using a base other than 10.

```julia
julia> round(pi, 2)
3.14

julia> round(pi, 3, 2)
3.125
```

> **Note**
>
> Rounding to specified digits in bases other than 2 can be inexact when operating on binary floating point numbers. For example, the `Float64` value represented by 1.15 is actually *less* than 1.15, yet will be rounded to 1.2.
>
> ```julia
> julia> x = 1.15
> 1.15
>
> julia> @sprintf "%.20f" x
> "1.14999999999999991118"
>
> julia> x < 115//100
> true
>
> julia> round(x, 1)
> 1.2
> ```

source

Base.Rounding.RoundingMode – Type.

```
RoundingMode
```

A type used for controlling the rounding mode of floating point operations (via `rounding`/`setrounding` functions), or as optional arguments for rounding to the nearest integer (via the `round` function).

Currently supported rounding modes are:

- `RoundNearest` (default)
- `RoundNearestTiesAway`
- `RoundNearestTiesUp`
- `RoundToZero`
- RoundFromZero (`BigFloat` only)
- `RoundUp`
- `RoundDown`

source

Base.Rounding.RoundNearest – Constant.

```
RoundNearest
```

The default rounding mode. Rounds to the nearest integer, with ties (fractional values of 0.5) being rounded to the nearest even integer.

source

Base.Rounding.RoundNearestTiesAway – Constant.

> RoundNearestTiesAway

Rounds to nearest integer, with ties rounded away from zero (C/C++ round behaviour).

source

Base.Rounding.RoundNearestTiesUp – Constant.

> RoundNearestTiesUp

Rounds to nearest integer, with ties rounded toward positive infinity (Java/JavaScript round behaviour).

source

Base.Rounding.RoundToZero – Constant.

> RoundToZero

round using this rounding mode is an alias for trunc.

source

Base.Rounding.RoundUp – Constant.

> RoundUp

round using this rounding mode is an alias for ceil.

source

Base.Rounding.RoundDown – Constant.

> RoundDown

round using this rounding mode is an alias for floor.

source

Base.round – Method.

> round(z, RoundingModeReal, RoundingModeImaginary)

Returns the nearest integral value of the same type as the complex-valued z to z, breaking ties using the specified RoundingModes. The first RoundingMode is used for rounding the real components while the second is used for rounding the imaginary components.

source

Base.ceil – Function.

> ceil([T,] x, [digits, [base]])

ceil(x) returns the nearest integral value of the same type as x that is greater than or equal to x.

ceil(T, x) converts the result to type T, throwing an InexactError if the value is not representable.

digits and base work as for round.

source

Base.floor – Function.

```
floor([T,] x, [digits, [base]])
```

floor(x) returns the nearest integral value of the same type as x that is less than or equal to x.

floor(T, x) converts the result to type T, throwing an InexactError if the value is not representable.

digits and base work as for round.

source

Base.trunc – Function.

```
trunc([T,] x, [digits, [base]])
```

trunc(x) returns the nearest integral value of the same type as x whose absolute value is less than or equal to x.

trunc(T, x) converts the result to type T, throwing an InexactError if the value is not representable.

digits and base work as for round.

source

Base.unsafe_trunc – Function.

```
unsafe_trunc(T, x)
```

unsafe_trunc(T, x) returns the nearest integral value of type T whose absolute value is less than or equal to x. If the value is not representable by T, an arbitrary value will be returned.

source

Base.signif – Function.

```
signif(x, digits, [base])
```

Rounds (in the sense of round) x so that there are digits significant digits, under a base base representation, default 10. E.g., signif(123.456, 2) is 120.0, and signif(357.913, 4, 2) is 352.0.

source

Base.min – Function.

```
min(x, y, ...)
```

Return the minimum of the arguments. See also the minimum function to take the minimum element from a collection.

```
julia> min(2, 5, 1)
1
```

source

Base.max – Function.

```
max(x, y, ...)
```

Return the maximum of the arguments. See also the maximum function to take the maximum element from a collection.

```
julia> max(2, 5, 1)
5
```

source

Base.minmax – Function.

```
minmax(x, y)
```

Return (min(x,y), max(x,y)). See also: extrema that returns (minimum(x), maximum(x)).

```
julia> minmax('c','b')
('b', 'c')
```

source

Base.Math.clamp – Function.

```
clamp(x, lo, hi)
```

Return x if lo <= x <= hi. If x < lo, return lo. If x > hi, return hi. Arguments are promoted to a common type.

```
julia> clamp.([pi, 1.0, big(10.)], 2., 9.)
3-element Array{BigFloat,1}:
 3.141592653589793238462643383279502884197169399375105820974944592307816406286198
 2.000000000000000000000000000000000000000000000000000000000000000000000000000000
 9.000000000000000000000000000000000000000000000000000000000000000000000000000000
```

source

Base.Math.clamp! – Function.

```
clamp!(array::AbstractArray, lo, hi)
```

Restrict values in array to the specified range, in-place. See also clamp.

source

Base.abs – Function.

```
abs(x)
```

The absolute value of x.

When abs is applied to signed integers, overflow may occur, resulting in the return of a negative value. This overflow occurs only when abs is applied to the minimum representable value of a signed integer. That is, when x == typemin(typeof(x)), abs(x) == x < 0, not -x as might be expected.

```
julia> abs(-3)
3

julia> abs(1 + im)
1.4142135623730951

julia> abs(typemin(Int64))
-9223372036854775808
```

source

Base.Checked.checked_abs – Function.

> Base.checked_abs(x)

Calculates abs(x), checking for overflow errors where applicable. For example, standard two's complement signed integers (e.g. Int) cannot represent abs(typemin(Int)), thus leading to an overflow.

The overflow protection may impose a perceptible performance penalty.

source

Base.Checked.checked_neg – Function.

> Base.checked_neg(x)

Calculates -x, checking for overflow errors where applicable. For example, standard two's complement signed integers (e.g. Int) cannot represent -typemin(Int), thus leading to an overflow.

The overflow protection may impose a perceptible performance penalty.

source

Base.Checked.checked_add – Function.

> Base.checked_add(x, y)

Calculates x+y, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

Base.Checked.checked_sub – Function.

> Base.checked_sub(x, y)

Calculates x-y, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

Base.Checked.checked_mul – Function.

> Base.checked_mul(x, y)

Calculates x*y, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

Base.Checked.checked_div – Function.

> Base.checked_div(x, y)

Calculates div(x,y), checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

`Base.Checked.checked_rem` – Function.

```
Base.checked_rem(x, y)
```

Calculates x%y, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

`Base.Checked.checked_fld` – Function.

```
Base.checked_fld(x, y)
```

Calculates `fld(x,y)`, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

`Base.Checked.checked_mod` – Function.

```
Base.checked_mod(x, y)
```

Calculates `mod(x,y)`, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

`Base.Checked.checked_cld` – Function.

```
Base.checked_cld(x, y)
```

Calculates `cld(x,y)`, checking for overflow errors where applicable.

The overflow protection may impose a perceptible performance penalty.

source

`Base.Checked.add_with_overflow` – Function.

```
Base.add_with_overflow(x, y) -> (r, f)
```

Calculates `r = x+y`, with the flag f indicating whether overflow has occurred.

source

`Base.Checked.sub_with_overflow` – Function.

```
Base.sub_with_overflow(x, y) -> (r, f)
```

Calculates `r = x-y`, with the flag f indicating whether overflow has occurred.

source

`Base.Checked.mul_with_overflow` – Function.

```
Base.mul_with_overflow(x, y) -> (r, f)
```

Calculates `r = x*y`, with the flag f indicating whether overflow has occurred.

source

Base.abs2 – Function.

```
abs2(x)
```

Squared absolute value of x.

source

Base.copysign – Function.

```
copysign(x, y) -> z
```

Return z which has the magnitude of x and the same sign as y.

```
julia> copysign(1, -2)
-1

julia> copysign(-1, 2)
1
```

source

Base.sign – Function.

```
sign(x)
```

Return zero if x==0 and $x/|x|$ otherwise (i.e., ±1 for real x).

source

Base.signbit – Function.

```
signbit(x)
```

Returns true if the value of the sign of x is negative, otherwise false.

```
julia> signbit(-4)
true

julia> signbit(5)
false

julia> signbit(5.5)
false

julia> signbit(-4.1)
true
```

source

Base.flipsign – Function.

```
flipsign(x, y)
```

Return x with its sign flipped if y is negative. For example abs(x) = flipsign(x,x).

source

Base.sqrt – Function.

```
sqrt(x)
```

Return $\sqrt{x}$. Throws DomainError for negative Real arguments. Use complex negative arguments instead. The prefix operator √ is equivalent to sqrt.

source

Base.isqrt – Function.

```
isqrt(n::Integer)
```

Integer square root: the largest integer m such that m*m <= n.

```
julia> isqrt(5)
2
```

source

Base.Math.cbrt – Function.

```
cbrt(x::Real)
```

Return the cube root of x, i.e. $x^{1/3}$. Negative values are accepted (returning the negative real root when $x < 0$).

The prefix operator is equivalent to cbrt.

```
julia> cbrt(big(27))
3.000000000000000000000000000000000000000000000000000000000000000000000000000000
```

source

Base.real – Method.

```
real(z)
```

Return the real part of the complex number z.

```
julia> real(1 + 3im)
1
```

source

Base.imag – Function.

```
imag(z)
```

Return the imaginary part of the complex number z.

```
julia> imag(1 + 3im)
3
```

source

Base.reim – Function.

```
reim(z)
```

Return both the real and imaginary parts of the complex number z.

```julia
julia> reim(1 + 3im)
(1, 3)
```

source

Base.conj – Function.

```
conj(v::RowVector)
```

Returns a `ConjArray` lazy view of the input, where each element is conjugated.

Example

```julia
julia> v = [1+im, 1-im].'
1×2 RowVector{Complex{Int64},Array{Complex{Int64},1}}:
 1+1im  1-1im

julia> conj(v)
1×2 RowVector{Complex{Int64},ConjArray{Complex{Int64},1,Array{Complex{Int64},1}}}:
 1-1im  1+1im
```

source

```
conj(z)
```

Compute the complex conjugate of a complex number z.

```julia
julia> conj(1 + 3im)
1 - 3im
```

source

Base.angle – Function.

```
angle(z)
```

Compute the phase angle in radians of a complex number z.

source

Base.cis – Function.

```
cis(z)
```

Return $\exp(iz)$.

source

Base.binomial – Function.

```
binomial(n,k)
```

Number of ways to choose k out of n items.

source

Base.`factorial` – Function.

```
factorial(n)
```

Factorial of n. If n is an `Integer`, the factorial is computed as an integer (promoted to at least 64 bits). Note that this may overflow if n is not small, but you can use `factorial(big(n))` to compute the result exactly in arbitrary precision. If n is not an `Integer`, `factorial(n)` is equivalent to `gamma(n+1)`.

```julia
julia> factorial(6)
720

julia> factorial(21)
ERROR: OverflowError()
[...]

julia> factorial(21.0)
5.109094217170944e19

julia> factorial(big(21))
51090942171709440000
```

source

Base.`gcd` – Function.

```
gcd(x,y)
```

Greatest common (positive) divisor (or zero if x and y are both zero).

```julia
julia> gcd(6,9)
3

julia> gcd(6,-9)
3
```

source

Base.`lcm` – Function.

```
lcm(x,y)
```

Least common (non-negative) multiple.

```julia
julia> lcm(2,3)
6

julia> lcm(-2,3)
6
```

source

Base.`gcdx` – Function.

```
gcdx(x,y)
```

Computes the greatest common (positive) divisor of x and y and their Bézout coefficients, i.e. the integer coefficients u and v that satisfy $ux + vy = d = gcd(x, y)$. $gcdx(x, y)$ returns (d, u, v).

```julia
julia> gcdx(12, 42)
(6, -3, 1)
```

```julia
julia> gcdx(240, 46)
(2, -9, 47)
```

> **Note**
>
> Bézout coefficients are *not* uniquely defined. gcdx returns the minimal Bézout coefficients that are computed by the extended Euclidean algorithm. (Ref: D. Knuth, TAoCP, 2/e, p. 325, Algorithm X.) For signed integers, these coefficients u and v are minimal in the sense that $|u| < |y/d|$ and $|v| < |x/d|$. Furthermore, the signs of u and v are chosen so that d is positive. For unsigned integers, the coefficients u and v might be near their typemax, and the identity then holds only via the unsigned integers' modulo arithmetic.

source

Base.ispow2 – Function.

```
ispow2(n::Integer) -> Bool
```

Test whether n is a power of two.

```julia
julia> ispow2(4)
true
```

```julia
julia> ispow2(5)
false
```

source

Base.nextpow2 – Function.

```
nextpow2(n::Integer)
```

The smallest power of two not less than n. Returns 0 for n==0, and returns -nextpow2(-n) for negative arguments.

```julia
julia> nextpow2(16)
16
```

```julia
julia> nextpow2(17)
32
```

source

Base.prevpow2 – Function.

```
prevpow2(n::Integer)
```

The largest power of two not greater than n. Returns 0 for n==0, and returns -prevpow2(-n) for negative arguments.

```julia
julia> prevpow2(5)
4
```

source

Base.nextpow – Function.

```
nextpow(a, x)
```

The smallest a^n not less than x, where n is a non-negative integer. a must be greater than 1, and x must be greater than 0.

source

Base.prevpow – Function.

```
prevpow(a, x)
```

The largest a^n not greater than x, where n is a non-negative integer. a must be greater than 1, and x must not be less than 1.

source

Base.nextprod – Function.

```
nextprod([k_1, k_2,...], n)
```

Next integer greater than or equal to n that can be written as $\prod k_i^{p_i}$ for integers p_1, p_2, etc.

```
julia> nextprod([2, 3], 105)
108

julia> 2^2 * 3^3
108
```

source

Base.invmod – Function.

```
invmod(x,m)
```

Take the inverse of x modulo m: y such that $xy = 1 \pmod m$, with $div(x, y) = 0$. This is undefined for $m = 0$, or if $gcd(x, m) \neq 1$.

```
julia> invmod(2,5)
3

julia> invmod(2,3)
2

julia> invmod(5,6)
5
```

source

Base.powermod – Function.

```
powermod(x::Integer, p::Integer, m)
```

Compute $x^p \pmod m$.

source

Base.Math.gamma – Function.

```
gamma(x)
```

Compute the gamma function of x.

source

Base.Math.lgamma – Function.

```
lgamma(x)
```

Compute the logarithm of the absolute value of gamma for Real x, while for Complex x compute the principal branch cut of the logarithm of gamma(x) (defined for negative real(x) by analytic continuation from positive real(x)).

source

Base.Math.lfact – Function.

```
lfact(x)
```

Compute the logarithmic factorial of a nonnegative integer x. Equivalent to lgamma of x + 1, but lgamma extends this function to non-integer x.

source

Base.Math.beta – Function.

```
beta(x, y)
```

Euler integral of the first kind $\mathrm{B}(x,y) = \Gamma(x)\Gamma(y)/\Gamma(x+y)$.

source

Base.Math.lbeta – Function.

```
lbeta(x, y)
```

Natural logarithm of the absolute value of the beta function $\log(|\mathrm{B}(x,y)|)$.

source

Base.ndigits – Function.

```
ndigits(n::Integer, b::Integer=10)
```

Compute the number of digits in integer n written in base b.

source

Base.widemul – Function.

```
widemul(x, y)
```

Multiply x and y, giving the result as a larger type.

```
julia> widemul(Float32(3.), 4.)
1.200000000000000000000000000000000000000000000000000000000000000000000000000000e+01
```

source

Base.Math.@evalpoly – Macro.

```
@evalpoly(z, c...)
```

Evaluate the polynomial $\sum_k c[k]z^{k-1}$ for the coefficients c[1], c[2], ...; that is, the coefficients are given in ascending order by power of z. This macro expands to efficient inline code that uses either Horner's method or, for complex z, a more efficient Goertzel-like algorithm.

```
julia> @evalpoly(3, 1, 0, 1)
10

julia> @evalpoly(2, 1, 0, 1)
5

julia> @evalpoly(2, 1, 1, 1)
7
```

source

47.3   Statistics

Base.mean – Function.

```
mean(f::Function, v)
```

Apply the function f to each element of v and take the mean.

```
julia> mean(√, [1, 2, 3])
1.3820881233139908

julia> mean([√1, √2, √3])
1.3820881233139908
```

source

```
mean(v[, region])
```

Compute the mean of whole array v, or optionally along the dimensions in `region`.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

Base.mean! – Function.

```
mean!(r, v)
```

Compute the mean of v over the singleton dimensions of r, and write results to r.

source

Base.std – Function.

```
std(v[, region]; corrected::Bool=true, mean=nothing)
```

Compute the sample standard deviation of a vector or array `v`, optionally along dimensions in `region`. The algorithm returns an estimator of the generative distribution's standard deviation under the assumption that each entry of `v` is an IID drawn from that generative distribution. This computation is equivalent to calculating `sqrt(sum((v - mean(v)).^2) / (length(v) - 1))`. A pre-computed mean may be provided. If `corrected` is `true`, then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `length(x)`.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

`Base.stdm` – Function.

```
stdm(v, m::Number; corrected::Bool=true)
```

Compute the sample standard deviation of a vector `v` with known mean `m`. If `corrected` is `true`, then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `length(x)`.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

`Base.var` – Function.

```
var(v[, region]; corrected::Bool=true, mean=nothing)
```

Compute the sample variance of a vector or array `v`, optionally along dimensions in `region`. The algorithm will return an estimator of the generative distribution's variance under the assumption that each entry of `v` is an IID drawn from that generative distribution. This computation is equivalent to calculating `sum(abs2, v - mean(v)) / (length(v) - 1)`. If `corrected` is `true`, then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `length(x)`. The mean `mean` over the region may be provided.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

`Base.varm` – Function.

```
varm(v, m[, region]; corrected::Bool=true)
```

Compute the sample variance of a collection `v` with known mean(s) `m`, optionally over `region`. `m` may contain means for each dimension of `v`. If `corrected` is `true`, then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `length(x)`.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

Base.middle – Function.

```
middle(x)
```

Compute the middle of a scalar value, which is equivalent to x itself, but of the type of `middle(x, x)` for consistency.

source

```
middle(x, y)
```

Compute the middle of two reals x and y, which is equivalent in both value and type to computing their mean ((`x + y) / 2`).

source

```
middle(range)
```

Compute the middle of a range, which consists of computing the mean of its extrema. Since a range is sorted, the mean is performed with the first and last element.

```julia
julia> middle(1:10)
5.5
```

source

```
middle(a)
```

Compute the middle of an array a, which consists of finding its extrema and then computing their mean.

```julia
julia> a = [1,2,3.6,10.9]
4-element Array{Float64,1}:
  1.0
  2.0
  3.6
 10.9

julia> middle(a)
5.95
```

source

Base.median – Function.

```
median(v[, region])
```

Compute the median of an entire array v, or, optionally, along the dimensions in `region`. For an even number of elements no exact median element exists, so the result is equivalent to calculating mean of two median elements.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended.

source

Base.median! – Function.

```
median!(v)
```

Like median, but may overwrite the input vector.

source

Base.quantile – Function.

```
quantile(v, p; sorted=false)
```

Compute the quantile(s) of a vector v at a specified probability or vector p. The keyword argument `sorted` indicates whether v can be assumed to be sorted.

The p should be on the interval [0,1], and v should not have any NaN values.

Quantiles are computed via linear interpolation between the points `((k-1)/(n-1), v[k])`, for `k = 1:n` where `n = length(v)`. This corresponds to Definition 7 of Hyndman and Fan (1996), and is the same as the R default.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended. `quantile` will throw an `ArgumentError` in the presence of NaN values in the data array.

- Hyndman, R.J and Fan, Y. (1996) "Sample Quantiles in Statistical Packages", *The American Statistician*, Vol. 50, No. 4, pp. 361-365

source

Base.quantile! – Function.

```
quantile!([q, ] v, p; sorted=false)
```

Compute the quantile(s) of a vector v at the probabilities p, with optional output into array q (if not provided, a new output array is created). The keyword argument `sorted` indicates whether v can be assumed to be sorted; if `false` (the default), then the elements of v may be partially sorted.

The elements of p should be on the interval [0,1], and v should not have any NaN values.

Quantiles are computed via linear interpolation between the points `((k-1)/(n-1), v[k])`, for `k = 1:n` where `n = length(v)`. This corresponds to Definition 7 of Hyndman and Fan (1996), and is the same as the R default.

> **Note**
>
> Julia does not ignore NaN values in the computation. For applications requiring the handling of missing data, the `DataArrays.jl` package is recommended. `quantile!` will throw an `ArgumentError` in the presence of NaN values in the data array.

- Hyndman, R.J and Fan, Y. (1996) "Sample Quantiles in Statistical Packages", *The American Statistician*, Vol. 50, No. 4, pp. 361-365

source

Base.cov – Function.

```
cov(x[, corrected=true])
```

Compute the variance of the vector x. If `corrected` is `true` (the default) then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `length(x)`.

source

```
cov(X[, vardim=1, corrected=true])
```

Compute the covariance matrix of the matrix X along the dimension `vardim`. If `corrected` is `true` (the default) then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `size(X, vardim)`.

source

```
cov(x, y[, corrected=true])
```

Compute the covariance between the vectors x and y. If `corrected` is `true` (the default), computes $\frac{1}{n-1}\sum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})^*$ where $*$ denotes the complex conjugate and n = `length(x)` = `length(y)`. If `corrected` is `false`, computes $rac1nsum_{i=1}^{n}(x_i - \bar{x})(y_i - \bar{y})^*$.

source

```
cov(X, Y[, vardim=1, corrected=true])
```

Compute the covariance between the vectors or matrices X and Y along the dimension `vardim`. If `corrected` is `true` (the default) then the sum is scaled with n-1, whereas the sum is scaled with n if `corrected` is `false` where n = `size(X, vardim)` = `size(Y, vardim)`.

source

`Base.cor` – Function.

```
cor(x)
```

Return the number one.

source

```
cor(X[, vardim=1])
```

Compute the Pearson correlation matrix of the matrix X along the dimension `vardim`.

source

```
cor(x, y)
```

Compute the Pearson correlation between the vectors x and y.

source

```
cor(X, Y[, vardim=1])
```

Compute the Pearson correlation between the vectors or matrices X and Y along the dimension `vardim`.

source

47.4 Signal Processing

Fast Fourier transform (FFT) functions in Julia are implemented by calling functions from FFTW.

`Base.DFT.fft` – Function.

```
fft(A [, dims])
```

Performs a multidimensional FFT of the array A. The optional `dims` argument specifies an iterable subset of dimensions (e.g. an integer, range, tuple, or array) to transform along. Most efficient if the size of A along the transformed dimensions is a product of small primes; see `nextprod()`. See also `plan_fft()` for even greater efficiency.

A one-dimensional FFT computes the one-dimensional discrete Fourier transform (DFT) as defined by

$$\mathrm{DFT}(A)[k] = \sum_{n=1}^{\mathrm{length}(A)} \exp\left(-i\frac{2\pi(n-1)(k-1)}{\mathrm{length}(A)}\right) A[n].$$

A multidimensional FFT simply performs this operation along each transformed dimension of A.

> **Note**
> - Julia starts FFTW up with 1 thread by default. Higher performance is usually possible by increasing number of threads. Use `FFTW.set_num_threads(Sys.CPU_CORES)` to use as many threads as cores on your system.
> - This performs a multidimensional FFT by default. FFT libraries in other languages such as Python and Octave perform a one-dimensional FFT along the first non-singleton dimension of the array. This is worth noting while performing comparisons. For more details, refer to the Noteworthy Differences from other Languages section of the manual.

source

Base.DFT.fft! – Function.

```
fft!(A [, dims])
```

Same as `fft`, but operates in-place on A, which must be an array of complex floating-point numbers.

source

Base.DFT.ifft – Function.

```
ifft(A [, dims])
```

Multidimensional inverse FFT.

A one-dimensional inverse FFT computes

$$\mathrm{IDFT}(A)[k] = \frac{1}{\mathrm{length}(A)} \sum_{n=1}^{\mathrm{length}(A)} \exp\left(+i\frac{2\pi(n-1)(k-1)}{\mathrm{length}(A)}\right) A[n].$$

A multidimensional inverse FFT simply performs this operation along each transformed dimension of A.

source

Base.DFT.ifft! – Function.

```
ifft!(A [, dims])
```

Same as `ifft`, but operates in-place on A.

source

`Base.DFT.bfft` – Function.

```
bfft(A [, dims])
```

Similar to `ifft`, but computes an unnormalized inverse (backward) transform, which must be divided by the product of the sizes of the transformed dimensions in order to obtain the inverse. (This is slightly more efficient than `ifft` because it omits a scaling step, which in some applications can be combined with other computational steps elsewhere.)

$$\mathrm{BDFT}(A)[k] = \mathrm{length}(A)\,\mathrm{IDFT}(A)[k]$$

source

`Base.DFT.bfft!` – Function.

```
bfft!(A [, dims])
```

Same as `bfft`, but operates in-place on A.

source

`Base.DFT.plan_fft` – Function.

```
plan_fft(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Pre-plan an optimized FFT along given dimensions (`dims`) of arrays matching the shape and type of A. (The first two arguments have the same meaning as for `fft`.) Returns an object P which represents the linear operator computed by the FFT, and which contains all of the information needed to compute `fft(A, dims)` quickly.

To apply P to an array A, use P * A; in general, the syntax for applying plans is much like that of matrices. (A plan can only be applied to arrays of the same size as the A for which the plan was created.) You can also apply a plan with a preallocated output array Â by calling `A_mul_B!(Â, plan, A)`. (For `A_mul_B!`, however, the input array A must be a complex floating-point array like the output Â.) You can compute the inverse-transform plan by `inv(P)` and apply the inverse plan with P \ Â (the inverse plan is cached and reused for subsequent calls to `inv` or `\`), and apply the inverse plan to a pre-allocated output array A with `A_ldiv_B!(A, P, Â)`.

The `flags` argument is a bitwise-or of FFTW planner flags, defaulting to `FFTW.ESTIMATE`. e.g. passing `FFTW.MEASURE` or `FFTW.PATIENT` will instead spend several seconds (or more) benchmarking different possible FFT algorithms and picking the fastest one; see the FFTW manual for more information on planner flags. The optional `timelimit` argument specifies a rough upper bound on the allowed planning time, in seconds. Passing `FFTW.MEASURE` or `FFTW.PATIENT` may cause the input array A to be overwritten with zeros during plan creation.

`plan_fft!` is the same as `plan_fft` but creates a plan that operates in-place on its argument (which must be an array of complex floating-point numbers). `plan_ifft` and so on are similar but produce plans that perform the equivalent of the inverse transforms `ifft` and so on.

source

`Base.DFT.plan_ifft` – Function.

```
plan_ifft(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Same as `plan_fft`, but produces a plan that performs inverse transforms `ifft`.

> source

Base.DFT.plan_bfft – Function.

```
plan_bfft(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Same as `plan_fft`, but produces a plan that performs an unnormalized backwards transform `bfft`.

> source

Base.DFT.plan_fft! – Function.

```
plan_fft!(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Same as `plan_fft`, but operates in-place on A.

> source

Base.DFT.plan_ifft! – Function.

```
plan_ifft!(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Same as `plan_ifft`, but operates in-place on A.

> source

Base.DFT.plan_bfft! – Function.

```
plan_bfft!(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Same as `plan_bfft`, but operates in-place on A.

> source

Base.DFT.rfft – Function.

```
rfft(A [, dims])
```

Multidimensional FFT of a real array A, exploiting the fact that the transform has conjugate symmetry in order to save roughly half the computational time and storage costs compared with `fft`. If A has size `(n_1,  ..., n_d)`, the result has size `(div(n_1,2)+1,  ..., n_d)`.

The optional `dims` argument specifies an iterable subset of one or more dimensions of A to transform, similar to `fft`. Instead of (roughly) halving the first dimension of A in the result, the `dims[1]` dimension is (roughly) halved in the same way.

> source

Base.DFT.irfft – Function.

```
irfft(A, d [, dims])
```

Inverse of `rfft`: for a complex array A, gives the corresponding real array whose FFT yields A in the first half. As for `rfft`, dims is an optional subset of dimensions to transform, defaulting to `1:ndims(A)`.

d is the length of the transformed real array along the `dims[1]` dimension, which must satisfy `div(d,2)+1 == size(A,dims[1])`. (This parameter cannot be inferred from `size(A)` since both `2*size(A,dims[1])-2` as well as `2*size(A,dims[1])-1` are valid sizes for the transformed real array.)

> source

Base.DFT.brfft – Function.

```
brfft(A, d [, dims])
```

Similar to irfft but computes an unnormalized inverse transform (similar to bfft), which must be divided by the product of the sizes of the transformed dimensions (of the real output array) in order to obtain the inverse transform.

source

Base.DFT.plan_rfft – Function.

```
plan_rfft(A [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Pre-plan an optimized real-input FFT, similar to plan_fft except for rfft instead of fft. The first two arguments, and the size of the transformed result, are the same as for rfft.

source

Base.DFT.plan_brfft – Function.

```
plan_brfft(A, d [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Pre-plan an optimized real-input unnormalized transform, similar to plan_rfft except for brfft instead of rfft. The first two arguments and the size of the transformed result, are the same as for brfft.

source

Base.DFT.plan_irfft – Function.

```
plan_irfft(A, d [, dims]; flags=FFTW.ESTIMATE;  timelimit=Inf)
```

Pre-plan an optimized inverse real-input FFT, similar to plan_rfft except for irfft and brfft, respectively. The first three arguments have the same meaning as for irfft.

source

Base.DFT.FFTW.dct – Function.

```
dct(A [, dims])
```

Performs a multidimensional type-II discrete cosine transform (DCT) of the array A, using the unitary normalization of the DCT. The optional dims argument specifies an iterable subset of dimensions (e.g. an integer, range, tuple, or array) to transform along. Most efficient if the size of A along the transformed dimensions is a product of small primes; see nextprod. See also plan_dct for even greater efficiency.

source

Base.DFT.FFTW.dct! – Function.

```
dct!(A [, dims])
```

Same as dct!, except that it operates in-place on A, which must be an array of real or complex floating-point values.

source

Base.DFT.FFTW.idct – Function.

```
idct(A [, dims])
```

Computes the multidimensional inverse discrete cosine transform (DCT) of the array A (technically, a type-III DCT with the unitary normalization). The optional `dims` argument specifies an iterable subset of dimensions (e.g. an integer, range, tuple, or array) to transform along. Most efficient if the size of A along the transformed dimensions is a product of small primes; see `nextprod`. See also `plan_idct` for even greater efficiency.

source

`Base.DFT.FFTW.idct!` – Function.

```
idct!(A [, dims])
```

Same as `idct!`, but operates in-place on A.

source

`Base.DFT.FFTW.plan_dct` – Function.

```
plan_dct(A [, dims [, flags [, timelimit]]])
```

Pre-plan an optimized discrete cosine transform (DCT), similar to `plan_fft` except producing a function that computes `dct`. The first two arguments have the same meaning as for `dct`.

source

`Base.DFT.FFTW.plan_dct!` – Function.

```
plan_dct!(A [, dims [, flags [, timelimit]]])
```

Same as `plan_dct`, but operates in-place on A.

source

`Base.DFT.FFTW.plan_idct` – Function.

```
plan_idct(A [, dims [, flags [, timelimit]]])
```

Pre-plan an optimized inverse discrete cosine transform (DCT), similar to `plan_fft` except producing a function that computes `idct`. The first two arguments have the same meaning as for `idct`.

source

`Base.DFT.FFTW.plan_idct!` – Function.

```
plan_idct!(A [, dims [, flags [, timelimit]]])
```

Same as `plan_idct`, but operates in-place on A.

source

`Base.DFT.fftshift` – Method.

```
fftshift(x)
```

Swap the first and second halves of each dimension of x.

source

`Base.DFT.fftshift` – Method.

```
fftshift(x,dim)
```

Swap the first and second halves of the given dimension or iterable of dimensions of array x.

source

Base.DFT.ifftshift – Function.

```
ifftshift(x, [dim])
```

Undoes the effect of `fftshift`.

source

Base.DSP.filt – Function.

```
filt(b, a, x, [si])
```

Apply filter described by vectors a and b to vector x, with an optional initial filter state vector si (defaults to zeros).

source

Base.DSP.filt! – Function.

```
filt!(out, b, a, x, [si])
```

Same as `filt` but writes the result into the `out` argument, which may alias the input x to modify it in-place.

source

Base.DSP.deconv – Function.

```
deconv(b,a) -> c
```

Construct vector c such that `b = conv(a,c) + r`. Equivalent to polynomial division.

source

Base.DSP.conv – Function.

```
conv(u,v)
```

Convolution of two vectors. Uses FFT algorithm.

source

Base.DSP.conv2 – Function.

```
conv2(u,v,A)
```

2-D convolution of the matrix A with the 2-D separable kernel generated by the vectors u and v. Uses 2-D FFT algorithm.

source

```
conv2(B,A)
```

2-D convolution of the matrix B with the matrix A. Uses 2-D FFT algorithm.

source

Base.DSP.xcorr – Function.

```
xcorr(u,v)
```

Compute the cross-correlation of two vectors.

source

The following functions are defined within the `Base.FFTW` module.

`Base.DFT.FFTW.r2r` – Function.

```
r2r(A, kind [, dims])
```

Performs a multidimensional real-input/real-output (r2r) transform of type `kind` of the array `A`, as defined in the FFTW manual. `kind` specifies either a discrete cosine transform of various types (`FFTW.REDFT00`, `FFTW.REDFT01`, `FFTW.REDFT10`, or `FFTW.REDFT11`), a discrete sine transform of various types (`FFTW.RODFT00`, `FFTW.RODFT01`, `FFTW.RODFT10`, or `FFTW.RODFT11`), a real-input DFT with halfcomplex-format output (`FFTW.R2HC` and its inverse `FFTW.HC2R`), or a discrete Hartley transform (`FFTW.DHT`). The `kind` argument may be an array or tuple in order to specify different transform types along the different dimensions of `A`; `kind[end]` is used for any unspecified dimensions. See the FFTW manual for precise definitions of these transform types, at http://www.fftw.org/doc.

The optional `dims` argument specifies an iterable subset of dimensions (e.g. an integer, range, tuple, or array) to transform along. `kind[i]` is then the transform type for `dims[i]`, with `kind[end]` being used for `i > length(kind)`.

See also `plan_r2r` to pre-plan optimized r2r transforms.

source

`Base.DFT.FFTW.r2r!` – Function.

```
r2r!(A, kind [, dims])
```

Same as `r2r`, but operates in-place on A, which must be an array of real or complex floating-point numbers.

source

`Base.DFT.FFTW.plan_r2r` – Function.

```
plan_r2r(A, kind [, dims [, flags [, timelimit]]])
```

Pre-plan an optimized r2r transform, similar to `plan_fft` except that the transforms (and the first three arguments) correspond to `r2r` and `r2r!`, respectively.

source

`Base.DFT.FFTW.plan_r2r!` – Function.

```
plan_r2r!(A, kind [, dims [, flags [, timelimit]]])
```

Similar to `plan_fft`, but corresponds to `r2r!`.

source

Chapter 48

Numbers

48.1   Standard Numeric Types

Abstract number types

Core.Number – Type.

> Number

Abstract supertype for all number types.

source

Core.Real – Type.

> Real <: Number

Abstract supertype for all real numbers.

source

Core.AbstractFloat – Type.

> AbstractFloat <: Real

Abstract supertype for all floating point numbers.

source

Core.Integer – Type.

> Integer <: Real

Abstract supertype for all integers.

source

Core.Signed – Type.

> Signed <: Integer

Abstract supertype for all signed integers.

source

`Core.Unsigned` – Type.

> `Unsigned <: Integer`

Abstract supertype for all unsigned integers.

> source

Concrete number types

`Core.Float16` – Type.

> `Float16 <: AbstractFloat`

16-bit floating point number type.

> source

`Core.Float32` – Type.

> `Float32 <: AbstractFloat`

32-bit floating point number type.

> source

`Core.Float64` – Type.

> `Float64 <: AbstractFloat`

64-bit floating point number type.

> source

`Base.MPFR.BigFloat` – Type.

> `BigFloat <: AbstractFloat`

Arbitrary precision floating point number type.

> source

`Core.Bool` – Type.

> `Bool <: Integer`

Boolean type.

> source

`Core.Int8` – Type.

> `Int8 <: Signed`

8-bit signed integer type.

> source

`Core.UInt8` – Type.

> `UInt8 <: Unsigned`

8-bit unsigned integer type.

source

Core.Int16 – Type.

> Int16 <: Signed

16-bit signed integer type.

source

Core.UInt16 – Type.

> UInt16 <: Unsigned

16-bit unsigned integer type.

source

Core.Int32 – Type.

> Int32 <: Signed

32-bit signed integer type.

source

Core.UInt32 – Type.

> UInt32 <: Unsigned

32-bit unsigned integer type.

source

Core.Int64 – Type.

> Int64 <: Signed

64-bit signed integer type.

source

Core.UInt64 – Type.

> UInt64 <: Unsigned

64-bit unsigned integer type.

source

Core.Int128 – Type.

> Int128 <: Signed

128-bit signed integer type.

source

Core.UInt128 – Type.

> UInt128 <: Unsigned

128-bit unsigned integer type.

source

Base.GMP.BigInt – Type.

> BigInt <: Integer

Arbitrary precision integer type.

source

Base.Complex – Type.

> Complex{T<:Real} <: Number

Complex number type with real and imaginary part of type T.

Complex32, Complex64 and Complex128 are aliases for Complex{Float16}, Complex{Float32} and Complex{Float64} respectively.

source

Base.Rational – Type.

> Rational{T<:Integer} <: Real

Rational number type, with numerator and denominator of type T.

source

Base.Irrational – Type.

> Irrational <: Real

Irrational number type.

source

48.2   Data Formats

Base.bin – Function.

> bin(n, pad::Int=1)

Convert an integer to a binary string, optionally specifying a number of digits to pad to.

```julia
julia> bin(10,2)
"1010"

julia> bin(10,8)
"00001010"
```

source

Base.hex – Function.

```
hex(n, pad::Int=1)
```

Convert an integer to a hexadecimal string, optionally specifying a number of digits to pad to.

source

Base.dec – Function.

```
dec(n, pad::Int=1)
```

Convert an integer to a decimal string, optionally specifying a number of digits to pad to.

source

Base.oct – Function.

```
oct(n, pad::Int=1)
```

Convert an integer to an octal string, optionally specifying a number of digits to pad to.

source

Base.base – Function.

```
base(base::Integer, n::Integer, pad::Integer=1)
```

Convert an integer n to a string in the given base, optionally specifying a number of digits to pad to.

```
julia> base(13,5,4)
"0005"

julia> base(5,13,4)
"0023"
```

source

Base.digits – Function.

```
digits([T<:Integer], n::Integer, base::T=10, pad::Integer=1)
```

Returns an array with element type T (default Int) of the digits of n in the given base, optionally padded with zeros to a specified size. More significant digits are at higher indexes, such that n == sum([digits[k]*base^(k-1) for k=1:length(digits)]).

source

Base.digits! – Function.

```
digits!(array, n::Integer, base::Integer=10)
```

Fills an array of the digits of n in the given base. More significant digits are at higher indexes. If the array length is insufficient, the least significant digits are filled up to the array length. If the array length is excessive, the excess portion is filled with zeros.

source

Base.bits – Function.

```
bits(n)
```

A string giving the literal bit representation of a number.

```
julia> bits(4)
"0000000000000000000000000000000000000000000000000000000000000100"

julia> bits(2.2)
"0100000000000001100110011001100110011001100110011001100110011010"
```

source

Base.parse – Method.

```
parse(type, str, [base])
```

Parse a string as a number. If the type is an integer type, then a base can be specified (the default is 10). If the type is a floating point type, the string is parsed as a decimal floating point number. If the string does not contain a valid number, an error is raised.

source

Base.tryparse – Function.

```
tryparse(type, str, [base])
```

Like parse, but returns a Nullable of the requested type. The result will be null if the string does not contain a valid number.

source

Base.big – Function.

```
big(x)
```

Convert a number to a maximum precision representation (typically BigInt or BigFloat). See BigFloat for information about some pitfalls with floating-point numbers.

source

Base.signed – Function.

```
signed(x)
```

Convert a number to a signed integer. If the argument is unsigned, it is reinterpreted as signed without checking for overflow.

source

Base.unsigned – Function.

```
unsigned(x) -> Unsigned
```

Convert a number to an unsigned integer. If the argument is signed, it is reinterpreted as unsigned without checking for negative values.

source

Base.float – Method.

```
float(x)
```

Convert a number or array to a floating point data type. When passed a string, this function is equivalent to `parse(Float64, x)`.

source

`Base.Math.significand` – Function.

```
significand(x)
```

Extract the `significand(s)` (a.k.a. mantissa), in binary representation, of a floating-point number. If x is a non-zero finite number, then the result will be a number of the same type on the interval $[1, 2)$. Otherwise x is returned.

Examples

```julia
julia> significand(15.2)/15.2
0.125

julia> significand(15.2)*8
15.2
```

source

`Base.Math.exponent` – Function.

```
exponent(x) -> Int
```

Get the exponent of a normalized floating-point number.

source

`Base.complex` – Method.

```
complex(r, [i])
```

Convert real numbers or arrays to complex. `i` defaults to zero.

source

`Base.bswap` – Function.

```
bswap(n)
```

Byte-swap an integer. Flip the bits of its binary representation.

```julia
julia> a = bswap(4)
288230376151711744

julia> bswap(a)
4

julia> bin(1)
"1"

julia> bin(bswap(1))
"1000000000000000000000000000000000000000000000000000000000000000"
```

source

Base.num2hex – Function.

```
num2hex(f)
```

Get a hexadecimal string of the binary representation of a floating point number.

```
julia> num2hex(2.2)
"400199999999999a"
```

source

Base.hex2num – Function.

```
hex2num(str)
```

Convert a hexadecimal string to the floating point number it represents.

source

Base.hex2bytes – Function.

```
hex2bytes(s::AbstractString)
```

Convert an arbitrarily long hexadecimal string to its binary representation. Returns an `Array{UInt8,1}`, i.e. an array of bytes.

```
julia> a = hex(12345)
"3039"

julia> hex2bytes(a)
2-element Array{UInt8,1}:
 0x30
 0x39
```

source

Base.bytes2hex – Function.

```
bytes2hex(bin_arr::Array{UInt8, 1}) -> String
```

Convert an array of bytes to its hexadecimal representation. All characters are in lower-case.

```
julia> a = hex(12345)
"3039"

julia> b = hex2bytes(a)
2-element Array{UInt8,1}:
 0x30
 0x39

julia> bytes2hex(b)
"3039"
```

source

48.3 General Number Functions and Constants

`Base.one` – Function.

```
one(x)
one(T::type)
```

Return a multiplicative identity for x: a value such that `one(x)*x == x*one(x) == x`. Alternatively `one(T)` can take a type T, in which case one returns a multiplicative identity for any x of type T.

If possible, `one(x)` returns a value of the same type as x, and `one(T)` returns a value of type T. However, this may not be the case for types representing dimensionful quantities (e.g. time in days), since the multiplicative identity must be dimensionless. In that case, `one(x)` should return an identity value of the same precision (and shape, for matrices) as x.

If you want a quantity that is of the same type as x, or of type T, even if x is dimensionful, use `oneunit` instead.

```julia
julia> one(3.7)
1.0

julia> one(Int)
1

julia> one(Dates.Day(1))
1
```

source

`Base.oneunit` – Function.

```
oneunit(x::T)
oneunit(T::Type)
```

Returns `T(one(x))`, where T is either the type of the argument or (if a type is passed) the argument. This differs from `one` for dimensionful quantities: one is dimensionless (a multiplicative identity) while `oneunit` is dimensionful (of the same type as x, or of type T).

```julia
julia> oneunit(3.7)
1.0

julia> oneunit(Dates.Day)
1 day
```

source

`Base.zero` – Function.

```
zero(x)
```

Get the additive identity element for the type of x (x can also specify the type itself).

source

`Base.pi` – Constant.

```
piπ
```

The constant pi.

```julia
julia> pi
π = 3.1415926535897...
```

source

Base.im – Constant.

```
im
```

The imaginary unit.

source

Base.eu – Constant.

```
e
eu
```

The constant e.

```julia
julia> e
e = 2.7182818284590...
```

source

Base.catalan – Constant.

```
catalan
```

Catalan's constant.

```julia
julia> catalan
catalan = 0.9159655941772...
```

source

Base.eulergamma – Constant.

```
γ
eulergamma
```

Euler's constant.

```julia
julia> eulergamma
γ = 0.5772156649015...
```

source

Base.golden – Constant.

```
φ
golden
```

The golden ratio.

```julia
julia> golden
φ = 1.6180339887498...
```

source

Base.Inf – Constant.

```
Inf
```

Positive infinity of type Float64.

source

Base.Inf32 – Constant.

```
Inf32
```

Positive infinity of type Float32.

source

Base.Inf16 – Constant.

```
Inf16
```

Positive infinity of type Float16.

source

Base.NaN – Constant.

```
NaN
```

A not-a-number value of type Float64.

source

Base.NaN32 – Constant.

```
NaN32
```

A not-a-number value of type Float32.

source

Base.NaN16 – Constant.

```
NaN16
```

A not-a-number value of type Float16.

source

Base.issubnormal – Function.

```
issubnormal(f) -> Bool
```

Test whether a floating point number is subnormal.

source

Base.isfinite – Function.

```
isfinite(f) -> Bool
```

Test whether a number is finite.

```
julia> isfinite(5)
true

julia> isfinite(NaN32)
false
```

source

Base.isinf – Function.

```
isinf(f) -> Bool
```

Test whether a number is infinite.

source

Base.isnan – Function.

```
isnan(f) -> Bool
```

Test whether a floating point number is not a number (NaN).

source

Base.iszero – Function.

```
iszero(x)
```

Return `true` if `x == zero(x)`; if x is an array, this checks whether all of the elements of x are zero.

source

Base.nextfloat – Function.

```
nextfloat(x::AbstractFloat, n::Integer)
```

The result of n iterative applications of `nextfloat` to x if `n >= 0`, or -n applications of `prevfloat` if `n < 0`.

source

```
nextfloat(x::AbstractFloat)
```

Returns the smallest floating point number y of the same type as x such `x < y`. If no such y exists (e.g. if x is Inf or NaN), then returns x.

source

Base.prevfloat – Function.

```
prevfloat(x::AbstractFloat)
```

Returns the largest floating point number y of the same type as x such `y < x`. If no such y exists (e.g. if x is -Inf or NaN), then returns x.

source

Base.isinteger – Function.

```
isinteger(x) -> Bool
```

Test whether x is numerically equal to some integer.

```
julia> isinteger(4.0)
true
```

source

Base.isreal – Function.

```
isreal(x) -> Bool
```

Test whether x or all its elements are numerically equal to some real number.

```
julia> isreal(5.)
true

julia> isreal([4.; complex(0,1)])
false
```

source

Core.Float32 – Method.

```
Float32(x [, mode::RoundingMode])
```

Create a Float32 from x. If x is not exactly representable then mode determines how x is rounded.

```
julia> Float32(1/3, RoundDown)
0.3333333f0

julia> Float32(1/3, RoundUp)
0.33333334f0
```

See RoundingMode for available rounding modes.

source

Core.Float64 – Method.

```
Float64(x [, mode::RoundingMode])
```

Create a Float64 from x. If x is not exactly representable then mode determines how x is rounded.

```
julia> Float64(pi, RoundDown)
3.141592653589793

julia> Float64(pi, RoundUp)
3.1415926535897936
```

See RoundingMode for available rounding modes.

source

Base.GMP.BigInt – Method.

```
BigInt(x)
```

Create an arbitrary precision integer. x may be an Int (or anything that can be converted to an Int). The usual mathematical operators are defined for this type, and results are promoted to a BigInt.

Instances can be constructed from strings via parse, or using the big string literal.

```
julia> parse(BigInt, "42")
42

julia> big"313"
313
```

source

Base.MPFR.BigFloat – Method.

```
BigFloat(x)
```

Create an arbitrary precision floating point number. x may be an Integer, a Float64 or a BigInt. The usual mathematical operators are defined for this type, and results are promoted to a BigFloat.

Note that because decimal literals are converted to floating point numbers when parsed, BigFloat(2.1) may not yield what you expect. You may instead prefer to initialize constants from strings via parse, or using the big string literal.

```
julia> BigFloat(2.1)
2.100000000000000088817841970012523233890533447265625000000000000000000000000000

julia> big"2.1"
2.099999999999999999999999999999999999999999999999999999999999999999999999999986
```

source

Base.Rounding.rounding – Function.

```
rounding(T)
```

Get the current floating point rounding mode for type T, controlling the rounding of basic arithmetic functions (+, -, *, / and sqrt) and type conversion.

See RoundingMode for available modes.

source

Base.Rounding.setrounding – Method.

```
setrounding(T, mode)
```

Set the rounding mode of floating point type T, controlling the rounding of basic arithmetic functions (+, -, *, / and sqrt) and type conversion. Other numerical functions may give incorrect or invalid values when using rounding modes other than the default RoundNearest.

Note that this may affect other types, for instance changing the rounding mode of Float64 will change the rounding mode of Float32. See RoundingMode for available modes.

> **Warning**
>
> This feature is still experimental, and may give unexpected or incorrect values.

source

Base.Rounding.setrounding – Method.

```
setrounding(f::Function, T, mode)
```

Change the rounding mode of floating point type T for the duration of f. It is logically equivalent to:

```
old = rounding(T)
setrounding(T, mode)
f()
setrounding(T, old)
```

See RoundingMode for available rounding modes.

> **Warning**
>
> This feature is still experimental, and may give unexpected or incorrect values. A known problem is the interaction with compiler optimisations, e.g.
>
> ```
> julia> setrounding(Float64,RoundDown) do
>            1.1 + 0.1
>        end
> 1.2000000000000002
> ```
>
> Here the compiler is *constant folding*, that is evaluating a known constant expression at compile time, however the rounding mode is only changed at runtime, so this is not reflected in the function result. This can be avoided by moving constants outside the expression, e.g.
>
> ```
> julia> x = 1.1; y = 0.1;
>
> julia> setrounding(Float64,RoundDown) do
>            x + y
>        end
> 1.2
> ```

source

Base.Rounding.get_zero_subnormals – Function.

```
get_zero_subnormals() -> Bool
```

Returns `false` if operations on subnormal floating-point values ("denormals") obey rules for IEEE arithmetic, and `true` if they might be converted to zeros.

source

Base.Rounding.set_zero_subnormals – Function.

```
set_zero_subnormals(yes::Bool) -> Bool
```

If yes is `false`, subsequent floating-point operations follow rules for IEEE arithmetic on subnormal values ("denormals"). Otherwise, floating-point operations are permitted (but not required) to convert subnormal inputs or outputs to zero. Returns `true` unless yes==true but the hardware does not support zeroing of subnormal numbers.

`set_zero_subnormals(true)` can speed up some computations on some hardware. However, it can break identities such as (x-y==0) == (x==y).

source

Integers

Base.count_ones – Function.

```
count_ones(x::Integer) -> Integer
```

Number of ones in the binary representation of x.

```julia
julia> count_ones(7)
3
```

source

Base.count_zeros – Function.

```
count_zeros(x::Integer) -> Integer
```

Number of zeros in the binary representation of x.

```julia
julia> count_zeros(Int32(2 ^ 16 - 1))
16
```

source

Base.leading_zeros – Function.

```
leading_zeros(x::Integer) -> Integer
```

Number of zeros leading the binary representation of x.

```julia
julia> leading_zeros(Int32(1))
31
```

source

Base.leading_ones – Function.

```
leading_ones(x::Integer) -> Integer
```

Number of ones leading the binary representation of x.

```julia
julia> leading_ones(UInt32(2 ^ 32 - 2))
31
```

source

Base.trailing_zeros – Function.

```
trailing_zeros(x::Integer) -> Integer
```

Number of zeros trailing the binary representation of x.

```julia
julia> trailing_zeros(2)
1
```

source

Base.trailing_ones – Function.

```
trailing_ones(x::Integer) -> Integer
```

Number of ones trailing the binary representation of x.

```
julia> trailing_ones(3)
2
```

source

Base.isodd – Function.

```
isodd(x::Integer) -> Bool
```

Returns `true` if x is odd (that is, not divisible by 2), and `false` otherwise.

```
julia> isodd(9)
true

julia> isodd(10)
false
```

source

Base.iseven – Function.

```
iseven(x::Integer) -> Bool
```

Returns `true` is x is even (that is, divisible by 2), and `false` otherwise.

```
julia> iseven(9)
false

julia> iseven(10)
true
```

source

48.4   BigFloats

The BigFloat type implements arbitrary-precision floating-point arithmetic using the GNU MPFR library.

Base.precision – Function.

```
precision(num::AbstractFloat)
```

Get the precision of a floating point number, as defined by the effective number of bits in the mantissa.

source

Base.precision – Method.

```
precision(BigFloat)
```

Get the precision (in bits) currently used for BigFloat arithmetic.

source

Base.MPFR.setprecision – Function.

```
setprecision([T=BigFloat,] precision::Int)
```

Set the precision (in bits) to be used for T arithmetic.

source

```
setprecision(f::Function, [T=BigFloat,] precision::Integer)
```

Change the T arithmetic precision (in bits) for the duration of f. It is logically equivalent to:

```
old = precision(BigFloat)
setprecision(BigFloat, precision)
f()
setprecision(BigFloat, old)
```

Often used as `setprecision(T, precision) do ... end`

source

Base.MPFR.BigFloat – Method.

```
BigFloat(x, prec::Int)
```

Create a representation of x as a BigFloat with precision `prec`.

source

Base.MPFR.BigFloat – Method.

```
BigFloat(x, rounding::RoundingMode)
```

Create a representation of x as a BigFloat with the current global precision and rounding mode `rounding`.

source

Base.MPFR.BigFloat – Method.

```
BigFloat(x, prec::Int, rounding::RoundingMode)
```

Create a representation of x as a BigFloat with precision `prec` and rounding mode `rounding`.

source

Base.MPFR.BigFloat – Method.

```
BigFloat(x::String)
```

Create a representation of the string x as a BigFloat.

source

48.5   Random Numbers

Random number generation in Julia uses the Mersenne Twister library via MersenneTwister objects. Julia has a global RNG, which is used by default. Other RNG types can be plugged in by inheriting the AbstractRNG type; they can then be used to have multiple streams of random numbers. Besides MersenneTwister, Julia also provides the RandomDevice RNG type, which is a wrapper over the OS provided entropy.

Most functions related to random generation accept an optional `AbstractRNG` as the first argument, `rng` , which defaults to the global one if not provided. Morever, some of them accept optionally dimension specifications `dims...` (which can be given as a tuple) to generate arrays of random values.

A `MersenneTwister` or RandomDevice RNG can generate random numbers of the following types: `Float16`, `Float32`, `Float64`, `Bool`, `Int8`, `UInt8`, `Int16`, `UInt16`, `Int32`, `UInt32`, `Int64`, `UInt64`, `Int128`, `UInt128`, `BigInt` (or complex numbers of those types). Random floating point numbers are generated uniformly in $[0, 1)$. As `BigInt` represents unbounded integers, the interval must be specified (e.g. `rand(big(1:6))`).

Base.Random.srand – Function.

```
srand([rng=GLOBAL_RNG], [seed]) -> rng
srand([rng=GLOBAL_RNG], filename, n=4) -> rng
```

Reseed the random number generator. If a `seed` is provided, the RNG will give a reproducible sequence of numbers, otherwise Julia will get entropy from the system. For `MersenneTwister`, the `seed` may be a non-negative integer, a vector of `UInt32` integers or a filename, in which case the seed is read from a file (4n bytes are read from the file, where n is an optional argument). RandomDevice does not support seeding.

source

Base.Random.MersenneTwister – Type.

```
MersenneTwister(seed)
```

Create a `MersenneTwister` RNG object. Different RNG objects can have their own seeds, which may be useful for generating different streams of random numbers.

source

Base.Random.RandomDevice – Type.

```
RandomDevice()
```

Create a `RandomDevice` RNG object. Two such objects will always generate different streams of random numbers.

source

Base.Random.rand – Function.

```
rand([rng=GLOBAL_RNG], [S], [dims...])
```

Pick a random element or array of random elements from the set of values specified by S; S can be

- an indexable collection (for example 1:n or `['x','y','z']`), or
- a type: the set of values to pick from is then equivalent to `typemin(S):typemax(S)` for integers (this is not applicable to `BigInt`), and to $[0, 1)$ for floating point numbers;

S defaults to `Float64`.

source

Base.Random.rand! – Function.

```
rand!([rng=GLOBAL_RNG], A, [coll])
```

Populate the array A with random values. If the indexable collection `coll` is specified, the values are picked randomly from `coll`. This is equivalent to `copy!(A, rand(rng, coll, size(A)))` or `copy!(A, rand(rng, eltype(A), size(A)))` but without allocating a new array.

source

Base.Random.bitrand – Function.

```
bitrand([rng=GLOBAL_RNG], [dims...])
```

Generate a `BitArray` of random boolean values.

source

Base.Random.randn – Function.

```
randn([rng=GLOBAL_RNG], [T=Float64], [dims...])
```

Generate a normally-distributed random number of type T with mean 0 and standard deviation 1. Optionally generate an array of normally-distributed random numbers. The `Base` module currently provides an implementation for the types `Float16`, `Float32`, and `Float64` (the default).

source

Base.Random.randn! – Function.

```
randn!([rng=GLOBAL_RNG], A::AbstractArray) -> A
```

Fill the array A with normally-distributed (mean 0, standard deviation 1) random numbers. Also see the `rand` function.

source

Base.Random.randexp – Function.

```
randexp([rng=GLOBAL_RNG], [T=Float64], [dims...])
```

Generate a random number of type T according to the exponential distribution with scale 1. Optionally generate an array of such random numbers. The `Base` module currently provides an implementation for the types `Float16`, `Float32`, and `Float64` (the default).

source

Base.Random.randexp! – Function.

```
randexp!([rng=GLOBAL_RNG], A::AbstractArray) -> A
```

Fill the array A with random numbers following the exponential distribution (with scale 1).

source

Base.Random.randjump – Function.

```
randjump(r::MersenneTwister, jumps::Integer, [jumppoly::AbstractString=dSFMT.JPOLY1e21]) ->
    Vector{MersenneTwister}
```

Create an array of the size `jumps` of initialized `MersenneTwister` RNG objects. The first RNG object given as a parameter and following `MersenneTwister` RNGs in the array are initialized such that a state of the RNG object in the array would be moved forward (without generating numbers) from a previous RNG object array element on a particular number of steps encoded by the jump polynomial `jumppoly`.

Default jump polynomial moves forward `MersenneTwister` RNG state by `10^20` steps.

source

Chapter 49

Strings

Base.length – Method.

```
length(s::AbstractString)
```

The number of characters in string s.

```
julia> length("jμΛΙα")
5
```

source

Base.sizeof – Method.

```
sizeof(s::AbstractString)
```

The number of bytes in string s.

```
julia> sizeof("")
3
```

source

Base.:* – Method.

```
*(x, y...)
```

Multiplication operator. x*y*z*... calls this function with all arguments, i.e. *(x, y, z, ...).

source

Base.:^ – Method.

```
^(s::AbstractString, n::Integer)
```

Repeat n times the string s. The repeat function is an alias to this operator.

```
julia> "Test "^3
"Test Test Test "
```

source

`Base.string` – Function.

```
string(xs...)
```

Create a string from any values using the `print` function.

```
julia> string("a", 1, true)
"a1true"
```

source

`Base.repr` – Function.

```
repr(x)
```

Create a string from any value using the `showall` function.

source

`Core.String` – Method.

```
String(s::AbstractString)
```

Convert a string to a contiguous byte array representation encoded as UTF-8 bytes. This representation is often appropriate for passing strings to C.

source

`Base.transcode` – Function.

```
transcode(T, src)
```

Convert string data between Unicode encodings. `src` is either a `String` or a `Vector{UIntXX}` of UTF-XX code units, where XX is 8, 16, or 32. T indicates the encoding of the return value: `String` to return a (UTF-8 encoded) `String` or `UIntXX` to return a `Vector{UIntXX}` of UTF-XX data. (The alias `Cwchar_t` can also be used as the integer type, for converting `wchar_t*` strings used by external C libraries.)

The `transcode` function succeeds as long as the input data can be reasonably represented in the target encoding; it always succeeds for conversions between UTF-XX encodings, even for invalid Unicode data.

Only conversion to/from UTF-8 is currently supported.

source

`Base.unsafe_string` – Function.

```
unsafe_string(p::Ptr{UInt8}, [length::Integer])
```

Copy a string from the address of a C-style (NUL-terminated) string encoded as UTF-8. (The pointer can be safely freed afterwards.) If `length` is specified (the length of the data in bytes), the string does not have to be NUL-terminated.

This function is labelled "unsafe" because it will crash if p is not a valid memory address to data of the requested length.

source

`Base.codeunit` – Method.

```
codeunit(s::AbstractString, i::Integer)
```

Get the `ith` code unit of an encoded string. For example, returns the `ith` byte of the representation of a UTF-8 string.

`Base.ascii` – Function.

```
ascii(s::AbstractString)
```

Convert a string to `String` type and check that it contains only ASCII data, otherwise throwing an `ArgumentError` indicating the position of the first non-ASCII byte.

```
julia> ascii("abcdeγfgh")
ERROR: ArgumentError: invalid ASCII at index 6 in "abcdeγfgh"
Stacktrace:
 [1] ascii(::String) at ./strings/util.jl:479

julia> ascii("abcdefgh")
"abcdefgh"
```

`Base.@r_str` – Macro.

```
@r_str -> Regex
```

Construct a regex, such as `r"^[a-z]*$"`. The regex also accepts one or more flags, listed after the ending quote, to change its behaviour:

- `i` enables case-insensitive matching
- `m` treats the `^` and `$` tokens as matching the start and end of individual lines, as opposed to the whole string.
- `s` allows the `.` modifier to match newlines.
- `x` enables "comment mode": whitespace is enabled except when escaped with `\`, and `#` is treated as starting a comment.

For example, this regex has all three flags enabled:

```
julia> match(r"a+.*b+.*?d$"ism, "Goodbye,\nOh, angry,\nBad world\n")
RegexMatch("angry,\nBad world")
```

`Base.Docs.@html_str` – Macro.

```
@html_str -> Docs.HTML
```

Create an HTML object from a literal string.

`Base.Docs.@text_str` – Macro.

```
@text_str -> Docs.Text
```

Create a `Text` object from a literal string.

`Base.UTF8proc.normalize_string` – Function.

```
normalize_string(s::AbstractString, normalform::Symbol)
```

Normalize the string s according to one of the four "normal forms" of the Unicode standard: `normalform` can be `:NFC`, `:NFD`, `:NFKC`, or `:NFKD`. Normal forms C (canonical composition) and D (canonical decomposition) convert different visually identical representations of the same abstract string into a single canonical form, with form C being more compact. Normal forms KC and KD additionally canonicalize "compatibility equivalents": they convert characters that are abstractly similar but visually distinct into a single canonical choice (e.g. they expand ligatures into the individual characters), with form KC being more compact.

Alternatively, finer control and additional transformations may be be obtained by calling `normalize_string(s; keywords...)`, where any number of the following boolean keywords options (which all default to `false` except for `compose`) are specified:

- `compose=false`: do not perform canonical composition
- `decompose=true`: do canonical decomposition instead of canonical composition (`compose=true` is ignored if present)
- `compat=true`: compatibility equivalents are canonicalized
- `casefold=true`: perform Unicode case folding, e.g. for case-insensitive string comparison
- `newline2lf=true`, `newline2ls=true`, or `newline2ps=true`: convert various newline sequences (LF, CRLF, CR, NEL) into a linefeed (LF), line-separation (LS), or paragraph-separation (PS) character, respectively
- `stripmark=true`: strip diacritical marks (e.g. accents)
- `stripignore=true`: strip Unicode's "default ignorable" characters (e.g. the soft hyphen or the left-to-right marker)
- `stripcc=true`: strip control characters; horizontal tabs and form feeds are converted to spaces; newlines are also converted to spaces unless a newline-conversion flag was specified
- `rejectna=true`: throw an error if unassigned code points are found
- `stable=true`: enforce Unicode Versioning Stability

For example, NFKC corresponds to the options `compose=true`, `compat=true`, `stable=true`.

source

`Base.UTF8proc.graphemes` – Function.

```
graphemes(s::AbstractString) -> GraphemeIterator
```

Returns an iterator over substrings of s that correspond to the extended graphemes in the string, as defined by Unicode UAX #29. (Roughly, these are what users would perceive as single characters, even though they may contain more than one codepoint; for example a letter combined with an accent mark is a single grapheme.)

source

`Base.isvalid` – Method.

```
isvalid(value) -> Bool
```

Returns `true` if the given value is valid for its type, which currently can be either `Char` or `String`.

source

`Base.isvalid` – Method.

```
isvalid(T, value) -> Bool
```

Returns `true` if the given value is valid for that type. Types currently can be either `Char` or `String`. Values for `Char` can be of type `Char` or `UInt32`. Values for `String` can be of that type, or `Vector{UInt8}`.

source

`Base.isvalid` – Method.

```
isvalid(str::AbstractString, i::Integer)
```

Tells whether index `i` is valid for the given string.

```julia
julia> str = "αβγdef";

julia> isvalid(str, 1)
true

julia> str[1]
'α': Unicode U+03b1 (category Ll: Letter, lowercase)

julia> isvalid(str, 2)
false

julia> str[2]
ERROR: UnicodeError: invalid character index
[...]
```

source

`Base.UTF8proc.is_assigned_char` – Function.

```
is_assigned_char(c) -> Bool
```

Returns `true` if the given char or integer is an assigned Unicode code point.

source

`Base.ismatch` – Function.

```
ismatch(r::Regex, s::AbstractString) -> Bool
```

Test whether a string contains a match of the given regular expression.

source

`Base.match` – Function.

```
match(r::Regex, s::AbstractString[, idx::Integer[, addopts]])
```

Search for the first match of the regular expression r in s and return a `RegexMatch` object containing the match, or nothing if the match failed. The matching substring can be retrieved by accessing `m.match` and the captured sequences can be retrieved by accessing `m.captures` The optional `idx` argument specifies an index at which to start the search.

source

`Base.eachmatch` – Function.

```
eachmatch(r::Regex, s::AbstractString[, overlap::Bool=false])
```

Search for all matches of a the regular expression r in s and return a iterator over the matches. If overlap is `true`, the matching sequences are allowed to overlap indices in the original string, otherwise they must be from distinct character ranges.

source

Base.matchall – Function.

```
matchall(r::Regex, s::AbstractString[, overlap::Bool=false]) -> Vector{AbstractString}
```

Return a vector of the matching substrings from eachmatch.

source

Base.lpad – Function.

```
lpad(s, n::Integer, p::AbstractString=" ")
```

Make a string at least n columns wide when printed by padding s on the left with copies of p.

```julia
julia> lpad("March",10)
"     March"
```

source

Base.rpad – Function.

```
rpad(s, n::Integer, p::AbstractString=" ")
```

Make a string at least n columns wide when printed by padding s on the right with copies of p.

```julia
julia> rpad("March",20)
"March               "
```

source

Base.search – Function.

```
search(string::AbstractString, chars::Chars, [start::Integer])
```

Search for the first occurrence of the given characters within the given string. The second argument may be a single character, a vector or a set of characters, a string, or a regular expression (though regular expressions are only allowed on contiguous strings, such as ASCII or UTF-8 strings). The third argument optionally specifies a starting index. The return value is a range of indexes where the matching sequence is found, such that `s[search(s,x)] == x`:

`search(string, "substring")` = start:end such that `string[start:end] == "substring"`, or `0:-1` if unmatched.

`search(string, 'c')` = index such that `string[index] == 'c'`, or 0 if unmatched.

```julia
julia> search("Hello to the world", "z")
0:-1

julia> search("JuliaLang","Julia")
1:5
```

Base.rsearch – Function.

```
rsearch(s::AbstractString, chars::Chars, [start::Integer])
```

Similar to search, but returning the last occurrence of the given characters within the given string, searching in reverse from start.

```
julia> rsearch("aaabbb","b")
6:6
```

Base.searchindex – Function.

```
searchindex(s::AbstractString, substring, [start::Integer])
```

Similar to search, but return only the start index at which the substring is found, or 0 if it is not.

```
julia> searchindex("Hello to the world", "z")
0

julia> searchindex("JuliaLang","Julia")
1

julia> searchindex("JuliaLang","Lang")
6
```

Base.rsearchindex – Function.

```
rsearchindex(s::AbstractString, substring, [start::Integer])
```

Similar to rsearch, but return only the start index at which the substring is found, or 0 if it is not.

```
julia> rsearchindex("aaabbb","b")
6

julia> rsearchindex("aaabbb","a")
3
```

Base.contains – Method.

```
contains(haystack::AbstractString, needle::AbstractString)
```

Determine whether the second argument is a substring of the first.

```
julia> contains("JuliaLang is pretty cool!", "Julia")
true
```

Base.reverse – Method.

```
reverse(s::AbstractString) -> AbstractString
```

Reverses a string.

```
julia> reverse("JuliaLang")
"gnaLailuJ"
```

source

Base.replace – Function.

```
replace(string::AbstractString, pat, r[, n::Integer=0])
```

Search for the given pattern pat, and replace each occurrence with r. If n is provided, replace at most n occurrences. As with search, the second argument may be a single character, a vector or a set of characters, a string, or a regular expression. If r is a function, each occurrence is replaced with r(s) where s is the matched substring. If pat is a regular expression and r is a SubstitutionString, then capture group references in r are replaced with the corresponding matched text.

source

Base.split – Function.

```
split(s::AbstractString, [chars]; limit::Integer=0, keep::Bool=true)
```

Return an array of substrings by splitting the given string on occurrences of the given character delimiters, which may be specified in any of the formats allowed by search's second argument (i.e. a single character, collection of characters, string, or regular expression). If chars is omitted, it defaults to the set of all space characters, and keep is taken to be false. The two keyword arguments are optional: they are a maximum size for the result and a flag determining whether empty fields should be kept in the result.

```
julia> a = "Ma.rch"
"Ma.rch"

julia> split(a,".")
2-element Array{SubString{String},1}:
 "Ma"
 "rch"
```

source

Base.rsplit – Function.

```
rsplit(s::AbstractString, [chars]; limit::Integer=0, keep::Bool=true)
```

Similar to split, but starting from the end of the string.

```
julia> a = "M.a.r.c.h"
"M.a.r.c.h"

julia> rsplit(a,".")
5-element Array{SubString{String},1}:
 "M"
 "a"
 "r"
 "c"
```

```
 "h"

julia> rsplit(a,".";limit=1)
1-element Array{SubString{String},1}:
 "M.a.r.c.h"

julia> rsplit(a,".";limit=2)
2-element Array{SubString{String},1}:
 "M.a.r.c"
 "h"
```

source

`Base.strip` – Function.

```
strip(s::AbstractString, [chars::Chars])
```

Return s with any leading and trailing whitespace removed. If `chars` (a character, or vector or set of characters) is provided, instead remove characters contained in it.

```
julia> strip("{3, 5}\n", ['{', '}', '\n'])
"3, 5"
```

source

`Base.lstrip` – Function.

```
lstrip(s::AbstractString[, chars::Chars])
```

Return s with any leading whitespace and delimiters removed. The default delimiters to remove are ' ', \t, \n, \v, \f, and \r. If `chars` (a character, or vector or set of characters) is provided, instead remove characters contained in it.

```
julia> a = lpad("March", 20)
"               March"

julia> lstrip(a)
"March"
```

source

`Base.rstrip` – Function.

```
rstrip(s::AbstractString[, chars::Chars])
```

Return s with any trailing whitespace and delimiters removed. The default delimiters to remove are ' ', \t, \n, \v, \f, and \r. If `chars` (a character, or vector or set of characters) is provided, instead remove characters contained in it.

```
julia> a = rpad("March", 20)
"March               "

julia> rstrip(a)
"March"
```

source

Base.startswith – Function.

> startswith(s::AbstractString, prefix::AbstractString)

Returns true if s starts with prefix. If prefix is a vector or set of characters, tests whether the first character of s belongs to that set.

See also endswith.

```
julia> startswith("JuliaLang", "Julia")
true
```

source

Base.endswith – Function.

> endswith(s::AbstractString, suffix::AbstractString)

Returns true if s ends with suffix. If suffix is a vector or set of characters, tests whether the last character of s belongs to that set.

See also startswith.

```
julia> endswith("Sunday", "day")
true
```

source

Base.uppercase – Function.

> uppercase(s::AbstractString)

Returns s with all characters converted to uppercase.

```
julia> uppercase("Julia")
"JULIA"
```

source

Base.lowercase – Function.

> lowercase(s::AbstractString)

Returns s with all characters converted to lowercase.

```
julia> lowercase("STRINGS AND THINGS")
"strings and things"
```

source

Base.titlecase – Function.

> titlecase(s::AbstractString)

Capitalizes the first character of each word in s.

```
julia> titlecase("the julia programming language")
"The Julia Programming Language"
```

`Base.ucfirst` – Function.

```
ucfirst(s::AbstractString)
```

Returns `string` with the first character converted to uppercase.

```
julia> ucfirst("python")
"Python"
```

`Base.lcfirst` – Function.

```
lcfirst(s::AbstractString)
```

Returns `string` with the first character converted to lowercase.

```
julia> lcfirst("Julia")
"julia"
```

`Base.join` – Function.

```
join(io::IO, strings, delim, [last])
```

Join an array of `strings` into a single string, inserting the given delimiter between adjacent strings. If `last` is given, it will be used instead of `delim` between the last two strings. For example,

```
julia> join(["apples", "bananas", "pineapples"], ", ", " and ")
"apples, bananas and pineapples"
```

`strings` can be any iterable over elements x which are convertible to strings via `print(io::IOBuffer, x)`. `strings` will be printed to `io`.

`Base.chop` – Function.

```
chop(s::AbstractString)
```

Remove the last character from s.

```
julia> a = "March"
"March"

julia> chop(a)
"Marc"
```

`Base.chomp` – Function.

```
chomp(s::AbstractString)
```

Remove a single trailing newline from a string.

```julia
julia> chomp("Hello\n")
"Hello"
```

source

Base.ind2chr – Function.

```julia
ind2chr(s::AbstractString, i::Integer)
```

Convert a byte index i to a character index with respect to string s.

See also chr2ind.

```julia
julia> str = "αβγdef";

julia> ind2chr(str, 3)
2

julia> chr2ind(str, 2)
3
```

source

Base.chr2ind – Function.

```julia
chr2ind(s::AbstractString, i::Integer)
```

Convert a character index i to a byte index.

See also ind2chr.

```julia
julia> str = "αβγdef";

julia> chr2ind(str, 2)
3

julia> ind2chr(str, 3)
2
```

source

Base.nextind – Function.

```julia
nextind(str::AbstractString, i::Integer)
```

Get the next valid string index after i. Returns a value greater than endof(str) at or after the end of the string.

```julia
julia> str = "αβγdef";

julia> nextind(str, 1)
3

julia> endof(str)
9

julia> nextind(str, 9)
10
```

Base.prevind – Function.

```
prevind(str::AbstractString, i::Integer)
```

Get the previous valid string index before `i`. Returns a value less than 1 at the beginning of the string.

```
julia> prevind("αβγdef", 3)
1

julia> prevind("αβγdef", 1)
0
```

Base.Random.randstring – Function.

```
randstring([rng,] len=8)
```

Create a random ASCII string of length `len`, consisting of upper- and lower-case letters and the digits 0-9. The optional `rng` argument specifies a random number generator, see Random Numbers.

Base.UTF8proc.charwidth – Function.

```
charwidth(c)
```

Gives the number of columns needed to print a character.

Base.strwidth – Function.

```
strwidth(s::AbstractString)
```

Gives the number of columns needed to print a string.

```
julia> strwidth("March")
5
```

Base.UTF8proc.isalnum – Function.

```
isalnum(c::Char) -> Bool
```

Tests whether a character is alphanumeric. A character is classified as alphabetic if it belongs to the Unicode general category Letter or Number, i.e. a character whose category code begins with 'L' or 'N'.

Base.UTF8proc.isalpha – Function.

```
isalpha(c::Char) -> Bool
```

Tests whether a character is alphabetic. A character is classified as alphabetic if it belongs to the Unicode general category Letter, i.e. a character whose category code begins with 'L'.

   source

Base.isascii – Function.

```
isascii(c::Union{Char,AbstractString}) -> Bool
```

Tests whether a character belongs to the ASCII character set, or whether this is true for all elements of a string.

   source

Base.UTF8proc.iscntrl – Function.

```
iscntrl(c::Char) -> Bool
```

Tests whether a character is a control character. Control characters are the non-printing characters of the Latin-1 subset of Unicode.

   source

Base.UTF8proc.isdigit – Function.

```
isdigit(c::Char) -> Bool
```

Tests whether a character is a numeric digit (0-9).

   source

Base.UTF8proc.isgraph – Function.

```
isgraph(c::Char) -> Bool
```

Tests whether a character is printable, and not a space. Any character that would cause a printer to use ink should be classified with `isgraph(c)==true`.

   source

Base.UTF8proc.islower – Function.

```
islower(c::Char) -> Bool
```

Tests whether a character is a lowercase letter. A character is classified as lowercase if it belongs to Unicode category Ll, Letter: Lowercase.

   source

Base.UTF8proc.isnumber – Function.

```
isnumber(c::Char) -> Bool
```

Tests whether a character is numeric. A character is classified as numeric if it belongs to the Unicode general category Number, i.e. a character whose category code begins with 'N'.

   source

Base.UTF8proc.isprint – Function.

```
isprint(c::Char) -> Bool
```

Tests whether a character is printable, including spaces, but not a control character.

source

Base.UTF8proc.ispunct – Function.

```
ispunct(c::Char) -> Bool
```

Tests whether a character belongs to the Unicode general category Punctuation, i.e. a character whose category code begins with 'P'.

source

Base.UTF8proc.isspace – Function.

```
isspace(c::Char) -> Bool
```

Tests whether a character is any whitespace character. Includes ASCII characters '\t', '\n', '\v', '\f', '\r', and ' ', Latin-1 character U+0085, and characters in Unicode category Zs.

source

Base.UTF8proc.isupper – Function.

```
isupper(c::Char) -> Bool
```

Tests whether a character is an uppercase letter. A character is classified as uppercase if it belongs to Unicode category Lu, Letter: Uppercase, or Lt, Letter: Titlecase.

source

Base.isxdigit – Function.

```
isxdigit(c::Char) -> Bool
```

Tests whether a character is a valid hexadecimal digit. Note that this does not include x (as in the standard 0x prefix).

```
julia> isxdigit('a')
true

julia> isxdigit('x')
false
```

source

Core.Symbol – Type.

```
Symbol(x...) -> Symbol
```

Create a Symbol by concatenating the string representations of the arguments together.

source

Base.escape_string – Function.

```
escape_string([io,] str::AbstractString[, esc::AbstractString]) -> AbstractString
```

General escaping of traditional C and Unicode escape sequences. Any characters in `esc` are also escaped (with a backslash). See also `unescape_string`.

source

`Base.unescape_string` – Function.

```
unescape_string([io,] s::AbstractString) -> AbstractString
```

General unescaping of traditional C and Unicode escape sequences. Reverse of `escape_string`.

source

Chapter 50

Arrays

50.1 Constructors and Types

Core.AbstractArray – Type.

```
AbstractArray{T, N}
```

Abstract array supertype which arrays inherit from.

source

Core.Array – Type.

```
Array{T}(dims)
Array{T,N}(dims)
```

Construct an uninitialized N-dimensional dense array with element type T, where N is determined from the length or number of dims. dims may be a tuple or a series of integer arguments corresponding to the lengths in each dimension. If the rank N is supplied explicitly as in Array{T,N}(dims), then it must match the length or number of dims.

Example

```julia
julia> A = Array{Float64, 2}(2, 2);

julia> ndims(A)
2

julia> eltype(A)
Float64
```

source

Base.getindex – Method.

```
getindex(type[, elements...])
```

Construct a 1-d array of the specified type. This is usually called with the syntax Type[]. Element values can be specified using Type[a,b,c,...].

```julia
julia> Int8[1, 2, 3]
3-element Array{Int8,1}:
```

```
 1
 2
 3

julia> getindex(Int8, 1, 2, 3)
3-element Array{Int8,1}:
 1
 2
 3
```

source

Base.zeros – Function.

```
zeros([A::AbstractArray,] [T=eltype(A)::Type,] [dims=size(A)::Tuple])
```

Create an array of all zeros with the same layout as A, element type T and size dims. The A argument can be skipped, which behaves like Array{Float64,0}() was passed. For convenience dims may also be passed in variadic form.

```
julia> zeros(1)
1-element Array{Float64,1}:
 0.0

julia> zeros(Int8, 2, 3)
2×3 Array{Int8,2}:
 0  0  0
 0  0  0

julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> zeros(A)
2×2 Array{Int64,2}:
 0  0
 0  0

julia> zeros(A, Float64)
2×2 Array{Float64,2}:
 0.0  0.0
 0.0  0.0

julia> zeros(A, Bool, (3,))
3-element Array{Bool,1}:
 false
 false
 false
```

See also ones, similar.

source

Base.ones – Function.

```
ones([A::AbstractArray,] [T=eltype(A)::Type,] [dims=size(A)::Tuple])
```

Create an array of all ones with the same layout as A, element type T and size `dims`. The A argument can be skipped, which behaves like `Array{Float64,0}()` was passed. For convenience `dims` may also be passed in variadic form.

```julia
julia> ones(Complex128, 2, 3)
2×3 Array{Complex{Float64},2}:
 1.0+0.0im  1.0+0.0im  1.0+0.0im
 1.0+0.0im  1.0+0.0im  1.0+0.0im

julia> ones(1,2)
1×2 Array{Float64,2}:
 1.0  1.0

julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> ones(A)
2×2 Array{Int64,2}:
 1  1
 1  1

julia> ones(A, Float64)
2×2 Array{Float64,2}:
 1.0  1.0
 1.0  1.0

julia> ones(A, Bool, (3,))
3-element Array{Bool,1}:
 true
 true
 true
```

See also `zeros`, `similar`.

source

Base.BitArray – Type.

```
BitArray(dims::Integer...)
BitArray{N}(dims::NTuple{N,Int})
```

Construct an uninitialized `BitArray` with the given dimensions. Behaves identically to the `Array` constructor.

```julia
julia> BitArray(2, 2)
2×2 BitArray{2}:
 false  false
 false   true

julia> BitArray((3, 1))
3×1 BitArray{2}:
 false
  true
 false
```

source

```
BitArray(itr)
```

Construct a `BitArray` generated by the given iterable object. The shape is inferred from the `itr` object.

```julia
julia> BitArray([1 0; 0 1])
2×2 BitArray{2}:
  true  false
 false   true

julia> BitArray(x+y == 3 for x = 1:2, y = 1:3)
2×3 BitArray{2}:
 false   true  false
  true  false  false

julia> BitArray(x+y == 3 for x = 1:2 for y = 1:3)
6-element BitArray{1}:
 false
  true
 false
  true
 false
 false
```

source

Base.`trues` – Function.

```
trues(dims)
```

Create a `BitArray` with all values set to `true`.

```julia
julia> trues(2,3)
2×3 BitArray{2}:
 true  true  true
 true  true  true
```

source

```
trues(A)
```

Create a `BitArray` with all values set to `true` of the same shape as A.

```julia
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> trues(A)
2×2 BitArray{2}:
 true  true
 true  true
```

source

Base.`falses` – Function.

```
falses(dims)
```

Create a `BitArray` with all values set to `false`.

```julia
julia> falses(2,3)
2×3 BitArray{2}:
 false  false  false
 false  false  false
```

source

```
falses(A)
```

Create a `BitArray` with all values set to `false` of the same shape as A.

```julia
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> falses(A)
2×2 BitArray{2}:
 false  false
 false  false
```

source

`Base.fill` – Function.

```
fill(x, dims)
```

Create an array filled with the value x. For example, `fill(1.0, (5,5))` returns a 5×5 array of floats, with each element initialized to `1.0`.

```julia
julia> fill(1.0, (5,5))
5×5 Array{Float64,2}:
 1.0  1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0  1.0
```

If x is an object reference, all elements will refer to the same object. `fill(Foo(), dims)` will return an array filled with the result of evaluating Foo() once.

source

`Base.fill!` – Function.

```
fill!(A, x)
```

Fill array A with the value x. If x is an object reference, all elements will refer to the same object. `fill!(A, Foo())` will return A filled with the result of evaluating Foo() once.

```julia
julia> A = zeros(2,3)
2×3 Array{Float64,2}:
 0.0  0.0  0.0
 0.0  0.0  0.0

julia> fill!(A, 2.)
2×3 Array{Float64,2}:
 2.0  2.0  2.0
 2.0  2.0  2.0

julia> a = [1, 1, 1]; A = fill!(Vector{Vector{Int}}(3), a); a[1] = 2; A
3-element Array{Array{Int64,1},1}:
 [2, 1, 1]
 [2, 1, 1]
 [2, 1, 1]

julia> x = 0; f() = (global x += 1; x); fill!(Vector{Int}(3), f())
3-element Array{Int64,1}:
 1
 1
 1
```

source

Base.similar – Method.

```
similar(array, [element_type=eltype(array)], [dims=size(array)])
```

Create an uninitialized mutable array with the given element type and size, based upon the given source array. The second and third arguments are both optional, defaulting to the given array's `eltype` and `size`. The dimensions may be specified either as a single tuple argument or as a series of integer arguments.

Custom AbstractArray subtypes may choose which specific array type is best-suited to return for the given element type and dimensionality. If they do not specialize this method, the default is an `Array{element_type}(dims...)`.

For example, `similar(1:10, 1, 4)` returns an uninitialized `Array{Int,2}` since ranges are neither mutable nor support 2 dimensions:

```julia
julia> similar(1:10, 1, 4)
1×4 Array{Int64,2}:
 4419743872  4374413872  4419743888  0
```

Conversely, `similar(trues(10,10), 2)` returns an uninitialized `BitVector` with two elements since `BitArrays` are both mutable and can support 1-dimensional arrays:

```julia
julia> similar(trues(10,10), 2)
2-element BitArray{1}:
 false
 false
```

Since `BitArrays` can only store elements of type `Bool`, however, if you request a different element type it will create a regular `Array` instead:

```julia
julia> similar(falses(10), Float64, 2, 4)
2×4 Array{Float64,2}:
 2.18425e-314  2.18425e-314  2.18425e-314  2.18425e-314
 2.18425e-314  2.18425e-314  2.18425e-314  2.18425e-314
```

Base.similar – Method.

```
similar(storagetype, indices)
```

Create an uninitialized mutable array analogous to that specified by `storagetype`, but with `indices` specified by the last argument. `storagetype` might be a type or a function.

Examples:

```
similar(Array{Int}, indices(A))
```

creates an array that "acts like" an `Array{Int}` (and might indeed be backed by one), but which is indexed identically to A. If A has conventional indexing, this will be identical to `Array{Int}(size(A))`, but if A has unconventional indexing then the indices of the result will match A.

```
similar(BitArray, (indices(A, 2),))
```

would create a 1-dimensional logical array whose indices match those of the columns of A.

```
similar(dims->zeros(Int, dims), indices(A))
```

would create an array of `Int`, initialized to zero, matching the indices of A.

Base.eye – Function.

```
eye([T::Type=Float64,] m::Integer, n::Integer)
```

m-by-n identity matrix. The default element type is Float64.

```
eye(m, n)
```

m-by-n identity matrix.

```
eye([T::Type=Float64,] n::Integer)
```

n-by-n identity matrix. The default element type is Float64.

```
eye(A)
```

Constructs an identity matrix of the same dimensions and type as A.

```
julia> A = [1 2 3; 4 5 6; 7 8 9]
3×3 Array{Int64,2}:
 1  2  3
 4  5  6
 7  8  9

julia> eye(A)
3×3 Array{Int64,2}:
 1  0  0
 0  1  0
 0  0  1
```

Note the difference from ones.

> source

Base.linspace – Function.

```
linspace(start, stop, n=50)
```

Construct a range of n linearly spaced elements from start to stop.

```
julia> linspace(1.3,2.9,9)
1.3:0.2:2.9
```

> source

Base.logspace – Function.

```
logspace(start::Real, stop::Real, n::Integer=50)
```

Construct a vector of n logarithmically spaced numbers from 10^start to 10^stop.

```
julia> logspace(1.,10.,5)
5-element Array{Float64,1}:
   10.0
 1778.28
    3.16228e5
    5.62341e7
    1.0e10
```

> source

Base.Random.randsubseq – Function.

```
randsubseq(A, p) -> Vector
```

Return a vector consisting of a random subsequence of the given array A, where each element of A is included (in order) with independent probability p. (Complexity is linear in p*length(A), so this function is efficient even if p is small and A is large.) Technically, this process is known as "Bernoulli sampling" of A.

> source

Base.Random.randsubseq! – Function.

```
randsubseq!(S, A, p)
```

Like randsubseq, but the results are stored in S (which is resized as needed).

> source

50.2   Basic functions

Base.ndims – Function.

```
ndims(A::AbstractArray) -> Integer
```

Returns the number of dimensions of A.

```julia
julia> A = ones(3,4,5);

julia> ndims(A)
3
```

source

Base.size – Function.

```julia
size(A::AbstractArray, [dim...])
```

Returns a tuple containing the dimensions of A. Optionally you can specify the dimension(s) you want the length of, and get the length of that dimension, or a tuple of the lengths of dimensions you asked for.

```julia
julia> A = ones(2,3,4);

julia> size(A, 2)
3

julia> size(A,3,2)
(4, 3)
```

source

Base.indices – Method.

```julia
indices(A)
```

Returns the tuple of valid indices for array A.

```julia
julia> A = ones(5,6,7);

julia> indices(A)
(Base.OneTo(5), Base.OneTo(6), Base.OneTo(7))
```

source

Base.indices – Method.

```julia
indices(A, d)
```

Returns the valid range of indices for array A along dimension d.

```julia
julia> A = ones(5,6,7);

julia> indices(A,2)
Base.OneTo(6)
```

source

Base.length – Method.

```julia
length(A::AbstractArray) -> Integer
```

Returns the number of elements in A.

```
julia> A = ones(3,4,5);

julia> length(A)
60
```

source

`Base.eachindex` – Function.

```
eachindex(A...)
```

Creates an iterable object for visiting each index of an AbstractArray A in an efficient manner. For array types that have opted into fast linear indexing (like `Array`), this is simply the range `1:length(A)`. For other array types, this returns a specialized Cartesian range to efficiently index into the array with indices specified for every dimension. For other iterables, including strings and dictionaries, this returns an iterator object supporting arbitrary index types (e.g. unevenly spaced or non-integer indices).

Example for a sparse 2-d array:

```
julia> A = sparse([1, 1, 2], [1, 3, 1], [1, 2, -5])
2×3 SparseMatrixCSC{Int64,Int64} with 3 stored entries:
  [1, 1]  =  1
  [2, 1]  =  -5
  [1, 3]  =  2

julia> for iter in eachindex(A)
           @show iter.I[1], iter.I[2]
           @show A[iter]
       end
(iter.I[1], iter.I[2]) = (1, 1)
A[iter] = 1
(iter.I[1], iter.I[2]) = (2, 1)
A[iter] = -5
(iter.I[1], iter.I[2]) = (1, 2)
A[iter] = 0
(iter.I[1], iter.I[2]) = (2, 2)
A[iter] = 0
(iter.I[1], iter.I[2]) = (1, 3)
A[iter] = 2
(iter.I[1], iter.I[2]) = (2, 3)
A[iter] = 0
```

If you supply more than one `AbstractArray` argument, `eachindex` will create an iterable object that is fast for all arguments (a `UnitRange` if all inputs have fast linear indexing, a `CartesianRange` otherwise). If the arrays have different sizes and/or dimensionalities, `eachindex` returns an iterable that spans the largest range along each dimension.

source

`Base.linearindices` – Function.

```
linearindices(A)
```

Returns a `UnitRange` specifying the valid range of indices for `A[i]` where `i` is an `Int`. For arrays with conventional indexing (indices start at 1), or any multidimensional array, this is `1:length(A)`; however, for one-dimensional arrays with unconventional indices, this is `indices(A, 1)`.

Calling this function is the "safe" way to write algorithms that exploit linear indexing.

```julia
julia> A = ones(5,6,7);

julia> b = linearindices(A);

julia> extrema(b)
(1, 210)
```

source

`Base.IndexStyle` – Type.

```
IndexStyle(A)
IndexStyle(typeof(A))
```

`IndexStyle` specifies the "native indexing style" for array A. When you define a new `AbstractArray` type, you can choose to implement either linear indexing or cartesian indexing. If you decide to implement linear indexing, then you must set this trait for your array type:

```
Base.IndexStyle(::Type{<:MyArray}) = IndexLinear()
```

The default is `IndexCartesian()`.

Julia's internal indexing machinery will automatically (and invisibly) convert all indexing operations into the preferred style using sub2ind or ind2sub. This allows users to access elements of your array using any indexing style, even when explicit methods have not been provided.

If you define both styles of indexing for your `AbstractArray`, this trait can be used to select the most performant indexing style. Some methods check this trait on their inputs, and dispatch to different algorithms depending on the most efficient access pattern. In particular, eachindex creates an iterator whose type depends on the setting of this trait.

source

`Base.countnz` – Function.

```
countnz(A) -> Integer
```

Counts the number of nonzero values in array A (dense or sparse). Note that this is not a constant-time operation. For sparse matrices, one should usually use nnz, which returns the number of stored values.

```julia
julia> A = [1 2 4; 0 0 1; 1 1 0]
3×3 Array{Int64,2}:
 1  2  4
 0  0  1
 1  1  0

julia> countnz(A)
6
```

source

`Base.conj!` – Function.

```
conj!(A)
```

Transform an array to its complex conjugate in-place.

See also conj.

```julia
julia> A = [1+im 2-im; 2+2im 3+im]
2×2 Array{Complex{Int64},2}:
 1+1im  2-1im
 2+2im  3+1im

julia> conj!(A);

julia> A
2×2 Array{Complex{Int64},2}:
 1-1im  2+1im
 2-2im  3-1im
```

source

Base.stride – Function.

```julia
stride(A, k::Integer)
```

Returns the distance in memory (in number of elements) between adjacent elements in dimension k.

```julia
julia> A = ones(3,4,5);

julia> stride(A,2)
3

julia> stride(A,3)
12
```

source

Base.strides – Function.

```julia
strides(A)
```

Returns a tuple of the memory strides in each dimension.

```julia
julia> A = ones(3,4,5);

julia> strides(A)
(1, 3, 12)
```

source

Base.ind2sub – Function.

```julia
ind2sub(a, index) -> subscripts
```

Returns a tuple of subscripts into array a corresponding to the linear index index.

```julia
julia> A = ones(5,6,7);

julia> ind2sub(A,35)
(5, 1, 2)

julia> ind2sub(A,70)
(5, 2, 3)
```

source

```
ind2sub(dims, index) -> subscripts
```

Returns a tuple of subscripts into an array with dimensions `dims`, corresponding to the linear index `index`.

Example:

```
i, j, ... = ind2sub(size(A), indmax(A))
```

provides the indices of the maximum element.

```
julia> ind2sub((3,4),2)
(2, 1)

julia> ind2sub((3,4),3)
(3, 1)

julia> ind2sub((3,4),4)
(1, 2)
```

source

`Base.sub2ind` – Function.

```
sub2ind(dims, i, j, k...) -> index
```

The inverse of `ind2sub`, returns the linear index corresponding to the provided subscripts.

```
julia> sub2ind((5,6,7),1,2,3)
66

julia> sub2ind((5,6,7),1,6,3)
86
```

source

`Base.LinAlg.checksquare` – Function.

```
LinAlg.checksquare(A)
```

Check that a matrix is square, then return its common dimension. For multiple arguments, return a vector.

Example

```
julia> A = ones(4,4); B = zeros(5,5);

julia> LinAlg.checksquare(A, B)
2-element Array{Int64,1}:
 4
 5
```

source

50.3 Broadcast and vectorization

See also the dot syntax for vectorizing functions; for example, `f.(args...)` implicitly calls `broadcast(f, args...)`. Rather than relying on "vectorized" methods of functions like `sin` to operate on arrays, you should use `sin.(a)` to vectorize via `broadcast`.

`Base.broadcast` – Function.

```
broadcast(f, As...)
```

Broadcasts the arrays, tuples, `Refs`, nullables, and/or scalars `As` to a container of the appropriate type and dimensions. In this context, anything that is not a subtype of `AbstractArray`, `Ref` (except for `Ptrs`), `Tuple`, or `Nullable` is considered a scalar. The resulting container is established by the following rules:

- If all the arguments are scalars, it returns a scalar.

- If the arguments are tuples and zero or more scalars, it returns a tuple.

- If the arguments contain at least one array or `Ref`, it returns an array (expanding singleton dimensions), and treats `Refs` as 0-dimensional arrays, and tuples as 1-dimensional arrays.

The following additional rule applies to `Nullable` arguments: If there is at least one `Nullable`, and all the arguments are scalars or `Nullable`, it returns a `Nullable` treating `Nullables` as "containers".

A special syntax exists for broadcasting: `f.(args...)` is equivalent to `broadcast(f, args...)`, and nested `f.(g.(args...))` calls are fused into a single broadcast loop.

```
julia> A = [1, 2, 3, 4, 5]
5-element Array{Int64,1}:
 1
 2
 3
 4
 5

julia> B = [1 2; 3 4; 5 6; 7 8; 9 10]
5×2 Array{Int64,2}:
 1   2
 3   4
 5   6
 7   8
 9  10

julia> broadcast(+, A, B)
5×2 Array{Int64,2}:
  2   3
  5   6
  8   9
 11  12
 14  15

julia> parse.(Int, ["1", "2"])
2-element Array{Int64,1}:
 1
 2

julia> abs.((1, -2))
```

```
(1, 2)

julia> broadcast(+, 1.0, (0, -2.0))
(1.0, -1.0)

julia> broadcast(+, 1.0, (0, -2.0), Ref(1))
2-element Array{Float64,1}:
 2.0
 0.0

julia> (+).([[0,2], [1,3]], Ref{Vector{Int}}([1,-1]))
2-element Array{Array{Int64,1},1}:
 [1, 1]
 [2, 2]

julia> string.(("one","two","three","four"), ": ", 1:4)
4-element Array{String,1}:
 "one: 1"
 "two: 2"
 "three: 3"
 "four: 4"

julia> Nullable("X") .* "Y"
Nullable{String}("XY")

julia> broadcast(/, 1.0, Nullable(2.0))
Nullable{Float64}(0.5)

julia> (1 + im) ./ Nullable{Int}()
Nullable{Complex{Float64}}()
```

source

Base.broadcast! – Function.

```
broadcast!(f, dest, As...)
```

Like broadcast, but store the result of broadcast(f, As...) in the dest array. Note that dest is only used to store the result, and does not supply arguments to f unless it is also listed in the As, as in broadcast!(f, A, A, B) to perform A[:] = broadcast(f, A, B).

source

Base.Broadcast.@__dot__ – Macro.

```
@. expr
```

Convert every function call or operator in expr into a "dot call" (e.g. convert f(x) to f.(x)), and convert every assignment in expr to a "dot assignment" (e.g. convert += to .+=).

If you want to *avoid* adding dots for selected function calls in expr, splice those function calls in with $. For example, @. sqrt(abs($sort(x))) is equivalent to sqrt.(abs.(sort(x))) (no dot for sort).

(@. is equivalent to a call to @__dot__.)

source

Base.Broadcast.broadcast_getindex – Function.

```
broadcast_getindex(A, inds...)
```

Broadcasts the `inds` arrays to a common size like `broadcast` and returns an array of the results `A[ks...]`, where `ks` goes over the positions in the broadcast result A.

```julia
julia> A = [1, 2, 3, 4, 5]
5-element Array{Int64,1}:
 1
 2
 3
 4
 5

julia> B = [1 2; 3 4; 5 6; 7 8; 9 10]
5×2 Array{Int64,2}:
 1   2
 3   4
 5   6
 7   8
 9  10

julia> C = broadcast(+,A,B)
5×2 Array{Int64,2}:
  2   3
  5   6
  8   9
 11  12
 14  15

julia> broadcast_getindex(C,[1,2,10])
3-element Array{Int64,1}:
  2
  5
 15
```

source

`Base.Broadcast.broadcast_setindex!` – Function.

```
broadcast_setindex!(A, X, inds...)
```

Broadcasts the X and `inds` arrays to a common size and stores the value from each position in X at the indices in A given by the same positions in `inds`.

source

50.4   Indexing and assignment

`Base.getindex` – Method.

```
getindex(A, inds...)
```

Returns a subset of array A as specified by `inds`, where each `ind` may be an `Int`, a `Range`, or a `Vector`. See the manual section on array indexing for details.

```julia
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> getindex(A, 1)
1

julia> getindex(A, [2, 1])
2-element Array{Int64,1}:
 3
 1

julia> getindex(A, 2:4)
3-element Array{Int64,1}:
 3
 2
 4
```

source

Base.setindex! – Method.

```
setindex!(A, X, inds...)
```

Store values from array X within some subset of A as specified by inds.

source

Base.copy! – Method.

```
copy!(dest, Rdest::CartesianRange, src, Rsrc::CartesianRange) -> dest
```

Copy the block of src in the range of Rsrc to the block of dest in the range of Rdest. The sizes of the two regions must match.

source

Base.isassigned – Function.

```
isassigned(array, i) -> Bool
```

Tests whether the given array has a value associated with index i. Returns false if the index is out of bounds, or has an undefined reference.

```julia
julia> isassigned(rand(3, 3), 5)
true

julia> isassigned(rand(3, 3), 3 * 3 + 1)
false

julia> mutable struct Foo end

julia> v = similar(rand(3), Foo)
3-element Array{Foo,1}:
 #undef
 #undef
```

```
  #undef

julia> isassigned(v, 1)
  false
```

source

`Base.Colon` – Type.

```
Colon()
```

Colons (:) are used to signify indexing entire objects or dimensions at once.

Very few operations are defined on Colons directly; instead they are converted by `to_indices` to an internal vector type (`Base.Slice`) to represent the collection of indices they span before being used.

source

`Base.IteratorsMD.CartesianIndex` – Type.

```
CartesianIndex(i, j, k...)   -> I
CartesianIndex((i, j, k...)) -> I
```

Create a multidimensional index I, which can be used for indexing a multidimensional array A. In particular, `A[I]` is equivalent to `A[i,j,k...]`. One can freely mix integer and `CartesianIndex` indices; for example, `A[Ipre, i, Ipost]` (where Ipre and Ipost are `CartesianIndex` indices and i is an Int) can be a useful expression when writing algorithms that work along a single dimension of an array of arbitrary dimensionality.

A `CartesianIndex` is sometimes produced by eachindex, and always when iterating with an explicit CartesianRange.

source

`Base.IteratorsMD.CartesianRange` – Type.

```
CartesianRange(Istart::CartesianIndex, Istop::CartesianIndex) -> R
CartesianRange(sz::Dims) -> R
CartesianRange(istart:istop, jstart:jstop, ...) -> R
```

Define a region R spanning a multidimensional rectangular range of integer indices. These are most commonly encountered in the context of iteration, where `for I in R ... end` will return CartesianIndex indices I equivalent to the nested loops

```
for j = jstart:jstop
    for i = istart:istop
        ...
    end
end
```

Consequently these can be useful for writing algorithms that work in arbitrary dimensions.

source

`Base.to_indices` – Function.

```
to_indices(A, I::Tuple)
```

Convert the tuple `I` to a tuple of indices for use in indexing into array A.

The returned tuple must only contain either `Ints` or `AbstractArrays` of scalar indices that are supported by array A. It will error upon encountering a novel index type that it does not know how to process.

For simple index types, it defers to the unexported `Base.to_index(A, i)` to process each index `i`. While this internal function is not intended to be called directly, `Base.to_index` may be extended by custom array or index types to provide custom indexing behaviors.

More complicated index types may require more context about the dimension into which they index. To support those cases, `to_indices(A, I)` calls `to_indices(A, indices(A), I)`, which then recursively walks through both the given tuple of indices and the dimensional indices of A in tandem. As such, not all index types are guaranteed to propagate to `Base.to_index`.

source

`Base.checkbounds` – Function.

```
checkbounds(Bool, A, I...)
```

Return `true` if the specified indices `I` are in bounds for the given array A. Subtypes of `AbstractArray` should specialize this method if they need to provide custom bounds checking behaviors; however, in many cases one can rely on A's indices and `checkindex`.

See also `checkindex`.

```
julia> A = rand(3, 3);

julia> checkbounds(Bool, A, 2)
true

julia> checkbounds(Bool, A, 3, 4)
false

julia> checkbounds(Bool, A, 1:3)
true

julia> checkbounds(Bool, A, 1:3, 2:4)
false
```

source

```
checkbounds(A, I...)
```

Throw an error if the specified indices `I` are not in bounds for the given array A.

source

`Base.checkindex` – Function.

```
checkindex(Bool, inds::AbstractUnitRange, index)
```

Return `true` if the given `index` is within the bounds of `inds`. Custom types that would like to behave as indices for all arrays can extend this method in order to provide a specialized bounds checking implementation.

```
julia> checkindex(Bool,1:20,8)
true

julia> checkindex(Bool,1:20,21)
false
```

50.5   Views (SubArrays and other view types)

Base.view – Function.

```
view(A, inds...)
```

Like getindex, but returns a view into the parent array A with the given indices instead of making a copy. Calling getindex or setindex! on the returned SubArray computes the indices to the parent array on the fly without checking bounds.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> b = view(A, :, 1)
2-element SubArray{Int64,1,Array{Int64,2},Tuple{Base.Slice{Base.OneTo{Int64}},Int64},true}:
 1
 3

julia> fill!(b, 0)
2-element SubArray{Int64,1,Array{Int64,2},Tuple{Base.Slice{Base.OneTo{Int64}},Int64},true}:
 0
 0

julia> A # Note A has changed even though we modified b
2×2 Array{Int64,2}:
 0  2
 0  4
```

Base.@view – Macro.

```
@view A[inds...]
```

Creates a SubArray from an indexing expression. This can only be applied directly to a reference expression (e.g. @view A[1,2:end]), and should *not* be used as the target of an assignment (e.g. @view(A[1,2:end]) = ...). See also @views to switch an entire block of code to use views for slicing.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> b = @view A[:, 1]
2-element SubArray{Int64,1,Array{Int64,2},Tuple{Base.Slice{Base.OneTo{Int64}},Int64},true}:
 1
 3

julia> fill!(b, 0)
2-element SubArray{Int64,1,Array{Int64,2},Tuple{Base.Slice{Base.OneTo{Int64}},Int64},true}:
 0
```

```
 0

julia> A
2×2 Array{Int64,2}:
 0  2
 0  4
```

source

Base.@views – Macro.

```
@views expression
```

Convert every array-slicing operation in the given expression (which may be a `begin`/`end` block, loop, function, etc.) to return a view. Scalar indices, non-array types, and explicit `getindex` calls (as opposed to `array[...]`) are unaffected.

Note that the `@views` macro only affects `array[...]` expressions that appear explicitly in the given expression, not array slicing that occurs in functions called by that code.

source

Base.parent – Function.

```
parent(A)
```

Returns the "parent array" of an array view type (e.g., `SubArray`), or the array itself if it is not a view.

source

Base.parentindexes – Function.

```
parentindexes(A)
```

From an array view A, returns the corresponding indexes in the parent.

source

Base.slicedim – Function.

```
slicedim(A, d::Integer, i)
```

Return all the data of A where the index for dimension d equals i. Equivalent to `A[:, :, ..., i, :, :, ...]` where i is in position d.

```
julia> A = [1 2 3 4; 5 6 7 8]
2×4 Array{Int64,2}:
 1  2  3  4
 5  6  7  8

julia> slicedim(A,2,3)
2-element Array{Int64,1}:
 3
 7
```

source

Base.reinterpret – Function.

```
reinterpret(type, A)
```

Change the type-interpretation of a block of memory. For arrays, this constructs an array with the same binary data as the given array, but with the specified element type. For example, `reinterpret(Float32, UInt32(7))` interprets the 4 bytes corresponding to UInt32(7) as a Float32.

```
julia> reinterpret(Float32, UInt32(7))
1.0f-44

julia> reinterpret(Float32, UInt32[1 2 3 4 5])
1×5 Array{Float32,2}:
 1.4013f-45  2.8026f-45  4.2039f-45  5.60519f-45  7.00649f-45
```

source

Base.reshape – Function.

```
reshape(A, dims...) -> R
reshape(A, dims) -> R
```

Return an array R with the same data as A, but with different dimension sizes or number of dimensions. The two arrays share the same underlying data, so that setting elements of R alters the values of A and vice versa.

The new dimensions may be specified either as a list of arguments or as a shape tuple. At most one dimension may be specified with a `:`, in which case its length is computed such that its product with all the specified dimensions is equal to the length of the original array A. The total number of elements must not change.

```
julia> A = collect(1:16)
16-element Array{Int64,1}:
  1
  2
  3
  4
  5
  6
  7
  8
  9
 10
 11
 12
 13
 14
 15
 16

julia> reshape(A, (4, 4))
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> reshape(A, 2, :)
2×8 Array{Int64,2}:
 1  3  5  7   9  11  13  15
 2  4  6  8  10  12  14  16
```

`Base.squeeze` – Function.

```
squeeze(A, dims)
```

Remove the dimensions specified by `dims` from array `A`. Elements of `dims` must be unique and within the range `1:ndims(A)`. `size(A,i)` must equal `1` for all `i` in `dims`.

```
julia> a = reshape(collect(1:4),(2,2,1,1))
2×2×1×1 Array{Int64,4}:
[:, :, 1, 1] =
 1  3
 2  4

julia> squeeze(a,3)
2×2×1 Array{Int64,3}:
[:, :, 1] =
 1  3
 2  4
```

`Base.vec` – Function.

```
vec(a::AbstractArray) -> Vector
```

Reshape the array `a` as a one-dimensional column vector. The resulting array shares the same underlying data as `a`, so modifying one will also modify the other.

```
julia> a = [1 2 3; 4 5 6]
2×3 Array{Int64,2}:
 1  2  3
 4  5  6

julia> vec(a)
6-element Array{Int64,1}:
 1
 4
 2
 5
 3
 6
```

See also `reshape`.

50.6   Concatenation and permutation

`Base.cat` – Function.

```
cat(dims, A...)
```

Concatenate the input arrays along the specified dimensions in the iterable `dims`. For dimensions not in `dims`, all input arrays should have the same size, which will also be the size of the output array along that dimension. For dimensions in `dims`, the size of the output array is the sum of the sizes of the input arrays along that dimension. If `dims` is a single number, the different arrays are tightly stacked along that dimension. If `dims` is an iterable containing several dimensions, this allows one to construct block diagonal matrices and their higher-dimensional analogues by simultaneously increasing several dimensions for every new input array and putting zero blocks elsewhere. For example, `cat([1,2], matrices...)` builds a block diagonal matrix, i.e. a block matrix with `matrices[1]`, `matrices[2]`, ... as diagonal blocks and matching zero blocks away from the diagonal.

source

Base.vcat – Function.

```
vcat(A...)
```

Concatenate along dimension 1.

```
julia> a = [1 2 3 4 5]
1×5 Array{Int64,2}:
 1  2  3  4  5

julia> b = [6 7 8 9 10; 11 12 13 14 15]
2×5 Array{Int64,2}:
  6   7   8   9  10
 11  12  13  14  15

julia> vcat(a,b)
3×5 Array{Int64,2}:
  1   2   3   4   5
  6   7   8   9  10
 11  12  13  14  15

julia> c = ([1 2 3], [4 5 6])
([1 2 3], [4 5 6])

julia> vcat(c...)
2×3 Array{Int64,2}:
 1  2  3
 4  5  6
```

source

Base.hcat – Function.

```
hcat(A...)
```

Concatenate along dimension 2.

```
julia> a = [1; 2; 3; 4; 5]
5-element Array{Int64,1}:
 1
 2
 3
 4
 5
```

```julia
julia> b = [6 7; 8 9; 10 11; 12 13; 14 15]
5×2 Array{Int64,2}:
  6   7
  8   9
 10  11
 12  13
 14  15

julia> hcat(a,b)
5×3 Array{Int64,2}:
 1   6   7
 2   8   9
 3  10  11
 4  12  13
 5  14  15

julia> c = ([1; 2; 3], [4; 5; 6])
([1, 2, 3], [4, 5, 6])

julia> hcat(c...)
3×2 Array{Int64,2}:
 1  4
 2  5
 3  6
```

source

Base.hvcat – Function.

```
hvcat(rows::Tuple{Vararg{Int}}, values...)
```

Horizontal and vertical concatenation in one call. This function is called for block matrix syntax. The first argument specifies the number of arguments to concatenate in each block row.

```julia
julia> a, b, c, d, e, f = 1, 2, 3, 4, 5, 6
(1, 2, 3, 4, 5, 6)

julia> [a b c; d e f]
2×3 Array{Int64,2}:
 1  2  3
 4  5  6

julia> hvcat((3,3), a,b,c,d,e,f)
2×3 Array{Int64,2}:
 1  2  3
 4  5  6

julia> [a b;c d; e f]
3×2 Array{Int64,2}:
 1  2
 3  4
 5  6

julia> hvcat((2,2,2), a,b,c,d,e,f)
3×2 Array{Int64,2}:
 1  2
```

```
 3   4
 5   6
```

If the first argument is a single integer n, then all block rows are assumed to have n block columns.

source

Base.flipdim – Function.

```
flipdim(A, d::Integer)
```

Reverse A in dimension d.

```
julia> b = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> flipdim(b,2)
2×2 Array{Int64,2}:
 2  1
 4  3
```

source

Base.circshift – Function.

```
circshift(A, shifts)
```

Circularly shift the data in an array. The second argument is a vector giving the amount to shift in each dimension.

```
julia> b = reshape(collect(1:16), (4,4))
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> circshift(b, (0,2))
4×4 Array{Int64,2}:
  9  13  1  5
 10  14  2  6
 11  15  3  7
 12  16  4  8

julia> circshift(b, (-1,0))
4×4 Array{Int64,2}:
 2  6  10  14
 3  7  11  15
 4  8  12  16
 1  5   9  13
```

See also circshift!.

source

Base.circshift! – Function.

```
circshift!(dest, src, shifts)
```

Circularly shift the data in `src`, storing the result in `dest`. `shifts` specifies the amount to shift in each dimension.

The `dest` array must be distinct from the `src` array (they cannot alias each other).

See also `circshift`.

source

Base.circcopy! – Function.

```
circcopy!(dest, src)
```

Copy `src` to `dest`, indexing each dimension modulo its length. `src` and `dest` must have the same size, but can be offset in their indices; any offset results in a (circular) wraparound. If the arrays have overlapping indices, then on the domain of the overlap `dest` agrees with `src`.

```julia
julia> src = reshape(collect(1:16), (4,4))
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> dest = OffsetArray{Int}((0:3,2:5))

julia> circcopy!(dest, src)
OffsetArrays.OffsetArray{Int64,2,Array{Int64,2}} with indices 0:3×2:5:
 8  12  16  4
 5   9  13  1
 6  10  14  2
 7  11  15  3

julia> dest[1:3,2:4] == src[1:3,2:4]
true
```

source

Base.contains – Method.

```
contains(fun, itr, x) -> Bool
```

Returns `true` if there is at least one element y in `itr` such that `fun(y, x)` is `true`.

```julia
julia> vec = [10, 100, 200]
3-element Array{Int64,1}:
  10
 100
 200

julia> contains(==, vec, 200)
true

julia> contains(==, vec, 300)
false
```

```
julia> contains(>, vec, 100)
true

julia> contains(>, vec, 200)
false
```

source

Base.find – Method.

```
find(A)
```

Return a vector of the linear indexes of the non-zeros in A (determined by A[i]!=0). A common use of this is to convert a boolean array to an array of indexes of the true elements. If there are no non-zero elements of A, find returns an empty array.

```
julia> A = [true false; false true]
2×2 Array{Bool,2}:
  true  false
 false   true

julia> find(A)
2-element Array{Int64,1}:
 1
 4
```

source

Base.find – Method.

```
find(f::Function, A)
```

Return a vector I of the linear indexes of A where f(A[I]) returns true. If there are no such elements of A, find returns an empty array.

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> find(isodd,A)
2-element Array{Int64,1}:
 1
 2
```

source

Base.findn – Function.

```
findn(A)
```

Return a vector of indexes for each dimension giving the locations of the non-zeros in A (determined by A[i]!=0). If there are no non-zero elements of A, findn returns a 2-tuple of empty arrays.

```julia
julia> A = [1 2 0; 0 0 3; 0 4 0]
3×3 Array{Int64,2}:
 1  2  0
 0  0  3
 0  4  0

julia> findn(A)
([1, 1, 3, 2], [1, 2, 2, 3])

julia> A = zeros(2,2)
2×2 Array{Float64,2}:
 0.0  0.0
 0.0  0.0

julia> findn(A)
(Int64[], Int64[])
```

source

`Base.findnz` – Function.

```
findnz(A)
```

Return a tuple (I, J, V) where I and J are the row and column indexes of the non-zero values in matrix A, and V is a vector of the non-zero values.

```julia
julia> A = [1 2 0; 0 0 3; 0 4 0]
3×3 Array{Int64,2}:
 1  2  0
 0  0  3
 0  4  0

julia> findnz(A)
([1, 1, 3, 2], [1, 2, 2, 3], [1, 2, 4, 3])
```

source

`Base.findfirst` – Method.

```
findfirst(A)
```

Return the linear index of the first non-zero value in A (determined by A[i]!=0). Returns 0 if no such value is found.

```julia
julia> A = [0 0; 1 0]
2×2 Array{Int64,2}:
 0  0
 1  0

julia> findfirst(A)
2
```

source

`Base.findfirst` – Method.

```
findfirst(A, v)
```

Return the linear index of the first element equal to v in A. Returns 0 if v is not found.

```julia
julia> A = [4 6; 2 2]
2×2 Array{Int64,2}:
 4  6
 2  2

julia> findfirst(A,2)
2

julia> findfirst(A,3)
0
```

source

Base.findfirst – Method.

```julia
findfirst(predicate::Function, A)
```

Return the linear index of the first element of A for which predicate returns true. Returns 0 if there is no such element.

```julia
julia> A = [1 4; 2 2]
2×2 Array{Int64,2}:
 1  4
 2  2

julia> findfirst(iseven, A)
2

julia> findfirst(x -> x>10, A)
0
```

source

Base.findlast – Method.

```julia
findlast(A)
```

Return the linear index of the last non-zero value in A (determined by A[i]!=0). Returns 0 if there is no non-zero value in A.

```julia
julia> A = [1 0; 1 0]
2×2 Array{Int64,2}:
 1  0
 1  0

julia> findlast(A)
2

julia> A = zeros(2,2)
2×2 Array{Float64,2}:
 0.0  0.0
 0.0  0.0

julia> findlast(A)
0
```

```
source
```

Base.findlast – Method.

```
findlast(A, v)
```

Return the linear index of the last element equal to v in A. Returns 0 if there is no element of A equal to v.

```julia
julia> A = [1 2; 2 1]
2×2 Array{Int64,2}:
 1  2
 2  1

julia> findlast(A,1)
4

julia> findlast(A,2)
3

julia> findlast(A,3)
0
```

```
source
```

Base.findlast – Method.

```
findlast(predicate::Function, A)
```

Return the linear index of the last element of A for which predicate returns true. Returns 0 if there is no such element.

```julia
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> findlast(isodd, A)
2

julia> findlast(x -> x > 5, A)
0
```

```
source
```

Base.findnext – Method.

```
findnext(A, i::Integer)
```

Find the next linear index >= i of a non-zero element of A, or 0 if not found.

```julia
julia> A = [0 0; 1 0]
2×2 Array{Int64,2}:
 0  0
 1  0

julia> findnext(A,1)
2
```

```
julia> findnext(A,3)
0
```

source

Base.findnext – Method.

```
findnext(predicate::Function, A, i::Integer)
```

Find the next linear index >= i of an element of A for which predicate returns true, or 0 if not found.

```
julia> A = [1 4; 2 2]
2×2 Array{Int64,2}:
 1  4
 2  2

julia> findnext(isodd, A, 1)
1

julia> findnext(isodd, A, 2)
0
```

source

Base.findnext – Method.

```
findnext(A, v, i::Integer)
```

Find the next linear index >= i of an element of A equal to v (using ==), or 0 if not found.

```
julia> A = [1 4; 2 2]
2×2 Array{Int64,2}:
 1  4
 2  2

julia> findnext(A,4,4)
0

julia> findnext(A,4,3)
3
```

source

Base.findprev – Method.

```
findprev(A, i::Integer)
```

Find the previous linear index <= i of a non-zero element of A, or 0 if not found.

```
julia> A = [0 0; 1 2]
2×2 Array{Int64,2}:
 0  0
 1  2

julia> findprev(A,2)
```

```
2

julia> findprev(A,1)
0
```

source

Base.findprev – Method.

```
findprev(predicate::Function, A, i::Integer)
```

Find the previous linear index <= i of an element of A for which predicate returns true, or 0 if not found.

```
julia> A = [4 6; 1 2]
2×2 Array{Int64,2}:
 4  6
 1  2

julia> findprev(isodd, A, 1)
0

julia> findprev(isodd, A, 3)
2
```

source

Base.findprev – Method.

```
findprev(A, v, i::Integer)
```

Find the previous linear index <= i of an element of A equal to v (using ==), or 0 if not found.

```
julia> A = [0 0; 1 2]
2×2 Array{Int64,2}:
 0  0
 1  2

julia> findprev(A, 1, 4)
2

julia> findprev(A, 1, 1)
0
```

source

Base.permutedims – Function.

```
permutedims(A, perm)
```

Permute the dimensions of array A. perm is a vector specifying a permutation of length ndims(A). This is a generalization of transpose for multi-dimensional arrays. Transpose is equivalent to permutedims(A, [2,1]).

See also: PermutedDimsArray.

```
julia> A = reshape(collect(1:8), (2,2,2))
2×2×2 Array{Int64,3}:
[:, :, 1] =
```

```
 1  3
 2  4

[:, :, 2] =
 5  7
 6  8

julia> permutedims(A, [3, 2, 1])
2×2×2 Array{Int64,3}:
[:, :, 1] =
 1  3
 5  7

[:, :, 2] =
 2  4
 6  8
```

source

Base.permutedims! – Function.

```
permutedims!(dest, src, perm)
```

Permute the dimensions of array `src` and store the result in the array `dest`. `perm` is a vector specifying a permutation of length `ndims(src)`. The preallocated array `dest` should have `size(dest) == size(src)[perm]` and is completely overwritten. No in-place permutation is supported and unexpected results will happen if `src` and `dest` have overlapping memory regions.

source

Base.PermutedDimsArrays.PermutedDimsArray – Type.

```
PermutedDimsArray(A, perm) -> B
```

Given an AbstractArray A, create a view B such that the dimensions appear to be permuted. Similar to `permutedims`, except that no copying occurs (B shares storage with A).

See also: permutedims.

Example

```
julia> A = rand(3,5,4);

julia> B = PermutedDimsArray(A, (3,1,2));

julia> size(B)
(4, 3, 5)

julia> B[3,1,2] == A[1,2,3]
true
```

source

Base.promote_shape – Function.

```
promote_shape(s1, s2)
```

Check two array shapes for compatibility, allowing trailing singleton dimensions, and return whichever shape has more dimensions.

```julia
julia> a = ones(3,4,1,1,1);

julia> b = ones(3,4);

julia> promote_shape(a,b)
(Base.OneTo(3), Base.OneTo(4), Base.OneTo(1), Base.OneTo(1), Base.OneTo(1))

julia> promote_shape((2,3,1,4), (2, 3, 1, 4, 1))
(2, 3, 1, 4, 1)
```

source

50.7   Array functions

Base.accumulate – Method.

```
accumulate(op, A, dim=1)
```

Cumulative operation op along a dimension dim (defaults to 1). See also accumulate! to use a preallocated output array, both for performance and to control the precision of the output (e.g. to avoid overflow). For common operations there are specialized variants of accumulate, see: cumsum, cumprod

```julia
julia> accumulate(+, [1,2,3])
3-element Array{Int64,1}:
 1
 3
 6

julia> accumulate(*, [1,2,3])
3-element Array{Int64,1}:
 1
 2
 6
```

source

```
accumulate(op, v0, A)
```

Like accumulate, but using a starting element v0. The first entry of the result will be op(v0, first(A)). For example:

```julia
julia> accumulate(+, 100, [1,2,3])
3-element Array{Int64,1}:
 101
 103
 106

julia> accumulate(min, 0, [1,2,-1])
3-element Array{Int64,1}:
  0
  0
 -1
```

source

Base.accumulate! – Function.

```
accumulate!(op, B, A, dim=1)
```

Cumulative operation op on A along a dimension, storing the result in B. The dimension defaults to 1. See also accumulate.

source

Base.cumprod – Function.

```
cumprod(A, dim=1)
```

Cumulative product along a dimension dim (defaults to 1). See also cumprod! to use a preallocated output array, both for performance and to control the precision of the output (e.g. to avoid overflow).

```
julia> a = [1 2 3; 4 5 6]
2×3 Array{Int64,2}:
 1  2  3
 4  5  6

julia> cumprod(a,1)
2×3 Array{Int64,2}:
 1   2   3
 4  10  18

julia> cumprod(a,2)
2×3 Array{Int64,2}:
 1   2    6
 4  20  120
```

source

Base.cumprod! – Function.

```
cumprod!(B, A, dim::Integer=1)
```

Cumulative product of A along a dimension, storing the result in B. The dimension defaults to 1. See also cumprod.

source

Base.cumsum – Function.

```
cumsum(A, dim=1)
```

Cumulative sum along a dimension dim (defaults to 1). See also cumsum! to use a preallocated output array, both for performance and to control the precision of the output (e.g. to avoid overflow).

```
julia> a = [1 2 3; 4 5 6]
2×3 Array{Int64,2}:
 1  2  3
 4  5  6

julia> cumsum(a,1)
2×3 Array{Int64,2}:
 1  2  3
```

```
 5  7  9

julia> cumsum(a,2)
2×3 Array{Int64,2}:
 1  3   6
 4  9  15
```

source

`Base.cumsum!` – Function.

```
cumsum!(B, A, dim::Integer=1)
```

Cumulative sum of A along a dimension, storing the result in B. The dimension defaults to 1. See also cumsum.

source

`Base.cumsum_kbn` – Function.

```
cumsum_kbn(A, [dim::Integer=1])
```

Cumulative sum along a dimension, using the Kahan-Babuska-Neumaier compensated summation algorithm for additional accuracy. The dimension defaults to 1.

source

`Base.LinAlg.diff` – Function.

```
diff(A, [dim::Integer=1])
```

Finite difference operator of matrix or vector A. If A is a matrix, compute the finite difference over a dimension dim (default 1).

Example

```
julia> a = [2 4; 6 16]
2×2 Array{Int64,2}:
 2   4
 6  16

julia> diff(a,2)
2×1 Array{Int64,2}:
  2
 10
```

source

`Base.LinAlg.gradient` – Function.

```
gradient(F::AbstractVector, [h::Real])
```

Compute differences along vector F, using h as the spacing between points. The default spacing is one.

Example

```
julia> a = [2,4,6,8];

julia> gradient(a)
```

```
4-element Array{Float64,1}:
 2.0
 2.0
 2.0
 2.0
```

source

Base.rot180 – Function.

```
rot180(A)
```

Rotate matrix A 180 degrees.

```julia
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> rot180(a)
2×2 Array{Int64,2}:
 4  3
 2  1
```

source

```
rot180(A, k)
```

Rotate matrix A 180 degrees an integer k number of times. If k is even, this is equivalent to a copy.

```julia
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> rot180(a,1)
2×2 Array{Int64,2}:
 4  3
 2  1

julia> rot180(a,2)
2×2 Array{Int64,2}:
 1  2
 3  4
```

source

Base.rotl90 – Function.

```
rotl90(A)
```

Rotate matrix A left 90 degrees.

```julia
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
```

```
 3  4
```

```
julia> rotl90(a)
2×2 Array{Int64,2}:
 2  4
 1  3
```

source

```
rotl90(A, k)
```

Rotate matrix A left 90 degrees an integer k number of times. If k is zero or a multiple of four, this is equivalent to a copy.

```
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4
```

```
julia> rotl90(a,1)
2×2 Array{Int64,2}:
 2  4
 1  3
```

```
julia> rotl90(a,2)
2×2 Array{Int64,2}:
 4  3
 2  1
```

```
julia> rotl90(a,3)
2×2 Array{Int64,2}:
 3  1
 4  2
```

```
julia> rotl90(a,4)
2×2 Array{Int64,2}:
 1  2
 3  4
```

source

Base.rotr90 – Function.

```
rotr90(A)
```

Rotate matrix A right 90 degrees.

```
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4
```

```
julia> rotr90(a)
2×2 Array{Int64,2}:
 3  1
 4  2
```

```
rotr90(A, k)
```

Rotate matrix A right 90 degrees an integer k number of times. If k is zero or a multiple of four, this is equivalent to a copy.

```
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> rotr90(a,1)
2×2 Array{Int64,2}:
 3  1
 4  2

julia> rotr90(a,2)
2×2 Array{Int64,2}:
 4  3
 2  1

julia> rotr90(a,3)
2×2 Array{Int64,2}:
 2  4
 1  3

julia> rotr90(a,4)
2×2 Array{Int64,2}:
 1  2
 3  4
```

`Base.reducedim` – Function.

```
reducedim(f, A, region[, v0])
```

Reduce 2-argument function `f` along dimensions of `A`. `region` is a vector specifying the dimensions to reduce, and `v0` is the initial value to use in the reductions. For `+`, `*`, `max` and `min` the `v0` argument is optional.

The associativity of the reduction is implementation-dependent; if you need a particular associativity, e.g. left-to-right, you should write your own loop. See documentation for `reduce`.

```
julia> a = reshape(collect(1:16), (4,4))
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> reducedim(max, a, 2)
4×1 Array{Int64,2}:
 13
 14
 15
 16
```

```
julia> reducedim(max, a, 1)
1×4 Array{Int64,2}:
 4  8  12  16
```

source

Base.mapreducedim – Function.

```
mapreducedim(f, op, A, region[, v0])
```

Evaluates to the same as `reducedim(op, map(f, A), region, f(v0))`, but is generally faster because the intermediate array is avoided.

```
julia> a = reshape(collect(1:16), (4,4))
4×4 Array{Int64,2}:
 1  5   9  13
 2  6  10  14
 3  7  11  15
 4  8  12  16

julia> mapreducedim(isodd, *, a, 1)
1×4 Array{Bool,2}:
 false  false  false  false

julia> mapreducedim(isodd, |, a, 1, true)
1×4 Array{Bool,2}:
 true  true  true  true
```

source

Base.mapslices – Function.

```
mapslices(f, A, dims)
```

Transform the given dimensions of array A using function f. f is called on each slice of A of the form A[..., :, ..., :, ...]. dims is an integer vector specifying where the colons go in this expression. The results are concatenated along the remaining dimensions. For example, if dims is [1,2] and A is 4-dimensional, f is called on A[:,:,i,j] for all i and j.

```
julia> a = reshape(collect(1:16),(2,2,2,2))
2×2×2×2 Array{Int64,4}:
[:, :, 1, 1] =
 1  3
 2  4

[:, :, 2, 1] =
 5  7
 6  8

[:, :, 1, 2] =
  9  11
 10  12

[:, :, 2, 2] =
 13  15
```

```
  14  16

julia> mapslices(sum, a, [1,2])
1×1×2×2 Array{Int64,4}:
[:, :, 1, 1] =
 10

[:, :, 2, 1] =
 26

[:, :, 1, 2] =
 42

[:, :, 2, 2] =
 58
```

source

Base.sum_kbn – *Function.*

```
sum_kbn(A)
```

Returns the sum of all elements of A, using the Kahan-Babuska-Neumaier compensated summation algorithm for additional accuracy.

source

50.8   Combinatorics

Base.Random.randperm – *Function.*

```
randperm([rng=GLOBAL_RNG,] n::Integer)
```

Construct a random permutation of length n. The optional `rng` argument specifies a random number generator (see Random Numbers). To randomly permute a arbitrary vector, see `shuffle` or `shuffle!`.

source

Base.invperm – *Function.*

```
invperm(v)
```

Return the inverse permutation of v. If B = A[v], then A == B[invperm(v)].

```
julia> v = [2; 4; 3; 1];

julia> invperm(v)
4-element Array{Int64,1}:
 4
 1
 3
 2

julia> A = ['a','b','c','d'];

julia> B = A[v]
4-element Array{Char,1}:
```

```
  'b'
  'd'
  'c'
  'a'

julia> B[invperm(v)]
4-element Array{Char,1}:
  'a'
  'b'
  'c'
  'd'
```

source

Base.isperm – Function.

```
isperm(v) -> Bool
```

Returns `true` if v is a valid permutation.

```
julia> isperm([1; 2])
true

julia> isperm([1; 3])
false
```

source

Base.permute! – Method.

```
permute!(v, p)
```

Permute vector v in-place, according to permutation p. No checking is done to verify that p is a permutation.

To return a new permutation, use `v[p]`. Note that this is generally faster than `permute!(v,p)` for large vectors.

See also `ipermute!`

```
julia> A = [1, 1, 3, 4];

julia> perm = [2, 4, 3, 1];

julia> permute!(A, perm);

julia> A
4-element Array{Int64,1}:
 1
 4
 3
 1
```

source

Base.ipermute! – Function.

```
ipermute!(v, p)
```

Like `permute!`, but the inverse of the given permutation is applied.

```julia
julia> A = [1, 1, 3, 4];

julia> perm = [2, 4, 3, 1];

julia> ipermute!(A, perm);

julia> A
4-element Array{Int64,1}:
 4
 1
 3
 1
```

source

Base.Random.randcycle – Function.

```
randcycle([rng=GLOBAL_RNG,] n::Integer)
```

Construct a random cyclic permutation of length n. The optional rng argument specifies a random number generator, see Random Numbers.

source

Base.Random.shuffle – Function.

```
shuffle([rng=GLOBAL_RNG,] v)
```

Return a randomly permuted copy of v. The optional rng argument specifies a random number generator (see Random Numbers). To permute v in-place, see shuffle!. To obtain randomly permuted indices, see randperm.

source

Base.Random.shuffle! – Function.

```
shuffle!([rng=GLOBAL_RNG,] v)
```

In-place version of shuffle: randomly permute the array v in-place, optionally supplying the random-number generator rng.

source

Base.reverse – Function.

```
reverse(v [, start=1 [, stop=length(v) ]] )
```

Return a copy of v reversed from start to stop.

source

Base.reverseind – Function.

```
reverseind(v, i)
```

Given an index i in reverse(v), return the corresponding index in v so that v[reverseind(v,i)] == reverse(v)[i]. (This can be nontrivial in the case where v is a Unicode string.)

source

Base.reverse! – Function.

```
reverse!(v [, start=1 [, stop=length(v) ]]) -> v
```

In-place version of reverse.

source

50.9   BitArrays

BitArrays are space-efficient "packed" boolean arrays, which store one bit per boolean value. They can be used similarly to Array{Bool} arrays (which store one byte per boolean value), and can be converted to/from the latter via Array(bitarray) and BitArray(array), respectively.

Base.flipbits! – Function.

```
flipbits!(B::BitArray{N}) -> BitArray{N}
```

Performs a bitwise not operation on B. See ~.

```julia
julia> A = trues(2,2)
2×2 BitArray{2}:
 true   true
 true   true

julia> flipbits!(A)
2×2 BitArray{2}:
 false   false
 false   false
```

source

Base.rol! – Function.

```
rol!(dest::BitVector, src::BitVector, i::Integer) -> BitVector
```

Performs a left rotation operation on src and puts the result into dest. i controls how far to rotate the bits.

source

```
rol!(B::BitVector, i::Integer) -> BitVector
```

Performs a left rotation operation in-place on B. i controls how far to rotate the bits.

source

Base.rol – Function.

```
rol(B::BitVector, i::Integer) -> BitVector
```

Performs a left rotation operation, returning a new BitVector. i controls how far to rotate the bits. See also rol!.

```julia
julia> A = BitArray([true, true, false, false, true])
5-element BitArray{1}:
   true
   true
  false
  false
```

```
   true

julia> rol(A,1)
5-element BitArray{1}:
   true
  false
  false
   true
   true

julia> rol(A,2)
5-element BitArray{1}:
  false
  false
   true
   true
   true

julia> rol(A,5)
5-element BitArray{1}:
   true
   true
  false
  false
   true
```

source

Base.ror! – Function.

```
ror!(dest::BitVector, src::BitVector, i::Integer) -> BitVector
```

Performs a right rotation operation on `src` and puts the result into `dest`. `i` controls how far to rotate the bits.

source

```
ror!(B::BitVector, i::Integer) -> BitVector
```

Performs a right rotation operation in-place on B. `i` controls how far to rotate the bits.

source

Base.ror – Function.

```
ror(B::BitVector, i::Integer) -> BitVector
```

Performs a right rotation operation on B, returning a new `BitVector`. `i` controls how far to rotate the bits. See also `ror!`.

```
julia> A = BitArray([true, true, false, false, true])
5-element BitArray{1}:
   true
   true
  false
  false
   true

julia> ror(A,1)
```

```
5-element BitArray{1}:
  true
  true
  true
 false
 false

julia> ror(A,2)
5-element BitArray{1}:
 false
  true
  true
  true
 false

julia> ror(A,5)
5-element BitArray{1}:
  true
  true
 false
 false
  true
```

source

50.10 Sparse Vectors and Matrices

Sparse vectors and matrices largely support the same set of operations as their dense counterparts. The following functions are specific to sparse arrays.

Base.SparseArrays.sparse – Function.

```
sparse(A)
```

Convert an AbstractMatrix A into a sparse matrix.

```
julia> A = eye(3)
3×3 Array{Float64,2}:
 1.0  0.0  0.0
 0.0  1.0  0.0
 0.0  0.0  1.0

julia> sparse(A)
3×3 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0
```

source

```
sparse(I, J, V,[ m, n, combine])
```

Create a sparse matrix S of dimensions m x n such that S[I[k], J[k]] = V[k]. The combine function is used to combine duplicates. If m and n are not specified, they are set to maximum(I) and maximum(J) respectively. If the combine function is not supplied, combine defaults to + unless the elements of V are Booleans in which case combine defaults to |. All elements of I must satisfy 1 <= I[k] <= m, and all elements of J must satisfy

`1 <= J[k] <= n`. Numerical zeros in (I, J, V) are retained as structural nonzeros; to drop numerical zeros, use
`dropzeros!`.

For additional documentation and an expert driver, see `Base.SparseArrays.sparse!`.

```julia
julia> Is = [1; 2; 3];

julia> Js = [1; 2; 3];

julia> Vs = [1; 2; 3];

julia> sparse(Is, Js, Vs)
3×3 SparseMatrixCSC{Int64,Int64} with 3 stored entries:
  [1, 1]  =  1
  [2, 2]  =  2
  [3, 3]  =  3
```

source

`Base.SparseArrays.sparsevec` – Function.

```
sparsevec(I, V, [m, combine])
```

Create a sparse vector S of length `m` such that `S[I[k]] = V[k]`. Duplicates are combined using the `combine`
function, which defaults to `+` if no `combine` argument is provided, unless the elements of V are Booleans in which
case `combine` defaults to `|`.

```julia
julia> II = [1, 3, 3, 5]; V = [0.1, 0.2, 0.3, 0.2];

julia> sparsevec(II, V)
5-element SparseVector{Float64,Int64} with 3 stored entries:
  [1]  =  0.1
  [3]  =  0.5
  [5]  =  0.2

julia> sparsevec(II, V, 8, -)
8-element SparseVector{Float64,Int64} with 3 stored entries:
  [1]  =  0.1
  [3]  =  -0.1
  [5]  =  0.2

julia> sparsevec([1, 3, 1, 2, 2], [true, true, false, false, false])
3-element SparseVector{Bool,Int64} with 3 stored entries:
  [1]  =  true
  [2]  =  false
  [3]  =  true
```

source

```
sparsevec(d::Dict, [m])
```

Create a sparse vector of length `m` where the nonzero indices are keys from the dictionary, and the nonzero values
are the values from the dictionary.

```julia
julia> sparsevec(Dict(1 => 3, 2 => 2))
2-element SparseVector{Int64,Int64} with 2 stored entries:
  [1]  =  3
  [2]  =  2
```

source

```
sparsevec(A)
```

Convert a vector A into a sparse vector of length m.

```julia
julia> sparsevec([1.0, 2.0, 0.0, 0.0, 3.0, 0.0])
6-element SparseVector{Float64,Int64} with 3 stored entries:
  [1]  =  1.0
  [2]  =  2.0
  [5]  =  3.0
```

source

`Base.SparseArrays.issparse` – Function.

```
issparse(S)
```

Returns `true` if S is sparse, and `false` otherwise.

source

`Base.full` – Function.

```
full(S)
```

Convert a sparse matrix or vector S into a dense matrix or vector.

```julia
julia> A = speye(3)
3×3 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0

julia> full(A)
3×3 Array{Float64,2}:
 1.0  0.0  0.0
 0.0  1.0  0.0
 0.0  0.0  1.0
```

source

`Base.SparseArrays.nnz` – Function.

```
nnz(A)
```

Returns the number of stored (filled) elements in a sparse array.

```julia
julia> A = speye(3)
3×3 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0

julia> nnz(A)
3
```

source

`Base.SparseArrays.spzeros` – Function.

```
spzeros([type,]m[,n])
```

Create a sparse vector of length m or sparse matrix of size m x n. This sparse array will not contain any nonzero values. No storage will be allocated for nonzero values during construction. The type defaults to `Float64` if not specified.

```julia
julia> spzeros(3, 3)
3×3 SparseMatrixCSC{Float64,Int64} with 0 stored entries

julia> spzeros(Float32, 4)
4-element SparseVector{Float32,Int64} with 0 stored entries
```

source

`Base.SparseArrays.spones` – Function.

```
spones(S)
```

Create a sparse array with the same structure as that of S, but with every nonzero element having the value `1.0`.

```julia
julia> A = sparse([1,2,3,4],[2,4,3,1],[5.,4.,3.,2.])
4×4 SparseMatrixCSC{Float64,Int64} with 4 stored entries:
  [4, 1]  =  2.0
  [1, 2]  =  5.0
  [3, 3]  =  3.0
  [2, 4]  =  4.0

julia> spones(A)
4×4 SparseMatrixCSC{Float64,Int64} with 4 stored entries:
  [4, 1]  =  1.0
  [1, 2]  =  1.0
  [3, 3]  =  1.0
  [2, 4]  =  1.0
```

Note the difference from speye.

source

`Base.SparseArrays.speye` – Method.

```
speye([type,]m[,n])
```

Create a sparse identity matrix of size m x m. When n is supplied, create a sparse identity matrix of size m x n. The type defaults to `Float64` if not specified.

`sparse(I, m, n)` is equivalent to `speye(Int, m, n)`, and `sparse(α*I, m, n)` can be used to efficiently create a sparse multiple α of the identity matrix.

source

`Base.SparseArrays.speye` – Method.

```
speye(S)
```

Create a sparse identity matrix with the same size as S.

```
julia> A = sparse([1,2,3,4],[2,4,3,1],[5.,4.,3.,2.])
4×4 SparseMatrixCSC{Float64,Int64} with 4 stored entries:
  [4, 1]  =  2.0
  [1, 2]  =  5.0
  [3, 3]  =  3.0
  [2, 4]  =  4.0

julia> speye(A)
4×4 SparseMatrixCSC{Float64,Int64} with 4 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0
  [4, 4]  =  1.0
```

Note the difference from spones.

source

```
speye([type,]m[,n])
```

Create a sparse identity matrix of size m x m. When n is supplied, create a sparse identity matrix of size m x n. The type defaults to Float64 if not specified.

sparse(I, m, n) is equivalent to speye(Int, m, n), and sparse(α*I, m, n) can be used to efficiently create a sparse multiple α of the identity matrix.

source

Base.SparseArrays.spdiagm – Function.

```
spdiagm(B, d[, m, n])
```

Construct a sparse diagonal matrix. B is a tuple of vectors containing the diagonals and d is a tuple containing the positions of the diagonals. In the case the input contains only one diagonal, B can be a vector (instead of a tuple) and d can be the diagonal position (instead of a tuple), defaulting to 0 (diagonal). Optionally, m and n specify the size of the resulting sparse matrix.

```
julia> spdiagm(([1,2,3,4],[4,3,2,1]),(-1,1))
5×5 SparseMatrixCSC{Int64,Int64} with 8 stored entries:
  [2, 1]  =  1
  [1, 2]  =  4
  [3, 2]  =  2
  [2, 3]  =  3
  [4, 3]  =  3
  [3, 4]  =  2
  [5, 4]  =  4
  [4, 5]  =  1
```

source

Base.SparseArrays.sprand – Function.

```
sprand([rng],[type],m,[n],p::AbstractFloat,[rfn])
```

Create a random length m sparse vector or m by n sparse matrix, in which the probability of any element being nonzero is independently given by p (and hence the mean density of nonzeros is also exactly p). Nonzero values are sampled from the distribution specified by rfn and have the type type. The uniform distribution is used in case rfn is not specified. The optional rng argument specifies a random number generator, see Random Numbers.

```
julia> rng = MersenneTwister(1234);

julia> sprand(rng, Bool, 2, 2, 0.5)
2×2 SparseMatrixCSC{Bool,Int64} with 2 stored entries:
  [1, 1]  =  true
  [2, 1]  =  true

julia> sprand(rng, Float64, 3, 0.75)
3-element SparseVector{Float64,Int64} with 1 stored entry:
  [3]  =  0.298614
```

source

Base.SparseArrays.sprandn – Function.

```
sprandn([rng], m[,n],p::AbstractFloat)
```

Create a random sparse vector of length m or sparse matrix of size m by n with the specified (independent) proba-bility p of any entry being nonzero, where nonzero values are sampled from the normal distribution. The optional rng argument specifies a random number generator, see Random Numbers.

```
julia> rng = MersenneTwister(1234);

julia> sprandn(rng, 2, 2, 0.75)
2×2 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =   0.532813
  [2, 1]  =  -0.271735
  [2, 2]  =   0.502334
```

source

Base.SparseArrays.nonzeros – Function.

```
nonzeros(A)
```

Return a vector of the structural nonzero values in sparse array A. This includes zeros that are explicitly stored in the sparse array. The returned vector points directly to the internal nonzero storage of A, and any modifications to the returned vector will mutate A as well. See rowvals and nzrange.

```
julia> A = speye(3)
3×3 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0

julia> nonzeros(A)
3-element Array{Float64,1}:
 1.0
 1.0
 1.0
```

source

Base.SparseArrays.rowvals – Function.

```
rowvals(A::SparseMatrixCSC)
```

Return a vector of the row indices of A. Any modifications to the returned vector will mutate A as well. Providing access to how the row indices are stored internally can be useful in conjunction with iterating over structural nonzero values. See also nonzeros and nzrange.

```julia
julia> A = speye(3)
3×3 SparseMatrixCSC{Float64,Int64} with 3 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0

julia> rowvals(A)
3-element Array{Int64,1}:
 1
 2
 3
```

source

Base.SparseArrays.nzrange – Function.

```
nzrange(A::SparseMatrixCSC, col::Integer)
```

Return the range of indices to the structural nonzero values of a sparse matrix column. In conjunction with nonzeros and rowvals, this allows for convenient iterating over a sparse matrix :

```
A = sparse(I,J,V)
rows = rowvals(A)
vals = nonzeros(A)
m, n = size(A)
for i = 1:n
   for j in nzrange(A, i)
      row = rows[j]
      val = vals[j]
      # perform sparse wizardry...
   end
end
```

source

Base.SparseArrays.dropzeros! – Method.

```
dropzeros!(A::SparseMatrixCSC, trim::Bool = true)
```

Removes stored numerical zeros from A, optionally trimming resulting excess space from A.rowval and A.nzval when trim is true.

For an out-of-place version, see dropzeros. For algorithmic information, see fkeep!.

source

Base.SparseArrays.dropzeros – Method.

```
dropzeros(A::SparseMatrixCSC, trim::Bool = true)
```

Generates a copy of A and removes stored numerical zeros from that copy, optionally trimming excess space from the result's rowval and nzval arrays when trim is true.

For an in-place version and algorithmic information, see dropzeros!.

source

Base.SparseArrays.dropzeros! – Method.

```
dropzeros!(x::SparseVector, trim::Bool = true)
```

Removes stored numerical zeros from x, optionally trimming resulting excess space from x.nzind and x.nzval when trim is true.

For an out-of-place version, see dropzeros. For algorithmic information, see fkeep!.

source

Base.SparseArrays.dropzeros – Method.

```
dropzeros(x::SparseVector, trim::Bool = true)
```

Generates a copy of x and removes numerical zeros from that copy, optionally trimming excess space from the result's nzind and nzval arrays when trim is true.

For an in-place version and algorithmic information, see dropzeros!.

source

Base.SparseArrays.permute – Function.

```
permute{Tv,Ti}(A::SparseMatrixCSC{Tv,Ti}, p::AbstractVector{<:Integer},
    q::AbstractVector{<:Integer})
```

Bilaterally permute A, returning PAQ (A[p,q]). Column-permutation q's length must match A's column count (length(q) == A.n). Row-permutation p's length must match A's row count (length(p) == A.m).

For expert drivers and additional information, see permute!.

source

Base.permute! – Method.

```
permute!{Tv,Ti}(X::SparseMatrixCSC{Tv,Ti}, A::SparseMatrixCSC{Tv,Ti},
    p::AbstractVector{<:Integer}, q::AbstractVector{<:Integer}[, C::SparseMatrixCSC{Tv,Ti}])
```

Bilaterally permute A, storing result PAQ (A[p,q]) in X. Stores intermediate result (AQ)^T (transpose(A[:,q])) in optional argument C if present. Requires that none of X, A, and, if present, C alias each other; to store result PAQ back into A, use the following method lacking X:

```
permute!{Tv,Ti}(A::SparseMatrixCSC{Tv,Ti}, p::AbstractVector{<:Integer},
    q::AbstractVector{<:Integer}[, C::SparseMatrixCSC{Tv,Ti}[, workcolptr::Vector{Ti}]])
```

X's dimensions must match those of A (X.m == A.m and X.n == A.n), and X must have enough storage to accommodate all allocated entries in A (length(X.rowval) >= nnz(A) and length(X.nzval) >= nnz(A)). Column-permutation q's length must match A's column count (length(q) == A.n). Row-permutation p's length must match A's row count (length(p) == A.m).

C's dimensions must match those of transpose(A) (C.m == A.n and C.n == A.m), and C must have enough storage to accommodate all allocated entries in A (length(C.rowval) >= nnz(A) and length(C.nzval) >= nnz(A)).

For additional (algorithmic) information, and for versions of these methods that forgo argument checking, see (unexported) parent methods unchecked_noalias_permute! and unchecked_aliasing_permute!.

See also: permute

source

Chapter 51

Tasks and Parallel Computing

51.1 Tasks

Core.Task – Type.

```
Task(func)
```

Create a Task (i.e. coroutine) to execute the given function (which must be callable with no arguments). The task exits when this function returns.

```julia
julia> a() = det(rand(1000, 1000));

julia> b = Task(a);
```

In this example, b is a runnable Task that hasn't started yet.

source

Base.current_task – Function.

```
current_task()
```

Get the currently running Task.

source

Base.istaskdone – Function.

```
istaskdone(t::Task) -> Bool
```

Determine whether a task has exited.

```julia
julia> a2() = det(rand(1000, 1000));

julia> b = Task(a2);

julia> istaskdone(b)
false

julia> schedule(b);

julia> yield();
```

```
julia> istaskdone(b)
true
```

source

Base.istaskstarted – Function.

```
istaskstarted(t::Task) -> Bool
```

Determine whether a task has started executing.

```
julia> a3() = det(rand(1000, 1000));

julia> b = Task(a3);

julia> istaskstarted(b)
false
```

source

Base.yield – Function.

```
yield()
```

Switch to the scheduler to allow another scheduled task to run. A task that calls this function is still runnable, and will be restarted immediately if there are no other runnable tasks.

source

```
yield(t::Task, arg = nothing)
```

A fast, unfair-scheduling version of `schedule(t, arg); yield()` which immediately yields to `t` before calling the scheduler.

source

Base.yieldto – Function.

```
yieldto(t::Task, arg = nothing)
```

Switch to the given task. The first time a task is switched to, the task's function is called with no arguments. On subsequent switches, `arg` is returned from the task's last call to `yieldto`. This is a low-level call that only switches tasks, not considering states or scheduling in any way. Its use is discouraged.

source

Base.task_local_storage – Method.

```
task_local_storage(key)
```

Look up the value of a key in the current task's task-local storage.

source

Base.task_local_storage – Method.

```
task_local_storage(key, value)
```

Assign a value to a key in the current task's task-local storage.

source

Base.task_local_storage – Method.

```
task_local_storage(body, key, value)
```

Call the function body with a modified task-local storage, in which `value` is assigned to `key`; the previous value of `key`, or lack thereof, is restored afterwards. Useful for emulating dynamic scoping.

source

Base.Condition – Type.

```
Condition()
```

Create an edge-triggered event source that tasks can wait for. Tasks that call `wait` on a `Condition` are suspended and queued. Tasks are woken up when `notify` is later called on the `Condition`. Edge triggering means that only tasks waiting at the time `notify` is called can be woken up. For level-triggered notifications, you must keep extra state to keep track of whether a notification has happened. The `Channel` type does this, and so can be used for level-triggered events.

source

Base.notify – Function.

```
notify(condition, val=nothing; all=true, error=false)
```

Wake up tasks waiting for a condition, passing them `val`. If `all` is `true` (the default), all waiting tasks are woken, otherwise only one is. If `error` is `true`, the passed value is raised as an exception in the woken tasks.

Returns the count of tasks woken up. Returns 0 if no tasks are waiting on `condition`.

source

Base.schedule – Function.

```
schedule(t::Task, [val]; error=false)
```

Add a `Task` to the scheduler's queue. This causes the task to run constantly when the system is otherwise idle, unless the task performs a blocking operation such as `wait`.

If a second argument `val` is provided, it will be passed to the task (via the return value of `yieldto`) when it runs again. If `error` is `true`, the value is raised as an exception in the woken task.

```julia
julia> a5() = det(rand(1000, 1000));

julia> b = Task(a5);

julia> istaskstarted(b)
false

julia> schedule(b);

julia> yield();

julia> istaskstarted(b)
true

julia> istaskdone(b)
true
```

source

Base.@schedule – Macro.

```
@schedule
```

Wrap an expression in a Task and add it to the local machine's scheduler queue. Similar to @async except that an enclosing @sync does NOT wait for tasks started with an @schedule.

source

Base.@task – Macro.

```
@task
```

Wrap an expression in a Task without executing it, and return the Task. This only creates a task, and does not run it.

```julia
julia> a1() = det(rand(1000, 1000));

julia> b = @task a1();

julia> istaskstarted(b)
false

julia> schedule(b);

julia> yield();

julia> istaskdone(b)
true
```

source

Base.sleep – Function.

```
sleep(seconds)
```

Block the current task for a specified number of seconds. The minimum sleep time is 1 millisecond or input of 0.001.

source

Base.Channel – Type.

```
Channel{T}(sz::Int)
```

Constructs a Channel with an internal buffer that can hold a maximum of sz objects of type T. put! calls on a full channel block until an object is removed with take!.

Channel(0) constructs an unbuffered channel. put! blocks until a matching take! is called. And vice-versa.

Other constructors:

- Channel(Inf): equivalent to Channel{Any}(typemax(Int))
- Channel(sz): equivalent to Channel{Any}(sz)

source

Base.put! – Method.

```
put!(c::Channel, v)
```

Appends an item v to the channel c. Blocks if the channel is full.

For unbuffered channels, blocks until a take! is performed by a different task.

source

Base.take! – Method.

```
take!(c::Channel)
```

Removes and returns a value from a Channel. Blocks until data is available.

For unbuffered channels, blocks until a put! is performed by a different task.

source

Base.isready – Method.

```
isready(c::Channel)
```

Determine whether a Channel has a value stored to it. Returns immediately, does not block.

For unbuffered channels returns true if there are tasks waiting on a put!.

source

Base.fetch – Method.

```
fetch(c::Channel)
```

Waits for and gets the first available item from the channel. Does not remove the item. fetch is unsupported on an unbuffered (0-size) channel.

source

Base.close – Method.

```
close(c::Channel)
```

Closes a channel. An exception is thrown by:

- put! on a closed channel.
- take! and fetch on an empty, closed channel.

source

Base.bind – Method.

```
bind(chnl::Channel, task::Task)
```

Associates the lifetime of chnl with a task. Channel chnl is automatically closed when the task terminates. Any uncaught exception in the task is propagated to all waiters on chnl.

The chnl object can be explicitly closed independent of task termination. Terminating tasks have no effect on already closed Channel objects.

When a channel is bound to multiple tasks, the first task to terminate will close the channel. When multiple channels are bound to the same task, termination of the task will close all of the bound channels.

```julia
julia> c = Channel(0);

julia> task = @schedule foreach(i->put!(c, i), 1:4);

julia> bind(c,task);

julia> for i in c
            @show i
       end;
i = 1
i = 2
i = 3
i = 4

julia> isopen(c)
false
```

```julia
julia> c = Channel(0);

julia> task = @schedule (put!(c,1);error("foo"));

julia> bind(c,task);

julia> take!(c)
1

julia> put!(c,1);
ERROR: foo
Stacktrace:
 [1] check_channel_state(::Channel{Any}) at ./channels.jl:131
 [2] put!(::Channel{Any}, ::Int64) at ./channels.jl:261
```

source

Base.asyncmap – Function.

```julia
asyncmap(f, c...; ntasks=0, batch_size=nothing)
```

Uses multiple concurrent tasks to map f over a collection (or multiple equal length collections). For multiple collection arguments, f is applied elementwise.

ntasks specifies the number of tasks to run concurrently. Depending on the length of the collections, if ntasks is unspecified, up to 100 tasks will be used for concurrent mapping.

ntasks can also be specified as a zero-arg function. In this case, the number of tasks to run in parallel is checked before processing every element and a new task started if the value of ntasks_func() is less than the current number of tasks.

If batch_size is specified, the collection is processed in batch mode. f must then be a function that must accept a Vector of argument tuples and must return a vector of results. The input vector will have a length of batch_size or less.

The following examples highlight execution in different tasks by returning the object_id of the tasks in which the mapping function is executed.

First, with ntasks undefined, each element is processed in a different task.

```
julia> tskoid() = object_id(current_task());

julia> asyncmap(x->tskoid(), 1:5)
5-element Array{UInt64,1}:
 0x6e15e66c75c75853
 0x440f8819a1baa682
 0x9fb3eeadd0c83985
 0xebd3e35fe90d4050
 0x29efc93edce2b961

julia> length(unique(asyncmap(x->tskoid(), 1:5)))
5
```

With `ntasks=2` all elements are processed in 2 tasks.

```
julia> asyncmap(x->tskoid(), 1:5; ntasks=2)
5-element Array{UInt64,1}:
 0x027ab1680df7ae94
 0xa23d2f80cd7cf157
 0x027ab1680df7ae94
 0xa23d2f80cd7cf157
 0x027ab1680df7ae94

julia> length(unique(asyncmap(x->tskoid(), 1:5; ntasks=2)))
2
```

With `batch_size` defined, the mapping function needs to be changed to accept an array of argument tuples and return an array of results. map is used in the modified mapping function to achieve this.

```
julia> batch_func(input) = map(x->string("args_tuple: ", x, ", element_val: ", x[1], ", task:
        ", tskoid()), input)
batch_func (generic function with 1 method)

julia> asyncmap(batch_func, 1:5; ntasks=2, batch_size=2)
5-element Array{String,1}:
 "args_tuple: (1,), element_val: 1, task: 9118321258196414413"
 "args_tuple: (2,), element_val: 2, task: 4904288162898683522"
 "args_tuple: (3,), element_val: 3, task: 9118321258196414413"
 "args_tuple: (4,), element_val: 4, task: 4904288162898683522"
 "args_tuple: (5,), element_val: 5, task: 9118321258196414413"
```

> **Note**
>
> Currently, all tasks in Julia are executed in a single OS thread co-operatively. Consequently, ayncmap is beneficial only when the mapping function involves any I/O - disk, network, remote worker invocation, etc.

source

Base.asyncmap! – Function.

```
asyncmap!(f, results, c...; ntasks=0, batch_size=nothing)
```

Like asyncmap(), but stores output in `results` rather than returning a collection.

source

51.2   General Parallel Computing Support

`Base.Distributed.addprocs` – Function.

```
addprocs(manager::ClusterManager; kwargs...) -> List of process identifiers
```

Launches worker processes via the specified cluster manager.

For example, Beowulf clusters are supported via a custom cluster manager implemented in the package `Clus-terManagers.jl`.

The number of seconds a newly launched worker waits for connection establishment from the master can be specified via variable `JULIA_WORKER_TIMEOUT` in the worker process's environment. Relevant only when using TCP/IP as transport.

source

```
addprocs(machines; tunnel=false, sshflags=``, max_parallel=10, kwargs...) -> List of process
    identifiers
```

Add processes on remote machines via SSH. Requires `julia` to be installed in the same location on each node, or to be available via a shared file system.

`machines` is a vector of machine specifications. Workers are started for each specification.

A machine specification is either a string `machine_spec` or a tuple - `(machine_spec, count)`.

`machine_spec` is a string of the form `[user@]host[:port] [bind_addr[:port]]`. `user` defaults to current user, `port` to the standard ssh port. If `[bind_addr[:port]]` is specified, other workers will connect to this worker at the specified `bind_addr` and `port`.

`count` is the number of workers to be launched on the specified host. If specified as `:auto` it will launch as many workers as the number of cores on the specific host.

Keyword arguments:

- `tunnel`: if `true` then SSH tunneling will be used to connect to the worker from the master process. Default is `false`.

- `sshflags`: specifies additional ssh options, e.g. sshflags=`` `-i /home/foo/bar.pem` ``

- `max_parallel`: specifies the maximum number of workers connected to in parallel at a host. Defaults to 10.

- `dir`: specifies the working directory on the workers. Defaults to the host's current directory (as found by `pwd()`)

- `enable_threaded_blas`: if `true` then BLAS will run on multiple threads in added processes. Default is `false`.

- `exename`: name of the `julia` executable. Defaults to `"$JULIA_HOME/julia"` or `"$JULIA_HOME/julia-debug"` as the case may be.

- `exeflags`: additional flags passed to the worker processes.

- `topology`: Specifies how the workers connect to each other. Sending a message between unconnected workers results in an error.

    - `topology=:all_to_all`: All processes are connected to each other. The default.

    - `topology=:master_slave`: Only the driver process, i.e. `pid` 1 connects to the workers. The workers do not connect to each other.

      – `topology=:custom`: The `launch` method of the cluster manager specifies the connection topology via fields `ident` and `connect_idents` in `WorkerConfig`. A worker with a cluster manager identity `ident` will connect to all workers specified in `connect_idents`.

Environment variables :

If the master process fails to establish a connection with a newly launched worker within 60.0 seconds, the worker treats it as a fatal situation and terminates. This timeout can be controlled via environment variable JU-LIA_WORKER_TIMEOUT. The value of `JULIA_WORKER_TIMEOUT` on the master process specifies the number of seconds a newly launched worker waits for connection establishment.

[source](#)

```
addprocs(; kwargs...) -> List of process identifiers
```

Equivalent to `addprocs(Sys.CPU_CORES; kwargs...)`

Note that workers do not run a `.juliarc.jl` startup script, nor do they synchronize their global state (such as global variables, new method definitions, and loaded modules) with any of the other running processes.

[source](#)

```
addprocs(np::Integer; restrict=true, kwargs...) -> List of process identifiers
```

Launches workers using the in-built `LocalManager` which only launches workers on the local host. This can be used to take advantage of multiple cores. `addprocs(4)` will add 4 processes on the local machine. If `restrict` is `true`, binding is restricted to `127.0.0.1`. Keyword args `dir`, `exename`, `exeflags`, `topology`, and `enable_threaded_blas` have the same effect as documented for `addprocs(machines)`.

[source](#)

`Base.Distributed.nprocs` – Function.

```
nprocs()
```

Get the number of available processes.

[source](#)

`Base.Distributed.nworkers` – Function.

```
nworkers()
```

Get the number of available worker processes. This is one less than `nprocs()`. Equal to `nprocs()` if `nprocs() == 1`.

[source](#)

`Base.Distributed.procs` – Method.

```
procs()
```

Returns a list of all process identifiers.

[source](#)

`Base.Distributed.procs` – Method.

```
procs(pid::Integer)
```

Returns a list of all process identifiers on the same physical node. Specifically all workers bound to the same ip-address as `pid` are returned.

source

`Base.Distributed.workers` – Function.

```
workers()
```

Returns a list of all worker process identifiers.

source

`Base.Distributed.rmprocs` – Function.

```
rmprocs(pids...; waitfor=typemax(Int))
```

Removes the specified workers. Note that only process 1 can add or remove workers.

Argument `waitfor` specifies how long to wait for the workers to shut down: - If unspecified, `rmprocs` will wait until all requested `pids` are removed. - An `ErrorException` is raised if all workers cannot be terminated before the requested `waitfor` seconds. - With a `waitfor` value of 0, the call returns immediately with the workers scheduled for removal in a different task. The scheduled `Task` object is returned. The user should call `wait` on the task before invoking any other parallel calls.

source

`Base.Distributed.interrupt` – Function.

```
interrupt(pids::Integer...)
```

Interrupt the current executing task on the specified workers. This is equivalent to pressing Ctrl-C on the local machine. If no arguments are given, all workers are interrupted.

source

```
interrupt(pids::AbstractVector=workers())
```

Interrupt the current executing task on the specified workers. This is equivalent to pressing Ctrl-C on the local machine. If no arguments are given, all workers are interrupted.

source

`Base.Distributed.myid` – Function.

```
myid()
```

Get the id of the current process.

source

`Base.Distributed.pmap` – Function.

```
pmap([::AbstractWorkerPool], f, c...; distributed=true, batch_size=1, on_error=nothing,
    retry_delays=[]), retry_check=nothing) -> collection
```

Transform collection c by applying f to each element using available workers and tasks.

For multiple collection arguments, apply f elementwise.

Note that f must be made available to all worker processes; see Code Availability and Loading Packages for details.

If a worker pool is not specified, all available workers, i.e., the default worker pool is used.

By default, pmap distributes the computation over all specified workers. To use only the local process and distribute over tasks, specify `distributed=false`. This is equivalent to using asyncmap. For example, `pmap(f, c; distributed=false)` is equivalent to `asyncmap(f,c; ntasks=()->nworkers())`

pmap can also use a mix of processes and tasks via the `batch_size` argument. For batch sizes greater than 1, the collection is processed in multiple batches, each of length `batch_size` or less. A batch is sent as a single request to a free worker, where a local asyncmap processes elements from the batch using multiple concurrent tasks.

Any error stops pmap from processing the remainder of the collection. To override this behavior you can specify an error handling function via argument `on_error` which takes in a single argument, i.e., the exception. The function can stop the processing by rethrowing the error, or, to continue, return any value which is then returned inline with the results to the caller.

Consider the following two examples. The first one returns the exception object inline, the second a 0 in place of any exception:

```julia
julia> pmap(x->iseven(x) ? error("foo") : x, 1:4; on_error=identity)
4-element Array{Any,1}:
 1
  ErrorException("foo")
 3
  ErrorException("foo")

julia> pmap(x->iseven(x) ? error("foo") : x, 1:4; on_error=ex->0)
4-element Array{Int64,1}:
 1
 0
 3
 0
```

Errors can also be handled by retrying failed computations. Keyword arguments `retry_delays` and `retry_check` are passed through to retry as keyword arguments `delays` and `check` respectively. If batching is specified, and an entire batch fails, all items in the batch are retried.

Note that if both `on_error` and `retry_delays` are specified, the `on_error` hook is called before retrying. If `on_error` does not throw (or rethrow) an exception, the element will not be retried.

Example: On errors, retry f on an element a maximum of 3 times without any delay between retries.

```julia
pmap(f, c; retry_delays = zeros(3))
```

Example: Retry f only if the exception is not of type `InexactError`, with exponentially increasing delays up to 3 times. Return a NaN in place for all `InexactError` occurrences.

```julia
pmap(f, c; on_error = e->(isa(e, InexactError) ? NaN : rethrow(e)), retry_delays =
↪    ExponentialBackOff(n = 3))
```

source

`Base.Distributed.RemoteException` – Type.

```julia
RemoteException(captured)
```

Exceptions on remote computations are captured and rethrown locally. A `RemoteException` wraps the `pid` of the worker and a captured exception. A `CapturedException` captures the remote exception and a serializable form of the call stack when the exception was raised.

source

`Base.Distributed.Future` – Type.

```
Future(pid::Integer=myid())
```

Create a `Future` on process `pid`. The default `pid` is the current process.

    source

`Base.Distributed.RemoteChannel` – Method.

```
RemoteChannel(pid::Integer=myid())
```

Make a reference to a `Channel{Any}(1)` on process `pid`. The default `pid` is the current process.

    source

`Base.Distributed.RemoteChannel` – Method.

```
RemoteChannel(f::Function, pid::Integer=myid())
```

Create references to remote channels of a specific size and type. `f()` is a function that when executed on `pid` must return an implementation of an `AbstractChannel`.

For example, `RemoteChannel(()->Channel{Int}(10), pid)`, will return a reference to a channel of type `Int` and size 10 on `pid`.

The default `pid` is the current process.

    source

`Base.wait` – Function.

```
wait([x])
```

Block the current task until some event occurs, depending on the type of the argument:

- `RemoteChannel` : Wait for a value to become available on the specified remote channel.
- `Future` : Wait for a value to become available for the specified future.
- `Channel`: Wait for a value to be appended to the channel.
- `Condition`: Wait for `notify` on a condition.
- `Process`: Wait for a process or process chain to exit. The `exitcode` field of a process can be used to determine success or failure.
- `Task`: Wait for a `Task` to finish, returning its result value. If the task fails with an exception, the exception is propagated (re-thrown in the task that called `wait`).
- RawFD: Wait for changes on a file descriptor (see `poll_fd` for keyword arguments and return code)

If no argument is passed, the task blocks for an undefined period. A task can only be restarted by an explicit call to `schedule` or `yieldto`.

Often `wait` is called within a `while` loop to ensure a waited-for condition is met before proceeding.

    source

`Base.fetch` – Method.

```
fetch(x)
```

Waits and fetches a value from x depending on the type of x:

- `Future`: Wait for and get the value of a `Future`. The fetched value is cached locally. Further calls to `fetch` on the same reference return the cached value. If the remote value is an exception, throws a `RemoteException` which captures the remote exception and backtrace.
- `RemoteChannel`: Wait for and get the value of a remote reference. Exceptions raised are same as for a `Future`.

Does not remove the item fetched.

source

`Base.Distributed.remotecall` – Method.

```
remotecall(f, id::Integer, args...; kwargs...) -> Future
```

Call a function f asynchronously on the given arguments on the specified process. Returns a `Future`. Keyword arguments, if any, are passed through to f.

source

`Base.Distributed.remotecall_wait` – Method.

```
remotecall_wait(f, id::Integer, args...; kwargs...)
```

Perform a faster `wait(remotecall(...))` in one message on the `Worker` specified by worker id `id`. Keyword arguments, if any, are passed through to f.

See also `wait` and `remotecall`.

source

`Base.Distributed.remotecall_fetch` – Method.

```
remotecall_fetch(f, id::Integer, args...; kwargs...)
```

Perform `fetch(remotecall(...))` in one message. Keyword arguments, if any, are passed through to f. Any remote exceptions are captured in a `RemoteException` and thrown.

See also `fetch` and `remotecall`.

source

`Base.Distributed.remote_do` – Method.

```
remote_do(f, id::Integer, args...; kwargs...) -> nothing
```

Executes f on worker `id` asynchronously. Unlike `remotecall`, it does not store the result of computation, nor is there a way to wait for its completion.

A successful invocation indicates that the request has been accepted for execution on the remote node.

While consecutive `remotecall`s to the same worker are serialized in the order they are invoked, the order of executions on the remote worker is undetermined. For example, `remote_do(f1, 2); remotecall(f2, 2); remote_do(f3, 2)` will serialize the call to f1, followed by f2 and f3 in that order. However, it is not guaranteed that f1 is executed before f3 on worker 2.

Any exceptions thrown by f are printed to `STDERR` on the remote worker.

Keyword arguments, if any, are passed through to f.

source

Base.put! – Method.

```
put!(rr::RemoteChannel, args...)
```

Store a set of values to the RemoteChannel. If the channel is full, blocks until space is available. Returns its first argument.

source

Base.put! – Method.

```
put!(rr::Future, v)
```

Store a value to a Future rr. Futures are write-once remote references. A put! on an already set Future throws an Exception. All asynchronous remote calls return Futures and set the value to the return value of the call upon completion.

source

Base.take! – Method.

```
take!(rr::RemoteChannel, args...)
```

Fetch value(s) from a RemoteChannel rr, removing the value(s) in the processs.

source

Base.isready – Method.

```
isready(rr::RemoteChannel, args...)
```

Determine whether a RemoteChannel has a value stored to it. Note that this function can cause race conditions, since by the time you receive its result it may no longer be true. However, it can be safely used on a Future since they are assigned only once.

source

Base.isready – Method.

```
isready(rr::Future)
```

Determine whether a Future has a value stored to it.

If the argument Future is owned by a different node, this call will block to wait for the answer. It is recommended to wait for rr in a separate task instead or to use a local Channel as a proxy:

```
c = Channel(1)
@async put!(c, remotecall_fetch(long_computation, p))
isready(c)  # will not block
```

source

Base.Distributed.WorkerPool – Type.

```
WorkerPool(workers::Vector{Int})
```

Create a WorkerPool from a vector of worker ids.

source

Base.Distributed.CachingPool – Type.

```
CachingPool(workers::Vector{Int})
```

An implementation of an `AbstractWorkerPool`. `remote`, `remotecall_fetch`, `pmap` (and other remote calls which execute functions remotely) benefit from caching the serialized/deserialized functions on the worker nodes, especially closures (which may capture large amounts of data).

The remote cache is maintained for the lifetime of the returned `CachingPool` object. To clear the cache earlier, use `clear!(pool)`.

For global variables, only the bindings are captured in a closure, not the data. `let` blocks can be used to capture global data.

For example:

```
const foo=rand(10^8);
wp=CachingPool(workers())
let foo=foo
    pmap(wp, i->sum(foo)+i, 1:100);
end
```

The above would transfer `foo` only once to each worker.

source

Base.Distributed.default_worker_pool – Function.

```
default_worker_pool()
```

`WorkerPool` containing idle `workers()` - used by `remote(f)` and `pmap` (by default).

source

Base.Distributed.clear! – Method.

```
clear!(pool::CachingPool) -> pool
```

Removes all cached functions from all participating workers.

source

Base.Distributed.remote – Function.

```
remote([::AbstractWorkerPool], f) -> Function
```

Returns an anonymous function that executes function `f` on an available worker using `remotecall_fetch`.

source

Base.Distributed.remotecall – Method.

```
remotecall(f, pool::AbstractWorkerPool, args...; kwargs...) -> Future
```

`WorkerPool` variant of `remotecall(f, pid, ....)`. Waits for and takes a free worker from `pool` and performs a `remotecall` on it.

source

Base.Distributed.remotecall_wait – Method.

```
remotecall_wait(f, pool::AbstractWorkerPool, args...; kwargs...) -> Future
```

`WorkerPool` variant of `remotecall_wait(f, pid, ....)`. Waits for and takes a free worker from `pool` and performs a `remotecall_wait` on it.

> source

`Base.Distributed.remotecall_fetch` – Method.

```
remotecall_fetch(f, pool::AbstractWorkerPool, args...; kwargs...) -> result
```

`WorkerPool` variant of `remotecall_fetch(f, pid, ....)`. Waits for and takes a free worker from `pool` and performs a `remotecall_fetch` on it.

> source

`Base.Distributed.remote_do` – Method.

```
remote_do(f, pool::AbstractWorkerPool, args...; kwargs...) -> nothing
```

`WorkerPool` variant of `remote_do(f, pid, ....)`. Waits for and takes a free worker from `pool` and performs a `remote_do` on it.

> source

`Base.timedwait` – Function.

```
timedwait(testcb::Function, secs::Float64; pollint::Float64=0.1)
```

Waits until `testcb` returns `true` or for `secs` seconds, whichever is earlier. `testcb` is polled every `pollint` seconds.

> source

`Base.Distributed.@spawn` – Macro.

```
@spawn
```

Creates a closure around an expression and runs it on an automatically-chosen process, returning a `Future` to the result.

> source

`Base.Distributed.@spawnat` – Macro.

```
@spawnat
```

Accepts two arguments, p and an expression. A closure is created around the expression and run asynchronously on process p. Returns a `Future` to the result.

> source

`Base.Distributed.@fetch` – Macro.

```
@fetch
```

Equivalent to `fetch(@spawn expr)`. See `fetch` and `@spawn`.

> source

`Base.Distributed.@fetchfrom` – Macro.

```
@fetchfrom
```

Equivalent to `fetch(@spawnat p expr)`. See `fetch` and `@spawnat`.

source

`Base.@async` – Macro.

```
@async
```

Like `@schedule`, `@async` wraps an expression in a `Task` and adds it to the local machine's scheduler queue. Additionally it adds the task to the set of items that the nearest enclosing `@sync` waits for.

source

`Base.@sync` – Macro.

```
@sync
```

Wait until all dynamically-enclosed uses of `@async`, `@spawn`, `@spawnat` and `@parallel` are complete. All exceptions thrown by enclosed async operations are collected and thrown as a `CompositeException`.

source

`Base.Distributed.@parallel` – Macro.

```
@parallel
```

A parallel for loop of the form :

```
@parallel [reducer] for var = range
    body
end
```

The specified range is partitioned and locally executed across all workers. In case an optional reducer function is specified, `@parallel` performs local reductions on each worker with a final reduction on the calling process.

Note that without a reducer function, `@parallel` executes asynchronously, i.e. it spawns independent tasks on all available workers and returns immediately without waiting for completion. To wait for completion, prefix the call with `@sync`, like :

```
@sync @parallel for var = range
    body
end
```

source

`Base.Distributed.@everywhere` – Macro.

```
@everywhere expr
```

Execute an expression under `Main` everywhere. Equivalent to calling `eval(Main, expr)` on all processes. Errors on any of the processes are collected into a `CompositeException` and thrown. For example :

```
@everywhere bar=1
```

will define `Main.bar` on all processes.

Unlike `@spawn` and `@spawnat`, `@everywhere` does not capture any local variables. Prefixing `@everywhere` with `@eval` allows us to broadcast local variables using interpolation :

```
foo = 1
@eval @everywhere bar=$foo
```

The expression is evaluated under `Main` irrespective of where `@everywhere` is called from. For example :

```
module FooBar
    foo() = @everywhere bar()=myid()
end
FooBar.foo()
```

will result in `Main.bar` being defined on all processes and not `FooBar.bar`.

source

`Base.Distributed.clear!` – Method.

```
clear!(syms, pids=workers(); mod=Main)
```

Clears global bindings in modules by initializing them to `nothing`. `syms` should be of type `Symbol` or a collection of `Symbols`. `pids` and `mod` identify the processes and the module in which global variables are to be reinitialized. Only those names found to be defined under `mod` are cleared.

An exception is raised if a global constant is requested to be cleared.

source

`Base.Distributed.remoteref_id` – Function.

```
Base.remoteref_id(r::AbstractRemoteRef) -> RRID
```

`Futures` and `RemoteChannels` are identified by fields:

- `where` - refers to the node where the underlying object/storage referred to by the reference actually exists.

- `whence` - refers to the node the remote reference was created from. Note that this is different from the node where the underlying object referred to actually exists. For example calling `RemoteChannel(2)` from the master process would result in a `where` value of 2 and a `whence` value of 1.

- `id` is unique across all references created from the worker specified by `whence`.

Taken together, `whence` and `id` uniquely identify a reference across all workers.

`Base.remoteref_id` is a low-level API which returns a `Base.RRID` object that wraps `whence` and `id` values of a remote reference.

source

`Base.Distributed.channel_from_id` – Function.

```
Base.channel_from_id(id) -> c
```

A low-level API which returns the backing `AbstractChannel` for an `id` returned by `remoteref_id`. The call is valid only on the node where the backing channel exists.

source

`Base.Distributed.worker_id_from_socket` – Function.

```
Base.worker_id_from_socket(s) -> pid
```

A low-level API which given a `IO` connection or a `Worker`, returns the `pid` of the worker it is connected to. This is useful when writing custom `serialize` methods for a type, which optimizes the data written out depending on the receiving process id.

source

`Base.Distributed.cluster_cookie` – Method.

```
Base.cluster_cookie() -> cookie
```

Returns the cluster cookie.

source

`Base.Distributed.cluster_cookie` – Method.

```
Base.cluster_cookie(cookie) -> cookie
```

Sets the passed cookie as the cluster cookie, then returns it.

source

51.3 Shared Arrays

`Base.SharedArray` – Type.

```
SharedArray{T}(dims::NTuple; init=false, pids=Int[])
SharedArray{T,N}(...)
```

Construct a `SharedArray` of a bits type `T` and size `dims` across the processes specified by `pids` - all of which have to be on the same host. If `N` is specified by calling `SharedArray{T,N}(dims)`, then `N` must match the length of `dims`.

If `pids` is left unspecified, the shared array will be mapped across all processes on the current host, including the master. But, `localindexes` and `indexpids` will only refer to worker processes. This facilitates work distribution code to use workers for actual computation with the master process acting as a driver.

If an `init` function of the type `initfn(S::SharedArray)` is specified, it is called on all the participating workers.

The shared array is valid as long as a reference to the `SharedArray` object exists on the node which created the mapping.

```
SharedArray{T}(filename::AbstractString, dims::NTuple, [offset=0]; mode=nothing, init=false,
    pids=Int[])
SharedArray{T,N}(...)
```

Construct a `SharedArray` backed by the file `filename`, with element type T (must be a bits type) and size `dims`, across the processes specified by `pids` - all of which have to be on the same host. This file is mmapped into the host memory, with the following consequences:

- The array data must be represented in binary format (e.g., an ASCII format like CSV cannot be supported)

- Any changes you make to the array values (e.g., `A[3] = 0`) will also change the values on disk

If `pids` is left unspecified, the shared array will be mapped across all processes on the current host, including the master. But, `localindexes` and `indexpids` will only refer to worker processes. This facilitates work distribution code to use workers for actual computation with the master process acting as a driver.

mode must be one of `"r"`, `"r+"`, `"w+"`, or `"a+"`, and defaults to `"r+"` if the file specified by `filename` already exists, or `"w+"` if not. If an `init` function of the type `initfn(S::SharedArray)` is specified, it is called on all the participating workers. You cannot specify an `init` function if the file is not writable.

`offset` allows you to skip the specified number of bytes at the beginning of the file.

source

Base.Distributed.procs – Method.

```
procs(S::SharedArray)
```

Get the vector of processes mapping the shared array.

source

Base.sdata – Function.

```
sdata(S::SharedArray)
```

Returns the actual `Array` object backing S.

source

Base.indexpids – Function.

```
indexpids(S::SharedArray)
```

Returns the current worker's index in the list of workers mapping the `SharedArray` (i.e. in the same list returned by `procs(S)`), or 0 if the `SharedArray` is not mapped locally.

source

Base.localindexes – Function.

```
localindexes(S::SharedArray)
```

Returns a range describing the "default" indexes to be handled by the current process. This range should be interpreted in the sense of linear indexing, i.e., as a sub-range of `1:length(S)`. In multi-process contexts, returns an empty range in the parent process (or any process for which `indexpids` returns 0).

It's worth emphasizing that `localindexes` exists purely as a convenience, and you can partition work on the array among workers any way you wish. For a `SharedArray`, all indexes should be equally fast for each worker process.

source

51.4   Multi-Threading

This experimental interface supports Julia's multi-threading capabilities. Types and functions described here might (and likely will) change in the future.

Base.Threads.threadid – Function.

```
Threads.threadid()
```

Get the ID number of the current thread of execution. The master thread has ID 1.

source

Base.Threads.nthreads – Function.

```
Threads.nthreads()
```

Get the number of threads available to the Julia process. This is the inclusive upper bound on `threadid()`.

source

Base.Threads.@threads – Macro.

```
Threads.@threads
```

A macro to parallelize a for-loop to run with multiple threads. This spawns `nthreads()` number of threads, splits the iteration space amongst them, and iterates in parallel. A barrier is placed at the end of the loop which waits for all the threads to finish execution, and the loop returns.

source

Base.Threads.Atomic – Type.

```
Threads.Atomic{T}
```

Holds a reference to an object of type T, ensuring that it is only accessed atomically, i.e. in a thread-safe manner.

Only certain "simple" types can be used atomically, namely the primitive integer and float-point types. These are `Int8...Int128`, `UInt8...UInt128`, and `Float16...Float64`.

New atomic objects can be created from a non-atomic values; if none is specified, the atomic object is initialized with zero.

Atomic objects can be accessed using the `[ ]` notation:

```julia
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> x[] = 1
1

julia> x[]
1
```

Atomic operations use an `atomic_` prefix, such as `atomic_add!`, `atomic_xchg!`, etc.

source

Base.Threads.atomic_cas! – Function.

```
Threads.atomic_cas!{T}(x::Atomic{T}, cmp::T, newval::T)
```

Atomically compare-and-set x

Atomically compares the value in x with cmp. If equal, write newval to x. Otherwise, leaves x unmodified. Returns the old value in x. By comparing the returned value to cmp (via ===) one knows whether x was modified and now holds the new value newval.

For further details, see LLVM's cmpxchg instruction.

This function can be used to implement transactional semantics. Before the transaction, one records the value in x. After the transaction, the new value is stored only if x has not been modified in the mean time.

```
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_cas!(x, 4, 2);

julia> x
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_cas!(x, 3, 2);

julia> x
Base.Threads.Atomic{Int64}(2)
```

source

Base.Threads.atomic_xchg! – Function.

```
Threads.atomic_xchg!{T}(x::Atomic{T}, newval::T)
```

Atomically exchange the value in x

Atomically exchanges the value in x with newval. Returns the **old** value.

For further details, see LLVM's `atomicrmw xchg` instruction.

```
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_xchg!(x, 2)
3

julia> x[]
2
```

source

Base.Threads.atomic_add! – Function.

```
Threads.atomic_add!{T}(x::Atomic{T}, val::T)
```

Atomically add val to x

Performs `x[] += val` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw add` instruction.

```
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_add!(x, 2)
3

julia> x[]
5
```

source

Base.Threads.atomic_sub! – Function.

```
Threads.atomic_sub!{T}(x::Atomic{T}, val::T)
```

Atomically subtract `val` from `x`

Performs `x[] -= val` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw sub` instruction.

```julia
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_sub!(x, 2)
3

julia> x[]
1
```

source

Base.Threads.atomic_and! – Function.

```
Threads.atomic_and!{T}(x::Atomic{T}, val::T)
```

Atomically bitwise-and `x` with `val`

Performs `x[] &= val` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw and` instruction.

```julia
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_and!(x, 2)
3

julia> x[]
2
```

source

Base.Threads.atomic_nand! – Function.

```
Threads.atomic_nand!{T}(x::Atomic{T}, val::T)
```

Atomically bitwise-nand (not-and) `x` with `val`

Performs `x[] = ~(x[] & val)` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw nand` instruction.

```julia
julia> x = Threads.Atomic{Int}(3)
Base.Threads.Atomic{Int64}(3)

julia> Threads.atomic_nand!(x, 2)
3

julia> x[]
-3
```

source

Base.Threads.atomic_or! – Function.

```
Threads.atomic_or!{T}(x::Atomic{T}, val::T)
```

Atomically bitwise-or x with `val`

Performs `x[] |= val` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw or` instruction.

```julia
julia> x = Threads.Atomic{Int}(5)
Base.Threads.Atomic{Int64}(5)

julia> Threads.atomic_or!(x, 7)
5

julia> x[]
7
```

source

Base.Threads.atomic_xor! – Function.

```
Threads.atomic_xor!{T}(x::Atomic{T}, val::T)
```

Atomically bitwise-xor (exclusive-or) x with `val`

Performs `x[] $= val` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw xor` instruction.

```julia
julia> x = Threads.Atomic{Int}(5)
Base.Threads.Atomic{Int64}(5)

julia> Threads.atomic_xor!(x, 7)
5

julia> x[]
2
```

source

Base.Threads.atomic_max! – Function.

```
Threads.atomic_max!{T}(x::Atomic{T}, val::T)
```

Atomically store the maximum of x and `val` in x

Performs `x[] = max(x[], val)` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw max` instruction.

```julia
julia> x = Threads.Atomic{Int}(5)
Base.Threads.Atomic{Int64}(5)

julia> Threads.atomic_max!(x, 7)
5

julia> x[]
7
```

source

Base.Threads.atomic_min! – Function.

```
Threads.atomic_min!{T}(x::Atomic{T}, val::T)
```

Atomically store the minimum of x and val in x

Performs `x[] = min(x[], val)` atomically. Returns the **old** value.

For further details, see LLVM's `atomicrmw min` instruction.

```julia
julia> x = Threads.Atomic{Int}(7)
Base.Threads.Atomic{Int64}(7)

julia> Threads.atomic_min!(x, 5)
7

julia> x[]
5
```

source

Base.Threads.atomic_fence – Function.

```
Threads.atomic_fence()
```

Insert a sequential-consistency memory fence

Inserts a memory fence with sequentially-consistent ordering semantics. There are algorithms where this is needed, i.e. where an acquire/release ordering is insufficient.

This is likely a very expensive operation. Given that all other atomic operations in Julia already have acquire/release semantics, explicit fences should not be necessary in most cases.

For further details, see LLVM's `fence` instruction.

source

51.5 ccall using a threadpool (Experimental)

Base.@threadcall – Macro.

```
@threadcall((cfunc, clib), rettype, (argtypes...), argvals...)
```

The @threadcall macro is called in the same way as `ccall` but does the work in a different thread. This is useful when you want to call a blocking C function without causing the main `julia` thread to become blocked. Concurrency is limited by size of the libuv thread pool, which defaults to 4 threads but can be increased by setting the UV_THREADPOOL_SIZE environment variable and restarting the `julia` process.

Note that the called function should never call back into Julia.

source

51.6   Synchronization Primitives

`Base.Threads.AbstractLock` – Type.

> `AbstractLock`

Abstract supertype describing types that implement the thread-safe synchronization primitives: `lock`, `trylock`, `unlock`, and `islocked`

> source

`Base.lock` – Function.

> `lock(the_lock)`

Acquires the lock when it becomes available. If the lock is already locked by a different task/thread, it waits for it to become available.

Each `lock` must be matched by an `unlock`.

> source

`Base.unlock` – Function.

> `unlock(the_lock)`

Releases ownership of the lock.

If this is a recursive lock which has been acquired before, it just decrements an internal counter and returns immediately.

> source

`Base.trylock` – Function.

> `trylock(the_lock) -> Success (Boolean)`

Acquires the lock if it is available, returning `true` if successful. If the lock is already locked by a different task/thread, returns `false`.

Each successful `trylock` must be matched by an `unlock`.

> source

`Base.islocked` – Function.

> `islocked(the_lock) -> Status (Boolean)`

Check whether the lock is held by any task/thread. This should not be used for synchronization (see instead `trylock`).

> source

`Base.ReentrantLock` – Type.

> `ReentrantLock()`

Creates a reentrant lock for synchronizing Tasks. The same task can acquire the lock as many times as required. Each `lock` must be matched with an `unlock`.

This lock is NOT threadsafe. See `Threads.Mutex` for a threadsafe lock.

> source

Base.Threads.Mutex – Type.

> Mutex()

These are standard system mutexes for locking critical sections of logic.

On Windows, this is a critical section object, on pthreads, this is a `pthread_mutex_t`.

See also SpinLock for a lighter-weight lock.

source

Base.Threads.SpinLock – Type.

> SpinLock()

Creates a non-reentrant lock. Recursive use will result in a deadlock. Each `lock` must be matched with an `unlock`.

Test-and-test-and-set spin locks are quickest up to about 30ish contending threads. If you have more contention than that, perhaps a lock is the wrong way to synchronize.

See also RecursiveSpinLock for a version that permits recursion.

See also Mutex for a more efficient version on one core or if the lock may be held for a considerable length of time.

source

Base.Threads.RecursiveSpinLock – Type.

> RecursiveSpinLock()

Creates a reentrant lock. The same thread can acquire the lock as many times as required. Each `lock` must be matched with an `unlock`.

See also SpinLock for a slightly faster version.

See also Mutex for a more efficient version on one core or if the lock may be held for a considerable length of time.

source

Base.Semaphore – Type.

> Semaphore(sem_size)

Creates a counting semaphore that allows at most `sem_size` acquires to be in use at any time. Each acquire must be mached with a release.

This construct is NOT threadsafe.

source

Base.acquire – Function.

> acquire(s::Semaphore)

Wait for one of the `sem_size` permits to be available, blocking until one can be acquired.

source

Base.release – Function.

> release(s::Semaphore)

Return one permit to the pool, possibly allowing another task to acquire it and resume execution.

source

51.7   Cluster Manager Interface

This interface provides a mechanism to launch and manage Julia workers on different cluster environments. There are two types of managers present in Base: `LocalManager`, for launching additional workers on the same host, and `SSHManager`, for launching on remote hosts via `ssh`. TCP/IP sockets are used to connect and transport messages between processes. It is possible for Cluster Managers to provide a different transport.

`Base.Distributed.launch` – Function.

```
launch(manager::ClusterManager, params::Dict, launched::Array, launch_ntfy::Condition)
```

Implemented by cluster managers. For every Julia worker launched by this function, it should append a `Worker-Config` entry to `launched` and notify `launch_ntfy`. The function MUST exit once all workers, requested by `manager` have been launched. `params` is a dictionary of all keyword arguments `addprocs` was called with.

source

`Base.Distributed.manage` – Function.

```
manage(manager::ClusterManager, id::Integer, config::WorkerConfig. op::Symbol)
```

Implemented by cluster managers. It is called on the master process, during a worker's lifetime, with appropriate op values:

- with `:register`/`:deregister` when a worker is added / removed from the Julia worker pool.
- with `:interrupt` when `interrupt(workers)` is called. The `ClusterManager` should signal the appropriate worker with an interrupt signal.
- with `:finalize` for cleanup purposes.

source

`Base.kill` – Method.

```
kill(manager::ClusterManager, pid::Int, config::WorkerConfig)
```

Implemented by cluster managers. It is called on the master process, by `rmprocs`. It should cause the remote worker specified by `pid` to exit. `Base.kill(manager::ClusterManager.....)` executes a remote `exit()` on `pid`.

source

`Base.Distributed.init_worker` – Function.

```
init_worker(cookie::AbstractString, manager::ClusterManager=DefaultClusterManager())
```

Called by cluster managers implementing custom transports. It initializes a newly launched process as a worker. Command line argument `--worker` has the effect of initializing a process as a worker using TCP/IP sockets for transport. `cookie` is a `cluster_cookie`.

source

`Base.connect` – Method.

```
connect(manager::ClusterManager, pid::Int, config::WorkerConfig) -> (instrm::IO, outstrm::IO)
```

Implemented by cluster managers using custom transports. It should establish a logical connection to worker with id `pid`, specified by `config` and return a pair of `IO` objects. Messages from `pid` to current process will be read off `instrm`, while messages to be sent to `pid` will be written to `outstrm`. The custom transport implementation must ensure that messages are delivered and received completely and in order. `Base.connect(manager::Cluster-Manager.....)` sets up TCP/IP socket connections in-between workers.

source

Base.Distributed.process_messages – Function.

```
Base.process_messages(r_stream::IO, w_stream::IO, incoming::Bool=true)
```

Called by cluster managers using custom transports. It should be called when the custom transport implementation receives the first message from a remote worker. The custom transport must manage a logical connection to the remote worker and provide two `IO` objects, one for incoming messages and the other for messages addressed to the remote worker. If `incoming` is `true`, the remote peer initiated the connection. Whichever of the pair initiates the connection sends the cluster cookie and its Julia version number to perform the authentication handshake.

See also `cluster_cookie`.

source

Chapter 52

Linear Algebra

52.1 Standard Functions

Linear algebra functions in Julia are largely implemented by calling functions from LAPACK. Sparse factorizations call functions from SuiteSparse.

`Base.:*` – Method.

```
*(x, y...)
```

Multiplication operator. `x*y*z*...` calls this function with all arguments, i.e. `*(x, y, z, ...)`.

source

`Base.:\` – Method.

```
\(x, y)
```

Left division operator: multiplication of y by the inverse of x on the left. Gives floating-point results for integer arguments.

```julia
julia> 3 \ 6
2.0

julia> inv(3) * 6
2.0

julia> A = [1 2; 3 4]; x = [5, 6];

julia> A \ x
2-element Array{Float64,1}:
 -4.0
  4.5

julia> inv(A) * x
2-element Array{Float64,1}:
 -4.0
  4.5
```

source

`Base.LinAlg.dot` – Function.

```
dot(n, X, incx, Y, incy)
```

Dot product of two vectors consisting of n elements of array X with stride `incx` and n elements of array Y with stride `incy`.

Example:

```
julia> dot(10, ones(10), 1, ones(20), 2)
10.0
```

source

Base.LinAlg.`vecdot` – Function.

```
vecdot(x, y)
```

For any iterable containers x and y (including arrays of any dimension) of numbers (or any element type for which `dot` is defined), compute the Euclidean dot product (the sum of `dot(x[i],y[i])`) as if they were vectors.

Examples

```
julia> vecdot(1:5, 2:6)
70

julia> x = fill(2., (5,5));

julia> y = fill(3., (5,5));

julia> vecdot(x, y)
150.0
```

source

Base.LinAlg.`cross` – Function.

```
cross(x, y)
×(x,y)
```

Compute the cross product of two 3-vectors.

Example

```
julia> a = [0;1;0]
3-element Array{Int64,1}:
 0
 1
 0

julia> b = [0;0;1]
3-element Array{Int64,1}:
 0
 0
 1

julia> cross(a,b)
3-element Array{Int64,1}:
 1
 0
 0
```

`Base.LinAlg.factorize` – Function.

```
factorize(A)
```

Compute a convenient factorization of A, based upon the type of the input matrix. `factorize` checks A to see if it is symmetric/triangular/etc. if A is passed as a generic matrix. `factorize` checks every element of A to verify/rule out each property. It will short-circuit as soon as it can rule out symmetry/triangular structure. The return value can be reused for efficient solving of multiple systems. For example: A=factorize(A); x=A\b; y=A\C.

Properties of A	type of factorization
Positive-definite	Cholesky (see `cholfact`)
Dense Symmetric/Hermitian	Bunch-Kaufman (see `bkfact`)
Sparse Symmetric/Hermitian	LDLt (see `ldltfact`)
Triangular	Triangular
Diagonal	Diagonal
Bidiagonal	Bidiagonal
Tridiagonal	LU (see `lufact`)
Symmetric real tridiagonal	LDLt (see `ldltfact`)
General square	LU (see `lufact`)
General non-square	QR (see `qrfact`)

If `factorize` is called on a Hermitian positive-definite matrix, for instance, then `factorize` will return a Cholesky factorization.

Example

```julia
julia> A = Array(Bidiagonal(ones(5, 5), true))
5×5 Array{Float64,2}:
 1.0  1.0  0.0  0.0  0.0
 0.0  1.0  1.0  0.0  0.0
 0.0  0.0  1.0  1.0  0.0
 0.0  0.0  0.0  1.0  1.0
 0.0  0.0  0.0  0.0  1.0

julia> factorize(A) # factorize will check to see that A is already factorized
5×5 Bidiagonal{Float64}:
 1.0  1.0
     1.0  1.0
         1.0  1.0
             1.0  1.0
                 1.0
```

This returns a 5×5 `Bidiagonal{Float64}`, which can now be passed to other linear algebra functions (e.g. eigensolvers) which will use specialized methods for `Bidiagonal` types.

`Base.LinAlg.Diagonal` – Type.

```
Diagonal(A::AbstractMatrix)
```

Constructs a matrix from the diagonal of A.

Example

```julia
julia> A = [1 2 3; 4 5 6; 7 8 9]
3×3 Array{Int64,2}:
 1  2  3
 4  5  6
 7  8  9

julia> Diagonal(A)
3×3 Diagonal{Int64}:
 1
    5
       9
```

source

```
Diagonal(V::AbstractVector)
```

Constructs a matrix with V as its diagonal.

Example

```julia
julia> V = [1; 2]
2-element Array{Int64,1}:
 1
 2

julia> Diagonal(V)
2×2 Diagonal{Int64}:
 1
    2
```

source

`Base.LinAlg.Bidiagonal` – Type.

```
Bidiagonal(dv, ev, isupper::Bool)
```

Constructs an upper (isupper=true) or lower (isupper=false) bidiagonal matrix using the given diagonal (dv) and off-diagonal (ev) vectors. The result is of type Bidiagonal and provides efficient specialized linear solvers, but may be converted into a regular matrix with convert(Array, _) (or Array(_) for short). ev's length must be one less than the length of dv.

Example

```julia
julia> dv = [1; 2; 3; 4]
4-element Array{Int64,1}:
 1
 2
 3
 4

julia> ev = [7; 8; 9]
3-element Array{Int64,1}:
 7
 8
 9

julia> Bu = Bidiagonal(dv, ev, true) # ev is on the first superdiagonal
```

```
4×4 Bidiagonal{Int64}:
 1  7
    2  8
       3  9
          4

julia> Bl = Bidiagonal(dv, ev, false) # ev is on the first subdiagonal
4×4 Bidiagonal{Int64}:
 1
 7  2
    8  3
       9  4
```

```
Bidiagonal(dv, ev, uplo::Char)
```

Constructs an upper (uplo='U') or lower (uplo='L') bidiagonal matrix using the given diagonal (dv) and off-diagonal (ev) vectors. The result is of type Bidiagonal and provides efficient specialized linear solvers, but may be converted into a regular matrix with convert(Array, _) (or Array(_) for short). ev's length must be one less than the length of dv.

Example

```
julia> dv = [1; 2; 3; 4]
4-element Array{Int64,1}:
 1
 2
 3
 4

julia> ev = [7; 8; 9]
3-element Array{Int64,1}:
 7
 8
 9

julia> Bu = Bidiagonal(dv, ev, 'U') #e is on the first superdiagonal
4×4 Bidiagonal{Int64}:
 1  7
    2  8
       3  9
          4

julia> Bl = Bidiagonal(dv, ev, 'L') #e is on the first subdiagonal
4×4 Bidiagonal{Int64}:
 1
 7  2
    8  3
       9  4
```

```
Bidiagonal(A, isupper::Bool)
```

Construct a Bidiagonal matrix from the main diagonal of A and its first super- (if isupper=true) or sub-diagonal (if isupper=false).

Example

```julia
julia> A = [1 1 1 1; 2 2 2 2; 3 3 3 3; 4 4 4 4]
4×4 Array{Int64,2}:
 1  1  1  1
 2  2  2  2
 3  3  3  3
 4  4  4  4

julia> Bidiagonal(A, true) #contains the main diagonal and first superdiagonal of A
4×4 Bidiagonal{Int64}:
 1  1
   2  2
     3  3
       4

julia> Bidiagonal(A, false) #contains the main diagonal and first subdiagonal of A
4×4 Bidiagonal{Int64}:
 1
 2  2
   3  3
     4  4
```

source

`Base.LinAlg.SymTridiagonal` – Type.

```julia
SymTridiagonal(dv, ev)
```

Construct a symmetric tridiagonal matrix from the diagonal and first sub/super-diagonal, respectively. The result is of type `SymTridiagonal` and provides efficient specialized eigensolvers, but may be converted into a regular matrix with `convert(Array, _)` (or `Array(_)` for short).

Example

```julia
julia> dv = [1; 2; 3; 4]
4-element Array{Int64,1}:
 1
 2
 3
 4

julia> ev = [7; 8; 9]
3-element Array{Int64,1}:
 7
 8
 9

julia> SymTridiagonal(dv, ev)
4×4 SymTridiagonal{Int64}:
 1  7
 7  2  8
    8  3  9
       9  4
```

source

Base.LinAlg.Tridiagonal – Type.

```
Tridiagonal(dl, d, du)
```

Construct a tridiagonal matrix from the first subdiagonal, diagonal, and first superdiagonal, respectively. The result is of type `Tridiagonal` and provides efficient specialized linear solvers, but may be converted into a regular matrix with `convert(Array, _)` (or `Array(_)` for short). The lengths of `dl` and `du` must be one less than the length of d.

Example

```
julia> dl = [1; 2; 3]
3-element Array{Int64,1}:
 1
 2
 3

julia> du = [4; 5; 6]
3-element Array{Int64,1}:
 4
 5
 6

julia> d = [7; 8; 9; 0]
4-element Array{Int64,1}:
 7
 8
 9
 0

julia> Tridiagonal(dl, d, du)
4×4 Tridiagonal{Int64}:
 7  4
 1  8  5
    2  9  6
       3  0
```

source

```
Tridiagonal(A)
```

returns a `Tridiagonal` array based on (abstract) matrix A, using its first lower diagonal, main diagonal, and first upper diagonal.

Example

```
julia> A = [1 2 3 4; 1 2 3 4; 1 2 3 4; 1 2 3 4]
4×4 Array{Int64,2}:
 1  2  3  4
 1  2  3  4
 1  2  3  4
 1  2  3  4

julia> Tridiagonal(A)
4×4 Tridiagonal{Int64}:
 1  2
 1  2  3
```

```
    2  3  4
       3  4
```

source

Base.LinAlg.Symmetric – Type.

```
Symmetric(A, uplo=:U)
```

Construct a Symmetric matrix from the upper (if uplo = :U) or lower (if uplo = :L) triangle of A.

Example

```julia
julia> A = [1 0 2 0 3; 0 4 0 5 0; 6 0 7 0 8; 0 9 0 1 0; 2 0 3 0 4]
5×5 Array{Int64,2}:
 1  0  2  0  3
 0  4  0  5  0
 6  0  7  0  8
 0  9  0  1  0
 2  0  3  0  4

julia> Supper = Symmetric(A)
5×5 Symmetric{Int64,Array{Int64,2}}:
 1  0  2  0  3
 0  4  0  5  0
 2  0  7  0  8
 0  5  0  1  0
 3  0  8  0  4

julia> Slower = Symmetric(A, :L)
5×5 Symmetric{Int64,Array{Int64,2}}:
 1  0  6  0  2
 0  4  0  9  0
 6  0  7  0  3
 0  9  0  1  0
 2  0  3  0  4
```

Note that Supper will not be equal to Slower unless A is itself symmetric (e.g. if A == A.').

source

Base.LinAlg.Hermitian – Type.

```
Hermitian(A, uplo=:U)
```

Construct a Hermitian matrix from the upper (if uplo = :U) or lower (if uplo = :L) triangle of A.

Example

```julia
julia> A = [1 0 2+2im 0 3-3im; 0 4 0 5 0; 6-6im 0 7 0 8+8im; 0 9 0 1 0; 2+2im 0 3-3im 0 4];

julia> Hupper = Hermitian(A)
5×5 Hermitian{Complex{Int64},Array{Complex{Int64},2}}:
 1+0im  0+0im  2+2im  0+0im  3-3im
 0+0im  4+0im  0+0im  5+0im  0+0im
 2-2im  0+0im  7+0im  0+0im  8+8im
 0+0im  5+0im  0+0im  1+0im  0+0im
```

```
 3+3im  0+0im  8-8im  0+0im  4+0im

julia> Hlower = Hermitian(A, :L)
5×5 Hermitian{Complex{Int64},Array{Complex{Int64},2}}:
 1+0im  0+0im  6+6im  0+0im  2-2im
 0+0im  4+0im  0+0im  9+0im  0+0im
 6-6im  0+0im  7+0im  0+0im  3+3im
 0+0im  9+0im  0+0im  1+0im  0+0im
 2+2im  0+0im  3-3im  0+0im  4+0im
```

Note that Hupper will not be equal to Hlower unless A is itself Hermitian (e.g. if A == A').

source

Base.LinAlg.lu – Function.

```
lu(A, pivot=Val{true}) -> L, U, p
```

Compute the LU factorization of A, such that A[p, :] = L*U. By default, pivoting is used. This can be overridden by passing Val{false} for the second argument.

See also lufact.

Example

```
julia> A = [4. 3.; 6. 3.]
2×2 Array{Float64,2}:
 4.0  3.0
 6.0  3.0

julia> L, U, p = lu(A)
([1.0 0.0; 0.666667 1.0], [6.0 3.0; 0.0 1.0], [2, 1])

julia> A[p, :] == L * U
true
```

source

Base.LinAlg.lufact – Function.

```
lufact(A::SparseMatrixCSC) -> F::UmfpackLU
```

Compute the LU factorization of a sparse matrix A.

For sparse A with real or complex element type, the return type of F is UmfpackLU{Tv, Ti}, with Tv = Float64 or Complex128 respectively and Ti is an integer type (Int32 or Int64).

The individual components of the factorization F can be accessed by indexing:

Component	Description
F[:L]	L (lower triangular) part of LU
F[:U]	U (upper triangular) part of LU
F[:p]	right permutation Vector
F[:q]	left permutation Vector
F[:Rs]	Vector of scaling factors
F[:(:)]	(L,U,p,q,Rs) components

The relation between F and A is

```
F[:L]*F[:U] == (F[:Rs] .* A)[F[:p], F[:q]]
```

F further supports the following functions:

- \
- cond
- det

> **Note**
>
> lufact(A::SparseMatrixCSC) uses the UMFPACK library that is part of SuiteSparse. As this library only supports sparse matrices with Float64 or Complex128 elements, lufact converts A into a copy that is of type SparseMatrixCSC{Float64} or SparseMatrixCSC{Complex128} as appropriate.

source

```
lufact(A [,pivot=Val{true}]) -> F::LU
```

Compute the LU factorization of A.

In most cases, if A is a subtype S of AbstractMatrix{T} with an element type T supporting +, -, * and /, the return type is LU{T,S{T}}. If pivoting is chosen (default) the element type should also support abs and <.

The individual components of the factorization F can be accessed by indexing:

Component	Description
F[:L]	L (lower triangular) part of LU
F[:U]	U (upper triangular) part of LU
F[:p]	(right) permutation Vector
F[:P]	(right) permutation Matrix

The relationship between F and A is

```
F[:L]*F[:U] == A[F[:p], :]
```

F further supports the following functions:

Supported function	LU	LU{T,Tridiagonal{T}}
/		
\		
cond		
inv		
det		
logdet		
logabsdet		
size		

Example

```
julia> A = [4 3; 6 3]
2×2 Array{Int64,2}:
 4  3
```

```
 6  3

julia> F = lufact(A)
Base.LinAlg.LU{Float64,Array{Float64,2}} with factors L and U:
[1.0 0.0; 1.5 1.0]
[4.0 3.0; 0.0 -1.5]

julia> F[:L] * F[:U] == A[F[:p], :]
true
```

source

Base.LinAlg.lufact! – Function.

```
lufact!(A, pivot=Val{true}) -> LU
```

lufact! is the same as lufact, but saves space by overwriting the input A, instead of creating a copy. An InexactError exception is thrown if the factorization produces a number not representable by the element type of A, e.g. for integer types.

source

Base.LinAlg.chol – Function.

```
chol(A) -> U
```

Compute the Cholesky factorization of a positive definite matrix A and return the UpperTriangular matrix U such that A = U'U.

Example

```
julia> A = [1. 2.; 2. 50.]
2×2 Array{Float64,2}:
 1.0   2.0
 2.0  50.0

julia> U = chol(A)
2×2 UpperTriangular{Float64,Array{Float64,2}}:
 1.0  2.0
      6.78233

julia> U'U
2×2 Array{Float64,2}:
 1.0   2.0
 2.0  50.0
```

source

```
chol(x::Number) -> y
```

Compute the square root of a non-negative number x.

Example

```
julia> chol(16)
4.0
```

source

`Base.LinAlg.cholfact` – Function.

```
cholfact(A; shift = 0.0, perm = Int[]) -> CHOLMOD.Factor
```

Compute the Cholesky factorization of a sparse positive definite matrix A. A must be a `SparseMatrixCSC`, `Symmetric{SparseMatrixCSC}`, or `Hermitian{SparseMatrixCSC}`. Note that even if A doesn't have the type tag, it must still be symmetric or Hermitian. A fill-reducing permutation is used. `F = cholfact(A)` is most frequently used to solve systems of equations with `F\b`, but also the methods `diag`, `det`, and `logdet` are defined for F. You can also extract individual factors from F, using `F[:L]`. However, since pivoting is on by default, the factorization is internally represented as `A == P'*L*L'*P` with a permutation matrix P; using just L without accounting for P will give incorrect answers. To include the effects of permutation, it's typically preferable to extract "combined" factors like `PtL = F[:PtL]` (the equivalent of `P'*L`) and `LtP = F[:UP]` (the equivalent of `L'*P`).

Setting the optional `shift` keyword argument computes the factorization of `A+shift*I` instead of A. If the `perm` argument is nonempty, it should be a permutation of `1:size(A,1)` giving the ordering to use (instead of CHOLMOD's default AMD ordering).

> **Note**
>
> This method uses the CHOLMOD library from SuiteSparse, which only supports doubles or complex doubles. Input matrices not of those element types will be converted to `SparseMatrixCSC{Float64}` or `SparseMatrixCSC{Complex128}` as appropriate.
>
> Many other functions from CHOLMOD are wrapped but not exported from the `Base.SparseArrays.CHOLMOD` module.

source

```
cholfact(A, [uplo::Symbol,] Val{false}) -> Cholesky
```

Compute the Cholesky factorization of a dense symmetric positive definite matrix A and return a `Cholesky` factorization. The matrix A can either be a `Symmetric` or `Hermitian` `StridedMatrix` or a *perfectly* symmetric or Hermitian `StridedMatrix`. In the latter case, the optional argument `uplo` may be `:L` for using the lower part or `:U` for the upper part of A. The default is to use `:U`. The triangular Cholesky factor can be obtained from the factorization F with: `F[:L]` and `F[:U]`. The following functions are available for `Cholesky` objects: `size`, `\`, `inv`, and `det`. A `PosDefException` exception is thrown in case the matrix is not positive definite.

Example

```
julia> A = [4. 12. -16.; 12. 37. -43.; -16. -43. 98.]
3×3 Array{Float64,2}:
   4.0   12.0  -16.0
  12.0   37.0  -43.0
 -16.0  -43.0   98.0

julia> C = cholfact(A)
Base.LinAlg.Cholesky{Float64,Array{Float64,2}} with factor:
[2.0 6.0 -8.0; 0.0 1.0 5.0; 0.0 0.0 3.0]

julia> C[:U]
3×3 UpperTriangular{Float64,Array{Float64,2}}:
 2.0  6.0  -8.0
      1.0   5.0
            3.0

julia> C[:L]
3×3 LowerTriangular{Float64,Array{Float64,2}}:
```

```
  2.0
  6.0  1.0
 -8.0  5.0  3.0

julia> C[:L] * C[:U] == A
true
```

source

```
cholfact(A, [uplo::Symbol,] Val{true}; tol = 0.0) -> CholeskyPivoted
```

Compute the pivoted Cholesky factorization of a dense symmetric positive semi-definite matrix A and return a CholeskyPivoted factorization. The matrix A can either be a Symmetric or Hermitian StridedMatrix or a *perfectly* symmetric or Hermitian StridedMatrix. In the latter case, the optional argument uplo may be :L for using the lower part or :U for the upper part of A. The default is to use :U. The triangular Cholesky factor can be obtained from the factorization F with: F[:L] and F[:U]. The following functions are available for PivotedC-holesky objects: size, \, inv, det, and rank. The argument tol determines the tolerance for determining the rank. For negative values, the tolerance is the machine precision.

source

Base.LinAlg.cholfact! – Function.

```
cholfact!(F::Factor, A; shift = 0.0) -> CHOLMOD.Factor
```

Compute the Cholesky (LL') factorization of A, reusing the symbolic factorization F. A must be a SparseMa-trixCSC, Symmetric{SparseMatrixCSC}, or Hermitian{SparseMatrixCSC}. Note that even if A doesn't have the type tag, it must still be symmetric or Hermitian.

See also cholfact.

> **Note**
>
> This method uses the CHOLMOD library from SuiteSparse, which only supports doubles or complex doubles. Input matrices not of those element types will be converted to SparseMatrixCSC{Float64} or SparseMatrixCSC{Complex128} as appropriate.

source

```
cholfact!(A, [uplo::Symbol,] Val{false}) -> Cholesky
```

The same as cholfact, but saves space by overwriting the input A, instead of creating a copy. An InexactError exception is thrown if the factorization produces a number not representable by the element type of A, e.g. for integer types.

Example

```
julia> A = [1 2; 2 50]
2×2 Array{Int64,2}:
 1   2
 2  50

julia> cholfact!(A)
ERROR: InexactError()
```

source

```
cholfact!(A, [uplo::Symbol,] Val{true}; tol = 0.0) -> CholeskyPivoted
```

The same as cholfact, but saves space by overwriting the input A, instead of creating a copy. An InexactError exception is thrown if the factorization produces a number not representable by the element type of A, e.g. for integer types.

source

Base.LinAlg.lowrankupdate – Function.

```
lowrankupdate(C::Cholesky, v::StridedVector) -> CC::Cholesky
```

Update a Cholesky factorization C with the vector v. If A = C[:U]'C[:U] then CC = cholfact(C[:U]'C[:U] + v*v') but the computation of CC only uses O(n^2) operations.

source

Base.LinAlg.lowrankdowndate – Function.

```
lowrankdowndate(C::Cholesky, v::StridedVector) -> CC::Cholesky
```

Downdate a Cholesky factorization C with the vector v. If A = C[:U]'C[:U] then CC = cholfact(C[:U]'C[:U] - v*v') but the computation of CC only uses O(n^2) operations.

source

Base.LinAlg.lowrankupdate! – Function.

```
lowrankupdate!(C::Cholesky, v::StridedVector) -> CC::Cholesky
```

Update a Cholesky factorization C with the vector v. If A = C[:U]'C[:U] then CC = cholfact(C[:U]'C[:U] + v*v') but the computation of CC only uses O(n^2) operations. The input factorization C is updated in place such that on exit C == CC. The vector v is destroyed during the computation.

source

Base.LinAlg.lowrankdowndate! – Function.

```
lowrankdowndate!(C::Cholesky, v::StridedVector) -> CC::Cholesky
```

Downdate a Cholesky factorization C with the vector v. If A = C[:U]'C[:U] then CC = cholfact(C[:U]'C[:U] - v*v') but the computation of CC only uses O(n^2) operations. The input factorization C is updated in place such that on exit C == CC. The vector v is destroyed during the computation.

source

Base.LinAlg.ldltfact – Function.

```
ldltfact(A; shift = 0.0, perm=Int[]) -> CHOLMOD.Factor
```

Compute the LDL' factorization of a sparse matrix A. A must be a SparseMatrixCSC, Symmetric{Sparse-MatrixCSC}, or Hermitian{SparseMatrixCSC}. Note that even if A doesn't have the type tag, it must still be symmetric or Hermitian. A fill-reducing permutation is used. F = ldltfact(A) is most frequently used to solve systems of equations A*x = b with F\b. The returned factorization object F also supports the methods diag, det, logdet, and inv. You can extract individual factors from F using F[:L]. However, since pivoting is on by default, the factorization is internally represented as A == P'*L*D*L'*P with a permutation matrix P; using just L without accounting for P will give incorrect answers. To include the effects of permutation, it is typically preferable to extract "combined" factors like PtL = F[:PtL] (the equivalent of P'*L) and LtP = F[:UP] (the equivalent of L'*P). The complete list of supported factors is :L, :PtL, :D, :UP, :U, :LD, :DU, :PtLD, :DUP.

Setting the optional `shift` keyword argument computes the factorization of `A+shift*I` instead of `A`. If the `perm` argument is nonempty, it should be a permutation of `1:size(A,1)` giving the ordering to use (instead of CHOLMOD's default AMD ordering).

> **Note**
>
> This method uses the CHOLMOD library from SuiteSparse, which only supports doubles or complex doubles. Input matrices not of those element types will be converted to `SparseMatrixCSC{Float64}` or `SparseMatrixCSC{Complex128}` as appropriate.
>
> Many other functions from CHOLMOD are wrapped but not exported from the `Base.SparseArrays.CHOLMOD` module.

source

```
ldltfact(S::SymTridiagonal) -> LDLt
```

Compute an LDLt factorization of a real symmetric tridiagonal matrix such that `A = L*Diagonal(d)*L'` where L is a unit lower triangular matrix and d is a vector. The main use of an LDLt factorization `F = ldltfact(A)` is to solve the linear system of equations `Ax = b` with `F\b`.

source

`Base.LinAlg.ldltfact!` – Function.

```
ldltfact!(F::Factor, A; shift = 0.0) -> CHOLMOD.Factor
```

Compute the LDL' factorization of A, reusing the symbolic factorization F. A must be a `SparseMatrixCSC`, `Symmetric{SparseMatrixCSC}`, or `Hermitian{SparseMatrixCSC}`. Note that even if A doesn't have the type tag, it must still be symmetric or Hermitian.

See also `ldltfact`.

> **Note**
>
> This method uses the CHOLMOD library from SuiteSparse, which only supports doubles or complex doubles. Input matrices not of those element types will be converted to `SparseMatrixCSC{Float64}` or `SparseMatrixCSC{Complex128}` as appropriate.

source

```
ldltfact!(S::SymTridiagonal) -> LDLt
```

Same as `ldltfact`, but saves space by overwriting the input A, instead of creating a copy.

source

`Base.LinAlg.qr` – Function.

```
qr(A, pivot=Val{false}; thin::Bool=true) -> Q, R, [p]
```

Compute the (pivoted) QR factorization of A such that either `A = Q*R` or `A[:,p] = Q*R`. Also see `qrfact`. The default is to compute a thin factorization. Note that R is not extended with zeros when the full Q is requested.

source

```
qr(v::AbstractVector) -> w, r
```

Computes the polar decomposition of a vector. Returns w, a unit vector in the direction of v, and r, the norm of
v.

See also `normalize`, `normalize!`, and `LinAlg.qr!`.

Example

```julia
julia> v = [1; 2]
2-element Array{Int64,1}:
 1
 2

julia> w, r = qr(v)
([0.447214, 0.894427], 2.23606797749979)

julia> w*r == v
true
```

source

`Base.LinAlg.qr!` – Function.

```
LinAlg.qr!(v::AbstractVector) -> w, r
```

Computes the polar decomposition of a vector. Instead of returning a new vector as `qr(v::AbstractVector)`,
this function mutates the input vector v in place. Returns w, a unit vector in the direction of v (this is a mutation
of v), and r, the norm of v.

See also `normalize`, `normalize!`, and `qr`.

source

`Base.LinAlg.qrfact` – Function.

```
qrfact(A) -> SPQR.Factorization
```

Compute the QR factorization of a sparse matrix A. A fill-reducing permutation is used. The main application of
this type is to solve least squares problems with `\`. The function calls the C library SPQR and a few additional
functions from the library are wrapped but not exported.

source

```
qrfact(A, pivot=Val{false}) -> F
```

Compute the QR factorization of the matrix A: an orthogonal (or unitary if A is complex-valued) matrix Q, and an
upper triangular matrix R such that

$$A = QR$$

The returned object F stores the factorization in a packed format:

- if `pivot == Val{true}` then F is a `QRPivoted` object,
- otherwise if the element type of A is a BLAS type (`Float32`, `Float64`, Complex64 or Complex128), then
  F is a `QRCompactWY` object,
- otherwise F is a `QR` object.

The individual components of the factorization F can be accessed by indexing with a symbol:

- `F[:Q]`: the orthogonal/unitary matrix Q
- `F[:R]`: the upper triangular matrix R
- `F[:p]`: the permutation vector of the pivot (`QRPivoted` only)
- `F[:P]`: the permutation matrix of the pivot (`QRPivoted` only)

The following functions are available for the QR objects: `inv`, `size`, and `\`. When A is rectangular, `\` will return a least squares solution and if the solution is not unique, the one with smallest norm is returned.

Multiplication with respect to either thin or full Q is allowed, i.e. both `F[:Q]*F[:R]` and `F[:Q]*A` are supported. A Q matrix can be converted into a regular matrix with `full` which has a named argument `thin`.

Example

```julia
julia> A = [3.0 -6.0; 4.0 -8.0; 0.0 1.0]
3×2 Array{Float64,2}:
 3.0  -6.0
 4.0  -8.0
 0.0   1.0

julia> F = qrfact(A)
Base.LinAlg.QRCompactWY{Float64,Array{Float64,2}} with factors Q and R:
[-0.6 0.0 0.8; -0.8 0.0 -0.6; 0.0 -1.0 0.0]
[-5.0 10.0; 0.0 -1.0]

julia> F[:Q] * F[:R] == A
true
```

> **Note**
>
> `qrfact` returns multiple types because LAPACK uses several representations that minimize the memory storage requirements of products of Householder elementary reflectors, so that the Q and R matrices can be stored compactly rather as two separate dense matrices.

source

Base.LinAlg.qrfact! – Function.

```
qrfact!(A, pivot=Val{false})
```

`qrfact!` is the same as `qrfact` when A is a subtype of `StridedMatrix`, but saves space by overwriting the input A, instead of creating a copy. An `InexactError` exception is thrown if the factorization produces a number not representable by the element type of A, e.g. for integer types.

source

Base.LinAlg.QR – Type.

```
QR <: Factorization
```

A QR matrix factorization stored in a packed format, typically obtained from `qrfact`. If A is an m×n matrix, then

$$A = QR$$

where Q is an orthogonal/unitary matrix and R is upper triangular. The matrix Q is stored as a sequence of Householder reflectors v_i and coefficients τ_i where:

$$Q = \prod_{i=1}^{\min(m,n)} (I - \tau_i v_i v_i^T).$$

The object has two fields:

- `factors` is an m×n matrix.
    - The upper triangular part contains the elements of R, that is R = `triu(F.factors)` for a QR object F.
    - The subdiagonal part contains the reflectors v_i stored in a packed format where v_i is the ith column of the matrix V = `eye(m,n) + tril(F.factors,-1)`.
- `τ` is a vector of length `min(m,n)` containing the coefficients au_i.

source

Base.LinAlg.QRCompactWY – Type.

```
QRCompactWY <: Factorization
```

A QR matrix factorization stored in a compact blocked format, typically obtained from `qrfact`. If A is an m×n matrix, then

$$A = QR$$

where Q is an orthogonal/unitary matrix and R is upper triangular. It is similar to the QR format except that the orthogonal/unitary matrix Q is stored in *Compact WY* format [1], as a lower trapezoidal matrix V and an upper triangular matrix T where

$$Q = \prod_{i=1}^{\min(m,n)} (I - \tau_i v_i v_i^T) = I - VTV^T$$

such that v_i is the ith column of V, and au_i is the ith diagonal element of T.

The object has two fields:

- `factors`, as in the QR type, is an m×n matrix.
    - The upper triangular part contains the elements of R, that is R = `triu(F.factors)` for a QR object F.
    - The subdiagonal part contains the reflectors v_i stored in a packed format such that V = `eye(m,n) + tril(F.factors,-1)`.
- `T` is a square matrix with `min(m,n)` columns, whose upper triangular part gives the matrix T above (the subdiagonal elements are ignored).

Note

This format should not to be confused with the older *WY* representation [2].

source

Base.LinAlg.QRPivoted – Type.

```
QRPivoted <: Factorization
```

A QR matrix factorization with column pivoting in a packed format, typically obtained from qrfact. If A is an m×n matrix, then

$$AP = QR$$

where P is a permutation matrix, Q is an orthogonal/unitary matrix and R is upper triangular. The matrix Q is stored as a sequence of Householder reflectors:

$$Q = \prod_{i=1}^{\min(m,n)} (I - \tau_i v_i v_i^T).$$

The object has three fields:

- factors is an m×n matrix.
  - The upper triangular part contains the elements of R, that is R = triu(F.factors) for a QR object F.
  - The subdiagonal part contains the reflectors v_i stored in a packed format where v_i is the ith column of the matrix V = eye(m,n) + tril(F.factors,-1).
- τ is a vector of length min(m,n) containing the coefficients au_i.
- jpvt is an integer vector of length n corresponding to the permutation P.

source

Base.LinAlg.lqfact! – Function.

```
lqfact!(A) -> LQ
```

Compute the LQ factorization of A, using the input matrix as a workspace. See also lq.

source

Base.LinAlg.lqfact – Function.

```
lqfact(A) -> LQ
```

Compute the LQ factorization of A. See also lq.

source

Base.LinAlg.lq – Function.

---

[2] C Bischof and C Van Loan, "The WY representation for products of Householder matrices", SIAM J Sci Stat Comput 8 (1987), s2-s13. doi:10.1137/0908009

[1] R Schreiber and C Van Loan, "A storage-efficient WY representation for products of Householder transformations", SIAM J Sci Stat Comput 10 (1989), 53-57. doi:10.1137/0910005

```
lq(A; [thin=true]) -> L, Q
```

Perform an LQ factorization of A such that A = L*Q. The default is to compute a thin factorization. The LQ factorization is the QR factorization of A.'. L is not extended with zeros if the full Q is requested.

source

Base.LinAlg.bkfact – Function.

```
bkfact(A, uplo::Symbol=:U, symmetric::Bool=issymmetric(A), rook::Bool=false) -> BunchKaufman
```

Compute the Bunch-Kaufman [3] factorization of a symmetric or Hermitian matrix A and return a BunchKaufman object. uplo indicates which triangle of matrix A to reference. If symmetric is true, A is assumed to be symmetric. If symmetric is false, A is assumed to be Hermitian. If rook is true, rook pivoting is used. If rook is false, rook pivoting is not used. The following functions are available for BunchKaufman objects: size, \, inv, issymmetric, ishermitian.

source

Base.LinAlg.bkfact! – Function.

```
bkfact!(A, uplo::Symbol=:U, symmetric::Bool=issymmetric(A), rook::Bool=false) -> BunchKaufman
```

bkfact! is the same as bkfact, but saves space by overwriting the input A, instead of creating a copy.

source

Base.LinAlg.eig – Function.

```
eig(A::Union{SymTridiagonal, Hermitian, Symmetric}, irange::UnitRange) -> D, V
eig(A::Union{SymTridiagonal, Hermitian, Symmetric}, vl::Real, vu::Real) -> D, V
eig(A, permute::Bool=true, scale::Bool=true) -> D, V
```

Computes eigenvalues (D) and eigenvectors (V) of A. See eigfact for details on the irange, vl, and vu arguments (for SymTridiagonal, Hermitian, and Symmetric matrices) and the permute and scale keyword arguments. The eigenvectors are returned columnwise.

Example

```
julia> eig([1.0 0.0 0.0; 0.0 3.0 0.0; 0.0 0.0 18.0])
([1.0, 3.0, 18.0], [1.0 0.0 0.0; 0.0 1.0 0.0; 0.0 0.0 1.0])
```

eig is a wrapper around eigfact, extracting all parts of the factorization to a tuple; where possible, using eigfact is recommended.

source

```
eig(A, B) -> D, V
```

Computes generalized eigenvalues (D) and vectors (V) of A with respect to B.

eig is a wrapper around eigfact, extracting all parts of the factorization to a tuple; where possible, using eigfact is recommended.

Example

---

[3] J R Bunch and L Kaufman, Some stable methods for calculating inertia and solving symmetric linear systems, Mathematics of Computation 31:137 (1977), 163-179. url.

```julia
julia> A = [1 0; 0 -1]
2×2 Array{Int64,2}:
 1   0
 0  -1

julia> B = [0 1; 1 0]
2×2 Array{Int64,2}:
 0  1
 1  0

julia> eig(A, B)
(Complex{Float64}[0.0+1.0im, 0.0-1.0im], Complex{Float64}[0.0-1.0im 0.0+1.0im; -1.0-0.0im
↪   -1.0+0.0im])
```

source

Base.LinAlg.eigvals – Function.

```
eigvals(A; permute::Bool=true, scale::Bool=true) -> values
```

Returns the eigenvalues of A.

For general non-symmetric matrices it is possible to specify how the matrix is balanced before the eigenvalue cal-culation. The option `permute=true` permutes the matrix to become closer to upper triangular, and `scale=true` scales the matrix by its diagonal elements to make rows and columns more equal in norm. The default is `true` for both options.

source

```
eigvals(A, B) -> values
```

Computes the generalized eigenvalues of A and B.

Example

```julia
julia> A = [1 0; 0 -1]
2×2 Array{Int64,2}:
 1   0
 0  -1

julia> B = [0 1; 1 0]
2×2 Array{Int64,2}:
 0  1
 1  0

julia> eigvals(A,B)
2-element Array{Complex{Float64},1}:
 0.0+1.0im
 0.0-1.0im
```

source

```
eigvals(A::Union{SymTridiagonal, Hermitian, Symmetric}, irange::UnitRange) -> values
```

Returns the eigenvalues of A. It is possible to calculate only a subset of the eigenvalues by specifying a `UnitRange` `irange` covering indices of the sorted eigenvalues, e.g. the 2nd to 8th eigenvalues.

```julia
julia> A = SymTridiagonal([1.; 2.; 1.], [2.; 3.])
3×3 SymTridiagonal{Float64}:
 1.0  2.0
 2.0  2.0  3.0
      3.0  1.0

julia> eigvals(A, 2:2)
1-element Array{Float64,1}:
 1.0

julia> eigvals(A)
3-element Array{Float64,1}:
 -2.14005
  1.0
  5.14005
```

source

```
eigvals(A::Union{SymTridiagonal, Hermitian, Symmetric}, vl::Real, vu::Real) -> values
```

Returns the eigenvalues of A. It is possible to calculate only a subset of the eigenvalues by specifying a pair vl and vu for the lower and upper boundaries of the eigenvalues.

```julia
julia> A = SymTridiagonal([1.; 2.; 1.], [2.; 3.])
3×3 SymTridiagonal{Float64}:
 1.0  2.0
 2.0  2.0  3.0
      3.0  1.0

julia> eigvals(A, -1, 2)
1-element Array{Float64,1}:
 1.0

julia> eigvals(A)
3-element Array{Float64,1}:
 -2.14005
  1.0
  5.14005
```

source

Base.LinAlg.eigvals! – Function.

```
eigvals!(A; permute::Bool=true, scale::Bool=true) -> values
```

Same as eigvals, but saves space by overwriting the input A, instead of creating a copy. The option permute=true permutes the matrix to become closer to upper triangular, and scale=true scales the matrix by its diagonal elements to make rows and columns more equal in norm.

source

```
eigvals!(A, B) -> values
```

Same as eigvals, but saves space by overwriting the input A (and B), instead of creating copies.

source

```
eigvals!(A::Union{SymTridiagonal, Hermitian, Symmetric}, irange::UnitRange) -> values
```

Same as `eigvals`, but saves space by overwriting the input A, instead of creating a copy. `irange` is a range of eigenvalue *indices* to search for - for instance, the 2nd to 8th eigenvalues.

source

```
eigvals!(A::Union{SymTridiagonal, Hermitian, Symmetric}, vl::Real, vu::Real) -> values
```

Same as `eigvals`, but saves space by overwriting the input A, instead of creating a copy. `vl` is the lower bound of the interval to search for eigenvalues, and `vu` is the upper bound.

source

Base.LinAlg.eigmax – Function.

```
eigmax(A; permute::Bool=true, scale::Bool=true)
```

Returns the largest eigenvalue of A. The option `permute=true` permutes the matrix to become closer to upper triangular, and `scale=true` scales the matrix by its diagonal elements to make rows and columns more equal in norm. Note that if the eigenvalues of A are complex, this method will fail, since complex numbers cannot be sorted.

Example

```
julia> A = [0 im; -im 0]
2×2 Array{Complex{Int64},2}:
 0+0im  0+1im
 0-1im  0+0im

julia> eigmax(A)
1.0

julia> A = [0 im; -1 0]
2×2 Array{Complex{Int64},2}:
  0+0im  0+1im
 -1+0im  0+0im

julia> eigmax(A)
ERROR: DomainError:
Stacktrace:
 [1] #eigmax#46(::Bool, ::Bool, ::Function, ::Array{Complex{Int64},2}) at
 ↪   ./linalg/eigen.jl:238
 [2] eigmax(::Array{Complex{Int64},2}) at ./linalg/eigen.jl:236
```

source

Base.LinAlg.eigmin – Function.

```
eigmin(A; permute::Bool=true, scale::Bool=true)
```

Returns the smallest eigenvalue of A. The option `permute=true` permutes the matrix to become closer to upper triangular, and `scale=true` scales the matrix by its diagonal elements to make rows and columns more equal in norm. Note that if the eigenvalues of A are complex, this method will fail, since complex numbers cannot be sorted.

Example

```
julia> A = [0 im; -im 0]
2×2 Array{Complex{Int64},2}:
```

```
   0+0im   0+1im
   0-1im   0+0im

julia> eigmin(A)
-1.0

julia> A = [0 im; -1 0]
2×2 Array{Complex{Int64},2}:
  0+0im   0+1im
 -1+0im   0+0im

julia> eigmin(A)
ERROR: DomainError:
Stacktrace:
 [1] #eigmin#47(::Bool, ::Bool, ::Function, ::Array{Complex{Int64},2}) at
 ↪   ./linalg/eigen.jl:280
 [2] eigmin(::Array{Complex{Int64},2}) at ./linalg/eigen.jl:278
```

source

Base.LinAlg.eigvecs – Function.

```
eigvecs(A::SymTridiagonal[, eigvals]) -> Matrix
```

Returns a matrix M whose columns are the eigenvectors of A. (The kth eigenvector can be obtained from the slice M[:, k].)

If the optional vector of eigenvalues eigvals is specified, eigvecs returns the specific corresponding eigenvectors.

Example

```
julia> A = SymTridiagonal([1.; 2.; 1.], [2.; 3.])
3×3 SymTridiagonal{Float64}:
 1.0  2.0
 2.0  2.0  3.0
      3.0  1.0

julia> eigvals(A)
3-element Array{Float64,1}:
 -2.14005
  1.0
  5.14005

julia> eigvecs(A)
3×3 Array{Float64,2}:
  0.418304  -0.83205      0.364299
 -0.656749  -7.39009e-16  0.754109
  0.627457   0.5547       0.546448

julia> eigvecs(A, [1.])
3×1 Array{Float64,2}:
  0.83205
  4.26351e-17
 -0.5547
```

source

```
eigvecs(A; permute::Bool=true, scale::Bool=true) -> Matrix
```

Returns a matrix M whose columns are the eigenvectors of A. (The kth eigenvector can be obtained from the slice M[:, k].) The permute and scale keywords are the same as for eigfact.

Example

```julia
julia> eigvecs([1.0 0.0 0.0; 0.0 3.0 0.0; 0.0 0.0 18.0])
3×3 Array{Float64,2}:
 1.0  0.0  0.0
 0.0  1.0  0.0
 0.0  0.0  1.0
```

source

```
eigvecs(A, B) -> Matrix
```

Returns a matrix M whose columns are the generalized eigenvectors of A and B. (The kth eigenvector can be obtained from the slice M[:, k].)

Example

```julia
julia> A = [1 0; 0 -1]
2×2 Array{Int64,2}:
 1   0
 0  -1

julia> B = [0 1; 1 0]
2×2 Array{Int64,2}:
 0  1
 1  0

julia> eigvecs(A, B)
2×2 Array{Complex{Float64},2}:
  0.0-1.0im   0.0+1.0im
 -1.0-0.0im  -1.0+0.0im
```

source

Base.LinAlg.eigfact – Function.

```
eigfact(A; permute::Bool=true, scale::Bool=true) -> Eigen
```

Computes the eigenvalue decomposition of A, returning an Eigen factorization object F which contains the eigenvalues in F[:values] and the eigenvectors in the columns of the matrix F[:vectors]. (The kth eigenvector can be obtained from the slice F[:vectors][:, k].)

The following functions are available for Eigen objects: inv, det, and isposdef.

For general nonsymmetric matrices it is possible to specify how the matrix is balanced before the eigenvector calculation. The option permute=true permutes the matrix to become closer to upper triangular, and scale=true scales the matrix by its diagonal elements to make rows and columns more equal in norm. The default is true for both options.

Example

```
julia> F = eigfact([1.0 0.0 0.0; 0.0 3.0 0.0; 0.0 0.0 18.0])
Base.LinAlg.Eigen{Float64,Float64,Array{Float64,2},Array{Float64,1}}([1.0, 3.0, 18.0], [1.0
↪    0.0 0.0; 0.0 1.0 0.0; 0.0 0.0 1.0])

julia> F[:values]
3-element Array{Float64,1}:
  1.0
  3.0
 18.0

julia> F[:vectors]
3×3 Array{Float64,2}:
 1.0  0.0  0.0
 0.0  1.0  0.0
 0.0  0.0  1.0
```

source

```
eigfact(A, B) -> GeneralizedEigen
```

Computes the generalized eigenvalue decomposition of A and B, returning a `GeneralizedEigen` factorization object F which contains the generalized eigenvalues in F[:values] and the generalized eigenvectors in the columns of the matrix F[:vectors]. (The kth generalized eigenvector can be obtained from the slice F[:vectors][:, k].)

source

```
eigfact(A::Union{SymTridiagonal, Hermitian, Symmetric}, irange::UnitRange) -> Eigen
```

Computes the eigenvalue decomposition of A, returning an `Eigen` factorization object F which contains the eigenvalues in F[:values] and the eigenvectors in the columns of the matrix F[:vectors]. (The kth eigenvector can be obtained from the slice F[:vectors][:, k].)

The following functions are available for `Eigen` objects: inv, det, and isposdef.

The `UnitRange` irange specifies indices of the sorted eigenvalues to search for.

> **Note**
> If irange is not 1:n, where n is the dimension of A, then the returned factorization will be a *truncated* factorization.

source

```
eigfact(A::Union{SymTridiagonal, Hermitian, Symmetric}, vl::Real, vu::Real) -> Eigen
```

Computes the eigenvalue decomposition of A, returning an `Eigen` factorization object F which contains the eigenvalues in F[:values] and the eigenvectors in the columns of the matrix F[:vectors]. (The kth eigenvector can be obtained from the slice F[:vectors][:, k].)

The following functions are available for `Eigen` objects: inv, det, and isposdef.

vl is the lower bound of the window of eigenvalues to search for, and vu is the upper bound.

> **Note**
> If [vl, vu] does not contain all eigenvalues of A, then the returned factorization will be a *truncated* factorization.

    source

Base.LinAlg.eigfact! – Function.

    eigfact!(A, [B])

Same as eigfact, but saves space by overwriting the input A (and B), instead of creating a copy.

    source

Base.LinAlg.hessfact – Function.

    hessfact(A) -> Hessenberg

Compute the Hessenberg decomposition of A and return a Hessenberg object. If F is the factorization object, the unitary matrix can be accessed with F[:Q] and the Hessenberg matrix with F[:H]. When Q is extracted, the resulting type is the HessenbergQ object, and may be converted to a regular matrix with convert(Array, _) (or Array(_) for short).

Example

```julia
julia> A = [4. 9. 7.; 4. 4. 1.; 4. 3. 2.]
3×3 Array{Float64,2}:
 4.0  9.0  7.0
 4.0  4.0  1.0
 4.0  3.0  2.0

julia> F = hessfact(A);

julia> F[:Q] * F[:H] * F[:Q]'
3×3 Array{Float64,2}:
 4.0  9.0  7.0
 4.0  4.0  1.0
 4.0  3.0  2.0
```

    source

Base.LinAlg.hessfact! – Function.

    hessfact!(A) -> Hessenberg

hessfact! is the same as hessfact, but saves space by overwriting the input A, instead of creating a copy.

    source

Base.LinAlg.schurfact – Function.

    schurfact(A::StridedMatrix) -> F::Schur

Computes the Schur factorization of the matrix A. The (quasi) triangular Schur factor can be obtained from the Schur object F with either F[:Schur] or F[:T] and the orthogonal/unitary Schur vectors can be obtained with F[:vectors] or F[:Z] such that A = F[:vectors]*F[:Schur]*F[:vectors]'. The eigenvalues of A can be obtained with F[:values].

Example

```
julia> A = [-2. 1. 3.; 2. 1. -1.; -7. 2. 7.]
3×3 Array{Float64,2}:
 -2.0  1.0   3.0
  2.0  1.0  -1.0
 -7.0  2.0   7.0

julia> F = schurfact(A)
Base.LinAlg.Schur{Float64,Array{Float64,2}} with factors T and Z:
[2.0 0.801792 6.63509; -8.55988e-11 2.0 8.08286; 0.0 0.0 1.99999]
[0.577351 0.154299 -0.801784; 0.577346 -0.77152 0.267262; 0.577354 0.617211 0.534522]
and values:
Complex{Float64}[2.0+8.28447e-6im, 2.0-8.28447e-6im, 1.99999+0.0im]

julia> F[:vectors] * F[:Schur] * F[:vectors]'
3×3 Array{Float64,2}:
 -2.0  1.0   3.0
  2.0  1.0  -1.0
 -7.0  2.0   7.0
```

[source]

```
schurfact(A::StridedMatrix, B::StridedMatrix) -> F::GeneralizedSchur
```

Computes the Generalized Schur (or QZ) factorization of the matrices A and B. The (quasi) triangular Schur factors can be obtained from the Schur object F with F[:S] and F[:T], the left unitary/orthogonal Schur vectors can be obtained with F[:left] or F[:Q] and the right unitary/orthogonal Schur vectors can be obtained with F[:right] or F[:Z] such that A=F[:left]*F[:S]*F[:right]' and B=F[:left]*F[:T]*F[:right]'. The generalized eigenvalues of A and B can be obtained with F[:alpha]./F[:beta].

[source]

Base.LinAlg.schurfact! – Function.

```
schurfact!(A::StridedMatrix) -> F::Schur
```

Same as schurfact but uses the input argument as workspace.

[source]

```
schurfact!(A::StridedMatrix, B::StridedMatrix) -> F::GeneralizedSchur
```

Same as schurfact but uses the input matrices A and B as workspace.

[source]

Base.LinAlg.schur – Function.

```
schur(A::StridedMatrix) -> T::Matrix, Z::Matrix, λ::Vector
```

Computes the Schur factorization of the matrix A. The methods return the (quasi) triangular Schur factor T and the orthogonal/unitary Schur vectors Z such that A = Z*T*Z'. The eigenvalues of A are returned in the vector λ.

See schurfact.

Example

```
julia> A = [-2. 1. 3.; 2. 1. -1.; -7. 2. 7.]
3×3 Array{Float64,2}:
 -2.0  1.0   3.0
```

```
  2.0  1.0  -1.0
 -7.0  2.0   7.0
```

julia> T, Z, lambda = schur(A)
([2.0 0.801792 6.63509; -8.55988e-11 2.0 8.08286; 0.0 0.0 1.99999], [0.577351 0.154299
↪   -0.801784; 0.577346 -0.77152 0.267262; 0.577354 0.617211 0.534522],
↪   Complex{Float64}[2.0+8.28447e-6im, 2.0-8.28447e-6im, 1.99999+0.0im])

julia> Z * T * Z'
3×3 Array{Float64,2}:
 -2.0 1.0   3.0
  2.0 1.0 -1.0
 -7.0 2.0   7.0

source

```
schur(A::StridedMatrix, B::StridedMatrix) -> S::StridedMatrix, T::StridedMatrix, Q::
    StridedMatrix, Z::StridedMatrix, α::Vector, β::Vector
```

See schurfact.

source

Base.LinAlg.ordschur – Function.

```
ordschur(F::Schur, select::Union{Vector{Bool},BitVector}) -> F::Schur
```

Reorders the Schur factorization F of a matrix A = Z*T*Z' according to the logical array `select` returning the reordered factorization F object. The selected eigenvalues appear in the leading diagonal of F[:Schur] and the corresponding leading columns of F[:vectors] form an orthogonal/unitary basis of the corresponding right invariant subspace. In the real case, a complex conjugate pair of eigenvalues must be either both included or both excluded via `select`.

source

```
ordschur(T::StridedMatrix, Z::StridedMatrix, select::Union{Vector{Bool},BitVector}) -> T::
    StridedMatrix, Z::StridedMatrix, λ::Vector
```

Reorders the Schur factorization of a real matrix A = Z*T*Z' according to the logical array `select` returning the reordered matrices T and Z as well as the vector of eigenvalues λ. The selected eigenvalues appear in the leading diagonal of T and the corresponding leading columns of Z form an orthogonal/unitary basis of the corresponding right invariant subspace. In the real case, a complex conjugate pair of eigenvalues must be either both included or both excluded via `select`.

source

```
ordschur(F::GeneralizedSchur, select::Union{Vector{Bool},BitVector}) -> F::GeneralizedSchur
```

Reorders the Generalized Schur factorization F of a matrix pair (A, B) = (Q*S*Z', Q*T*Z') according to the logical array `select` and returns a GeneralizedSchur object F. The selected eigenvalues appear in the leading diagonal of both F[:S] and F[:T], and the left and right orthogonal/unitary Schur vectors are also reordered such that (A, B) = F[:Q]*(F[:S], F[:T])*F[:Z]' still holds and the generalized eigenvalues of A and B can still be obtained with F[:alpha]./F[:beta].

source

```
ordschur(S::StridedMatrix, T::StridedMatrix, Q::StridedMatrix, Z::StridedMatrix, select) -> S
    ::StridedMatrix, T::StridedMatrix, Q::StridedMatrix, Z::StridedMatrix, α::Vector, β::
    Vector
```

Reorders the Generalized Schur factorization of a matrix pair (A, B) = (Q*S*Z', Q*T*Z') according to the logical array `select` and returns the matrices S, T, Q, Z and vectors α and β. The selected eigenvalues appear in the leading diagonal of both S and T, and the left and right unitary/orthogonal Schur vectors are also reordered such that (A, B) = Q*(S, T)*Z' still holds and the generalized eigenvalues of A and B can still be obtained with α./β.

source

`Base.LinAlg.ordschur!` – Function.

```
ordschur!(F::Schur, select::Union{Vector{Bool},BitVector}) -> F::Schur
```

Same as ordschur but overwrites the factorization F.

source

```
ordschur!(T::StridedMatrix, Z::StridedMatrix, select::Union{Vector{Bool},BitVector}) -> T::
    StridedMatrix, Z::StridedMatrix, λ::Vector
```

Same as ordschur but overwrites the input arguments.

source

```
ordschur!(F::GeneralizedSchur, select::Union{Vector{Bool},BitVector}) -> F::GeneralizedSchur
```

Same as ordschur but overwrites the factorization F.

source

```
ordschur!(S::StridedMatrix, T::StridedMatrix, Q::StridedMatrix, Z::StridedMatrix, select) ->
    S::StridedMatrix, T::StridedMatrix, Q::StridedMatrix, Z::StridedMatrix, α::Vector, β::
    Vector
```

Same as ordschur but overwrites the factorization the input arguments.

source

`Base.LinAlg.svdfact` – Function.

```
svdfact(A, thin::Bool=true) -> SVD
```

Compute the singular value decomposition (SVD) of A and return an SVD object.

U, S, V and Vt can be obtained from the factorization F with F[:U], F[:S], F[:V] and F[:Vt], such that A = U*diagm(S)*Vt. The algorithm produces Vt and hence Vt is more efficient to extract than V. The singular values in S are sorted in descending order.

If `thin=true` (default), a thin SVD is returned. For a $M \times N$ matrix A, U is $M \times M$ for a full SVD (thin=false) and $M \times \min(M, N)$ for a thin SVD.

Example

```julia
julia> A = [1. 0. 0. 0. 2.; 0. 0. 3. 0. 0.; 0. 0. 0. 0. 0.; 0. 2. 0. 0. 0.]
4×5 Array{Float64,2}:
 1.0  0.0  0.0  0.0  2.0
 0.0  0.0  3.0  0.0  0.0
 0.0  0.0  0.0  0.0  0.0
 0.0  2.0  0.0  0.0  0.0

julia> F = svdfact(A)
```

```
Base.LinAlg.SVD{Float64,Float64,Array{Float64,2}}([0.0 1.0 0.0 0.0; 1.0 0.0 0.0 0.0; 0.0 0.0
↪    0.0 -1.0; 0.0 0.0 1.0 0.0], [3.0, 2.23607, 2.0, 0.0], [-0.0 0.0 … -0.0 0.0; 0.447214 0.0
↪    … 0.0 0.894427; -0.0 1.0 … -0.0 0.0; 0.0 0.0 … 1.0 0.0])

julia> F[:U] * diagm(F[:S]) * F[:Vt]
4×5 Array{Float64,2}:
 1.0  0.0  0.0  0.0  2.0
 0.0  0.0  3.0  0.0  0.0
 0.0  0.0  0.0  0.0  0.0
 0.0  2.0  0.0  0.0  0.0
```

source

```
svdfact(A, B) -> GeneralizedSVD
```

Compute the generalized SVD of A and B, returning a `GeneralizedSVD` factorization object F, such that `A = F[:U]*F[:D1]*F[:R0]*F[:Q]'` and `B = F[:V]*F[:D2]*F[:R0]*F[:Q]'`.

For an M-by-N matrix A and P-by-N matrix B,

- `F[:U]` is a M-by-M orthogonal matrix,
- `F[:V]` is a P-by-P orthogonal matrix,
- `F[:Q]` is a N-by-N orthogonal matrix,
- `F[:R0]` is a (K+L)-by-N matrix whose rightmost (K+L)-by-(K+L) block is nonsingular upper block triangular,
- `F[:D1]` is a M-by-(K+L) diagonal matrix with 1s in the first K entries,
- `F[:D2]` is a P-by-(K+L) matrix whose top right L-by-L block is diagonal,

K+L is the effective numerical rank of the matrix `[A; B]`.

The entries of `F[:D1]` and `F[:D2]` are related, as explained in the LAPACK documentation for the generalized SVD and the xGGSVD3 routine which is called underneath (in LAPACK 3.6.0 and newer).

source

Base.LinAlg.svdfact! – Function.

```
svdfact!(A, thin::Bool=true) -> SVD
```

`svdfact!` is the same as `svdfact`, but saves space by overwriting the input A, instead of creating a copy.

source

```
svdfact!(A, B) -> GeneralizedSVD
```

`svdfact!` is the same as `svdfact`, but modifies the arguments A and B in-place, instead of making copies.

source

Base.LinAlg.svd – Function.

```
svd(A, thin::Bool=true) -> U, S, V
```

Computes the SVD of A, returning U, vector S, and V such that `A == U*diagm(S)*V'`. The singular values in S are sorted in descending order.

If `thin=true` (default), a thin SVD is returned. For a $M \times N$ matrix A, U is $M \times M$ for a full SVD (`thin=false`) and $M \times \min(M, N)$ for a thin SVD.

svd is a wrapper around `svdfact`, extracting all parts of the SVD factorization to a tuple. Direct use of `svdfact` is therefore more efficient.

Example

```julia
julia> A = [1. 0. 0. 0. 2.; 0. 0. 3. 0. 0.; 0. 0. 0. 0. 0.; 0. 2. 0. 0. 0.]
4×5 Array{Float64,2}:
 1.0  0.0  0.0  0.0  2.0
 0.0  0.0  3.0  0.0  0.0
 0.0  0.0  0.0  0.0  0.0
 0.0  2.0  0.0  0.0  0.0

julia> U, S, V = svd(A)
([0.0 1.0 0.0 0.0; 1.0 0.0 0.0 0.0; 0.0 0.0 0.0 -1.0; 0.0 0.0 1.0 0.0], [3.0, 2.23607, 2.0,
↪    0.0], [-0.0 0.447214 -0.0 0.0; 0.0 0.0 1.0 0.0; … ; -0.0 0.0 -0.0 1.0; 0.0 0.894427 0.0
↪    0.0])

julia> U*diagm(S)*V'
4×5 Array{Float64,2}:
 1.0  0.0  0.0  0.0  2.0
 0.0  0.0  3.0  0.0  0.0
 0.0  0.0  0.0  0.0  0.0
 0.0  2.0  0.0  0.0  0.0
```

source

```
svd(A, B) -> U, V, Q, D1, D2, R0
```

Wrapper around `svdfact` extracting all parts of the factorization to a tuple. Direct use of `svdfact` is therefore generally more efficient. The function returns the generalized SVD of A and B, returning U, V, Q, D1, D2, and R0 such that `A = U*D1*R0*Q'` and `B = V*D2*R0*Q'`.

source

`Base.LinAlg.svdvals` – Function.

```
svdvals(A)
```

Returns the singular values of A in descending order.

Example

```julia
julia> A = [1. 0. 0. 0. 2.; 0. 0. 3. 0. 0.; 0. 0. 0. 0. 0.; 0. 2. 0. 0. 0.]
4×5 Array{Float64,2}:
 1.0  0.0  0.0  0.0  2.0
 0.0  0.0  3.0  0.0  0.0
 0.0  0.0  0.0  0.0  0.0
 0.0  2.0  0.0  0.0  0.0

julia> svdvals(A)
4-element Array{Float64,1}:
 3.0
 2.23607
 2.0
 0.0
```

source

```
svdvals(A, B)
```

Return the generalized singular values from the generalized singular value decomposition of A and B. See also svdfact.

source

Base.LinAlg.Givens – Type.

```
LinAlg.Givens(i1,i2,c,s) -> G
```

A Givens rotation linear operator. The fields c and s represent the cosine and sine of the rotation angle, respectively. The Givens type supports left multiplication G*A and conjugated transpose right multiplication A*G'. The type doesn't have a size and can therefore be multiplied with matrices of arbitrary size as long as i2<=size(A,2) for G*A or i2<=size(A,1) for A*G'.

See also: givens

source

Base.LinAlg.givens – Function.

```
givens{T}(f::T, g::T, i1::Integer, i2::Integer) -> (G::Givens, r::T)
```

Computes the Givens rotation G and scalar r such that for any vector x where

```
x[i1] = f
x[i2] = g
```

the result of the multiplication

```
y = G*x
```

has the property that

```
y[i1] = r
y[i2] = 0
```

See also: LinAlg.Givens

source

```
givens(A::AbstractArray, i1::Integer, i2::Integer, j::Integer) -> (G::Givens, r)
```

Computes the Givens rotation G and scalar r such that the result of the multiplication

```
B = G*A
```

has the property that

```
B[i1,j] = r
B[i2,j] = 0
```

See also: LinAlg.Givens

source

```
givens(x::AbstractVector, i1::Integer, i2::Integer) -> (G::Givens, r)
```

Computes the Givens rotation G and scalar r such that the result of the multiplication

```
B = G*x
```

has the property that

```
B[i1] = r
B[i2] = 0
```

See also: `LinAlg.Givens`

`source`

`Base.LinAlg.triu` – Function.

```
triu(M)
```

Upper triangle of a matrix.

Example

```
julia> a = ones(4,4)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0

julia> triu(a)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 0.0  1.0  1.0  1.0
 0.0  0.0  1.0  1.0
 0.0  0.0  0.0  1.0
```

`source`

```
triu(M, k::Integer)
```

Returns the upper triangle of `M` starting from the `k`th superdiagonal.

Example

```
julia> a = ones(4,4)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0

julia> triu(a,3)
4×4 Array{Float64,2}:
 0.0  0.0  0.0  1.0
 0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0

julia> triu(a,-3)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
```

```
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
```

source

Base.LinAlg.triu! – Function.

```
triu!(M)
```

Upper triangle of a matrix, overwriting M in the process. See also triu.

source

```
triu!(M, k::Integer)
```

Returns the upper triangle of M starting from the kth superdiagonal, overwriting M in the process.

Example

```julia
julia> M = [1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5]
5×5 Array{Int64,2}:
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5

julia> triu!(M, 1)
5×5 Array{Int64,2}:
 0  2  3  4  5
 0  0  3  4  5
 0  0  0  4  5
 0  0  0  0  5
 0  0  0  0  0
```

source

Base.LinAlg.tril – Function.

```
tril(M)
```

Lower triangle of a matrix.

Example

```julia
julia> a = ones(4,4)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0

julia> tril(a)
4×4 Array{Float64,2}:
 1.0  0.0  0.0  0.0
 1.0  1.0  0.0  0.0
 1.0  1.0  1.0  0.0
 1.0  1.0  1.0  1.0
```

```
tril(M, k::Integer)
```

Returns the lower triangle of M starting from the kth superdiagonal.

Example

```julia
julia> a = ones(4,4)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0

julia> tril(a,3)
4×4 Array{Float64,2}:
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0
 1.0  1.0  1.0  1.0

julia> tril(a,-3)
4×4 Array{Float64,2}:
 0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0
 0.0  0.0  0.0  0.0
 1.0  0.0  0.0  0.0
```

`Base.LinAlg.tril!` – Function.

```
tril!(M)
```

Lower triangle of a matrix, overwriting M in the process. See also `tril`.

```
tril!(M, k::Integer)
```

Returns the lower triangle of M starting from the kth superdiagonal, overwriting M in the process.

Example

```julia
julia> M = [1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5; 1 2 3 4 5]
5×5 Array{Int64,2}:
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5

julia> tril!(M, 2)
5×5 Array{Int64,2}:
 1  2  3  0  0
 1  2  3  4  0
 1  2  3  4  5
 1  2  3  4  5
 1  2  3  4  5
```

source

`Base.LinAlg.diagind` – Function.

```
diagind(M, k::Integer=0)
```

A Range giving the indices of the kth diagonal of the matrix M.

Example

```julia
julia> A = [1 2 3; 4 5 6; 7 8 9]
3×3 Array{Int64,2}:
 1  2  3
 4  5  6
 7  8  9

julia> diagind(A,-1)
2:4:6
```

source

`Base.LinAlg.diag` – Function.

```
diag(M, k::Integer=0)
```

The kth diagonal of a matrix, as a vector. Use `diagm` to construct a diagonal matrix.

Example

```julia
julia> A = [1 2 3; 4 5 6; 7 8 9]
3×3 Array{Int64,2}:
 1  2  3
 4  5  6
 7  8  9

julia> diag(A,1)
2-element Array{Int64,1}:
 2
 6
```

source

`Base.LinAlg.diagm` – Function.

```
diagm(v, k::Integer=0)
```

Construct a matrix by placing v on the kth diagonal.

Example

```julia
julia> diagm([1,2,3],1)
4×4 Array{Int64,2}:
 0  1  0  0
 0  0  2  0
 0  0  0  3
 0  0  0  0
```

source

Base.LinAlg.scale! – Function.

```
scale!(A, b)
scale!(b, A)
```

Scale an array A by a scalar b overwriting A in-place.

If A is a matrix and b is a vector, then scale!(A,b) scales each column i of A by b[i] (similar to A*Diagonal(b)), while scale!(b,A) scales each row i of A by b[i] (similar to Diagonal(b)*A), again operating in-place on A. An InexactError exception is thrown if the scaling produces a number not representable by the element type of A, e.g. for integer types.

Example

```julia
julia> a = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> b = [1; 2]
2-element Array{Int64,1}:
 1
 2

julia> scale!(a,b)
2×2 Array{Int64,2}:
 1  4
 3  8

julia> a = [1 2; 3 4];

julia> b = [1; 2];

julia> scale!(b,a)
2×2 Array{Int64,2}:
 1  2
 6  8
```

source

Base.LinAlg.rank – Function.

```
rank(M[, tol::Real])
```

Compute the rank of a matrix by counting how many singular values of M have magnitude greater than tol. By default, the value of tol is the largest dimension of M multiplied by the eps of the eltype of M.

source

Base.LinAlg.norm – Function.

```
norm(A::AbstractArray, p::Real=2)
```

Compute the p-norm of a vector or the operator norm of a matrix A, defaulting to the 2-norm.

```
norm(A::AbstractVector, p::Real=2)
```

For vectors, this is equivalent to `vecnorm` and equal to:

$$\|A\|_p = \left(\sum_{i=1}^{n} |a_i|^p \right)^{1/p}$$

with a_i the entries of A and n its length.

p can assume any numeric value (even though not all values produce a mathematically valid vector norm). In particular, `norm(A, Inf)` returns the largest value in `abs(A)`, whereas `norm(A, -Inf)` returns the smallest.

Example

```
julia> v = [3, -2, 6]
3-element Array{Int64,1}:
  3
 -2
  6

julia> norm(v)
7.0

julia> norm(v, Inf)
6.0
```

source

```
norm(A::AbstractMatrix, p::Real=2)
```

For matrices, the matrix norm induced by the vector p-norm is used, where valid values of p are 1, 2, or `Inf`. (Note that for sparse matrices, p=2 is currently not implemented.) Use `vecnorm` to compute the Frobenius norm.

When p=1, the matrix norm is the maximum absolute column sum of A:

$$\|A\|_1 = \max_{1 \leq j \leq n} \sum_{i=1}^{m} |a_{ij}|$$

with a_{ij} the entries of A, and m and n its dimensions.

When p=2, the matrix norm is the spectral norm, equal to the largest singular value of A.

When p=`Inf`, the matrix norm is the maximum absolute row sum of A:

$$\|A\|_\infty = \max_{1 \leq i \leq m} \sum_{j=1}^{n} |a_{ij}|$$

Example

```
julia> A = [1 -2 -3; 2 3 -1]
2×3 Array{Int64,2}:
 1  -2  -3
 2   3  -1

julia> norm(A, Inf)
6.0
```

source

```
norm(x::Number, p::Real=2)
```

For numbers, return $(|x|^p)^{1/p}$. This is equivalent to `vecnorm`.

source

```
norm(A::RowVector, q::Real=2)
```

For row vectors, return the q-norm of A, which is equivalent to the p-norm with value `p = q/(q-1)`. They coincide at `p = q = 2`.

The difference in norm between a vector space and its dual arises to preserve the relationship between duality and the inner product, and the result is consistent with the p-norm of `1 × n` matrix.

source

Base.LinAlg.vecnorm – Function.

```
vecnorm(A, p::Real=2)
```

For any iterable container A (including arrays of any dimension) of numbers (or any element type for which `norm` is defined), compute the p-norm (defaulting to p=2) as if A were a vector of the corresponding length.

The p-norm is defined as:

$$\|A\|_p = \left(\sum_{i=1}^{n} |a_i|^p \right)^{1/p}$$

with a_i the entries of A and n its length.

p can assume any numeric value (even though not all values produce a mathematically valid vector norm). In particular, `vecnorm(A, Inf)` returns the largest value in `abs(A)`, whereas `vecnorm(A, -Inf)` returns the smallest. If A is a matrix and p=2, then this is equivalent to the Frobenius norm.

Example

```
julia> vecnorm([1 2 3; 4 5 6; 7 8 9])
16.881943016134134

julia> vecnorm([1 2 3 4 5 6 7 8 9])
16.881943016134134
```

source

```
vecnorm(x::Number, p::Real=2)
```

For numbers, return $(|x|^p)^{1/p}$.

source

Base.LinAlg.normalize! – Function.

```
normalize!(v::AbstractVector, p::Real=2)
```

Normalize the vector v in-place so that its p-norm equals unity, i.e. `norm(v, p) == 1`. See also `normalize` and `vecnorm`.

source

Base.LinAlg.normalize – Function.

```
normalize(v::AbstractVector, p::Real=2)
```

Normalize the vector v so that its p-norm equals unity, i.e. `norm(v, p) == vecnorm(v, p) == 1`. See also `normalize!` and `vecnorm`.

Example

```
julia> a = [1,2,4];

julia> b = normalize(a)
3-element Array{Float64,1}:
 0.218218
 0.436436
 0.872872

julia> norm(b)
1.0

julia> c = normalize(a, 1)
3-element Array{Float64,1}:
 0.142857
 0.285714
 0.571429

julia> norm(c, 1)
1.0
```

source

Base.LinAlg.cond – Function.

```
cond(M, p::Real=2)
```

Condition number of the matrix M, computed using the operator p-norm. Valid values for p are 1, 2 (default), or `Inf`.

source

Base.LinAlg.condskeel – Function.

```
condskeel(M, [x, p::Real=Inf])
```

$$\kappa_S(M, p) = \left\| \left| M \right| \left| M^{-1} \right| \right\|_p$$
$$\kappa_S(M, x, p) = \left\| \left| M \right| \left| M^{-1} \right| \left| x \right| \right\|_p$$

Skeel condition number κ_S of the matrix M, optionally with respect to the vector x, as computed using the operator p-norm. p is `Inf` by default, if not provided. Valid values for p are 1, 2, or `Inf`.

This quantity is also known in the literature as the Bauer condition number, relative condition number, or componentwise relative condition number.

source

Base.LinAlg.trace – Function.

```
trace(M)
```

Matrix trace. Sums the diagonal elements of M.

Example

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> trace(A)
5
```

source

Base.LinAlg.det – Function.

```
det(M)
```

Matrix determinant.

Example

```
julia> M = [1 0; 2 2]
2×2 Array{Int64,2}:
 1  0
 2  2

julia> det(M)
2.0
```

source

Base.LinAlg.logdet – Function.

```
logdet(M)
```

Log of matrix determinant. Equivalent to log(det(M)), but may provide increased accuracy and/or speed.

Example

```
julia> M = [1 0; 2 2]
2×2 Array{Int64,2}:
 1  0
 2  2

julia> logdet(M)
0.6931471805599453

julia> logdet(eye(3))
0.0
```

source

Base.LinAlg.logabsdet – Function.

```
logabsdet(M)
```

Log of absolute value of matrix determinant. Equivalent to (log(abs(det(M))), sign(det(M))), but may provide increased accuracy and/or speed.

source

`Base.inv` – Function.

```
inv(M)
```

Matrix inverse. Computes matrix N such that M * N = I, where I is the identity matrix. Computed by solving the left-division N = M \ I.

Example

```julia
julia> M = [2 5; 1 3]
2×2 Array{Int64,2}:
 2  5
 1  3

julia> N = inv(M)
2×2 Array{Float64,2}:
  3.0  -5.0
 -1.0   2.0

julia> M*N == N*M == eye(2)
true
```

source

`Base.LinAlg.pinv` – Function.

```
pinv(M[, tol::Real])
```

Computes the Moore-Penrose pseudoinverse.

For matrices M with floating point elements, it is convenient to compute the pseudoinverse by inverting only singular values above a given threshold, tol.

The optimal choice of tol varies both with the value of M and the intended application of the pseudoinverse. The default value of tol is eps(real(float(one(eltype(M)))))*maximum(size(A)), which is essentially machine epsilon for the real part of a matrix element multiplied by the larger matrix dimension. For inverting dense ill-conditioned matrices in a least-squares sense, tol = sqrt(eps(real(float(one(eltype(M)))))) is recommended.

For more information, see [4], [5], [6], [7].

Example

```julia
julia> M = [1.5 1.3; 1.2 1.9]
2×2 Array{Float64,2}:
 1.5  1.3
 1.2  1.9

julia> N = pinv(M)
2×2 Array{Float64,2}:
  1.47287   -1.00775
```

```
 -0.930233    1.16279

julia> M * N
2×2 Array{Float64,2}:
 1.0          -2.22045e-16
 4.44089e-16   1.0
```

source

`Base.LinAlg.nullspace` – Function.

```
nullspace(M)
```

Basis for nullspace of M.

Example

```
julia> M = [1 0 0; 0 1 0; 0 0 0]
3×3 Array{Int64,2}:
 1  0  0
 0  1  0
 0  0  0

julia> nullspace(M)
3×1 Array{Float64,2}:
 0.0
 0.0
 1.0
```

source

`Base.repmat` – Function.

```
repmat(A, m::Integer, n::Integer=1)
```

Construct a matrix by repeating the given matrix (or vector) m times in dimension 1 and n times in dimension 2.

```
julia> repmat([1, 2, 3], 2)
6-element Array{Int64,1}:
 1
 2
 3
 1
 2
 3

julia> repmat([1, 2, 3], 2, 3)
```

---

[4] Issue 8859, "Fix least squares", https://github.com/JuliaLang/julia/pull/8859

[5] Åke Björck, "Numerical Methods for Least Squares Problems", SIAM Press, Philadelphia, 1996, "Other Titles in Applied Mathematics", Vol. 51. doi:10.1137/1.9781611971484

[6] G. W. Stewart, "Rank Degeneracy", SIAM Journal on Scientific and Statistical Computing, 5(2), 1984, 403-413. doi:10.1137/0905030

[7] Konstantinos Konstantinides and Kung Yao, "Statistical analysis of effective singular values in matrix rank determination", IEEE Transactions on Acoustics, Speech and Signal Processing, 36(5), 1988, 757-763. doi:10.1109/29.1585

```
6×3 Array{Int64,2}:
 1  1  1
 2  2  2
 3  3  3
 1  1  1
 2  2  2
 3  3  3
```

source

Base.repeat – Function.

```
repeat(A::AbstractArray; inner=ntuple(x->1, ndims(A)), outer=ntuple(x->1, ndims(A)))
```

Construct an array by repeating the entries of A. The i-th element of `inner` specifies the number of times that the individual entries of the i-th dimension of A should be repeated. The i-th element of `outer` specifies the number of times that a slice along the i-th dimension of A should be repeated. If `inner` or `outer` are omitted, no repetition is performed.

```
julia> repeat(1:2, inner=2)
4-element Array{Int64,1}:
 1
 1
 2
 2

julia> repeat(1:2, outer=2)
4-element Array{Int64,1}:
 1
 2
 1
 2

julia> repeat([1 2; 3 4], inner=(2, 1), outer=(1, 3))
4×6 Array{Int64,2}:
 1  2  1  2  1  2
 1  2  1  2  1  2
 3  4  3  4  3  4
 3  4  3  4  3  4
```

source

Base.kron – Function.

```
kron(A, B)
```

Kronecker tensor product of two vectors or two matrices.

Example

```
julia> A = [1 2; 3 4]
2×2 Array{Int64,2}:
 1  2
 3  4

julia> B = [im 1; 1 -im]
```

```
2×2 Array{Complex{Int64},2}:
 0+1im  1+0im
 1+0im  0-1im

julia> kron(A, B)
4×4 Array{Complex{Int64},2}:
 0+1im  1+0im  0+2im  2+0im
 1+0im  0-1im  2+0im  0-2im
 0+3im  3+0im  0+4im  4+0im
 3+0im  0-3im  4+0im  0-4im
```

source

Base.SparseArrays.blkdiag – Function.

```
blkdiag(A...)
```

Concatenate matrices block-diagonally. Currently only implemented for sparse matrices.

Example

```
julia> blkdiag(speye(3), 2*speye(2))
5×5 SparseMatrixCSC{Float64,Int64} with 5 stored entries:
  [1, 1]  =  1.0
  [2, 2]  =  1.0
  [3, 3]  =  1.0
  [4, 4]  =  2.0
  [5, 5]  =  2.0
```

source

Base.LinAlg.linreg – Function.

```
linreg(x, y)
```

Perform simple linear regression using Ordinary Least Squares. Returns a and b such that a + b*x is the closest straight line to the given points (x, y), i.e., such that the squared error between y and a + b*x is minimized.

Examples:

```
using PyPlot
x = 1.0:12.0
y = [5.5, 6.3, 7.6, 8.8, 10.9, 11.79, 13.48, 15.02, 17.77, 20.81, 22.0, 22.99]
a, b = linreg(x, y)          # Linear regression
plot(x, y, "o")              # Plot (x, y) points
plot(x, a + b*x)             # Plot line determined by linear regression
```

See also:

\, cov, std, mean.

source

Base.LinAlg.expm – Function.

```
expm(A)
```

Compute the matrix exponential of A, defined by

$$e^A = \sum_{n=0}^{\infty} \frac{A^n}{n!}.$$

For symmetric or Hermitian A, an eigendecomposition (`eigfact`) is used, otherwise the scaling and squaring algorithm (see [8]) is chosen.

Example

```julia
julia> A = eye(2, 2)
2×2 Array{Float64,2}:
 1.0  0.0
 0.0  1.0

julia> expm(A)
2×2 Array{Float64,2}:
 2.71828  0.0
 0.0      2.71828
```

source

`Base.LinAlg.logm` – Function.

```julia
logm(A{T}::StridedMatrix{T})
```

If A has no negative real eigenvalue, compute the principal matrix logarithm of A, i.e. the unique matrix X such that $e^X = A$ and $-\pi < Im(\lambda) < \pi$ for all the eigenvalues λ of X. If A has nonpositive eigenvalues, a nonprincipal matrix function is returned whenever possible.

If A is symmetric or Hermitian, its eigendecomposition (`eigfact`) is used, if A is triangular an improved version of the inverse scaling and squaring method is employed (see [9] and [10]). For general matrices, the complex Schur form (`schur`) is computed and the triangular algorithm is used on the triangular factor.

Example

```julia
julia> A = 2.7182818 * eye(2)
2×2 Array{Float64,2}:
 2.71828  0.0
 0.0      2.71828

julia> logm(A)
2×2 Symmetric{Float64,Array{Float64,2}}:
 1.0  0.0
 0.0  1.0
```

source

---

[8] Nicholas J. Higham, "The squaring and scaling method for the matrix exponential revisited", SIAM Journal on Matrix Analysis and Applications, 26(4), 2005, 1179-1193. doi:10.1137/090768539

[9] Awad H. Al-Mohy and Nicholas J. Higham, "Improved inverse scaling and squaring algorithms for the matrix logarithm", SIAM Journal on Scientific Computing, 34(4), 2012, C153-C169. doi:10.1137/110852553

[10] Awad H. Al-Mohy, Nicholas J. Higham and Samuel D. Relton, "Computing the Fréchet derivative of the matrix logarithm and estimating the condition number", SIAM Journal on Scientific Computing, 35(4), 2013, C394-C410. doi:10.1137/120885991

Base.LinAlg.sqrtm – Function.

```
sqrtm(A)
```

If A has no negative real eigenvalues, compute the principal matrix square root of A, that is the unique matrix X with eigenvalues having positive real part such that $X^2 = A$. Otherwise, a nonprincipal square root is returned.

If A is symmetric or Hermitian, its eigendecomposition (`eigfact`) is used to compute the square root. Otherwise, the square root is determined by means of the Björck-Hammarling method [11], which computes the complex Schur form (`schur`) and then the complex square root of the triangular factor.

Example

```
julia> A = [4 0; 0 4]
2×2 Array{Int64,2}:
 4  0
 0  4

julia> sqrtm(A)
2×2 Array{Float64,2}:
 2.0  0.0
 0.0  2.0
```

source

Base.LinAlg.lyap – Function.

```
lyap(A, C)
```

Computes the solution X to the continuous Lyapunov equation AX + XA' + C = 0, where no eigenvalue of A has a zero real part and no two eigenvalues are negative complex conjugates of each other.

source

Base.LinAlg.sylvester – Function.

```
sylvester(A, B, C)
```

Computes the solution X to the Sylvester equation AX + XB + C = 0, where A, B and C have compatible dimensions and A and -B have no eigenvalues with equal real part.

source

Base.LinAlg.issymmetric – Function.

```
issymmetric(A) -> Bool
```

Test whether a matrix is symmetric.

Example

```
julia> a = [1 2; 2 -1]
2×2 Array{Int64,2}:
 1   2
 2  -1
```

---

[11]Åke Björck and Sven Hammarling, "A Schur method for the square root of a matrix", Linear Algebra and its Applications, 52-53, 1983, 127-140. doi:10.1016/0024-3795(83)80010-X

```
julia> issymmetric(a)
true

julia> b = [1 im; -im 1]
2×2 Array{Complex{Int64},2}:
 1+0im  0+1im
 0-1im  1+0im

julia> issymmetric(b)
false
```

source

Base.LinAlg.isposdef – Function.

```
isposdef(A) -> Bool
```

Test whether a matrix is positive definite.

Example

```
julia> A = [1 2; 2 50]
2×2 Array{Int64,2}:
 1   2
 2  50

julia> isposdef(A)
true
```

source

Base.LinAlg.isposdef! – Function.

```
isposdef!(A) -> Bool
```

Test whether a matrix is positive definite, overwriting A in the process.

Example

```
julia> A = [1. 2.; 2. 50.];

julia> isposdef!(A)
true

julia> A
2×2 Array{Float64,2}:
 1.0  2.0
 2.0  6.78233
```

source

Base.LinAlg.istril – Function.

```
istril(A) -> Bool
```

Test whether a matrix is lower triangular.

Example

```julia
julia> a = [1 2; 2 -1]
2×2 Array{Int64,2}:
 1   2
 2  -1

julia> istril(a)
false

julia> b = [1 0; -im -1]
2×2 Array{Complex{Int64},2}:
 1+0im   0+0im
 0-1im  -1+0im

julia> istril(b)
true
```

source

Base.LinAlg.istriu – Function.

```
istriu(A) -> Bool
```

Test whether a matrix is upper triangular.

Example

```julia
julia> a = [1 2; 2 -1]
2×2 Array{Int64,2}:
 1   2
 2  -1

julia> istriu(a)
false

julia> b = [1 im; 0 -1]
2×2 Array{Complex{Int64},2}:
 1+0im   0+1im
 0+0im  -1+0im

julia> istriu(b)
true
```

source

Base.LinAlg.isdiag – Function.

```
isdiag(A) -> Bool
```

Test whether a matrix is diagonal.

Example

```julia
julia> a = [1 2; 2 -1]
2×2 Array{Int64,2}:
 1   2
 2  -1
```

```
julia> isdiag(a)
false

julia> b = [im 0; 0 -im]
2×2 Array{Complex{Int64},2}:
 0+1im  0+0im
 0+0im  0-1im

julia> isdiag(b)
true
```

source

Base.LinAlg.ishermitian – Function.

```
ishermitian(A) -> Bool
```

Test whether a matrix is Hermitian.

Example

```
julia> a = [1 2; 2 -1]
2×2 Array{Int64,2}:
 1   2
 2  -1

julia> ishermitian(a)
true

julia> b = [1 im; -im 1]
2×2 Array{Complex{Int64},2}:
 1+0im  0+1im
 0-1im  1+0im

julia> ishermitian(b)
true
```

source

Base.LinAlg.RowVector – Type.

```
RowVector(vector)
```

A lazy-view wrapper of an `AbstractVector`, which turns a length-n vector into a 1×n shaped row vector and represents the transpose of a vector (the elements are also transposed recursively). This type is usually constructed (and unwrapped) via the `transpose` function or `.'` operator (or related `ctranspose` or `'` operator).

By convention, a vector can be multiplied by a matrix on its left (`A * v`) whereas a row vector can be multiplied by a matrix on its right (such that `v.' * A = (A.' * v).'`). It differs from a 1×n-sized matrix by the facts that its transpose returns a vector and the inner product `v1.' * v2` returns a scalar, but will otherwise behave similarly.

source

Base.LinAlg.ConjArray – Type.

```
ConjArray(array)
```

A lazy-view wrapper of an `AbstractArray`, taking the elementwise complex conjugate. This type is usually constructed (and unwrapped) via the `conj` function (or related `ctranspose`), but currently this is the default behavior for `RowVector` only. For other arrays, the `ConjArray` constructor can be used directly.

Examples

```julia
julia> [1+im, 1-im]'
1×2 RowVector{Complex{Int64},ConjArray{Complex{Int64},1,Array{Complex{Int64},1}}}:
 1-1im  1+1im

julia> ConjArray([1+im 0; 0 1-im])
2×2 ConjArray{Complex{Int64},2,Array{Complex{Int64},2}}:
 1-1im  0+0im
 0+0im  1+1im
```

source

`Base.transpose` – Function.

```
transpose(A::AbstractMatrix)
```

The transposition operator (`.'`).

Example

```julia
julia> A = [1 2 3; 4 5 6; 7 8 9]
3×3 Array{Int64,2}:
 1  2  3
 4  5  6
 7  8  9

julia> transpose(A)
3×3 Array{Int64,2}:
 1  4  7
 2  5  8
 3  6  9
```

source

```
transpose(v::AbstractVector)
```

The transposition operator (`.'`).

Example

```julia
julia> v = [1,2,3]
3-element Array{Int64,1}:
 1
 2
 3

julia> transpose(v)
1×3 RowVector{Int64,Array{Int64,1}}:
 1  2  3
```

source

`Base.LinAlg.transpose!` – Function.

```
transpose!(dest,src)
```

Transpose array `src` and store the result in the preallocated array `dest`, which should have a size corresponding to `(size(src,2),size(src,1))`. No in-place transposition is supported and unexpected results will happen if `src` and `dest` have overlapping memory regions.

source

`Base.ctranspose` – Function.

```
ctranspose(A)
```

The conjugate transposition operator (`'`).

Example

```julia
julia> A = [3+2im 9+2im; 8+7im  4+6im]
2×2 Array{Complex{Int64},2}:
 3+2im  9+2im
 8+7im  4+6im

julia> ctranspose(A)
2×2 Array{Complex{Int64},2}:
 3-2im  8-7im
 9-2im  4-6im
```

source

`Base.LinAlg.ctranspose!` – Function.

```
ctranspose!(dest,src)
```

Conjugate transpose array `src` and store the result in the preallocated array `dest`, which should have a size corresponding to `(size(src,2),size(src,1))`. No in-place transposition is supported and unexpected results will happen if `src` and `dest` have overlapping memory regions.

source

`Base.LinAlg.eigs` – Method.

```
eigs(A; nev=6, ncv=max(20,2*nev+1), which=:LM, tol=0.0, maxiter=300, sigma=nothing, ritzvec=
     true, v0=zeros((0,))) -> (d,[v,],nconv,niter,nmult,resid)
```

Computes eigenvalues d of A using implicitly restarted Lanczos or Arnoldi iterations for real symmetric or general nonsymmetric matrices respectively.

The following keyword arguments are supported:

- nev: Number of eigenvalues
- ncv: Number of Krylov vectors used in the computation; should satisfy nev+1 `<=` ncv `<=` n for real symmetric problems and nev+2 `<=` ncv `<=` n for other problems, where n is the size of the input matrix A. The default is ncv = max(20,2*nev+1). Note that these restrictions limit the input matrix A to be of dimension at least 2.
- which: type of eigenvalues to compute. See the note below.

which	type of eigenvalues
:LM	eigenvalues of largest magnitude (default)
:SM	eigenvalues of smallest magnitude
:LR	eigenvalues of largest real part
:SR	eigenvalues of smallest real part
:LI	eigenvalues of largest imaginary part (nonsymmetric or complex A only)
:SI	eigenvalues of smallest imaginary part (nonsymmetric or complex A only)
:BE	compute half of the eigenvalues from each end of the spectrum, biased in favor of the high end. (real symmetric A only)

- `tol`: parameter defining the relative tolerance for convergence of Ritz values (eigenvalue estimates). A Ritz value is considered converged when its associated residual is less than or equal to the product of `tol` and $max(^{2/3}, ||)$, where = `eps(real(eltype(A)))`/2 is LAPACK's machine epsilon. The residual associated with and its corresponding Ritz vector v is defined as the norm $||Av - v||$. The specified value of `tol` should be positive; otherwise, it is ignored and is used instead. Default: .

- `maxiter`: Maximum number of iterations (default = 300)

- `sigma`: Specifies the level shift used in inverse iteration. If `nothing` (default), defaults to ordinary (forward) iterations. Otherwise, find eigenvalues close to `sigma` using shift and invert iterations.

- `ritzvec`: Returns the Ritz vectors v (eigenvectors) if `true`

- `v0`: starting vector from which to start the iterations

`eigs` returns the nev requested eigenvalues in d, the corresponding Ritz vectors v (only if `ritzvec=true`), the number of converged eigenvalues nconv, the number of iterations `niter` and the number of matrix vector multiplications `nmult`, as well as the final residual vector `resid`.

Example

```julia
julia> A = spdiagm(1:4);

julia> λ,  = eigs(A, nev = 2);

julia> λ
2-element Array{Float64,1}:
 4.0
 3.0
```

Note

The `sigma` and `which` keywords interact: the description of eigenvalues searched for by `which` do *not* necessarily refer to the eigenvalues of A, but rather the linear operator constructed by the specification of the iteration mode implied by `sigma`.

sigma	iteration mode	which refers to eigenvalues of
nothing	ordinary (forward)	A
real or complex	inverse with level shift sigma	$(A - \sigma I)^{-1}$

Note

Although `tol` has a default value, the best choice depends strongly on the matrix A. We recommend that users _always_ specify a value for `tol` which suits their specific needs.

For details of how the errors in the computed eigenvalues are estimated, see:

- B. N. Parlett, "The Symmetric Eigenvalue Problem", SIAM: Philadelphia, 2/e (1998), Ch. 13.2, "Accessing Accuracy in Lanczos Problems", pp. 290-292 ff.

- R. B. Lehoucq and D. C. Sorensen, "Deflation Techniques for an Implicitly Restarted Arnoldi Iteration", SIAM Journal on Matrix Analysis and Applications (1996), 17(4), 789–821. doi:10.1137/S0895479895281484

source

Base.LinAlg.eigs – Method.

```
eigs(A, B; nev=6, ncv=max(20,2*nev+1), which=:LM, tol=0.0, maxiter=300, sigma=nothing,
     ritzvec=true, v0=zeros((0,))) -> (d,[v,],nconv,niter,nmult,resid)
```

Computes generalized eigenvalues d of A and B using implicitly restarted Lanczos or Arnoldi iterations for real symmetric or general nonsymmetric matrices respectively.

The following keyword arguments are supported:

- nev: Number of eigenvalues

- ncv: Number of Krylov vectors used in the computation; should satisfy nev+1 <= ncv <= n for real symmetric problems and nev+2 <= ncv <= n for other problems, where n is the size of the input matrices A and B. The default is ncv = max(20,2*nev+1). Note that these restrictions limit the input matrix A to be of dimension at least 2.

- which: type of eigenvalues to compute. See the note below.

which	type of eigenvalues
:LM	eigenvalues of largest magnitude (default)
:SM	eigenvalues of smallest magnitude
:LR	eigenvalues of largest real part
:SR	eigenvalues of smallest real part
:LI	eigenvalues of largest imaginary part (nonsymmetric or complex A only)
:SI	eigenvalues of smallest imaginary part (nonsymmetric or complex A only)
:BE	compute half of the eigenvalues from each end of the spectrum, biased in favor of the high end. (real symmetric A only)

- tol: relative tolerance used in the convergence criterion for eigenvalues, similar to tol in the eigs(A) method for the ordinary eigenvalue problem, but effectively for the eigenvalues of $B^{-1}A$ instead of A. See the documentation for the ordinary eigenvalue problem in eigs(A) and the accompanying note about tol.

- maxiter: Maximum number of iterations (default = 300)

- sigma: Specifies the level shift used in inverse iteration. If nothing (default), defaults to ordinary (forward) iterations. Otherwise, find eigenvalues close to sigma using shift and invert iterations.

- ritzvec: Returns the Ritz vectors v (eigenvectors) if true

- v0: starting vector from which to start the iterations

eigs returns the nev requested eigenvalues in d, the corresponding Ritz vectors v (only if ritzvec=true), the number of converged eigenvalues nconv, the number of iterations niter and the number of matrix vector multiplications nmult, as well as the final residual vector resid.

Example

```
julia> A = speye(4, 4); B = spdiagm(1:4);

julia> λ,  = eigs(A, B, nev = 2);

julia> λ
2-element Array{Float64,1}:
 1.0
 0.5
```

Note

The `sigma` and `which` keywords interact: the description of eigenvalues searched for by `which` do *not* necessarily refer to the eigenvalue problem $Av = Bv\lambda$, but rather the linear operator constructed by the specification of the iteration mode implied by `sigma`.

sigma	iteration mode	which refers to the problem
nothing	ordinary (forward)	$Av = Bv\lambda$
real or complex	inverse with level shift `sigma`	$(A - \sigma B)^{-1}B = v\nu$

source

`Base.LinAlg.svds` – Function.

```
svds(A; nsv=6, ritzvec=true, tol=0.0, maxiter=1000, ncv=2*nsv, u0=zeros((0,)), v0=zeros((0,))
     ) -> (SVD([left_sv,] s, [right_sv,]), nconv, niter, nmult, resid)
```

Computes the largest singular values s of A using implicitly restarted Lanczos iterations derived from `eigs`.

Inputs

- A: Linear operator whose singular values are desired. A may be represented as a subtype of `AbstractArray`, e.g., a sparse matrix, or any other type supporting the four methods `size(A)`, `eltype(A)`, `A * vector`, and `A' * vector`.
- nsv: Number of singular values. Default: 6.
- ritzvec: If `true`, return the left and right singular vectors `left_sv` and `right_sv`. If `false`, omit the singular vectors. Default: `true`.
- tol: tolerance, see `eigs`.
- maxiter: Maximum number of iterations, see `eigs`. Default: 1000.
- ncv: Maximum size of the Krylov subspace, see `eigs` (there called nev). Default: `2*nsv`.
- u0: Initial guess for the first left Krylov vector. It may have length m (the first dimension of A), or 0.
- v0: Initial guess for the first right Krylov vector. It may have length n (the second dimension of A), or 0.

Outputs

- svd: An SVD object containing the left singular vectors, the requested values, and the right singular vectors. If `ritzvec = false`, the left and right singular vectors will be empty.
- nconv: Number of converged singular values.
- niter: Number of iterations.
- nmult: Number of matrix–vector products used.
- resid: Final residual vector.

Example

```
julia> A = spdiagm(1:4);

julia> s = svds(A, nsv = 2)[1];

julia> s[:S]
2-element Array{Float64,1}:
 4.0
 3.0
```

Implementation

`svds(A)` is formally equivalent to calling `eigs` to perform implicitly restarted Lanczos tridiagonaliza-tion on the Hermitian matrix $\begin{pmatrix} 0 & A' \\ A & 0 \end{pmatrix}$, whose eigenvalues are plus and minus the singular values of A.

source

`Base.LinAlg.peakflops` – Function.

```
peakflops(n::Integer=2000; parallel::Bool=false)
```

`peakflops` computes the peak flop rate of the computer by using double precision `gemm!`. By default, if no arguments are specified, it multiplies a matrix of size n x n, where n = 2000. If the underlying BLAS is using multiple threads, higher flop rates are realized. The number of BLAS threads can be set with `BLAS.set_num_threads(n)`.

If the keyword argument `parallel` is set to `true`, `peakflops` is run in parallel on all the worker processors. The flop rate of the entire parallel computer is returned. When running in parallel, only 1 BLAS thread is used. The argument n still refers to the size of the problem that is solved on each processor.

source

52.2 Low-level matrix operations

Matrix operations involving transpositions operations like A' \ B are converted by the Julia parser into calls to specially named functions like `Ac_ldiv_B`. If you want to overload these operations for your own types, then it is useful to know the names of these functions.

Also, in many cases there are in-place versions of matrix operations that allow you to supply a pre-allocated output vector or matrix. This is useful when optimizing critical code in order to avoid the overhead of repeated allocations. These in-place operations are suffixed with ! below (e.g. `A_mul_B!`) according to the usual Julia convention.

`Base.LinAlg.A_ldiv_B!` – Function.

```
A_ldiv_B!([Y,] A, B) -> Y
```

Compute A \ B in-place and store the result in Y, returning the result. If only two arguments are passed, then `A_ldiv_B!(A, B)` overwrites B with the result.

The argument A should *not* be a matrix. Rather, instead of matrices it should be a factorization object (e.g. produced by `factorize` or `cholfact`). The reason for this is that factorization itself is both expensive and typically allocates memory (although it can also be done in-place via, e.g., `lufact!`), and performance-critical situations requiring A_ldiv_B! usually also require fine-grained control over the factorization of A.

source

Base.A_ldiv_Bc – Function.

```
A_ldiv_Bc(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.A_ldiv_Bt – Function.

```
A_ldiv_Bt(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.LinAlg.A_mul_B! – Function.

```
A_mul_B!(Y, A, B) -> Y
```

Calculates the matrix-matrix or matrix-vector product AB and stores the result in Y, overwriting the existing value of Y. Note that Y must not be aliased with either A or B.

Example

```julia
julia> A=[1.0 2.0; 3.0 4.0]; B=[1.0 1.0; 1.0 1.0]; Y = similar(B); A_mul_B!(Y, A, B);

julia> Y
2×2 Array{Float64,2}:
 3.0  3.0
 7.0  7.0
```

source

Base.A_mul_Bc – Function.

```
A_mul_Bc(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.A_mul_Bt – Function.

```
A_mul_Bt(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.A_rdiv_Bc – Function.

```
A_rdiv_Bc(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

Base.A_rdiv_Bt – Function.

```
A_rdiv_Bt(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

Base.Ac_ldiv_B – Function.

```
Ac_ldiv_B(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.LinAlg.Ac_ldiv_B! – Function.

```
Ac_ldiv_B!([Y,] A, B) -> Y
```

Similar to A_ldiv_B!, but return $A \setminus B$, computing the result in-place in Y (or overwriting B if Y is not supplied).

source

Base.Ac_ldiv_Bc – Function.

```
Ac_ldiv_Bc(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.Ac_mul_B – Function.

```
Ac_mul_B(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.Ac_mul_Bc – Function.

```
Ac_mul_Bc(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.Ac_rdiv_B – Function.

```
Ac_rdiv_B(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

Base.Ac_rdiv_Bc – Function.

```
Ac_rdiv_Bc(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

Base.`At_ldiv_B` – Function.

```
At_ldiv_B(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.LinAlg.`At_ldiv_B!` – Function.

```
At_ldiv_B!([Y,] A, B) -> Y
```

Similar to `A_ldiv_B!`, but return $A \setminus B$, computing the result in-place in Y (or overwriting B if Y is not supplied).

source

Base.`At_ldiv_Bt` – Function.

```
At_ldiv_Bt(A, B)
```

For matrices or vectors A and B, calculates $A \setminus B$.

source

Base.`At_mul_B` – Function.

```
At_mul_B(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.`At_mul_Bt` – Function.

```
At_mul_Bt(A, B)
```

For matrices or vectors A and B, calculates AB.

source

Base.`At_rdiv_B` – Function.

```
At_rdiv_B(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

Base.`At_rdiv_Bt` – Function.

```
At_rdiv_Bt(A, B)
```

For matrices or vectors A and B, calculates A/B.

source

52.3   BLAS Functions

In Julia (as in much of scientific computation), dense linear-algebra operations are based on the LAPACK library, which in turn is built on top of basic linear-algebra building-blocks known as the BLAS. There are highly optimized implementations of BLAS available for every computer architecture, and sometimes in high-performance linear algebra routines it is useful to call the BLAS functions directly.

`Base.LinAlg.BLAS` provides wrappers for some of the BLAS functions. Those BLAS functions that overwrite one of the input arrays have names ending in `'!'`. Usually, a BLAS function has four methods defined, for `Float64`, `Float32`, `Complex128`, and `Complex64` arrays.

BLAS Character Arguments

Many BLAS functions accept arguments that determine whether to transpose an argument (`trans`), which triangle of a matrix to reference (`uplo` or `ul`), whether the diagonal of a triangular matrix can be assumed to be all ones (`dA`) or which side of a matrix multiplication the input argument belongs on (`side`). The possiblities are:

Multplication Order

side	Meaning
`'L'`	The argument goes on the *left* side of a matrix-matrix operation.
`'R'`	The argument goes on the *right* side of a matrix-matrix operation.

Triangle Referencing

uplo/ul	Meaning
`'U'`	Only the *upper* triangle of the matrix will be used.
`'L'`	Only the *lower* triangle of the matrix will be used.

Transposition Operation

trans/tX	Meaning
`'N'`	The input matrix X is not transposed or conjugated.
`'T'`	The input matrix X will be transposed.
`'C'`	The input matrix X will be conjugated and transposed.

Unit Diagonal

diag/dX	Meaning
`'N'`	The diagonal values of the matrix X will be read.
`'U'`	The diagonal of the matrix X is assumed to be all ones.

Base.LinAlg.BLAS.dotu – Function.

```
dotu(n, X, incx, Y, incy)
```

Dot function for two complex vectors consisting of n elements of array X with stride `incx` and n elements of array Y with stride `incy`.

Example:

```
julia> Base.BLAS.dotu(10, im*ones(10), 1, complex.(ones(20), ones(20)), 2)
-10.0 + 10.0im
```

source

Base.LinAlg.BLAS.dotc – Function.

```
dotc(n, X, incx, U, incy)
```

Dot function for two complex vectors, consisting of n elements of array X with stride `incx` and n elements of array U with stride `incy`, conjugating the first vector.

Example:

```
julia> Base.BLAS.dotc(10, im*ones(10), 1, complex.(ones(20), ones(20)), 2)
10.0 - 10.0im
```

source

Base.LinAlg.BLAS.blascopy! – Function.

```
blascopy!(n, X, incx, Y, incy)
```

Copy n elements of array X with stride `incx` to array Y with stride `incy`. Returns Y.

source

Base.LinAlg.BLAS.nrm2 – Function.

```
nrm2(n, X, incx)
```

2-norm of a vector consisting of n elements of array X with stride `incx`.

Example:

```
julia> Base.BLAS.nrm2(4, ones(8), 2)
2.0

julia> Base.BLAS.nrm2(1, ones(8), 2)
1.0
```

source

Base.LinAlg.BLAS.asum – Function.

```
asum(n, X, incx)
```

Sum of the absolute values of the first n elements of array X with stride `incx`.

Example:

```
julia> Base.BLAS.asum(5, im*ones(10), 2)
5.0

julia> Base.BLAS.asum(2, im*ones(10), 5)
2.0
```

source

Base.LinAlg.axpy! – Function.

```
axpy!(a, X, Y)
```

Overwrite Y with a*X + Y, where a is a scalar. Returns Y.

Example:

```julia
julia> x = [1; 2; 3];

julia> y = [4; 5; 6];

julia> Base.BLAS.axpy!(2, x, y)
3-element Array{Int64,1}:
   6
   9
  12
```

source

Base.LinAlg.BLAS.scal! – Function.

```
scal!(n, a, X, incx)
```

Overwrite X with a*X for the first n elements of array X with stride incx. Returns X.

source

Base.LinAlg.BLAS.scal – Function.

```
scal(n, a, X, incx)
```

Returns X scaled by a for the first n elements of array X with stride incx.

source

Base.LinAlg.BLAS.ger! – Function.

```
ger!(alpha, x, y, A)
```

Rank-1 update of the matrix A with vectors x and y as alpha*x*y' + A.

source

Base.LinAlg.BLAS.syr! – Function.

```
syr!(uplo, alpha, x, A)
```

Rank-1 update of the symmetric matrix A with vector x as alpha*x*x.' + A. uplo controls which triangle of A is updated. Returns A.

source

Base.LinAlg.BLAS.syrk! – Function.

```
syrk!(uplo, trans, alpha, A, beta, C)
```

Rank-k update of the symmetric matrix C as alpha*A*A.' + beta*C or alpha*A.'*A + beta*C according to trans. Only the uplo triangle of C is used. Returns C.

source

Base.LinAlg.BLAS.syrk – Function.

```
syrk(uplo, trans, alpha, A)
```

Returns either the upper triangle or the lower triangle of A, according to uplo, of alpha*A*A.' or alpha*A.'*A, according to trans.

source

Base.LinAlg.BLAS.her! – Function.

```
her!(uplo, alpha, x, A)
```

Methods for complex arrays only. Rank-1 update of the Hermitian matrix A with vector x as alpha*x*x' + A. uplo controls which triangle of A is updated. Returns A.

source

Base.LinAlg.BLAS.herk! – Function.

```
herk!(uplo, trans, alpha, A, beta, C)
```

Methods for complex arrays only. Rank-k update of the Hermitian matrix C as alpha*A*A' + beta*C or alpha*A'*A + beta*C according to trans. Only the uplo triangle of C is updated. Returns C.

source

Base.LinAlg.BLAS.herk – Function.

```
herk(uplo, trans, alpha, A)
```

Methods for complex arrays only. Returns the uplo triangle of alpha*A*A' or alpha*A'*A, according to trans.

source

Base.LinAlg.BLAS.gbmv! – Function.

```
gbmv!(trans, m, kl, ku, alpha, A, x, beta, y)
```

Update vector y as alpha*A*x + beta*y or alpha*A'*x + beta*y according to trans. The matrix A is a general band matrix of dimension m by size(A,2) with kl sub-diagonals and ku super-diagonals. alpha and beta are scalars. Returns the updated y.

source

Base.LinAlg.BLAS.gbmv – Function.

```
gbmv(trans, m, kl, ku, alpha, A, x)
```

Returns alpha*A*x or alpha*A'*x according to trans. The matrix A is a general band matrix of dimension m by size(A,2) with kl sub-diagonals and ku super-diagonals, and alpha is a scalar.

source

Base.LinAlg.BLAS.sbmv! – Function.

```
sbmv!(uplo, k, alpha, A, x, beta, y)
```

Update vector y as `alpha*A*x + beta*y` where A is a a symmetric band matrix of order `size(A,2)` with k super-diagonals stored in the argument A. The storage layout for A is described the reference BLAS module, level-2 BLAS at http://www.netlib.org/lapack/explore-html/. Only the `uplo` triangle of A is used.

Returns the updated y.

source

Base.LinAlg.BLAS.sbmv – Method.

```
sbmv(uplo, k, alpha, A, x)
```

Returns `alpha*A*x` where A is a symmetric band matrix of order `size(A,2)` with k super-diagonals stored in the argument A. Only the `uplo` triangle of A is used.

source

Base.LinAlg.BLAS.sbmv – Method.

```
sbmv(uplo, k, A, x)
```

Returns `A*x` where A is a symmetric band matrix of order `size(A,2)` with k super-diagonals stored in the argument A. Only the `uplo` triangle of A is used.

source

Base.LinAlg.BLAS.gemm! – Function.

```
gemm!(tA, tB, alpha, A, B, beta, C)
```

Update C as `alpha*A*B + beta*C` or the other three variants according to `tA` and tB. Returns the updated C.

source

Base.LinAlg.BLAS.gemm – Method.

```
gemm(tA, tB, alpha, A, B)
```

Returns `alpha*A*B` or the other three variants according to `tA` and tB.

source

Base.LinAlg.BLAS.gemm – Method.

```
gemm(tA, tB, A, B)
```

Returns `A*B` or the other three variants according to `tA` and tB.

source

Base.LinAlg.BLAS.gemv! – Function.

```
gemv!(tA, alpha, A, x, beta, y)
```

Update the vector y as `alpha*A*x + beta*y` or `alpha*A'x + beta*y` according to `tA`. alpha and beta are scalars. Returns the updated y.

source

Base.LinAlg.BLAS.gemv – Method.

```
gemv(tA, alpha, A, x)
```

Returns `alpha*A*x` or `alpha*A'x` according to `tA`. `alpha` is a scalar.

> source

`Base.LinAlg.BLAS.gemv` – Method.

```
gemv(tA, A, x)
```

Returns `A*x` or `A'x` according to `tA`.

> source

`Base.LinAlg.BLAS.symm!` – Function.

```
symm!(side, ul, alpha, A, B, beta, C)
```

Update `C` as `alpha*A*B + beta*C` or `alpha*B*A + beta*C` according to `side`. A is assumed to be symmetric. Only the `ul` triangle of A is used. Returns the updated `C`.

> source

`Base.LinAlg.BLAS.symm` – Method.

```
symm(side, ul, alpha, A, B)
```

Returns `alpha*A*B` or `alpha*B*A` according to `side`. A is assumed to be symmetric. Only the `ul` triangle of A is used.

> source

`Base.LinAlg.BLAS.symm` – Method.

```
symm(side, ul, A, B)
```

Returns `A*B` or `B*A` according to `side`. A is assumed to be symmetric. Only the `ul` triangle of A is used.

> source

`Base.LinAlg.BLAS.symv!` – Function.

```
symv!(ul, alpha, A, x, beta, y)
```

Update the vector `y` as `alpha*A*x + beta*y`. A is assumed to be symmetric. Only the `ul` triangle of A is used. `alpha` and `beta` are scalars. Returns the updated `y`.

> source

`Base.LinAlg.BLAS.symv` – Method.

```
symv(ul, alpha, A, x)
```

Returns `alpha*A*x`. A is assumed to be symmetric. Only the `ul` triangle of A is used. `alpha` is a scalar.

> source

`Base.LinAlg.BLAS.symv` – Method.

```
symv(ul, A, x)
```

Returns A*x. A is assumed to be symmetric. Only the ul triangle of A is used.

Base.LinAlg.BLAS.trmm! – Function.

```
trmm!(side, ul, tA, dA, alpha, A, B)
```

Update B as alpha*A*B or one of the other three variants determined by side and tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones. Returns the updated B.

Base.LinAlg.BLAS.trmm – Function.

```
trmm(side, ul, tA, dA, alpha, A, B)
```

Returns alpha*A*B or one of the other three variants determined by side and tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones.

Base.LinAlg.BLAS.trsm! – Function.

```
trsm!(side, ul, tA, dA, alpha, A, B)
```

Overwrite B with the solution to A*X = alpha*B or one of the other three variants determined by side and tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones. Returns the updated B.

Base.LinAlg.BLAS.trsm – Function.

```
trsm(side, ul, tA, dA, alpha, A, B)
```

Returns the solution to A*X = alpha*B or one of the other three variants determined by determined by side and tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones.

Base.LinAlg.BLAS.trmv! – Function.

```
trmv!(ul, tA, dA, A, b)
```

Returns op(A)*b, where op is determined by tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones. The multiplication occurs in-place on b.

Base.LinAlg.BLAS.trmv – Function.

```
trmv(ul, tA, dA, A, b)
```

Returns op(A)*b, where op is determined by tA. Only the ul triangle of A is used. dA determines if the diagonal values are read or are assumed to be all ones.

Base.LinAlg.BLAS.trsv! – Function.

```
trsv!(ul, tA, dA, A, b)
```

Overwrite b with the solution to A*x = b or one of the other two variants determined by tA and ul. dA determines if the diagonal values are read or are assumed to be all ones. Returns the updated b.

source

Base.LinAlg.BLAS.trsv – Function.

```
trsv(ul, tA, dA, A, b)
```

Returns the solution to A*x = b or one of the other two variants determined by tA and ul. dA determines if the diagonal values are read or are assumed to be all ones.

source

Base.LinAlg.BLAS.set_num_threads – Function.

```
set_num_threads(n)
```

Set the number of threads the BLAS library should use.

source

Base.LinAlg.I – Constant.

```
I
```

An object of type UniformScaling, representing an identity matrix of any size.

Example

```
julia> ones(5, 6) * I == ones(5, 6)
true

julia> [1 2im 3; 1im 2 3] * I
2×3 Array{Complex{Int64},2}:
 1+0im  0+2im  3+0im
 0+1im  2+0im  3+0im
```

source

52.4   LAPACK Functions

Base.LinAlg.LAPACK provides wrappers for some of the LAPACK functions for linear algebra. Those functions that overwrite one of the input arrays have names ending in '!'.

Usually a function has 4 methods defined, one each for Float64, Float32, Complex128 and Complex64 arrays.

Note that the LAPACK API provided by Julia can and will change in the future. Since this API is not user-facing, there is no commitment to support/deprecate this specific set of functions in future releases.

Base.LinAlg.LAPACK.gbtrf! – Function.

```
gbtrf!(kl, ku, m, AB) -> (AB, ipiv)
```

Compute the LU factorization of a banded matrix AB. `kl` is the first subdiagonal containing a nonzero band, `ku` is the last superdiagonal containing one, and `m` is the first dimension of the matrix AB. Returns the LU factorization in-place and `ipiv`, the vector of pivots used.

source

Base.LinAlg.LAPACK.gbtrs! – Function.

```
gbtrs!(trans, kl, ku, m, AB, ipiv, B)
```

Solve the equation `AB * X = B`. `trans` determines the orientation of AB. It may be N (no transpose), T (transpose), or C (conjugate transpose). `kl` is the first subdiagonal containing a nonzero band, `ku` is the last superdiagonal containing one, and `m` is the first dimension of the matrix AB. `ipiv` is the vector of pivots returned from `gbtrf!`. Returns the vector or matrix X, overwriting B in-place.

source

Base.LinAlg.LAPACK.gebal! – Function.

```
gebal!(job, A) -> (ilo, ihi, scale)
```

Balance the matrix A before computing its eigensystem or Schur factorization. `job` can be one of N (A will not be permuted or scaled), P (A will only be permuted), S (A will only be scaled), or B (A will be both permuted and scaled). Modifies A in-place and returns `ilo`, `ihi`, and `scale`. If permuting was turned on, $A[i,j] = 0$ if $j > i$ and $1 < j < ilo$ or $j > ihi$. `scale` contains information about the scaling/permutations performed.

source

Base.LinAlg.LAPACK.gebak! – Function.

```
gebak!(job, side, ilo, ihi, scale, V)
```

Transform the eigenvectors V of a matrix balanced using `gebal!` to the unscaled/unpermuted eigenvectors of the original matrix. Modifies V in-place. `side` can be L (left eigenvectors are transformed) or R (right eigenvectors are transformed).

source

Base.LinAlg.LAPACK.gebrd! – Function.

```
gebrd!(A) -> (A, d, e, tauq, taup)
```

Reduce A in-place to bidiagonal form `A = QBP'`. Returns A, containing the bidiagonal matrix B; d, containing the diagonal elements of B; e, containing the off-diagonal elements of B; `tauq`, containing the elementary reflectors representing Q; and `taup`, containing the elementary reflectors representing P.

source

Base.LinAlg.LAPACK.gelqf! – Function.

```
gelqf!(A, tau)
```

Compute the LQ factorization of A, `A = LQ`. `tau` contains scalars which parameterize the elementary reflectors of the factorization. `tau` must have length greater than or equal to the smallest dimension of A.

Returns A and `tau` modified in-place.

source

```
gelqf!(A) -> (A, tau)
```

Compute the LQ factorization of A, A = LQ.

Returns A, modified in-place, and `tau`, which contains scalars which parameterize the elementary reflectors of the factorization.

source

`Base.LinAlg.LAPACK.geqlf!` – Function.

```
geqlf!(A, tau)
```

Compute the QL factorization of A, A = QL. `tau` contains scalars which parameterize the elementary reflectors of the factorization. `tau` must have length greater than or equal to the smallest dimension of A.

Returns A and `tau` modified in-place.

source

```
geqlf!(A) -> (A, tau)
```

Compute the QL factorization of A, A = QL.

Returns A, modified in-place, and `tau`, which contains scalars which parameterize the elementary reflectors of the factorization.

source

`Base.LinAlg.LAPACK.geqrf!` – Function.

```
geqrf!(A, tau)
```

Compute the QR factorization of A, A = QR. `tau` contains scalars which parameterize the elementary reflectors of the factorization. `tau` must have length greater than or equal to the smallest dimension of A.

Returns A and `tau` modified in-place.

source

```
geqrf!(A) -> (A, tau)
```

Compute the QR factorization of A, A = QR.

Returns A, modified in-place, and `tau`, which contains scalars which parameterize the elementary reflectors of the factorization.

source

`Base.LinAlg.LAPACK.geqp3!` – Function.

```
geqp3!(A, jpvt, tau)
```

Compute the pivoted QR factorization of A, AP = QR using BLAS level 3. P is a pivoting matrix, represented by `jpvt`. `tau` stores the elementary reflectors. `jpvt` must have length length greater than or equal to n if A is an (m x n) matrix. `tau` must have length greater than or equal to the smallest dimension of A.

A, `jpvt`, and `tau` are modified in-place.

source

```
geqp3!(A, jpvt) -> (A, jpvt, tau)
```

Compute the pivoted QR factorization of A, AP = QR using BLAS level 3. P is a pivoting matrix, represented by jpvt. jpvt must have length greater than or equal to n if A is an (m x n) matrix.

Returns A and jpvt, modified in-place, and tau, which stores the elementary reflectors.

```
geqp3!(A) -> (A, jpvt, tau)
```

Compute the pivoted QR factorization of A, AP = QR using BLAS level 3.

Returns A, modified in-place, jpvt, which represents the pivoting matrix P, and tau, which stores the elementary reflectors.

Base.LinAlg.LAPACK.gerqf! – Function.

```
gerqf!(A, tau)
```

Compute the RQ factorization of A, A = RQ. tau contains scalars which parameterize the elementary reflectors of the factorization. tau must have length greater than or equal to the smallest dimension of A.

Returns A and tau modified in-place.

```
gerqf!(A) -> (A, tau)
```

Compute the RQ factorization of A, A = RQ.

Returns A, modified in-place, and tau, which contains scalars which parameterize the elementary reflectors of the factorization.

Base.LinAlg.LAPACK.geqrt! – Function.

```
geqrt!(A, T)
```

Compute the blocked QR factorization of A, A = QR. T contains upper triangular block reflectors which parameterize the elementary reflectors of the factorization. The first dimension of T sets the block size and it must be between 1 and n. The second dimension of T must equal the smallest dimension of A.

Returns A and T modified in-place.

```
geqrt!(A, nb) -> (A, T)
```

Compute the blocked QR factorization of A, A = QR. nb sets the block size and it must be between 1 and n, the second dimension of A.

Returns A, modified in-place, and T, which contains upper triangular block reflectors which parameterize the elementary reflectors of the factorization.

Base.LinAlg.LAPACK.geqrt3! – Function.

```
geqrt3!(A, T)
```

Recursively computes the blocked QR factorization of A, A = QR. T contains upper triangular block reflectors which parameterize the elementary reflectors of the factorization. The first dimension of T sets the block size and it must be between 1 and n. The second dimension of T must equal the smallest dimension of A.

Returns A and T modified in-place.

source

```
geqrt3!(A) -> (A, T)
```

Recursively computes the blocked QR factorization of A, A = QR.

Returns A, modified in-place, and T, which contains upper triangular block reflectors which parameterize the elementary reflectors of the factorization.

source

Base.LinAlg.LAPACK.getrf! – Function.

```
getrf!(A) -> (A, ipiv, info)
```

Compute the pivoted LU factorization of A, A = LU.

Returns A, modified in-place, ipiv, the pivoting information, and an info code which indicates success (info = 0), a singular value in U (info = i, in which case U[i,i] is singular), or an error code (info < 0).

source

Base.LinAlg.LAPACK.tzrzf! – Function.

```
tzrzf!(A) -> (A, tau)
```

Transforms the upper trapezoidal matrix A to upper triangular form in-place. Returns A and tau, the scalar parameters for the elementary reflectors of the transformation.

source

Base.LinAlg.LAPACK.ormrz! – Function.

```
ormrz!(side, trans, A, tau, C)
```

Multiplies the matrix C by Q from the transformation supplied by tzrzf!. Depending on side or trans the multiplication can be left-sided (side = L, Q*C) or right-sided (side = R, C*Q) and Q can be unmodified (trans = N), transposed (trans = T), or conjugate transposed (trans = C). Returns matrix C which is modified in-place with the result of the multiplication.

source

Base.LinAlg.LAPACK.gels! – Function.

```
gels!(trans, A, B) -> (F, B, ssr)
```

Solves the linear equation A * X = B, A.' * X =B, or A' * X = B using a QR or LQ factorization. Modifies the matrix/vector B in place with the solution. A is overwritten with its QR or LQ factorization. trans may be one of N (no modification), T (transpose), or C (conjugate transpose). gels! searches for the minimum norm/least squares solution. A may be under or over determined. The solution is returned in B.

source

Base.LinAlg.LAPACK.gesv! – Function.

```
gesv!(A, B) -> (B, A, ipiv)
```

Solves the linear equation A * X = B where A is a square matrix using the LU factorization of A. A is overwritten with its LU factorization and B is overwritten with the solution X. `ipiv` contains the pivoting information for the LU factorization of A.

source

Base.LinAlg.LAPACK.getrs! – Function.

```
getrs!(trans, A, ipiv, B)
```

Solves the linear equation A * X = B, A.' * X =B, or A' * X = B for square A. Modifies the matrix/vector B in place with the solution. A is the LU factorization from `getrf!`, with `ipiv` the pivoting information. `trans` may be one of N (no modification), T (transpose), or C (conjugate transpose).

source

Base.LinAlg.LAPACK.getri! – Function.

```
getri!(A, ipiv)
```

Computes the inverse of A, using its LU factorization found by `getrf!`. `ipiv` is the pivot information output and A contains the LU factorization of `getrf!`. A is overwritten with its inverse.

source

Base.LinAlg.LAPACK.gesvx! – Function.

```
gesvx!(fact, trans, A, AF, ipiv, equed, R, C, B) -> (X, equed, R, C, B, rcond, ferr, berr,
    work)
```

Solves the linear equation A * X = B (trans = N), A.' * X =B (trans = T), or A' * X = B (trans = C) using the LU factorization of A. `fact` may be E, in which case A will be equilibrated and copied to AF; F, in which case AF and `ipiv` from a previous LU factorization are inputs; or N, in which case A will be copied to AF and then factored. If `fact` = F, `equed` may be N, meaning A has not been equilibrated; R, meaning A was multiplied by `diagm(R)` from the left; C, meaning A was multiplied by `diagm(C)` from the right; or B, meaning A was multiplied by `diagm(R)` from the left and `diagm(C)` from the right. If `fact` = F and `equed` = R or B the elements of R must all be positive. If `fact` = F and `equed` = C or B the elements of C must all be positive.

Returns the solution X; equed, which is an output if `fact` is not N, and describes the equilibration that was performed; R, the row equilibration diagonal; C, the column equilibration diagonal; B, which may be overwritten with its equilibrated form `diagm(R)*B` (if `trans` = N and `equed` = R, B) or `diagm(C)*B` (if `trans` = T, C and `equed` = C, B); rcond, the reciprocal condition number of A after equilbrating; `ferr`, the forward error bound for each solution vector in X; `berr`, the forward error bound for each solution vector in X; and `work`, the reciprocal pivot growth factor.

source

```
gesvx!(A, B)
```

The no-equilibration, no-transpose simplification of `gesvx!`.

source

Base.LinAlg.LAPACK.gelsd! – Function.

```
gelsd!(A, B, rcond) -> (B, rnk)
```

Computes the least norm solution of A * X = B by finding the SVD factorization of A, then dividing-and-conquering the problem. B is overwritten with the solution X. Singular values below `rcond` will be treated as zero. Returns the solution in B and the effective rank of A in `rnk`.

source

Base.LinAlg.LAPACK.gelsy! – Function.

```
gelsy!(A, B, rcond) -> (B, rnk)
```

Computes the least norm solution of A * X = B by finding the full QR factorization of A, then dividing-and-conquering the problem. B is overwritten with the solution X. Singular values below `rcond` will be treated as zero. Returns the solution in B and the effective rank of A in `rnk`.

source

Base.LinAlg.LAPACK.gglse! – Function.

```
gglse!(A, c, B, d) -> (X,res)
```

Solves the equation A * x = c where x is subject to the equality constraint B * x = d. Uses the formula ||c – A*x||^2 = 0 to solve. Returns X and the residual sum-of-squares.

source

Base.LinAlg.LAPACK.geev! – Function.

```
geev!(jobvl, jobvr, A) -> (W, VL, VR)
```

Finds the eigensystem of A. If `jobvl = N`, the left eigenvectors of A aren't computed. If `jobvr = N`, the right eigenvectors of A aren't computed. If `jobvl = V` or `jobvr = V`, the corresponding eigenvectors are computed. Returns the eigenvalues in W, the right eigenvectors in VR, and the left eigenvectors in VL.

source

Base.LinAlg.LAPACK.gesdd! – Function.

```
gesdd!(job, A) -> (U, S, VT)
```

Finds the singular value decomposition of A, A = U * S * V', using a divide and conquer approach. If `job = A`, all the columns of U and the rows of V' are computed. If `job = N`, no columns of U or rows of V' are computed. If `job = O`, A is overwritten with the columns of (thin) U and the rows of (thin) V'. If `job = S`, the columns of (thin) U and the rows of (thin) V' are computed and returned separately.

source

Base.LinAlg.LAPACK.gesvd! – Function.

```
gesvd!(jobu, jobvt, A) -> (U, S, VT)
```

Finds the singular value decomposition of A, A = U * S * V'. If `jobu = A`, all the columns of U are computed. If `jobvt = A` all the rows of V' are computed. If `jobu = N`, no columns of U are computed. If `jobvt = N` no rows of V' are computed. If `jobu = O`, A is overwritten with the columns of (thin) U. If `jobvt = O`, A is overwritten with the rows of (thin) V'. If `jobu = S`, the columns of (thin) U are computed and returned separately. If `jobvt = S` the rows of (thin) V' are computed and returned separately. `jobu` and `jobvt` can't both be O.

Returns U, S, and Vt, where S are the singular values of A.

source

Base.LinAlg.LAPACK.ggsvd! – Function.

```
ggsvd!(jobu, jobv, jobq, A, B) -> (U, V, Q, alpha, beta, k, l, R)
```

Finds the generalized singular value decomposition of A and B, U'*A*Q = D1*R and V'*B*Q = D2*R. D1 has `alpha` on its diagonal and D2 has `beta` on its diagonal. If `jobu = U`, the orthogonal/unitary matrix U is computed. If `jobv = V` the orthogonal/unitary matrix V is computed. If `jobq = Q`, the orthogonal/unitary matrix Q is computed. If `jobu`, `jobv` or `jobq` is N, that matrix is not computed. This function is only available in LAPACK versions prior to 3.6.0.

source

Base.LinAlg.LAPACK.ggsvd3! – Function.

```
ggsvd3!(jobu, jobv, jobq, A, B) -> (U, V, Q, alpha, beta, k, l, R)
```

Finds the generalized singular value decomposition of A and B, U'*A*Q = D1*R and V'*B*Q = D2*R. D1 has `alpha` on its diagonal and D2 has `beta` on its diagonal. If `jobu = U`, the orthogonal/unitary matrix U is computed. If `jobv = V` the orthogonal/unitary matrix V is computed. If `jobq = Q`, the orthogonal/unitary matrix Q is computed. If `jobu`, `jobv`, or `jobq` is N, that matrix is not computed. This function requires LAPACK 3.6.0.

source

Base.LinAlg.LAPACK.geevx! – Function.

```
geevx!(balanc, jobvl, jobvr, sense, A) -> (A, w, VL, VR, ilo, ihi, scale, abnrm, rconde,
    rcondv)
```

Finds the eigensystem of A with matrix balancing. If `jobvl = N`, the left eigenvectors of A aren't computed. If `jobvr = N`, the right eigenvectors of A aren't computed. If `jobvl = V` or `jobvr = V`, the corresponding eigenvectors are computed. If `balanc = N`, no balancing is performed. If `balanc = P`, A is permuted but not scaled. If `balanc = S`, A is scaled but not permuted. If `balanc = B`, A is permuted and scaled. If `sense = N`, no reciprocal condition numbers are computed. If `sense = E`, reciprocal condition numbers are computed for the eigenvalues only. If `sense = V`, reciprocal condition numbers are computed for the right eigenvectors only. If `sense = B`, reciprocal condition numbers are computed for the right eigenvectors and the eigenvectors. If `sense = E,B`, the right and left eigenvectors must be computed.

source

Base.LinAlg.LAPACK.ggev! – Function.

```
ggev!(jobvl, jobvr, A, B) -> (alpha, beta, vl, vr)
```

Finds the generalized eigendecomposition of A and B. If `jobvl = N`, the left eigenvectors aren't computed. If `jobvr = N`, the right eigenvectors aren't computed. If `jobvl = V` or `jobvr = V`, the corresponding eigenvectors are computed.

source

Base.LinAlg.LAPACK.gtsv! – Function.

```
gtsv!(dl, d, du, B)
```

Solves the equation A * X = B where A is a tridiagonal matrix with `dl` on the subdiagonal, `d` on the diagonal, and `du` on the superdiagonal.

Overwrites B with the solution X and returns it.

source

Base.LinAlg.LAPACK.gttrf! – Function.

```
gttrf!(dl, d, du) -> (dl, d, du, du2, ipiv)
```

Finds the LU factorization of a tridiagonal matrix with dl on the subdiagonal, d on the diagonal, and du on the superdiagonal.

Modifies dl, d, and du in-place and returns them and the second superdiagonal du2 and the pivoting vector ipiv.

source

Base.LinAlg.LAPACK.gttrs! – Function.

```
gttrs!(trans, dl, d, du, du2, ipiv, B)
```

Solves the equation A * X = B (trans = N), A.' * X = B (trans = T), or A' * X = B (trans = C) using the LU factorization computed by gttrf!. B is overwritten with the solution X.

source

Base.LinAlg.LAPACK.orglq! – Function.

```
orglq!(A, tau, k = length(tau))
```

Explicitly finds the matrix Q of a LQ factorization after calling gelqf! on A. Uses the output of gelqf!. A is overwritten by Q.

source

Base.LinAlg.LAPACK.orgqr! – Function.

```
orgqr!(A, tau, k = length(tau))
```

Explicitly finds the matrix Q of a QR factorization after calling geqrf! on A. Uses the output of geqrf!. A is overwritten by Q.

source

Base.LinAlg.LAPACK.orgql! – Function.

```
orgql!(A, tau, k = length(tau))
```

Explicitly finds the matrix Q of a QL factorization after calling geqlf! on A. Uses the output of geqlf!. A is overwritten by Q.

source

Base.LinAlg.LAPACK.orgrq! – Function.

```
orgrq!(A, tau, k = length(tau))
```

Explicitly finds the matrix Q of a RQ factorization after calling gerqf! on A. Uses the output of gerqf!. A is overwritten by Q.

source

Base.LinAlg.LAPACK.ormlq! – Function.

```
ormlq!(side, trans, A, tau, C)
```

Computes Q * C (trans = N), Q.' * C (trans = T), Q' * C (trans = C) for side = L or the equivalent right-sided multiplication for side = R using Q from a LQ factorization of A computed using gelqf!. C is overwritten.

source

Base.LinAlg.LAPACK.ormqr! – Function.

```
ormqr!(side, trans, A, tau, C)
```

Computes Q * C (trans = N), Q.' * C (trans = T), Q' * C (trans = C) for side = L or the equivalent right-sided multiplication for side = R using Q from a QR factorization of A computed using geqrf!. C is overwritten.

source

Base.LinAlg.LAPACK.ormql! – Function.

```
ormql!(side, trans, A, tau, C)
```

Computes Q * C (trans = N), Q.' * C (trans = T), Q' * C (trans = C) for side = L or the equivalent right-sided multiplication for side = R using Q from a QL factorization of A computed using geqlf!. C is overwritten.

source

Base.LinAlg.LAPACK.ormrq! – Function.

```
ormrq!(side, trans, A, tau, C)
```

Computes Q * C (trans = N), Q.' * C (trans = T), Q' * C (trans = C) for side = L or the equivalent right-sided multiplication for side = R using Q from a RQ factorization of A computed using gerqf!. C is overwritten.

source

Base.LinAlg.LAPACK.gemqrt! – Function.

```
gemqrt!(side, trans, V, T, C)
```

Computes Q * C (trans = N), Q.' * C (trans = T), Q' * C (trans = C) for side = L or the equivalent right-sided multiplication for side = R using Q from a QR factorization of A computed using geqrt!. C is overwritten.

source

Base.LinAlg.LAPACK.posv! – Function.

```
posv!(uplo, A, B) -> (A, B)
```

Finds the solution to A * X = B where A is a symmetric or Hermitian positive definite matrix. If uplo = U the upper Cholesky decomposition of A is computed. If uplo = L the lower Cholesky decomposition of A is computed. A is overwritten by its Cholesky decomposition. B is overwritten with the solution X.

source

Base.LinAlg.LAPACK.potrf! – Function.

```
potrf!(uplo, A)
```

Computes the Cholesky (upper if `uplo = U`, lower if `uplo = L`) decomposition of positive-definite matrix A. A is overwritten and returned with an info code.

> source

`Base.LinAlg.LAPACK.potri!` – Function.

```
potri!(uplo, A)
```

Computes the inverse of positive-definite matrix A after calling `potrf!` to find its (upper if `uplo = U`, lower if `uplo = L`) Cholesky decomposition.

A is overwritten by its inverse and returned.

> source

`Base.LinAlg.LAPACK.potrs!` – Function.

```
potrs!(uplo, A, B)
```

Finds the solution to A `* X = B` where A is a symmetric or Hermitian positive definite matrix whose Cholesky decomposition was computed by `potrf!`. If `uplo = U` the upper Cholesky decomposition of A was computed. If `uplo = L` the lower Cholesky decomposition of A was computed. B is overwritten with the solution X.

> source

`Base.LinAlg.LAPACK.pstrf!` – Function.

```
pstrf!(uplo, A, tol) -> (A, piv, rank, info)
```

Computes the (upper if `uplo = U`, lower if `uplo = L`) pivoted Cholesky decomposition of positive-definite matrix A with a user-set tolerance `tol`. A is overwritten by its Cholesky decomposition.

Returns A, the pivots `piv`, the rank of A, and an `info` code. If `info = 0`, the factorization succeeded. If `info = i > 0`, then A is indefinite or rank-deficient.

> source

`Base.LinAlg.LAPACK.ptsv!` – Function.

```
ptsv!(D, E, B)
```

Solves A `* X = B` for positive-definite tridiagonal A. D is the diagonal of A and E is the off-diagonal. B is overwritten with the solution X and returned.

> source

`Base.LinAlg.LAPACK.pttrf!` – Function.

```
pttrf!(D, E)
```

Computes the LDLt factorization of a positive-definite tridiagonal matrix with D as diagonal and E as off-diagonal. D and E are overwritten and returned.

> source

`Base.LinAlg.LAPACK.pttrs!` – Function.

```
pttrs!(D, E, B)
```

Solves A * X = B for positive-definite tridiagonal A with diagonal D and off-diagonal E after computing A's LDLt factorization using `pttrf!`. B is overwritten with the solution X.

source

Base.LinAlg.LAPACK.trtri! – Function.

```
trtri!(uplo, diag, A)
```

Finds the inverse of (upper if `uplo` = U, lower if `uplo` = L) triangular matrix A. If `diag` = N, A has non-unit diagonal elements. If `diag` = U, all diagonal elements of A are one. A is overwritten with its inverse.

source

Base.LinAlg.LAPACK.trtrs! – Function.

```
trtrs!(uplo, trans, diag, A, B)
```

Solves A * X = B (`trans` = N), A.' * X = B (`trans` = T), or A' * X = B (`trans` = C) for (upper if `uplo` = U, lower if `uplo` = L) triangular matrix A. If `diag` = N, A has non-unit diagonal elements. If `diag` = U, all diagonal elements of A are one. B is overwritten with the solution X.

source

Base.LinAlg.LAPACK.trcon! – Function.

```
trcon!(norm, uplo, diag, A)
```

Finds the reciprocal condition number of (upper if `uplo` = U, lower if `uplo` = L) triangular matrix A. If `diag` = N, A has non-unit diagonal elements. If `diag` = U, all diagonal elements of A are one. If `norm` = I, the condition number is found in the infinity norm. If `norm` = O or 1, the condition number is found in the one norm.

source

Base.LinAlg.LAPACK.trevc! – Function.

```
trevc!(side, howmny, select, T, VL = similar(T), VR = similar(T))
```

Finds the eigensystem of an upper triangular matrix T. If `side` = R, the right eigenvectors are computed. If `side` = L, the left eigenvectors are computed. If `side` = B, both sets are computed. If `howmny` = A, all eigenvectors are found. If `howmny` = B, all eigenvectors are found and backtransformed using VL and VR. If `howmny` = S, only the eigenvectors corresponding to the values in `select` are computed.

source

Base.LinAlg.LAPACK.trrfs! – Function.

```
trrfs!(uplo, trans, diag, A, B, X, Ferr, Berr) -> (Ferr, Berr)
```

Estimates the error in the solution to A * X = B (`trans` = N), A.' * X = B (`trans` = T), A' * X = B (`trans` = C) for `side` = L, or the equivalent equations a right-handed `side` = R X * A after computing X using `trtrs!`. If `uplo` = U, A is upper triangular. If `uplo` = L, A is lower triangular. If `diag` = N, A has non-unit diagonal elements. If `diag` = U, all diagonal elements of A are one. `Ferr` and `Berr` are optional inputs. `Ferr` is the forward error and `Berr` is the backward error, each component-wise.

source

Base.LinAlg.LAPACK.stev! – Function.

```
stev!(job, dv, ev) -> (dv, Zmat)
```

Computes the eigensystem for a symmetric tridiagonal matrix with dv as diagonal and ev as off-diagonal. If job = N only the eigenvalues are found and returned in dv. If job = V then the eigenvectors are also found and returned in Zmat.

source

Base.LinAlg.LAPACK.stebz! – Function.

```
stebz!(range, order, vl, vu, il, iu, abstol, dv, ev) -> (dv, iblock, isplit)
```

Computes the eigenvalues for a symmetric tridiagonal matrix with dv as diagonal and ev as off-diagonal. If range = A, all the eigenvalues are found. If range = V, the eigenvalues in the half-open interval (vl, vu] are found. If range = I, the eigenvalues with indices between il and iu are found. If order = B, eigvalues are ordered within a block. If order = E, they are ordered across all the blocks. abstol can be set as a tolerance for convergence.

source

Base.LinAlg.LAPACK.stegr! – Function.

```
stegr!(jobz, range, dv, ev, vl, vu, il, iu) -> (w, Z)
```

Computes the eigenvalues (jobz = N) or eigenvalues and eigenvectors (jobz = V) for a symmetric tridiagonal matrix with dv as diagonal and ev as off-diagonal. If range = A, all the eigenvalues are found. If range = V, the eigenvalues in the half-open interval (vl, vu] are found. If range = I, the eigenvalues with indices between il and iu are found. The eigenvalues are returned in w and the eigenvectors in Z.

source

Base.LinAlg.LAPACK.stein! – Function.

```
stein!(dv, ev_in, w_in, iblock_in, isplit_in)
```

Computes the eigenvectors for a symmetric tridiagonal matrix with dv as diagonal and ev_in as off-diagonal. w_in specifies the input eigenvalues for which to find corresponding eigenvectors. iblock_in specifies the submatrices corresponding to the eigenvalues in w_in. isplit_in specifies the splitting points between the submatrix blocks.

source

Base.LinAlg.LAPACK.syconv! – Function.

```
syconv!(uplo, A, ipiv) -> (A, work)
```

Converts a symmetric matrix A (which has been factorized into a triangular matrix) into two matrices L and D. If uplo = U, A is upper triangular. If uplo = L, it is lower triangular. ipiv is the pivot vector from the triangular factorization. A is overwritten by L and D.

source

Base.LinAlg.LAPACK.sysv! – Function.

```
sysv!(uplo, A, B) -> (B, A, ipiv)
```

Finds the solution to A * X = B for symmetric matrix A. If uplo = U, the upper half of A is stored. If uplo = L, the lower half is stored. B is overwritten by the solution X. A is overwritten by its Bunch-Kaufman factorization. ipiv contains pivoting information about the factorization.

source

`Base.LinAlg.LAPACK.sytrf!` – Function.

```
sytrf!(uplo, A) -> (A, ipiv, info)
```

Computes the Bunch-Kaufman factorization of a symmetric matrix A. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored.

Returns A, overwritten by the factorization, a pivot vector `ipiv`, and the error code `info` which is a non-negative integer. If `info` is positive the matrix is singular and the diagonal part of the factorization is exactly zero at position `info`.

> source

`Base.LinAlg.LAPACK.sytri!` – Function.

```
sytri!(uplo, A, ipiv)
```

Computes the inverse of a symmetric matrix A using the results of `sytrf!`. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored. A is overwritten by its inverse.

> source

`Base.LinAlg.LAPACK.sytrs!` – Function.

```
sytrs!(uplo, A, ipiv, B)
```

Solves the equation A * X = B for a symmetric matrix A using the results of `sytrf!`. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored. B is overwritten by the solution X.

> source

`Base.LinAlg.LAPACK.hesv!` – Function.

```
hesv!(uplo, A, B) -> (B, A, ipiv)
```

Finds the solution to A * X = B for Hermitian matrix A. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored. B is overwritten by the solution X. A is overwritten by its Bunch-Kaufman factorization. `ipiv` contains pivoting information about the factorization.

> source

`Base.LinAlg.LAPACK.hetrf!` – Function.

```
hetrf!(uplo, A) -> (A, ipiv, info)
```

Computes the Bunch-Kaufman factorization of a Hermitian matrix A. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored.

Returns A, overwritten by the factorization, a pivot vector `ipiv`, and the error code `info` which is a non-negative integer. If `info` is positive the matrix is singular and the diagonal part of the factorization is exactly zero at position `info`.

> source

`Base.LinAlg.LAPACK.hetri!` – Function.

```
hetri!(uplo, A, ipiv)
```

Computes the inverse of a Hermitian matrix A using the results of `sytrf!`. If `uplo = U`, the upper half of A is stored. If `uplo = L`, the lower half is stored. A is overwritten by its inverse.

> source

Base.LinAlg.LAPACK.hetrs! – Function.

```
hetrs!(uplo, A, ipiv, B)
```

Solves the equation A * X = B for a Hermitian matrix A using the results of sytrf!. If uplo = U, the upper half of A is stored. If uplo = L, the lower half is stored. B is overwritten by the solution X.

source

Base.LinAlg.LAPACK.syev! – Function.

```
syev!(jobz, uplo, A)
```

Finds the eigenvalues (jobz = N) or eigenvalues and eigenvectors (jobz = V) of a symmetric matrix A. If uplo = U, the upper triangle of A is used. If uplo = L, the lower triangle of A is used.

source

Base.LinAlg.LAPACK.syevr! – Function.

```
syevr!(jobz, range, uplo, A, vl, vu, il, iu, abstol) -> (W, Z)
```

Finds the eigenvalues (jobz = N) or eigenvalues and eigenvectors (jobz = V) of a symmetric matrix A. If uplo = U, the upper triangle of A is used. If uplo = L, the lower triangle of A is used. If range = A, all the eigenvalues are found. If range = V, the eigenvalues in the half-open interval (vl, vu] are found. If range = I, the eigenvalues with indices between il and iu are found. abstol can be set as a tolerance for convergence.

The eigenvalues are returned in W and the eigenvectors in Z.

source

Base.LinAlg.LAPACK.sygvd! – Function.

```
sygvd!(itype, jobz, uplo, A, B) -> (w, A, B)
```

Finds the generalized eigenvalues (jobz = N) or eigenvalues and eigenvectors (jobz = V) of a symmetric matrix A and symmetric positive-definite matrix B. If uplo = U, the upper triangles of A and B are used. If uplo = L, the lower triangles of A and B are used. If itype = 1, the problem to solve is A * x = lambda * B * x. If itype = 2, the problem to solve is A * B * x = lambda * x. If itype = 3, the problem to solve is B * A * x = lambda * x.

source

Base.LinAlg.LAPACK.bdsqr! – Function.

```
bdsqr!(uplo, d, e_, Vt, U, C) -> (d, Vt, U, C)
```

Computes the singular value decomposition of a bidiagonal matrix with d on the diagonal and e_ on the off-diagonal. If uplo = U, e_ is the superdiagonal. If uplo = L, e_ is the subdiagonal. Can optionally also compute the product Q' * C.

Returns the singular values in d, and the matrix C overwritten with Q' * C.

source

Base.LinAlg.LAPACK.bdsdc! – Function.

```
bdsdc!(uplo, compq, d, e_) -> (d, e, u, vt, q, iq)
```

Computes the singular value decomposition of a bidiagonal matrix with d on the diagonal and e_ on the off-diagonal using a divide and conqueq method. If `uplo = U`, e_ is the superdiagonal. If `uplo = L`, e_ is the subdiagonal. If `compq = N`, only the singular values are found. If `compq = I`, the singular values and vectors are found. If `compq = P`, the singular values and vectors are found in compact form. Only works for real types.

Returns the singular values in d, and if `compq = P`, the compact singular vectors in `iq`.

source

Base.LinAlg.LAPACK.gecon! – Function.

```
gecon!(normtype, A, anorm)
```

Finds the reciprocal condition number of matrix A. If `normtype = I`, the condition number is found in the infinity norm. If `normtype = 0` or 1, the condition number is found in the one norm. A must be the result of `getrf!` and anorm is the norm of A in the relevant norm.

source

Base.LinAlg.LAPACK.gehrd! – Function.

```
gehrd!(ilo, ihi, A) -> (A, tau)
```

Converts a matrix A to Hessenberg form. If A is balanced with `gebal!` then ilo and ihi are the outputs of `gebal!`. Otherwise they should be `ilo = 1` and `ihi = size(A,2)`. tau contains the elementary reflectors of the factorization.

source

Base.LinAlg.LAPACK.orghr! – Function.

```
orghr!(ilo, ihi, A, tau)
```

Explicitly finds Q, the orthogonal/unitary matrix from `gehrd!`. ilo, ihi, A, and tau must correspond to the input/output to `gehrd!`.

source

Base.LinAlg.LAPACK.gees! – Function.

```
gees!(jobvs, A) -> (A, vs, w)
```

Computes the eigenvalues (`jobvs = N`) or the eigenvalues and Schur vectors (`jobvs = V`) of matrix A. A is overwritten by its Schur form.

Returns A, vs containing the Schur vectors, and w, containing the eigenvalues.

source

Base.LinAlg.LAPACK.gges! – Function.

```
gges!(jobvsl, jobvsr, A, B) -> (A, B, alpha, beta, vsl, vsr)
```

Computes the generalized eigenvalues, generalized Schur form, left Schur vectors (`jobsvl = V`), or right Schur vectors (`jobvsr = V`) of A and B.

The generalized eigenvalues are returned in `alpha` and `beta`. The left Schur vectors are returned in `vsl` and the right Schur vectors are returned in `vsr`.

source

`Base.LinAlg.LAPACK.trexc!` – Function.

```
trexc!(compq, ifst, ilst, T, Q) -> (T, Q)
```

Reorder the Schur factorization of a matrix. If `compq = V`, the Schur vectors Q are reordered. If `compq = N` they are not modified. `ifst` and `ilst` specify the reordering of the vectors.

source

`Base.LinAlg.LAPACK.trsen!` – Function.

```
trsen!(compq, job, select, T, Q) -> (T, Q, w)
```

Reorder the Schur factorization of a matrix and optionally finds reciprocal condition numbers. If `job = N`, no condition numbers are found. If `job = E`, only the condition number for this cluster of eigenvalues is found. If `job = V`, only the condition number for the invariant subspace is found. If `job = B` then the condition numbers for the cluster and subspace are found. If `compq = V` the Schur vectors Q are updated. If `compq = N` the Schur vectors are not modified. `select` determines which eigenvalues are in the cluster.

Returns T, Q, and reordered eigenvalues in w.

source

`Base.LinAlg.LAPACK.tgsen!` – Function.

```
tgsen!(select, S, T, Q, Z) -> (S, T, alpha, beta, Q, Z)
```

Reorders the vectors of a generalized Schur decomposition. `select` specifices the eigenvalues in each cluster.

source

`Base.LinAlg.LAPACK.trsyl!` – Function.

```
trsyl!(transa, transb, A, B, C, isgn=1) -> (C, scale)
```

Solves the Sylvester matrix equation A * X +/- X * B = scale*C where A and B are both quasi-upper triangular. If `transa = N`, A is not modified. If `transa = T`, A is transposed. If `transa = C`, A is conjugate transposed. Similarly for `transb` and B. If `isgn = 1`, the equation A * X + X * B = scale * C is solved. If `isgn = -1`, the equation A * X - X * B = scale * C is solved.

Returns X (overwriting C) and `scale`.

source

Chapter 53

Constants

`Core.nothing` – Constant.

```
nothing
```

The singleton instance of type `Void`, used by convention when there is no value to return (as in a C `void` function). Can be converted to an empty `Nullable` value.

source

`Base.PROGRAM_FILE` – Constant.

```
PROGRAM_FILE
```

A string containing the script name passed to Julia from the command line. Note that the script name remains unchanged from within included files. Alternatively see `@__FILE__`.

source

`Base.ARGS` – Constant.

```
ARGS
```

An array of the command line arguments passed to Julia, as strings.

source

`Base.C_NULL` – Constant.

```
C_NULL
```

The C null pointer constant, sometimes used when calling external code.

source

`Base.VERSION` – Constant.

```
VERSION
```

A `VersionNumber` object describing which version of Julia is in use. For details see Version Number Literals.

source

`Base.LOAD_PATH` – Constant.

```
LOAD_PATH
```

An array of paths as strings or custom loader objects for the `require` function and `using` and `import` statements to consider when loading code. To create a custom loader type, define the type and then add appropriate methods to the `Base.load_hook` function with the following signature:

```
Base.load_hook(loader::Loader, name::String, found::Any)
```

The `loader` argument is the current value in `LOAD_PATH`, name is the name of the module to load, and `found` is the path of any previously found code to provide name. If no provider has been found earlier in `LOAD_PATH` then the value of `found` will be `nothing`. Custom loader functionality is experimental and may break or change in Julia 1.0.

> source

Base.JULIA_HOME – Constant.

```
JULIA_HOME
```

A string containing the full path to the directory containing the `julia` executable.

> source

Core.ANY – Constant.

```
ANY
```

Equivalent to Any for dispatch purposes, but signals the compiler to skip code generation specialization for that field.

> source

Base.Sys.CPU_CORES – Constant.

```
Sys.CPU_CORES
```

The number of logical CPU cores available in the system.

See the Hwloc.jl package for extended information, including number of physical cores.

> source

Base.Sys.WORD_SIZE – Constant.

```
Sys.WORD_SIZE
```

Standard word size on the current machine, in bits.

> source

Base.Sys.KERNEL – Constant.

```
Sys.KERNEL
```

A symbol representing the name of the operating system, as returned by uname of the build configuration.

> source

Base.Sys.ARCH – Constant.

> `Sys.ARCH`

A symbol representing the architecture of the build configuration.

source

`Base.Sys.MACHINE` – Constant.

> `Sys.MACHINE`

A string containing the build triple.

source

See also:

- STDIN

- STDOUT

- STDERR

- ENV

- ENDIAN_BOM

- `Libc.MS_ASYNC`

- `Libc.MS_INVALIDATE`

- `Libc.MS_SYNC`

- Libdl.DL_LOAD_PATH

- Libdl.RTLD_DEEPBIND

- Libdl.RTLD_LOCAL

- Libdl.RTLD_NOLOAD

- Libdl.RTLD_LAZY

- Libdl.RTLD_NOW

- Libdl.RTLD_GLOBAL

- Libdl.RTLD_NODELETE

- Libdl.RTLD_FIRST

Chapter 54

Filesystem

`Base.Filesystem.pwd` – Function.

```
pwd() -> AbstractString
```

Get the current working directory.

source

`Base.Filesystem.cd` – Method.

```
cd(dir::AbstractString=homedir())
```

Set the current working directory.

source

`Base.Filesystem.cd` – Method.

```
cd(f::Function, dir::AbstractString=homedir())
```

Temporarily changes the current working directory and applies function f before returning.

source

`Base.Filesystem.readdir` – Function.

```
readdir(dir::AbstractString=".") -> Vector{String}
```

Returns the files and directories in the directory dir (or the current working directory if not given).

source

`Base.Filesystem.walkdir` – Function.

```
walkdir(dir; topdown=true, follow_symlinks=false, onerror=throw)
```

The walkdir method returns an iterator that walks the directory tree of a directory. The iterator returns a tuple containing (`rootpath`, `dirs`, `files`). The directory tree can be traversed top-down or bottom-up. If walkdir encounters a SystemError it will rethrow the error by default. A custom error handling function can be provided through onerror keyword argument. onerror is called with a SystemError as argument.

```
for (root, dirs, files) in walkdir(".")
    println("Directories in $root")
    for dir in dirs
        println(joinpath(root, dir)) # path to directories
    end
    println("Files in $root")
    for file in files
        println(joinpath(root, file)) # path to files
    end
end
```

source

Base.Filesystem.mkdir – Function.

```
mkdir(path::AbstractString, mode::Unsigned=0o777)
```

Make a new directory with name path and permissions mode. mode defaults to 0o777, modified by the current file creation mask. This function never creates more than one directory. If the directory already exists, or some intermediate directories do not exist, this function throws an error. See mkpath for a function which creates all required intermediate directories.

source

Base.Filesystem.mkpath – Function.

```
mkpath(path::AbstractString, mode::Unsigned=0o777)
```

Create all directories in the given path, with permissions mode. mode defaults to 0o777, modified by the current file creation mask.

source

Base.Filesystem.symlink – Function.

```
symlink(target::AbstractString, link::AbstractString)
```

Creates a symbolic link to target with the name link.

> **Note**
> This function raises an error under operating systems that do not support soft symbolic links, such as Windows XP.

source

Base.Filesystem.readlink – Function.

```
readlink(path::AbstractString) -> AbstractString
```

Returns the target location a symbolic link path points to.

source

Base.Filesystem.chmod – Function.

```
chmod(path::AbstractString, mode::Integer; recursive::Bool=false)
```

Change the permissions mode of `path` to `mode`. Only integer modes (e.g. `0o777`) are currently supported. If `recursive=true` and the path is a directory all permissions in that directory will be recursively changed.

source

`Base.Filesystem.chown` – Function.

```
chown(path::AbstractString, owner::Integer, group::Integer=-1)
```

Change the owner and/or group of `path` to `owner` and/or `group`. If the value entered for `owner` or `group` is `-1` the corresponding ID will not change. Only integer `owners` and `groups` are currently supported.

source

`Base.stat` – Function.

```
stat(file)
```

Returns a structure whose fields contain information about the file. The fields of the structure are:

Name	Description
size	The size (in bytes) of the file
device	ID of the device that contains the file
inode	The inode number of the file
mode	The protection mode of the file
nlink	The number of hard links to the file
uid	The user id of the owner of the file
gid	The group id of the file owner
rdev	If this file refers to a device, the ID of the device it refers to
blksize	The file-system preferred block size for the file
blocks	The number of such blocks allocated
mtime	Unix timestamp of when the file was last modified
ctime	Unix timestamp of when the file was created

source

`Base.Filesystem.lstat` – Function.

```
lstat(file)
```

Like `stat`, but for symbolic links gets the info for the link itself rather than the file it refers to. This function must be called on a file path rather than a file object or a file descriptor.

source

`Base.Filesystem.ctime` – Function.

```
ctime(file)
```

Equivalent to `stat(file).ctime`

source

`Base.Filesystem.mtime` – Function.

```
mtime(file)
```

Equivalent to `stat(file).mtime`.

`Base.Filesystem.filemode` – Function.

```
filemode(file)
```

Equivalent to `stat(file).mode`

`Base.Filesystem.filesize` – Function.

```
filesize(path...)
```

Equivalent to `stat(file).size`.

`Base.Filesystem.uperm` – Function.

```
uperm(file)
```

Gets the permissions of the owner of the file as a bitfield of

Value	Description
01	Execute Permission
02	Write Permission
04	Read Permission

For allowed arguments, see `stat`.

`Base.Filesystem.gperm` – Function.

```
gperm(file)
```

Like `uperm` but gets the permissions of the group owning the file.

`Base.Filesystem.operm` – Function.

```
operm(file)
```

Like `uperm` but gets the permissions for people who neither own the file nor are a member of the group owning the file

`Base.Filesystem.cp` – Function.

```
cp(src::AbstractString, dst::AbstractString; remove_destination::Bool=false, follow_symlinks
    ::Bool=false)
```

Copy the file, link, or directory from `src` to `dest`. `remove_destination=true` will first remove an existing `dst`.

If `follow_symlinks=false`, and `src` is a symbolic link, `dst` will be created as a symbolic link. If `follow_symlinks=true` and `src` is a symbolic link, `dst` will be a copy of the file or directory `src` refers to.

`Base.download` – Function.

```
download(url::AbstractString, [localfile::AbstractString])
```

Download a file from the given url, optionally renaming it to the given local file name. Note that this function relies on the availability of external tools such as `curl`, `wget` or `fetch` to download the file and is provided for convenience. For production use or situations in which more options are needed, please use a package that provides the desired functionality instead.

source

`Base.Filesystem.mv` – Function.

```
mv(src::AbstractString, dst::AbstractString; remove_destination::Bool=false)
```

Move the file, link, or directory from `src` to `dst`. `remove_destination=true` will first remove an existing `dst`.

source

`Base.Filesystem.rm` – Function.

```
rm(path::AbstractString; force::Bool=false, recursive::Bool=false)
```

Delete the file, link, or empty directory at the given path. If `force=true` is passed, a non-existing path is not treated as error. If `recursive=true` is passed and the path is a directory, then all contents are removed recursively.

source

`Base.Filesystem.touch` – Function.

```
touch(path::AbstractString)
```

Update the last-modified timestamp on a file to the current time.

source

`Base.Filesystem.tempname` – Function.

```
tempname()
```

Generate a unique temporary file path.

source

`Base.Filesystem.tempdir` – Function.

```
tempdir()
```

Obtain the path of a temporary directory (possibly shared with other processes).

source

`Base.Filesystem.mktemp` – Method.

```
mktemp(parent=tempdir())
```

Returns `(path, io)`, where `path` is the path of a new temporary file in `parent` and `io` is an open file object for this path.

source

Base.Filesystem.mktemp – Method.

```
mktemp(f::Function, parent=tempdir())
```

Apply the function f to the result of `mktemp(parent)` and remove the temporary file upon completion.

source

Base.Filesystem.mktempdir – Method.

```
mktempdir(parent=tempdir())
```

Create a temporary directory in the `parent` directory and return its path. If `parent` does not exist, throw an error.

source

Base.Filesystem.mktempdir – Method.

```
mktempdir(f::Function, parent=tempdir())
```

Apply the function f to the result of `mktempdir(parent)` and remove the temporary directory upon completion.

source

Base.Filesystem.isblockdev – Function.

```
isblockdev(path) -> Bool
```

Returns `true` if `path` is a block device, `false` otherwise.

source

Base.Filesystem.ischardev – Function.

```
ischardev(path) -> Bool
```

Returns `true` if `path` is a character device, `false` otherwise.

source

Base.Filesystem.isdir – Function.

```
isdir(path) -> Bool
```

Returns `true` if `path` is a directory, `false` otherwise.

source

Base.Filesystem.isfifo – Function.

```
isfifo(path) -> Bool
```

Returns `true` if `path` is a FIFO, `false` otherwise.

source

Base.Filesystem.isfile – Function.

```
isfile(path) -> Bool
```

Returns `true` if `path` is a regular file, `false` otherwise.

source

`Base.Filesystem.islink` – Function.

```
islink(path) -> Bool
```

Returns `true` if `path` is a symbolic link, `false` otherwise.

source

`Base.Filesystem.ismount` – Function.

```
ismount(path) -> Bool
```

Returns `true` if `path` is a mount point, `false` otherwise.

source

`Base.Filesystem.ispath` – Function.

```
ispath(path) -> Bool
```

Returns `true` if `path` is a valid filesystem path, `false` otherwise.

source

`Base.Filesystem.issetgid` – Function.

```
issetgid(path) -> Bool
```

Returns `true` if `path` has the setgid flag set, `false` otherwise.

source

`Base.Filesystem.issetuid` – Function.

```
issetuid(path) -> Bool
```

Returns `true` if `path` has the setuid flag set, `false` otherwise.

source

`Base.Filesystem.issocket` – Function.

```
issocket(path) -> Bool
```

Returns `true` if `path` is a socket, `false` otherwise.

source

`Base.Filesystem.issticky` – Function.

```
issticky(path) -> Bool
```

Returns `true` if `path` has the sticky bit set, `false` otherwise.

source

`Base.Filesystem.homedir` – Function.

```
homedir() -> AbstractString
```

Return the current user's home directory.

> **Note**
> homedir determines the home directory via libuv's uv_os_homedir. For details (for example on
> how to specify the home directory via environment variables), see the uv_os_homedir documenta-
> tion.

source

Base.Filesystem.dirname – Function.

```
dirname(path::AbstractString) -> AbstractString
```

Get the directory part of a path.

```julia
julia> dirname("/home/myuser")
"/home"
```

See also: basename

source

Base.Filesystem.basename – Function.

```
basename(path::AbstractString) -> AbstractString
```

Get the file name part of a path.

```julia
julia> basename("/home/myuser/example.jl")
"example.jl"
```

See also: dirname

source

Base.@__FILE__ – Macro.

```
@__FILE__ -> AbstractString
```

@__FILE__ expands to a string with the absolute file path of the file containing the macro. Returns nothing if
run from a REPL or an empty string if evaluated by julia -e <expr>. Alternatively see PROGRAM_FILE.

source

Base.@__DIR__ – Macro.

```
@__DIR__ -> AbstractString
```

@__DIR__ expands to a string with the directory part of the absolute path of the file containing the macro. Returns
nothing if run from a REPL or an empty string if evaluated by julia -e <expr>.

source

@__LINE__ – Macro.

```
@__LINE__ -> Int
```

@__LINE__ expands to the line number of the call-site.

source

Base.Filesystem.isabspath – Function.

```
isabspath(path::AbstractString) -> Bool
```

Determines whether a path is absolute (begins at the root directory).

```
julia> isabspath("/home")
true

julia> isabspath("home")
false
```

source

Base.Filesystem.isdirpath – Function.

```
isdirpath(path::AbstractString) -> Bool
```

Determines whether a path refers to a directory (for example, ends with a path separator).

```
julia> isdirpath("/home")
false

julia> isdirpath("/home/")
true
```

source

Base.Filesystem.joinpath – Function.

```
joinpath(parts...) -> AbstractString
```

Join path components into a full path. If some argument is an absolute path, then prior components are dropped.

```
julia> joinpath("/home/myuser","example.jl")
"/home/myuser/example.jl"
```

source

Base.Filesystem.abspath – Function.

```
abspath(path::AbstractString) -> AbstractString
```

Convert a path to an absolute path by adding the current directory if necessary.

source

```
abspath(path::AbstractString, paths::AbstractString...) -> AbstractString
```

Convert a set of paths to an absolute path by joining them together and adding the current directory if necessary. Equivalent to abspath(joinpath(path, paths...)).

source

Base.Filesystem.normpath – Function.

```
normpath(path::AbstractString) -> AbstractString
```

Normalize a path, removing "." and ".." entries.

```
julia> normpath("/home/myuser/../example.jl")
"/home/example.jl"
```

source

Base.Filesystem.realpath – Function.

```
realpath(path::AbstractString) -> AbstractString
```

Canonicalize a path by expanding symbolic links and removing "." and ".." entries.

source

Base.Filesystem.relpath – Function.

```
relpath(path::AbstractString, startpath::AbstractString = ".") -> AbstractString
```

Return a relative filepath to `path` either from the current directory or from an optional start directory. This is a path computation: the filesystem is not accessed to confirm the existence or nature of `path` or `startpath`.

source

Base.Filesystem.expanduser – Function.

```
expanduser(path::AbstractString) -> AbstractString
```

On Unix systems, replace a tilde character at the start of a path with the current user's home directory.

source

Base.Filesystem.splitdir – Function.

```
splitdir(path::AbstractString) -> (AbstractString, AbstractString)
```

Split a path into a tuple of the directory name and file name.

```
julia> splitdir("/home/myuser")
("/home", "myuser")
```

source

Base.Filesystem.splitdrive – Function.

```
splitdrive(path::AbstractString) -> (AbstractString, AbstractString)
```

On Windows, split a path into the drive letter part and the path part. On Unix systems, the first component is always the empty string.

source

Base.Filesystem.splitext – Function.

```
splitext(path::AbstractString) -> (AbstractString, AbstractString)
```

If the last component of a path contains a dot, split the path into everything before the dot and everything including and after the dot. Otherwise, return a tuple of the argument unmodified and the empty string.

```julia
julia> splitext("/home/myuser/example.jl")
("/home/myuser/example", ".jl")

julia> splitext("/home/myuser/example")
("/home/myuser/example", "")
```

source

Chapter 55

I/O and Network

55.1   General I/O

Base.STDOUT – Constant.

> STDOUT

Global variable referring to the standard out stream.

> source

Base.STDERR – Constant.

> STDERR

Global variable referring to the standard error stream.

> source

Base.STDIN – Constant.

> STDIN

Global variable referring to the standard input stream.

> source

Base.open – Function.

```
open(filename::AbstractString, [read::Bool, write::Bool, create::Bool, truncate::Bool, append
    ::Bool]) -> IOStream
```

Open a file in a mode specified by five boolean arguments. The default is to open files for reading only. Returns a stream for accessing the file.

> source

```
open(filename::AbstractString, [mode::AbstractString]) -> IOStream
```

Alternate syntax for open, where a string-based mode specifier is used instead of the five booleans. The values of mode correspond to those from `fopen(3)` or Perl open, and are equivalent to setting the following boolean groups:

> source

Mode	Description
r	read
r+	read, write
w	write, create, truncate
w+	read, write, create, truncate
a	write, create, append
a+	read, write, create, append

```
open(f::Function, args...)
```

Apply the function f to the result of `open(args...)` and close the resulting file descriptor upon completion.

Example: `open(readstring, "file.txt")`

source

```
open(command, mode::AbstractString="r", stdio=DevNull)
```

Start running command asynchronously, and return a tuple (`stream`,`process`). If mode is `"r"`, then `stream` reads from the process's standard output and `stdio` optionally specifies the process's standard input stream. If mode is `"w"`, then `stream` writes to the process's standard input and `stdio` optionally specifies the process's standard output stream.

source

```
open(f::Function, command, mode::AbstractString="r", stdio=DevNull)
```

Similar to `open(command, mode, stdio)`, but calls `f(stream)` on the resulting read or write stream, then closes the stream and waits for the process to complete. Returns the value returned by f.

source

Base.IOBuffer – Type.

```
IOBuffer([data,],[readable::Bool=true, writable::Bool=true, [maxsize::Int=typemax(Int)]])
```

Create an `IOBuffer`, which may optionally operate on a pre-existing array. If the readable/writable arguments are given, they restrict whether or not the buffer may be read from or written to respectively. By default the buffer is readable but not writable. The last argument optionally specifies a size beyond which the buffer may not be grown.

source

```
IOBuffer() -> IOBuffer
```

Create an in-memory I/O stream.

source

```
IOBuffer(size::Int)
```

Create a fixed size IOBuffer. The buffer will not grow dynamically.

source

```
IOBuffer(string::String)
```

Create a read-only `IOBuffer` on the data underlying the given string.

```
julia> io = IOBuffer("Haho");

julia> String(take!(io))
"Haho"

julia> String(take!(io))
"Haho"
```

source

Base.take! – Method.

```
take!(b::IOBuffer)
```

Obtain the contents of an IOBuffer as an array, without copying. Afterwards, the IOBuffer is reset to its initial state.

source

Base.fdio – Function.

```
fdio([name::AbstractString, ]fd::Integer[, own::Bool=false]) -> IOStream
```

Create an IOStream object from an integer file descriptor. If own is true, closing this object will close the underlying descriptor. By default, an IOStream is closed when it is garbage collected. name allows you to associate the descriptor with a named file.

source

Base.flush – Function.

```
flush(stream)
```

Commit all currently buffered writes to the given stream.

source

Base.close – Function.

```
close(stream)
```

Close an I/O stream. Performs a flush first.

source

Base.write – Function.

```
write(stream::IO, x)
write(filename::AbstractString, x)
```

Write the canonical binary representation of a value to the given I/O stream or file. Returns the number of bytes written into the stream.

You can write multiple values with the same write call. i.e. the following are equivalent:

```
write(stream, x, y...)
write(stream, x) + write(stream, y...)
```

source

Base.read – Function.

```
read(filename::AbstractString, args...)
```

Open a file and read its contents. `args` is passed to `read`: this is equivalent to `open(io->read(io, args...),` `filename)`.

source

```
read(stream::IO, T, dims)
```

Read a series of values of type `T` from `stream`, in canonical binary representation. `dims` is either a tuple or a series of integer arguments specifying the size of the `Array{T}` to return.

source

```
read(s::IO, nb=typemax(Int))
```

Read at most nb bytes from s, returning a `Vector{UInt8}` of the bytes read.

source

```
read(s::IOStream, nb::Integer; all=true)
```

Read at most nb bytes from s, returning a `Vector{UInt8}` of the bytes read.

If `all` is `true` (the default), this function will block repeatedly trying to read all requested bytes, until an error or end-of-file occurs. If `all` is `false`, at most one `read` call is performed, and the amount of data returned is device-dependent. Note that not all stream types support the `all` option.

source

```
read(stream::IO, T)
```

Read a single value of type `T` from `stream`, in canonical binary representation.

source

Base.read! – Function.

```
read!(stream::IO, array::Union{Array, BitArray})
read!(filename::AbstractString, array::Union{Array, BitArray})
```

Read binary data from an I/O stream or file, filling in `array`.

source

Base.readbytes! – Function.

```
readbytes!(stream::IO, b::AbstractVector{UInt8}, nb=length(b))
```

Read at most nb bytes from `stream` into b, returning the number of bytes read. The size of b will be increased if needed (i.e. if nb is greater than `length(b)` and enough bytes could be read), but it will never be decreased.

source

```
readbytes!(stream::IOStream, b::AbstractVector{UInt8}, nb=length(b); all::Bool=true)
```

Read at most nb bytes from `stream` into b, returning the number of bytes read. The size of b will be increased if needed (i.e. if nb is greater than `length(b)` and enough bytes could be read), but it will never be decreased.

See read for a description of the `all` option.

source

Base.unsafe_read – Function.

> unsafe_read(io::IO, ref, nbytes::UInt)

Copy `nbytes` from the `IO` stream object into `ref` (converted to a pointer).

It is recommended that subtypes `T<:IO` override the following method signature to provide more efficient implementations: `unsafe_read(s::T, p::Ptr{UInt8}, n::UInt)`

> source

Base.unsafe_write – Function.

> unsafe_write(io::IO, ref, nbytes::UInt)

Copy `nbytes` from `ref` (converted to a pointer) into the `IO` object.

It is recommended that subtypes `T<:IO` override the following method signature to provide more efficient implementations: `unsafe_write(s::T, p::Ptr{UInt8}, n::UInt)`

> source

Base.position – Function.

> position(s)

Get the current position of a stream.

> source

Base.seek – Function.

> seek(s, pos)

Seek a stream to the given position.

> source

Base.seekstart – Function.

> seekstart(s)

Seek a stream to its beginning.

> source

Base.seekend – Function.

> seekend(s)

Seek a stream to its end.

> source

Base.skip – Function.

> skip(s, offset)

Seek a stream relative to the current position.

> source

Base.mark – Function.

```
mark(s)
```

Add a mark at the current position of stream s. Returns the marked position.

See also unmark, reset, ismarked.

source

Base.unmark – Function.

```
unmark(s)
```

Remove a mark from stream s. Returns `true` if the stream was marked, `false` otherwise.

See also mark, reset, ismarked.

source

Base.reset – Function.

```
reset(s)
```

Reset a stream s to a previously marked position, and remove the mark. Returns the previously marked position. Throws an error if the stream is not marked.

See also mark, unmark, ismarked.

source

Base.ismarked – Function.

```
ismarked(s)
```

Returns `true` if stream s is marked.

See also mark, unmark, reset.

source

Base.eof – Function.

```
eof(stream) -> Bool
```

Tests whether an I/O stream is at end-of-file. If the stream is not yet exhausted, this function will block to wait for more data if necessary, and then return `false`. Therefore it is always safe to read one byte after seeing eof return `false`. eof will return `false` as long as buffered data is still available, even if the remote end of a connection is closed.

source

Base.isreadonly – Function.

```
isreadonly(stream) -> Bool
```

Determine whether a stream is read-only.

source

Base.iswritable – Function.

```
iswritable(io) -> Bool
```

Returns `true` if the specified IO object is writable (if that can be determined).

source

Base.`isreadable` – Function.

```
isreadable(io) -> Bool
```

Returns `true` if the specified IO object is readable (if that can be determined).

source

Base.`isopen` – Function.

```
isopen(object) -> Bool
```

Determine whether an object - such as a stream, timer, or mmap – is not yet closed. Once an object is closed, it will never produce a new event. However, a closed stream may still have data to read in its buffer, use `eof` to check for the ability to read data. Use `poll_fd` to be notified when a stream might be writable or readable.

source

Base.`Serializer.serialize` – Function.

```
serialize(stream, value)
```

Write an arbitrary value to a stream in an opaque format, such that it can be read back by `deserialize`. The read-back value will be as identical as possible to the original. In general, this process will not work if the reading and writing are done by different versions of Julia, or an instance of Julia with a different system image. `Ptr` values are serialized as all-zero bit patterns (NULL).

source

Base.`Serializer.deserialize` – Function.

```
deserialize(stream)
```

Read a value written by `serialize`. `deserialize` assumes the binary data read from `stream` is correct and has been serialized by a compatible implementation of `serialize`. It has been designed with simplicity and performance as a goal and does not validate the data read. Malformed data can result in process termination. The caller has to ensure the integrity and correctness of data read from `stream`.

source

Base.`Grisu.print_shortest` – Function.

```
print_shortest(io, x)
```

Print the shortest possible representation, with the minimum number of consecutive non-zero digits, of number x, ensuring that it would parse to the exact same number.

source

Base.`fd` – Function.

```
fd(stream)
```

Returns the file descriptor backing the stream or file. Note that this function only applies to synchronous `File`'s and `IOStream`'s not to any of the asynchronous streams.

source

`Base.redirect_stdout` – Function.

```
redirect_stdout([stream]) -> (rd, wr)
```

Create a pipe to which all C and Julia level STDOUT output will be redirected. Returns a tuple (`rd, wr`) representing the pipe ends. Data written to STDOUT may now be read from the rd end of the pipe. The wr end is given for convenience in case the old STDOUT object was cached by the user and needs to be replaced elsewhere.

> **Note**
>
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.redirect_stdout` – Method.

```
redirect_stdout(f::Function, stream)
```

Run the function f while redirecting STDOUT to `stream`. Upon completion, STDOUT is restored to its prior setting.

> **Note**
>
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.redirect_stderr` – Function.

```
redirect_stderr([stream]) -> (rd, wr)
```

Like `redirect_stdout`, but for STDERR.

> **Note**
>
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.redirect_stderr` – Method.

```
redirect_stderr(f::Function, stream)
```

Run the function f while redirecting STDERR to `stream`. Upon completion, STDERR is restored to its prior setting.

> **Note**
>
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.redirect_stdin` – Function.

```
redirect_stdin([stream]) -> (rd, wr)
```

Like `redirect_stdout`, but for `STDIN`. Note that the order of the return tuple is still (`rd, wr`), i.e. data to be read from `STDIN` may be written to `wr`.

> **Note**
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.redirect_stdin` – Method.

```
redirect_stdin(f::Function, stream)
```

Run the function `f` while redirecting `STDIN` to `stream`. Upon completion, `STDIN` is restored to its prior setting.

> **Note**
> `stream` must be a TTY, a Pipe, or a TCPSocket.

source

`Base.readchomp` – Function.

```
readchomp(x)
```

Read the entirety of `x` as a string and remove a single trailing newline. Equivalent to `chomp!(readstring(x))`.

source

`Base.truncate` – Function.

```
truncate(file,n)
```

Resize the file or buffer given by the first argument to exactly n bytes, filling previously unallocated space with '\0' if the file or buffer is grown.

source

`Base.skipchars` – Function.

```
skipchars(stream, predicate; linecomment::Char)
```

Advance the stream until before the first character for which `predicate` returns `false`. For example `skipchars(stream, isspace)` will skip all whitespace. If keyword argument `linecomment` is specified, characters from that character through the end of a line will also be skipped.

source

`Base.DataFmt.countlines` – Function.

```
countlines(io::IO, eol::Char='\n')
```

Read `io` until the end of the stream/file and count the number of lines. To specify a file pass the filename as the first argument. EOL markers other than '\n' are supported by passing them as the second argument.

source

`Base.PipeBuffer` – Function.

```
PipeBuffer(data::Vector{UInt8}=UInt8[],[maxsize::Int=typemax(Int)])
```

An `IOBuffer` that allows reading and performs writes by appending. Seeking and truncating are not supported. See `IOBuffer` for the available constructors. If `data` is given, creates a `PipeBuffer` to operate on a data vector, optionally specifying a size beyond which the underlying `Array` may not be grown.

source

Base.readavailable – Function.

```
readavailable(stream)
```

Read all available data on the stream, blocking the task only if no data is available. The result is a `Vector{UInt8, 1}`.

source

Base.IOContext – Type.

```
IOContext
```

`IOContext` provides a mechanism for passing output configuration settings among `show` methods.

In short, it is an immutable dictionary that is a subclass of `IO`. It supports standard dictionary operations such as `getindex`, and can also be used as an I/O stream.

source

Base.IOContext – Method.

```
IOContext(io::IO, KV::Pair)
```

Create an `IOContext` that wraps a given stream, adding the specified `key=>value` pair to the properties of that stream (note that `io` can itself be an `IOContext`).

- use `(key => value) in dict` to see if this particular combination is in the properties set
- use `get(dict, key, default)` to retrieve the most recent value for a particular key

The following properties are in common use:

- `:compact`: Boolean specifying that small values should be printed more compactly, e.g. that numbers should be printed with fewer digits. This is set when printing array elements.
- `:limit`: Boolean specifying that containers should be truncated, e.g. showing `…` in place of most elements.
- `:displaysize`: A `Tuple{Int, Int}` giving the size in rows and columns to use for text output. This can be used to override the display size for called functions, but to get the size of the screen use the `displaysize` function.

```julia
julia> function f(io::IO)
           if get(io, :short, false)
               print(io, "short")
           else
               print(io, "loooooong")
           end
       end
f (generic function with 1 method)

julia> f(STDOUT)
loooooong
julia> f(IOContext(STDOUT, :short => true))
short
```

source

Base.IOContext – Method.

```
IOContext(io::IO, context::IOContext)
```

Create an `IOContext` that wraps an alternate IO but inherits the properties of `context`.

source

55.2 Text I/O

Base.show – Method.

```
show(x)
```

Write an informative text representation of a value to the current output stream. New types should overload `show(io, x)` where the first argument is a stream. The representation used by `show` generally includes Julia-specific formatting and type information.

source

Base.showcompact – Function.

```
showcompact(x)
```

Show a compact representation of a value.

This is used for printing array elements without repeating type information (which would be redundant with that printed once for the whole array), and without line breaks inside the representation of an element.

To offer a compact representation different from its standard one, a custom type should test `get(io, :compact, false)` in its normal show method.

source

Base.showall – Function.

```
showall(x)
```

Similar to show, except shows all elements of arrays.

source

Base.summary – Function.

```
summary(x)
```

Return a string giving a brief description of a value. By default returns `string(typeof(x))`, e.g. Int64.

For arrays, returns a string of size and type info, e.g. `10-element Array{Int64,1}`.

```
julia> summary(1)
"Int64"

julia> summary(zeros(2))
"2-element Array{Float64,1}"
```

source

Base.print – Function.

```
print(io::IO, x)
```

Write to io (or to the default output stream STDOUT if io is not given) a canonical (un-decorated) text representation of a value if there is one, otherwise call show. The representation used by print includes minimal formatting and tries to avoid Julia-specific details.

```julia
julia> print("Hello World!")
Hello World!
julia> io = IOBuffer();

julia> print(io, "Hello World!")

julia> String(take!(io))
"Hello World!"
```

source

Base.println – Function.

```
println(io::IO, xs...)
```

Print (using print) xs followed by a newline. If io is not supplied, prints to STDOUT.

source

Base.print_with_color – Function.

```
print_with_color(color::Union{Symbol, Int}, [io], xs...; bold::Bool = false)
```

Print xs in a color specified as a symbol.

color may take any of the values :normal, :default, :bold, :black, :blue, :cyan, :green, :light_black, :light_blue, :light_cyan, :light_green, :light_magenta, :light_red, :light_yellow, :magenta, :nothing, :red, :white, or :yellow or an integer between 0 and 255 inclusive. Note that not all terminals support 256 colors. If the keyword bold is given as true, the result will be printed in bold.

source

Base.info – Function.

```
info([io, ] msg..., [prefix="INFO: "])
```

Display an informational message. Argument msg is a string describing the information to be displayed. The prefix keyword argument can be used to override the default prepending of msg.

```julia
julia> info("hello world")
INFO: hello world

julia> info("hello world"; prefix="MY INFO: ")
MY INFO: hello world
```

See also logging.

source

Base.warn – Function.

```
warn([io, ] msg..., [prefix="WARNING: ", once=false, key=nothing, bt=nothing, filename=
    nothing, lineno::Int=0])
```

Display a warning. Argument `msg` is a string describing the warning to be displayed. Set `once` to true and specify a `key` to only display `msg` the first time `warn` is called. If `bt` is not `nothing` a backtrace is displayed. If `filename` is not `nothing` both it and `lineno` are displayed.

See also logging.

source

```
warn(msg)
```

Display a warning. Argument `msg` is a string describing the warning to be displayed.

```julia
julia> warn("Beep Beep")
WARNING: Beep Beep
```

source

`Base.logging` – Function.

```
logging(io [, m [, f]][; kind=:all])
logging([; kind=:all])
```

Stream output of informational, warning, and/or error messages to `io`, overriding what was otherwise specified. Optionally, divert stream only for module `m`, or specifically function `f` within `m`. `kind` can be `:all` (the default), `:info`, `:warn`, or `:error`. See `Base.log_{info,warn,error}_to` for the current set of redirections. Call `logging` with no arguments (or just the `kind`) to reset everything.

source

`Base.Printf.@printf` – Macro.

```
@printf([io::IOStream], "%Fmt", args...)
```

Print `args` using C `printf()` style format specification string, with some caveats: `Inf` and `NaN` are printed consistently as `Inf` and `NaN` for flags `%a`, `%A`, `%e`, `%E`, `%f`, `%F`, `%g`, and `%G`. Furthermore, if a floating point number is equally close to the numeric values of two possible output strings, the output string further away from zero is chosen.

Optionally, an `IOStream` may be passed as the first argument to redirect output.

Examples

```julia
julia> @printf("%f %F %f %F\n", Inf, Inf, NaN, NaN)
Inf Inf NaN NaN


julia> @printf "%.0f %.1f %f\n" 0.5 0.025 -0.0078125
1 0.0 -0.007813
```

source

`Base.Printf.@sprintf` – Macro.

```
@sprintf("%Fmt", args...)
```

Return `@printf` formatted output as string.

Examples

```
julia> s = @sprintf "this is a %s %15.1f" "test" 34.567;

julia> println(s)
this is a test            34.6
```

source

Base.sprint – Function.

```
sprint(f::Function, args...)
```

Call the given function with an I/O stream and the supplied extra arguments. Everything written to this I/O stream is returned as a string.

```
julia> sprint(showcompact, 66.66666)
"66.6667"
```

source

Base.showerror – Function.

```
showerror(io, e)
```

Show a descriptive representation of an exception object.

source

Base.dump – Function.

```
dump(x)
```

Show every part of the representation of a value.

source

Base.readstring – Function.

```
readstring(stream::IO)
readstring(filename::AbstractString)
```

Read the entire contents of an I/O stream or a file as a string. The text is assumed to be encoded in UTF-8.

source

Base.readline – Function.

```
readline(stream::IO=STDIN; chomp::Bool=true)
readline(filename::AbstractString; chomp::Bool=true)
```

Read a single line of text from the given I/O stream or file (defaults to STDIN). When reading from a file, the text is assumed to be encoded in UTF-8. Lines in the input end with `'\n'` or `"\r\n"` or the end of an input stream. When chomp is true (as it is by default), these trailing newline characters are removed from the line before it is returned. When chomp is false, they are returned as part of the line.

source

Base.readuntil – Function.

```
readuntil(stream::IO, delim)
readuntil(filename::AbstractString, delim)
```

Read a string from an I/O stream or a file, up to and including the given delimiter byte. The text is assumed to be encoded in UTF-8.

source

Base.readlines – Function.

```
readlines(stream::IO=STDIN; chomp::Bool=true)
readlines(filename::AbstractString; chomp::Bool=true)
```

Read all lines of an I/O stream or a file as a vector of strings. Behavior is equivalent to saving the result of reading `readline` repeatedly with the same arguments and saving the resulting lines as a vector of strings.

source

Base.eachline – Function.

```
eachline(stream::IO=STDIN; chomp::Bool=true)
eachline(filename::AbstractString; chomp::Bool=true)
```

Create an iterable `EachLine` object that will yield each line from an I/O stream or a file. Iteration calls `readline` on the stream argument repeatedly with chomp passed through, determining whether trailing end-of-line characters are removed. When called with a file name, the file is opened once at the beginning of iteration and closed at the end. If iteration is interrupted, the file will be closed when the `EachLine` object is garbage collected.

source

Base.DataFmt.readdlm – Method.

```
readdlm(source, delim::Char, T::Type, eol::Char; header=false, skipstart=0, skipblanks=true,
    use_mmap, quotes=true, dims, comments=true, comment_char='#')
```

Read a matrix from the source where each line (separated by `eol`) gives one row, with elements separated by the given delimiter. The source can be a text file, stream or byte array. Memory mapped files can be used by passing the byte array representation of the mapped segment as source.

If `T` is a numeric type, the result is an array of that type, with any non-numeric elements as NaN for floating-point types, or zero. Other useful values of `T` include `String`, `AbstractString`, and `Any`.

If `header` is `true`, the first row of data will be read as header and the tuple (`data_cells`, `header_cells`) is returned instead of only `data_cells`.

Specifying `skipstart` will ignore the corresponding number of initial lines from the input.

If `skipblanks` is `true`, blank lines in the input will be ignored.

If `use_mmap` is `true`, the file specified by `source` is memory mapped for potential speedups. Default is `true` except on Windows. On Windows, you may want to specify `true` if the file is large, and is only read once and not written to.

If `quotes` is `true`, columns enclosed within double-quote (") characters are allowed to contain new lines and column delimiters. Double-quote characters within a quoted field must be escaped with another double-quote. Specifying `dims` as a tuple of the expected rows and columns (including header, if any) may speed up reading of large files. If `comments` is `true`, lines beginning with `comment_char` and text following `comment_char` in any line are ignored.

source

`Base.DataFmt.readdlm` – Method.

```
readdlm(source, delim::Char, eol::Char; options...)
```

If all data is numeric, the result will be a numeric array. If some elements cannot be parsed as numbers, a heterogeneous array of numbers and strings is returned.

> source

`Base.DataFmt.readdlm` – Method.

```
readdlm(source, delim::Char, T::Type; options...)
```

The end of line delimiter is taken as \n.

> source

`Base.DataFmt.readdlm` – Method.

```
readdlm(source, delim::Char; options...)
```

The end of line delimiter is taken as \n. If all data is numeric, the result will be a numeric array. If some elements cannot be parsed as numbers, a heterogeneous array of numbers and strings is returned.

> source

`Base.DataFmt.readdlm` – Method.

```
readdlm(source, T::Type; options...)
```

The columns are assumed to be separated by one or more whitespaces. The end of line delimiter is taken as \n.

> source

`Base.DataFmt.readdlm` – Method.

```
readdlm(source; options...)
```

The columns are assumed to be separated by one or more whitespaces. The end of line delimiter is taken as \n. If all data is numeric, the result will be a numeric array. If some elements cannot be parsed as numbers, a heterogeneous array of numbers and strings is returned.

> source

`Base.DataFmt.writedlm` – Function.

```
writedlm(f, A, delim='\t'; opts)
```

Write A (a vector, matrix, or an iterable collection of iterable rows) as text to f (either a filename string or an IO stream) using the given delimiter delim (which defaults to tab, but can be any printable Julia object, typically a Char or AbstractString).

For example, two vectors x and y of the same length can be written as two columns of tab-delimited text to f by either writedlm(f, [x y]) or by writedlm(f, zip(x, y)).

> source

`Base.DataFmt.readcsv` – Function.

```
readcsv(source, [T::Type]; options...)
```

Equivalent to `readdlm` with `delim` set to comma, and type optionally defined by T.

source

Base.DataFmt.writecsv – Function.

```
writecsv(filename, A; opts)
```

Equivalent to `writedlm` with `delim` set to comma.

source

Base.Base64.Base64EncodePipe – Type.

```
Base64EncodePipe(ostream)
```

Returns a new write-only I/O stream, which converts any bytes written to it into base64-encoded ASCII bytes written to `ostream`. Calling `close` on the Base64EncodePipe stream is necessary to complete the encoding (but does not close `ostream`).

```julia
julia> io = IOBuffer();

julia> iob64_encode = Base64EncodePipe(io);

julia> write(iob64_encode, "Hello!")
6

julia> close(iob64_encode);

julia> str = String(take!(io))
"SGVsbG8h"

julia> String(base64decode(str))
"Hello!"
```

source

Base.Base64.Base64DecodePipe – Type.

```
Base64DecodePipe(istream)
```

Returns a new read-only I/O stream, which decodes base64-encoded data read from `istream`.

```julia
julia> io = IOBuffer();

julia> iob64_decode = Base64DecodePipe(io);

julia> write(io, "SGVsbG8h")
8

julia> seekstart(io);

julia> String(read(iob64_decode))
"Hello!"
```

source

Base.Base64.base64encode – Function.

```
base64encode(writefunc, args...)
base64encode(args...)
```

Given a `write`-like function `writefunc`, which takes an I/O stream as its first argument, `base64encode(write-func, args...)` calls `writefunc` to write `args...` to a base64-encoded string, and returns the string. `base64encode(args...)` is equivalent to `base64encode(write, args...)`: it converts its arguments into bytes using the standard `write` functions and returns the base64-encoded string.

See also `base64decode`.

source

Base.Base64.base64decode – Function.

```
base64decode(string)
```

Decodes the base64-encoded `string` and returns a `Vector{UInt8}` of the decoded bytes.

See also `base64encode`

```julia
julia> b = base64decode("SGVsbG8h")
6-element Array{UInt8,1}:
 0x48
 0x65
 0x6c
 0x6c
 0x6f
 0x21

julia> String(b)
"Hello!"
```

source

Base.displaysize – Function.

```
displaysize(io) -> (lines, columns)
```

Return the nominal size of the screen that may be used for rendering output to this io object

source

55.3   Multimedia I/O

Just as text output is performed by `print` and user-defined types can indicate their textual representation by overloading `show`, Julia provides a standardized mechanism for rich multimedia output (such as images, formatted text, or even audio and video), consisting of three parts:

- A function `display(x)` to request the richest available multimedia display of a Julia object x (with a plain-text fallback).

- Overloading `show` allows one to indicate arbitrary multimedia representations (keyed by standard MIME types) of user-defined types.

- Multimedia-capable display backends may be registered by subclassing a generic `Display` type and pushing them onto a stack of display backends via `pushdisplay`.

The base Julia runtime provides only plain-text display, but richer displays may be enabled by loading external modules or by using graphical Julia environments (such as the IPython-based IJulia notebook).

`Base.Multimedia.display` – Function.

```
display(x)
display(d::Display, x)
display(mime, x)
display(d::Display, mime, x)
```

Display x using the topmost applicable display in the display stack, typically using the richest supported multimedia output for x, with plain-text `STDOUT` output as a fallback. The `display(d, x)` variant attempts to display x on the given display d only, throwing a `MethodError` if d cannot display objects of this type.

There are also two variants with a `mime` argument (a MIME type string, such as `"image/png"`), which attempt to display x using the requested MIME type *only*, throwing a `MethodError` if this type is not supported by either the display(s) or by x. With these variants, one can also supply the "raw" data in the requested MIME type by passing `x::AbstractString` (for MIME types with text-based storage, such as text/html or application/postscript) or `x::Vector{UInt8}` (for binary MIME types).

source

`Base.Multimedia.redisplay` – Function.

```
redisplay(x)
redisplay(d::Display, x)
redisplay(mime, x)
redisplay(d::Display, mime, x)
```

By default, the `redisplay` functions simply call `display`. However, some display backends may override `redisplay` to modify an existing display of x (if any). Using `redisplay` is also a hint to the backend that x may be redisplayed several times, and the backend may choose to defer the display until (for example) the next interactive prompt.

source

`Base.Multimedia.displayable` – Function.

```
displayable(mime) -> Bool
displayable(d::Display, mime) -> Bool
```

Returns a boolean value indicating whether the given `mime` type (string) is displayable by any of the displays in the current display stack, or specifically by the display d in the second variant.

source

`Base.show` – Method.

```
show(stream, mime, x)
```

The `display` functions ultimately call `show` in order to write an object x as a given `mime` type to a given I/O `stream` (usually a memory buffer), if possible. In order to provide a rich multimedia representation of a user-defined type T, it is only necessary to define a new `show` method for T, via: `show(stream, ::MIME"mime", x::T) = ...`, where `mime` is a MIME-type string and the function body calls `write` (or similar) to write that representation of x to `stream`. (Note that the `MIME""` notation only supports literal strings; to construct MIME types in a more flexible manner use `MIME{Symbol("")}`.)

For example, if you define a `MyImage` type and know how to write it to a PNG file, you could define a function `show(stream, ::MIME"image/png", x::MyImage) = ...` to allow your images to be displayed on any

PNG-capable `Display` (such as IJulia). As usual, be sure to `import Base.show` in order to add new methods to the built-in Julia function show.

The default MIME type is `MIME"text/plain"`. There is a fallback definition for `text/plain` output that calls show with 2 arguments. Therefore, this case should be handled by defining a 2-argument show(`stream::IO`, `x::MyType`) method.

Technically, the `MIME"mime"` macro defines a singleton type for the given `mime` string, which allows us to exploit Julia's dispatch mechanisms in determining how to display objects of any given type.

The first argument to show can be an `IOContext` specifying output format properties. See `IOContext` for details.

source

Base.Multimedia.mimewritable – Function.

```
mimewritable(mime, x)
```

Returns a boolean value indicating whether or not the object x can be written as the given `mime` type. (By default, this is determined automatically by the existence of the corresponding show method for `typeof(x)`.)

source

Base.Multimedia.reprmime – Function.

```
reprmime(mime, x)
```

Returns an `AbstractString` or `Vector{UInt8}` containing the representation of x in the requested `mime` type, as written by show (throwing a `MethodError` if no appropriate show is available). An `AbstractString` is returned for MIME types with textual representations (such as `"text/html"` or `"application/postscript"`), whereas binary data is returned as `Vector{UInt8}`. (The function `istextmime(mime)` returns whether or not Julia treats a given `mime` type as text.)

As a special case, if x is an `AbstractString` (for textual MIME types) or a `Vector{UInt8}` (for binary MIME types), the `reprmime` function assumes that x is already in the requested `mime` format and simply returns x. This special case does not apply to the `"text/plain"` MIME type. This is useful so that raw data can be passed to display(`m::MIME`, x).

source

Base.Multimedia.stringmime – Function.

```
stringmime(mime, x)
```

Returns an `AbstractString` containing the representation of x in the requested `mime` type. This is similar to reprmime except that binary data is base64-encoded as an ASCII string.

source

As mentioned above, one can also define new display backends. For example, a module that can display PNG images in a window can register this capability with Julia, so that calling `display(x)` on types with PNG representations will automatically display the image using the module's window.

In order to define a new display backend, one should first create a subtype D of the abstract class `Display`. Then, for each MIME type (mime string) that can be displayed on D, one should define a function display(`d::D`, `::MIME"mime"`, x) = ... that displays x as that MIME type, usually by calling `reprmime(mime, x)`. A `MethodError` should be thrown if x cannot be displayed as that MIME type; this is automatic if one calls `reprmime`. Finally, one should define a function display(`d::D`, x) that queries `mimewritable(mime, x)` for the mime types supported by D and displays the "best" one; a `MethodError` should be thrown if no supported MIME types are found for x. Similarly, some

subtypes may wish to override `redisplay(d::D, ...)`. (Again, one should `import Base.display` to add new methods to `display`.) The return values of these functions are up to the implementation (since in some cases it may be useful to return a display "handle" of some type). The display functions for D can then be called directly, but they can also be invoked automatically from `display(x)` simply by pushing a new display onto the display-backend stack with:

`Base.Multimedia.pushdisplay` – Function.

```
pushdisplay(d::Display)
```

Pushes a new display d on top of the global display-backend stack. Calling `display(x)` or `display(mime, x)` will display x on the topmost compatible backend in the stack (i.e., the topmost backend that does not throw a `MethodError`).

source

`Base.Multimedia.popdisplay` – Function.

```
popdisplay()
popdisplay(d::Display)
```

Pop the topmost backend off of the display-backend stack, or the topmost copy of d in the second variant.

source

`Base.Multimedia.TextDisplay` – Type.

```
TextDisplay(io::IO)
```

Returns a `TextDisplay <: Display`, which displays any object as the text/plain MIME type (by default), writing the text representation to the given I/O stream. (This is how objects are printed in the Julia REPL.)

source

`Base.Multimedia.istextmime` – Function.

```
istextmime(m::MIME)
```

Determine whether a MIME type is text data. MIME types are assumed to be binary data except for a set of types known to be text data (possibly Unicode).

source

55.4   Memory-mapped I/O

`Base.Mmap.Anonymous` – Type.

```
Mmap.Anonymous(name, readonly, create)
```

Create an `IO`-like object for creating zeroed-out mmapped-memory that is not tied to a file for use in `Mmap.mmap`. Used by `SharedArray` for creating shared memory arrays.

source

`Base.Mmap.mmap` – Method.

```
Mmap.mmap(io::Union{IOStream,AbstractString,Mmap.AnonymousMmap}[, type::Type{Array{T,N}},
    dims, offset]; grow::Bool=true, shared::Bool=true)
      Mmap.mmap(type::Type{Array{T,N}}, dims)
```

Create an `Array` whose values are linked to a file, using memory-mapping. This provides a convenient way of working with data too large to fit in the computer's memory.

The type is an `Array{T,N}` with a bits-type element of `T` and dimension `N` that determines how the bytes of the array are interpreted. Note that the file must be stored in binary format, and no format conversions are possible (this is a limitation of operating systems, not Julia).

`dims` is a tuple or single `Integer` specifying the size or length of the array.

The file is passed via the stream argument, either as an open `IOStream` or filename string. When you initialize the stream, use `"r"` for a "read-only" array, and `"w+"` to create a new array used to write values to disk.

If no `type` argument is specified, the default is `Vector{UInt8}`.

Optionally, you can specify an offset (in bytes) if, for example, you want to skip over a header in the file. The default value for the offset is the current stream position for an `IOStream`.

The `grow` keyword argument specifies whether the disk file should be grown to accommodate the requested size of array (if the total file size is < requested array size). Write privileges are required to grow the file.

The `shared` keyword argument specifies whether the resulting `Array` and changes made to it will be visible to other processes mapping the same file.

For example, the following code

```
# Create a file for mmapping
# (you could alternatively use mmap to do this step, too)
A = rand(1:20, 5, 30)
s = open("/tmp/mmap.bin", "w+")
# We'll write the dimensions of the array as the first two Ints in the file
write(s, size(A,1))
write(s, size(A,2))
# Now write the data
write(s, A)
close(s)

# Test by reading it back in
s = open("/tmp/mmap.bin")   # default is read-only
m = read(s, Int)
n = read(s, Int)
A2 = Mmap.mmap(s, Matrix{Int}, (m,n))
```

creates a m-by-n `Matrix{Int}`, linked to the file associated with stream s.

A more portable file would need to encode the word size – 32 bit or 64 bit – and endianness information in the header. In practice, consider encoding binary data using standard formats like HDF5 (which can be used with memory-mapping).

source

`Base.Mmap.mmap` – Method.

```
Mmap.mmap(io, BitArray, [dims, offset])
```

Create a `BitArray` whose values are linked to a file, using memory-mapping; it has the same purpose, works in the same way, and has the same arguments, as mmap, but the byte representation is different.

Example: B = Mmap.mmap(s, BitArray, (25,30000))

This would create a 25-by-30000 `BitArray`, linked to the file associated with stream s.

source

`Base.Mmap.sync!` – Function.

```
Mmap.sync!(array)
```

Forces synchronization between the in-memory version of a memory-mapped `Array` or `BitArray` and the on-disk version.

source

55.5 Network I/O

`Base.connect` – Method.

```
connect([host], port::Integer) -> TCPSocket
```

Connect to the host `host` on port `port`.

source

`Base.connect` – Method.

```
connect(path::AbstractString) -> PipeEndpoint
```

Connect to the named pipe / UNIX domain socket at `path`.

source

`Base.listen` – Method.

```
listen([addr, ]port::Integer; backlog::Integer=BACKLOG_DEFAULT) -> TCPServer
```

Listen on port on the address specified by `addr`. By default this listens on `localhost` only. To listen on all interfaces pass `IPv4(0)` or `IPv6(0)` as appropriate. `backlog` determines how many connections can be pending (not having called `accept`) before the server will begin to reject them. The default value of `backlog` is 511.

source

`Base.listen` – Method.

```
listen(path::AbstractString) -> PipeServer
```

Create and listen on a named pipe / UNIX domain socket.

source

`Base.getaddrinfo` – Function.

```
getaddrinfo(host::AbstractString) -> IPAddr
```

Gets the IP address of the `host` (may have to do a DNS lookup)

source

`Base.getsockname` – Function.

```
getsockname(sock::Union{TCPServer, TCPSocket}) -> (IPAddr, UInt16)
```

Get the IP address and the port that the given `TCPSocket` is connected to (or bound to, in the case of `TCPServer`).

source

`Base.IPv4` – Type.

> IPv4(host::Integer) -> IPv4

Returns an IPv4 object from ip address `host` formatted as an `Integer`.

> **julia>** **IPv4**(3223256218)
> ip"192.30.252.154"

source

`Base.IPv6` – Type.

> IPv6(host::Integer) -> IPv6

Returns an IPv6 object from ip address `host` formatted as an `Integer`.

> **julia>** **IPv6**(3223256218)
> ip"::c01e:fc9a"

source

`Base.nb_available` – Function.

> nb_available(stream)

Returns the number of bytes available for reading before a read from this stream or buffer will block.

source

`Base.accept` – Function.

> accept(server[,client])

Accepts a connection on the given server and returns a connection to the client. An uninitialized client stream may be provided, in which case it will be used instead of creating a new stream.

source

`Base.listenany` – Function.

> listenany([host::IPAddr,] port_hint) -> (UInt16, TCPServer)

Create a `TCPServer` on any port, using hint as a starting point. Returns a tuple of the actual port that the server was created on and the server itself.

source

`Base.Filesystem.poll_fd` – Function.

> poll_fd(fd, timeout_s::Real=-1; readable=false, writable=false)

Monitor a file descriptor `fd` for changes in the read or write availability, and with a timeout given by `timeout_s` seconds.

The keyword arguments determine which of read and/or write status should be monitored; at least one of them must be set to `true`.

The returned value is an object with boolean fields `readable`, `writable`, and `timedout`, giving the result of the polling.

source

Base.Filesystem.poll_file – Function.

```
poll_file(path::AbstractString, interval_s::Real=5.007, timeout_s::Real=-1) -> (previous::
    StatStruct, current::StatStruct)
```

Monitor a file for changes by polling every `interval_s` seconds until a change occurs or `timeout_s` seconds have elapsed. The `interval_s` should be a long period; the default is 5.007 seconds.

Returns a pair of `StatStruct` objects (`previous`, `current`) when a change is detected.

To determine when a file was modified, compare `mtime(prev) != mtime(current)` to detect notification of changes. However, using `watch_file` for this operation is preferred, since it is more reliable and efficient, although in some situations it may not be available.

source

Base.Filesystem.watch_file – Function.

```
watch_file(path::AbstractString, timeout_s::Real=-1)
```

Watch file or directory `path` for changes until a change occurs or `timeout_s` seconds have elapsed.

The returned value is an object with boolean fields `changed`, `renamed`, and `timedout`, giving the result of watching the file.

This behavior of this function varies slightly across platforms. See https://nodejs.org/api/fs.html#fs_caveats for more detailed information.

source

Base.bind – Function.

```
bind(socket::Union{UDPSocket, TCPSocket}, host::IPAddr, port::Integer; ipv6only=false,
    reuseaddr=false, kws...)
```

Bind `socket` to the given `host:port`. Note that `0.0.0.0` will listen on all devices.

- The `ipv6only` parameter disables dual stack mode. If `ipv6only=true`, only an IPv6 stack is created.
- If `reuseaddr=true`, multiple threads or processes can bind to the same address without error if they all set `reuseaddr=true`, but only the last to bind will receive any traffic.

source

```
bind(chnl::Channel, task::Task)
```

Associates the lifetime of `chnl` with a task. Channel `chnl` is automatically closed when the task terminates. Any uncaught exception in the task is propagated to all waiters on `chnl`.

The `chnl` object can be explicitly closed independent of task termination. Terminating tasks have no effect on already closed Channel objects.

When a channel is bound to multiple tasks, the first task to terminate will close the channel. When multiple channels are bound to the same task, termination of the task will close all of the bound channels.

```
julia> c = Channel(0);

julia> task = @schedule foreach(i->put!(c, i), 1:4);

julia> bind(c,task);
```

```julia
julia> for i in c
           @show i
       end;
i = 1
i = 2
i = 3
i = 4

julia> isopen(c)
false

julia> c = Channel(0);

julia> task = @schedule (put!(c,1);error("foo"));

julia> bind(c,task);

julia> take!(c)
1

julia> put!(c,1);
ERROR: foo
Stacktrace:
 [1] check_channel_state(::Channel{Any}) at ./channels.jl:131
 [2] put!(::Channel{Any}, ::Int64) at ./channels.jl:261
```

source

Base.send – Function.

```julia
send(socket::UDPSocket, host, port::Integer, msg)
```

Send msg over socket to host:port.

source

Base.recv – Function.

```julia
recv(socket::UDPSocket)
```

Read a UDP packet from the specified socket, and return the bytes received. This call blocks.

source

Base.recvfrom – Function.

```julia
recvfrom(socket::UDPSocket) -> (address, data)
```

Read a UDP packet from the specified socket, returning a tuple of (address, data), where address will be either IPv4 or IPv6 as appropriate.

source

Base.setopt – Function.

```julia
setopt(sock::UDPSocket; multicast_loop = nothing, multicast_ttl=nothing, enable_broadcast=
    nothing, ttl=nothing)
```

Set UDP socket options.

- `multicast_loop`: loopback for multicast packets (default: `true`).
- `multicast_ttl`: TTL for multicast packets (default: `nothing`).
- `enable_broadcast`: flag must be set to `true` if socket will be used for broadcast messages, or else the UDP system will return an access error (default: `false`).
- `ttl`: Time-to-live of packets sent on the socket (default: `nothing`).

source

`Base.ntoh` – Function.

```
ntoh(x)
```

Converts the endianness of a value from Network byte order (big-endian) to that used by the Host.

source

`Base.hton` – Function.

```
hton(x)
```

Converts the endianness of a value from that used by the Host to Network byte order (big-endian).

source

`Base.ltoh` – Function.

```
ltoh(x)
```

Converts the endianness of a value from Little-endian to that used by the Host.

source

`Base.htol` – Function.

```
htol(x)
```

Converts the endianness of a value from that used by the Host to Little-endian.

source

`Base.ENDIAN_BOM` – Constant.

```
ENDIAN_BOM
```

The 32-bit byte-order-mark indicates the native byte order of the host machine. Little-endian machines will contain the value `0x04030201`. Big-endian machines will contain the value `0x01020304`.

source

Chapter 56

Punctuation

Extended documentation for mathematical symbols & functions is here.

symbol	meaning
@m	invoke macro m; followed by space-separated expressions
!	prefix "not" operator
a!( )	at the end of a function name, ! indicates that a function modifies its argument(s)
#	begin single line comment
#=	begin multi-line comment (these are nestable)
=#	end multi-line comment
$	string and expression interpolation
%	remainder operator
^	exponent operator
&	bitwise and
&&	short-circuiting boolean and
|	bitwise or
||	short-circuiting boolean or
	bitwise xor operator
*	multiply, or matrix multiply
( )	the empty tuple
~	bitwise not operator
\	backslash operator
'	complex transpose operator A^H
a[ ]	array indexing
[,]	vertical concatenation
[;]	also vertical concatenation
[ ]	with space-separated expressions, horizontal concatenation
T{ }	parametric type instantiation
;	statement separator
,	separate function arguments or tuple components
?	3-argument conditional operator (conditional ? if_true : if_false)
" "	delimit string literals
' '	delimit character literals
` `	delimit external process (command) specifications
. . .	splice arguments into a function call or declare a varargs function or type
.	access named fields in objects/modules, also prefixes elementwise operator/function calls
a:b	range a, a+1, a+2, ..., b
a:s:b	range a, a+s, a+2s, ..., b
:	index an entire dimension (1:end)
::	type annotation, depending on context
:( )	quoted expression
:a	symbol a
<:	subtype operator
>:	supertype operator (reverse of subtype operator)
===	egal comparison operator

Chapter 57

Sorting and Related Functions

Julia has an extensive, flexible API for sorting and interacting with already-sorted arrays of values. By default, Julia picks reasonable algorithms and sorts in standard ascending order:

```
julia> sort([2,3,1])
3-element Array{Int64,1}:
 1
 2
 3
```

You can easily sort in reverse order as well:

```
julia> sort([2,3,1], rev=true)
3-element Array{Int64,1}:
 3
 2
 1
```

To sort an array in-place, use the "bang" version of the sort function:

```
julia> a = [2,3,1];

julia> sort!(a);

julia> a
3-element Array{Int64,1}:
 1
 2
 3
```

Instead of directly sorting an array, you can compute a permutation of the array's indices that puts the array into sorted order:

```
julia> v = randn(5)
5-element Array{Float64,1}:
  0.297288
  0.382396
 -0.597634
```

```
 -0.0104452
 -0.839027

julia> p = sortperm(v)
5-element Array{Int64,1}:
 5
 3
 4
 1
 2

julia> v[p]
5-element Array{Float64,1}:
 -0.839027
 -0.597634
 -0.0104452
  0.297288
  0.382396
```

Arrays can easily be sorted according to an arbitrary transformation of their values:

```
julia> sort(v, by=abs)
5-element Array{Float64,1}:
 -0.0104452
  0.297288
  0.382396
 -0.597634
 -0.839027
```

Or in reverse order by a transformation:

```
julia> sort(v, by=abs, rev=true)
5-element Array{Float64,1}:
 -0.839027
 -0.597634
  0.382396
  0.297288
 -0.0104452
```

If needed, the sorting algorithm can be chosen:

```
julia> sort(v, alg=InsertionSort)
5-element Array{Float64,1}:
 -0.839027
 -0.597634
 -0.0104452
  0.297288
  0.382396
```

All the sorting and order related functions rely on a "less than" relation defining a total order on the values to be manipulated. The `isless` function is invoked by default, but the relation can be specified via the `lt` keyword.

57.1 Sorting Functions

Base.sort! – Function.

```
sort!(v; alg::Algorithm=defalg(v), lt=isless, by=identity, rev::Bool=false, order::Ordering=
    Forward)
```

Sort the vector v in place. QuickSort is used by default for numeric arrays while MergeSort is used for other arrays. You can specify an algorithm to use via the alg keyword (see Sorting Algorithms for available algorithms). The by keyword lets you provide a function that will be applied to each element before comparison; the lt keyword allows providing a custom "less than" function; use rev=true to reverse the sorting order. These options are independent and can be used together in all possible combinations: if both by and lt are specified, the lt function is applied to the result of the by function; rev=true reverses whatever ordering specified via the by and lt keywords.

```
julia> v = [3, 1, 2]; sort!(v); v
3-element Array{Int64,1}:
 1
 2
 3

julia> v = [3, 1, 2]; sort!(v, rev = true); v
3-element Array{Int64,1}:
 3
 2
 1

julia> v = [(1, "c"), (3, "a"), (2, "b")]; sort!(v, by = x -> x[1]); v
3-element Array{Tuple{Int64,String},1}:
 (1, "c")
 (2, "b")
 (3, "a")

julia> v = [(1, "c"), (3, "a"), (2, "b")]; sort!(v, by = x -> x[2]); v
3-element Array{Tuple{Int64,String},1}:
 (3, "a")
 (2, "b")
 (1, "c")
```

source

Base.sort – Function.

```
sort(v; alg::Algorithm=defalg(v), lt=isless, by=identity, rev::Bool=false, order::Ordering=
    Forward)
```

Variant of sort! that returns a sorted copy of v leaving v itself unmodified.

```
julia> v = [3, 1, 2];

julia> sort(v)
3-element Array{Int64,1}:
 1
 2
 3
```

```julia
julia> v
3-element Array{Int64,1}:
 3
 1
 2
```

```julia
sort(A, dim::Integer; alg::Algorithm=DEFAULT_UNSTABLE, lt=isless, by=identity, rev::Bool=
    false, order::Ordering=Forward, initialized::Bool=false)
```

Sort a multidimensional array A along the given dimension. See sort! for a description of possible keyword arguments.

```julia
julia> A = [4 3; 1 2]
2×2 Array{Int64,2}:
 4  3
 1  2

julia> sort(A, 1)
2×2 Array{Int64,2}:
 1  2
 4  3

julia> sort(A, 2)
2×2 Array{Int64,2}:
 3  4
 1  2
```

Base.sortperm – Function.

```julia
sortperm(v; alg::Algorithm=DEFAULT_UNSTABLE, lt=isless, by=identity, rev::Bool=false, order::
    Ordering=Forward)
```

Return a permutation vector of indices of v that puts it in sorted order. Specify alg to choose a particular sorting algorithm (see Sorting Algorithms). MergeSort is used by default, and since it is stable, the resulting permutation will be the lexicographically first one that puts the input array into sorted order – i.e. indices of equal elements appear in ascending order. If you choose a non-stable sorting algorithm such as QuickSort, a different permutation that puts the array into order may be returned. The order is specified using the same keywords as sort!.

See also sortperm!.

```julia
julia> v = [3, 1, 2];

julia> p = sortperm(v)
3-element Array{Int64,1}:
 2
 3
 1

julia> v[p]
3-element Array{Int64,1}:
 1
 2
 3
```

Base.Sort.sortperm! – Function.

```
sortperm!(ix, v; alg::Algorithm=DEFAULT_UNSTABLE, lt=isless, by=identity, rev::Bool=false,
    order::Ordering=Forward, initialized::Bool=false)
```

Like sortperm, but accepts a preallocated index vector ix. If initialized is false (the default), ix is initialized to contain the values 1:length(v).

```julia
julia> v = [3, 1, 2]; p = zeros(Int, 3);

julia> sortperm!(p, v); p
3-element Array{Int64,1}:
 2
 3
 1

julia> v[p]
3-element Array{Int64,1}:
 1
 2
 3
```

Base.Sort.sortrows – Function.

```
sortrows(A; alg::Algorithm=DEFAULT_UNSTABLE, lt=isless, by=identity, rev::Bool=false, order::
    Ordering=Forward)
```

Sort the rows of matrix A lexicographically. See sort! for a description of possible keyword arguments.

Examples

```julia
julia> sortrows([7 3 5; -1 6 4; 9 -2 8])
3×3 Array{Int64,2}:
 -1   6  4
  7   3  5
  9  -2  8

julia> sortrows([7 3 5; -1 6 4; 9 -2 8], lt=(x,y)->isless(x[2],y[2]))
3×3 Array{Int64,2}:
  9  -2  8
  7   3  5
 -1   6  4

julia> sortrows([7 3 5; -1 6 4; 9 -2 8], rev=true)
3×3 Array{Int64,2}:
  9  -2  8
  7   3  5
 -1   6  4
```

Base.Sort.sortcols – Function.

```
sortcols(A; alg::Algorithm=DEFAULT_UNSTABLE, lt=isless, by=identity, rev::Bool=false, order::
    Ordering=Forward)
```

Sort the columns of matrix A lexicographically. See `sort!` for a description of possible keyword arguments.

Examples

```
julia> sortcols([7 3 5; 6 -1 -4; 9 -2 8])
3×3 Array{Int64,2}:
  3   5  7
 -1  -4  6
 -2   8  9

julia> sortcols([7 3 5; 6 -1 -4; 9 -2 8], alg=InsertionSort, lt=(x,y)->isless(x[2],y[2]))
3×3 Array{Int64,2}:
  5   3  7
 -4  -1  6
  8  -2  9

julia> sortcols([7 3 5; 6 -1 -4; 9 -2 8], rev=true)
3×3 Array{Int64,2}:
 7   5   3
 6  -4  -1
 9   8  -2
```

source

57.2   Order-Related Functions

`Base.issorted` – Function.

```
issorted(v, lt=isless, by=identity, rev:Bool=false, order::Ordering=Forward)
```

Test whether a vector is in sorted order. The `lt`, `by` and `rev` keywords modify what order is considered to be sorted just as they do for `sort`.

```
julia> issorted([1, 2, 3])
true

julia> issorted([(1, "b"), (2, "a")], by = x -> x[1])
true

julia> issorted([(1, "b"), (2, "a")], by = x -> x[2])
false

julia> issorted([(1, "b"), (2, "a")], by = x -> x[2], rev=true)
true
```

source

`Base.Sort.searchsorted` – Function.

```
searchsorted(a, x, [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Returns the range of indices of a which compare as equal to x (using binary search) according to the order specified by the by, lt and rev keywords, assuming that a is already sorted in that order. Returns an empty range located at the insertion point if a does not contain values equal to x.

source

Base.Sort.searchsortedfirst – Function.

```
searchsortedfirst(a, x, [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Returns the index of the first value in a greater than or equal to x, according to the specified order. Returns length(a)+1 if x is greater than all values in a.

source

Base.Sort.searchsortedlast – Function.

```
searchsortedlast(a, x, [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Returns the index of the last value in a less than or equal to x, according to the specified order. Returns 0 if x is less than all values in a.

source

Base.Sort.select! – Function.

```
select!(v, k, [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Partially sort the vector v in place, according to the order specified by by, lt and rev so that the value at index k (or range of adjacent values if k is a range) occurs at the position where it would appear if the array were fully sorted via a non-stable algorithm. If k is a single index, that value is returned; if k is a range, an array of values at those indices is returned. Note that select! does not fully sort the input array.

source

Base.Sort.select – Function.

```
select(v, k, [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Variant of select! which copies v before partially sorting it, thereby returning the same thing as select! but leaving v unmodified.

source

Base.Sort.selectperm – Function.

```
selectperm(v, k, [alg=<algorithm>,] [by=<transform>,] [lt=<comparison>,] [rev=false])
```

Return a partial permutation of the vector v, according to the order specified by by, lt and rev, so that v[output] returns the first k (or range of adjacent values if k is a range) values of a fully sorted version of v. If k is a single index (Integer), an array of the first k indices is returned; if k is a range, an array of those indices is returned. Note that the handling of integer values for k is different from select in that it returns a vector of k elements instead of just the k th element. Also note that this is equivalent to, but more efficient than, calling sortperm(...)[k]

source

Base.Sort.selectperm! – Function.

```
selectperm!(ix, v, k, [alg=<algorithm>,] [by=<transform>,] [lt=<comparison>,] [rev=false,] [
    initialized=false])
```

Like `selectperm`, but accepts a preallocated index vector `ix`. If `initialized` is `false` (the default), ix is initialized to contain the values `1:length(ix)`.

source

57.3 Sorting Algorithms

There are currently four sorting algorithms available in base Julia:

- `InsertionSort`

- `QuickSort`

- `PartialQuickSort(k)`

- `MergeSort`

`InsertionSort` is an O(n^2) stable sorting algorithm. It is efficient for very small n, and is used internally by `QuickSort`.

`QuickSort` is an O(n log n) sorting algorithm which is in-place, very fast, but not stable – i.e. elements which are considered equal will not remain in the same order in which they originally appeared in the array to be sorted. `QuickSort` is the default algorithm for numeric values, including integers and floats.

`PartialQuickSort(k)` is similar to `QuickSort`, but the output array is only sorted up to index k if k is an integer, or in the range of k if k is an `OrdinalRange`. For example:

```
x = rand(1:500, 100)
k = 50
k2 = 50:100
s = sort(x; alg=QuickSort)
ps = sort(x; alg=PartialQuickSort(k))
qs = sort(x; alg=PartialQuickSort(k2))
map(issorted, (s, ps, qs))           # => (true, false, false)
map(x->issorted(x[1:k]), (s, ps, qs)) # => (true, true, false)
map(x->issorted(x[k2]), (s, ps, qs))  # => (true, false, true)
s[1:k] == ps[1:k]                     # => true
s[k2] == qs[k2]                       # => true
```

`MergeSort` is an O(n log n) stable sorting algorithm but is not in-place – it requires a temporary array of half the size of the input array – and is typically not quite as fast as `QuickSort`. It is the default algorithm for non-numeric data.

The default sorting algorithms are chosen on the basis that they are fast and stable, or *appear* to be so. For numeric types indeed, `QuickSort` is selected as it is faster and indistinguishable in this case from a stable sort (unless the array records its mutations in some way). The stability property comes at a non-negligible cost, so if you don't need it, you may want to explicitly specify your preferred algorithm, e.g. `sort!(v, alg=QuickSort)`.

The mechanism by which Julia picks default sorting algorithms is implemented via the `Base.Sort.defalg` function. It allows a particular algorithm to be registered as the default in all sorting functions for specific arrays. For example, here are the two default methods from `sort.jl`:

```
defalg(v::AbstractArray) = MergeSort
defalg{T<:Number}(v::AbstractArray{T}) = QuickSort
```

As for numeric arrays, choosing a non-stable default algorithm for array types for which the notion of a stable sort is meaningless (i.e. when two values comparing equal can not be distinguished) may make sense.

Chapter 58

Package Manager Functions

All package manager functions are defined in the Pkg module. None of the Pkg module's functions are exported; to use them, you'll need to prefix each function call with an explicit `Pkg.`, e.g. `Pkg.status()` or `Pkg.dir()`.

Functions for package development (e.g. `tag`, `publish`, etc.) have been moved to the PkgDev package. See PkgDev README for the documentation of those functions.

`Base.Pkg.dir` – Function.

```
dir() -> AbstractString
```

Returns the absolute path of the package directory. This defaults to `joinpath(homedir(),".julia","v$(VERSION.major).$(VERSION.minor)")` on all platforms (i.e. `~/.julia/v0.6` in UNIX shell syntax). If the JULIA_PKGDIR environment variable is set, then that path is used in the returned value as `joinpath(ENV["JULIA_PKGDIR"],"v$(VERSION.major).$(VERSION.minor)")`. If JULIA_PKGDIR is a relative path, it is interpreted relative to whatever the current working directory is.

source

```
dir(names...) -> AbstractString
```

Equivalent to `normpath(Pkg.dir(),names...)` – i.e. it appends path components to the package directory and normalizes the resulting path. In particular, `Pkg.dir(pkg)` returns the path to the package `pkg`.

source

`Base.Pkg.init` – Function.

```
init(meta::AbstractString=DEFAULT_META, branch::AbstractString=META_BRANCH)
```

Initialize `Pkg.dir()` as a package directory. This will be done automatically when the JULIA_PKGDIR is not set and `Pkg.dir()` uses its default value. As part of this process, clones a local METADATA git repository from the site and branch specified by its arguments, which are typically not provided. Explicit (non-default) arguments can be used to support a custom METADATA setup.

source

`Base.Pkg.resolve` – Function.

```
resolve()
```

Determines an optimal, consistent set of package versions to install or upgrade to. The optimal set of package versions is based on the contents of `Pkg.dir("REQUIRE")` and the state of installed packages in `Pkg.dir()`, Packages that are no longer required are moved into `Pkg.dir(".trash")`.

source

Base.Pkg.edit – Function.

```
edit()
```

Opens `Pkg.dir("REQUIRE")` in the editor specified by the `VISUAL` or `EDITOR` environment variables; when the editor command returns, it runs `Pkg.resolve()` to determine and install a new optimal set of installed package versions.

source

Base.Pkg.add – Function.

```
add(pkg, vers...)
```

Add a requirement entry for `pkg` to `Pkg.dir("REQUIRE")` and call `Pkg.resolve()`. If `vers` are given, they must be `VersionNumber` objects and they specify acceptable version intervals for `pkg`.

source

Base.Pkg.rm – Function.

```
rm(pkg)
```

Remove all requirement entries for `pkg` from `Pkg.dir("REQUIRE")` and call `Pkg.resolve()`.

source

Base.Pkg.clone – Function.

```
clone(pkg)
```

If `pkg` has a URL registered in `Pkg.dir("METADATA")`, clone it from that URL on the default branch. The package does not need to have any registered versions.

source

```
clone(url, [pkg])
```

Clone a package directly from the git URL `url`. The package does not need to be registered in `Pkg.dir("META-DATA")`. The package repo is cloned by the name `pkg` if provided; if not provided, `pkg` is determined automatically from `url`.

source

Base.Pkg.setprotocol! – Function.

```
setprotocol!(proto)
```

Set the protocol used to access GitHub-hosted packages. Defaults to 'https', with a blank `proto` delegating the choice to the package developer.

source

Base.Pkg.available – Function.

```
available() -> Vector{String}
```

Returns the names of available packages.

source

```
available(pkg) -> Vector{VersionNumber}
```

Returns the version numbers available for package pkg.

source

Base.Pkg.installed – Function.

```
installed() -> Dict{String,VersionNumber}
```

Returns a dictionary mapping installed package names to the installed version number of each package.

source

```
installed(pkg) -> Void | VersionNumber
```

If pkg is installed, return the installed version number. If pkg is registered, but not installed, return nothing.

source

Base.Pkg.status – Function.

```
status()
```

Prints out a summary of what packages are installed and what version and state they're in.

source

```
status(pkg)
```

Prints out a summary of what version and state pkg, specifically, is in.

source

Base.Pkg.update – Function.

```
update(pkgs...)
```

Update the metadata repo – kept in Pkg.dir("METADATA") – then update any fixed packages that can safely be pulled from their origin; then call Pkg.resolve() to determine a new optimal set of packages versions.

Without arguments, updates all installed packages. When one or more package names are provided as arguments, only those packages and their dependencies are updated.

source

Base.Pkg.checkout – Function.

```
checkout(pkg, [branch="master"]; merge=true, pull=true)
```

Checkout the Pkg.dir(pkg) repo to the branch branch. Defaults to checking out the "master" branch. To go back to using the newest compatible released version, use Pkg.free(pkg). Changes are merged (fast-forward only) if the keyword argument merge == true, and the latest version is pulled from the upstream repo if pull == true.

source

Base.Pkg.pin – Function.

```
pin(pkg)
```

Pin pkg at the current version. To go back to using the newest compatible released version, use `Pkg.free(pkg)`

source

```
pin(pkg, version)
```

Pin pkg at registered version `version`.

source

`Base.Pkg.free` – Function.

```
free(pkg)
```

Free the package pkg to be managed by the package manager again. It calls `Pkg.resolve()` to determine optimal package versions after. This is an inverse for both `Pkg.checkout` and `Pkg.pin`.

You can also supply an iterable collection of package names, e.g., `Pkg.free(("Pkg1", "Pkg2"))` to free multiple packages at once.

source

`Base.Pkg.build` – Function.

```
build()
```

Run the build scripts for all installed packages in depth-first recursive order.

source

```
build(pkgs...)
```

Run the build script in `deps/build.jl` for each package in `pkgs` and all of their dependencies in depth-first recursive order. This is called automatically by `Pkg.resolve()` on all installed or updated packages.

source

`Base.Pkg.test` – Function.

```
test(; coverage=false)
```

Run the tests for all installed packages ensuring that each package's test dependencies are installed for the duration of the test. A package is tested by running its `test/runtests.jl` file and test dependencies are specified in `test/REQUIRE`. Coverage statistics for the packages may be generated by passing `coverage=true`. The default behavior is not to run coverage.

source

```
test(pkgs...; coverage=false)
```

Run the tests for each package in `pkgs` ensuring that each package's test dependencies are installed for the duration of the test. A package is tested by running its `test/runtests.jl` file and test dependencies are specified in `test/REQUIRE`. Coverage statistics for the packages may be generated by passing `coverage=true`. The default behavior is not to run coverage.

source

`Base.Pkg.dependents` – Function.

```
dependents(pkg)
```

List the packages that have pkg as a dependency.

source

Chapter 59

Dates and Time

59.1   Dates and Time Types

Base.Dates.Period – Type.

```
Period
Year
Month
Week
Day
Hour
Minute
Second
Millisecond
Microsecond
Nanosecond
```

Period types represent discrete, human representations of time.

source

Base.Dates.CompoundPeriod – Type.

```
CompoundPeriod
```

A CompoundPeriod is useful for expressing time periods that are not a fixed multiple of smaller periods. For example, "a year and a day" is not a fixed number of days, but can be expressed using a CompoundPeriod. In fact, a CompoundPeriod is automatically generated by addition of different period types, e.g. Year(1) + Day(1) produces a CompoundPeriod result.

source

Base.Dates.Instant – Type.

```
Instant
```

Instant types represent integer-based, machine representations of time as continuous timelines starting from an epoch.

source

Base.Dates.UTInstant – Type.

> UTInstant{T}

The UTInstant represents a machine timeline based on UT time (1 day = one revolution of the earth). The T is a Period parameter that indicates the resolution or precision of the instant.

> source

Base.Dates.TimeType – Type.

> TimeType

TimeType types wrap Instant machine instances to provide human representations of the machine instant. Time, DateTime and Date are subtypes of TimeType.

> source

Base.Dates.DateTime – Type.

> DateTime

DateTime wraps a UTInstant{Millisecond} and interprets it according to the proleptic Gregorian calendar.

> source

Base.Dates.Date – Type.

> Date

Date wraps a UTInstant{Day} and interprets it according to the proleptic Gregorian calendar.

> source

Base.Dates.Time – Type.

> Time

Time wraps a Nanosecond and represents a specific moment in a 24-hour day.

> source

59.2   Dates Functions

All Dates functions are defined in the Dates module; note that only the Date, DateTime, and now functions are exported; to use all other Dates functions, you'll need to prefix each function call with an explicit Dates., e.g. Dates.dayofweek(dt). Alternatively, you can write using Base.Dates to bring all exported functions into Main to be used without the Dates. prefix.

Base.Dates.DateTime – Method.

> DateTime(y, [m, d, h, mi, s, ms]) -> DateTime

Construct a DateTime type by parts. Arguments must be convertible to Int64.

> source

Base.Dates.DateTime – Method.

> DateTime(periods::Period...) -> DateTime

Construct a `DateTime` type by `Period` type parts. Arguments may be in any order. `DateTime` parts not provided will default to the value of `Dates.default(period)`.

source

Base.Dates.DateTime – Method.

```
DateTime(f::Function, y[, m, d, h, mi, s]; step=Day(1), limit=10000) -> DateTime
```

Create a `DateTime` through the adjuster API. The starting point will be constructed from the provided `y`, `m`, `d...` arguments, and will be adjusted until `f::Function` returns `true`. The step size in adjusting can be provided manually through the `step` keyword. `limit` provides a limit to the max number of iterations the adjustment API will pursue before throwing an error (in the case that `f::Function` is never satisfied).

source

Base.Dates.DateTime – Method.

```
DateTime(dt::Date) -> DateTime
```

Converts a `Date` to a `DateTime`. The hour, minute, second, and millisecond parts of the new `DateTime` are assumed to be zero.

source

Base.Dates.DateTime – Method.

```
DateTime(dt::AbstractString, format::AbstractString; locale="english") -> DateTime
```

Construct a `DateTime` by parsing the `dt` date string following the pattern given in the `format` string.

This method creates a `DateFormat` object each time it is called. If you are parsing many date strings of the same format, consider creating a `DateFormat` object once and using that as the second argument instead.

source

Base.Dates.format – Function.

```
format(io::IO, tok::AbstractDateToken, dt::TimeType, locale)
```

Format the `tok` token from `dt` and write it to `io`. The formatting can be based on `locale`.

All subtypes of `AbstractDateToken` must define this method in order to be able to print a Date / DateTime object according to a `DateFormat` containing that token.

source

Base.Dates.DateFormat – Type.

```
DateFormat(format::AbstractString, locale="english") -> DateFormat
```

Construct a date formatting object that can be used for parsing date strings or formatting a date object as a string. The following character codes can be used to construct the `format` string:

Characters not listed above are normally treated as delimiters between date and time slots. For example a `dt` string of "1996-01-15T00:00:00.0" would have a `format` string like "y-m-dTH:M:S.s". If you need to use a code character as a delimiter you can escape it using backslash. The date "1995y01m" would have the format "y\ym\m".

Creating a DateFormat object is expensive. Whenever possible, create it once and use it many times or try the `dateformat""` string macro. Using this macro creates the DateFormat object once at macro expansion time and reuses it later. see `@dateformat_str`.

Code	Matches	Comment
y	1996, 96	Returns year of 1996, 0096
Y	1996, 96	Returns year of 1996, 0096. Equivalent to y
m	1, 01	Matches 1 or 2-digit months
u	Jan	Matches abbreviated months according to the `locale` keyword
U	January	Matches full month names according to the `locale` keyword
d	1, 01	Matches 1 or 2-digit days
H	00	Matches hours
M	00	Matches minutes
S	00	Matches seconds
s	.500	Matches milliseconds
e	Mon, Tues	Matches abbreviated days of the week
E	Monday	Matches full name days of the week
yyyymmdd	19960101	Matches fixed-width year, month, and day

See DateTime and format for how to use a DateFormat object to parse and write Date strings respectively.

source

Base.Dates.@dateformat_str – Macro.

```
dateformat"Y-m-d H:M:S"
```

Create a DateFormat object. Similar to `DateFormat("Y-m-d H:M:S")` but creates the DateFormat object once during macro expansion.

See DateFormat for details about format specifiers.

source

Base.Dates.DateTime – Method.

```
DateTime(dt::AbstractString, df::DateFormat) -> DateTime
```

Construct a `DateTime` by parsing the `dt` date string following the pattern given in the DateFormat object. Similar to `DateTime(::AbstractString, ::AbstractString)` but more efficient when repeatedly parsing similarly formatted date strings with a pre-created `DateFormat` object.

source

Base.Dates.Date – Method.

```
Date(y, [m, d]) -> Date
```

Construct a `Date` type by parts. Arguments must be convertible to Int64.

source

Base.Dates.Date – Method.

```
Date(period::Period...) -> Date
```

Construct a `Date` type by `Period` type parts. Arguments may be in any order. `Date` parts not provided will default to the value of `Dates.default(period)`.

source

Base.Dates.Date – Method.

```
Date(f::Function, y[, m, d]; step=Day(1), limit=10000) -> Date
```

Create a `Date` through the adjuster API. The starting point will be constructed from the provided `y`, `m`, `d` arguments, and will be adjusted until `f::Function` returns `true`. The step size in adjusting can be provided manually through the `step` keyword. `limit` provides a limit to the max number of iterations the adjustment API will pursue before throwing an error (given that `f::Function` is never satisfied).

source

Base.Dates.Date – Method.

```
Date(dt::DateTime) -> Date
```

Converts a `DateTime` to a `Date`. The hour, minute, second, and millisecond parts of the `DateTime` are truncated, so only the year, month and day parts are used in construction.

source

Base.Dates.Date – Method.

```
Date(dt::AbstractString, format::AbstractString; locale="english") -> Date
```

Construct a `Date` object by parsing a `dt` date string following the pattern given in the `format` string. Follows the same conventions as `DateTime(::AbstractString, ::AbstractString)`.

source

Base.Dates.Date – Method.

```
Date(dt::AbstractString, df::DateFormat) -> Date
```

Parse a date from a date string `dt` using a `DateFormat` object `df`.

source

Base.Dates.Time – Method.

```
Time(h, [mi, s, ms, us, ns]) -> Time
```

Construct a `Time` type by parts. Arguments must be convertible to Int64.

source

Base.Dates.Time – Method.

```
Time(period::TimePeriod...) -> Time
```

Construct a `Time` type by `Period` type parts. Arguments may be in any order. `Time` parts not provided will default to the value of `Dates.default(period)`.

source

Base.Dates.Time – Method.

```
Time(f::Function, h, mi=0; step::Period=Second(1), limit::Int=10000)
Time(f::Function, h, mi, s; step::Period=Millisecond(1), limit::Int=10000)
Time(f::Function, h, mi, s, ms; step::Period=Microsecond(1), limit::Int=10000)
Time(f::Function, h, mi, s, ms, us; step::Period=Nanosecond(1), limit::Int=10000)
```

Create a `Time` through the adjuster API. The starting point will be constructed from the provided `h, mi, s, ms, us` arguments, and will be adjusted until `f::Function` returns `true`. The step size in adjusting can be provided manually through the `step` keyword. `limit` provides a limit to the max number of iterations the adjustment API will pursue before throwing an error (in the case that `f::Function` is never satisfied). Note that the default step will adjust to allow for greater precision for the given arguments; i.e. if hour, minute, and second arguments are provided, the default step will be `Millisecond(1)` instead of `Second(1)`.

source

`Base.Dates.Time` – Method.

```
Time(dt::DateTime) -> Time
```

Converts a `DateTime` to a `Time`. The hour, minute, second, and millisecond parts of the `DateTime` are used to create the new `Time`. Microsecond and nanoseconds are zero by default.

source

`Base.Dates.now` – Method.

```
now() -> DateTime
```

Returns a `DateTime` corresponding to the user's system time including the system timezone locale.

source

`Base.Dates.now` – Method.

```
now(::Type{UTC}) -> DateTime
```

Returns a `DateTime` corresponding to the user's system time as UTC/GMT.

source

`Base.eps` – Function.

```
eps(::DateTime) -> Millisecond
eps(::Date) -> Day
eps(::Time) -> Nanosecond
```

Returns `Millisecond(1)` for `DateTime` values, `Day(1)` for `Date` values, and `Nanosecond(1)` for `Time` values.

source

Accessor Functions

`Base.Dates.year` – Function.

```
year(dt::TimeType) -> Int64
```

The year of a `Date` or `DateTime` as an `Int64`.

source

`Base.Dates.month` – Function.

```
month(dt::TimeType) -> Int64
```

The month of a `Date` or `DateTime` as an `Int64`.

source

Base.Dates.week – Function.

```
week(dt::TimeType) -> Int64
```

Return the ISO week date of a Date or DateTime as an Int64. Note that the first week of a year is the week that contains the first Thursday of the year which can result in dates prior to January 4th being in the last week of the previous year. For example week(Date(2005,1,1)) is the 53rd week of 2004.

source

Base.Dates.day – Function.

```
day(dt::TimeType) -> Int64
```

The day of month of a Date or DateTime as an Int64.

source

Base.Dates.hour – Function.

```
hour(dt::DateTime) -> Int64
```

The hour of day of a DateTime as an Int64.

source

```
hour(t::Time) -> Int64
```

The hour of a Time as an Int64.

source

Base.Dates.minute – Function.

```
minute(dt::DateTime) -> Int64
```

The minute of a DateTime as an Int64.

source

```
minute(t::Time) -> Int64
```

The minute of a Time as an Int64.

source

Base.Dates.second – Function.

```
second(dt::DateTime) -> Int64
```

The second of a DateTime as an Int64.

source

```
second(t::Time) -> Int64
```

The second of a Time as an Int64.

source

Base.Dates.millisecond – Function.

```
millisecond(dt::DateTime) -> Int64
```

The millisecond of a `DateTime` as an `Int64`.

source

```
millisecond(t::Time) -> Int64
```

The millisecond of a `Time` as an `Int64`.

source

`Base.Dates.microsecond` – Function.

```
microsecond(t::Time) -> Int64
```

The microsecond of a `Time` as an `Int64`.

source

`Base.Dates.nanosecond` – Function.

```
nanosecond(t::Time) -> Int64
```

The nanosecond of a `Time` as an `Int64`.

source

`Base.Dates.Year` – Method.

```
Year(v)
```

Construct a `Year` object with the given v value. Input must be losslessly convertible to an `Int64`.

source

`Base.Dates.Month` – Method.

```
Month(v)
```

Construct a `Month` object with the given v value. Input must be losslessly convertible to an `Int64`.

source

`Base.Dates.Week` – Method.

```
Week(v)
```

Construct a `Week` object with the given v value. Input must be losslessly convertible to an `Int64`.

source

`Base.Dates.Day` – Method.

```
Day(v)
```

Construct a `Day` object with the given v value. Input must be losslessly convertible to an `Int64`.

source

`Base.Dates.Hour` – Method.

```
Hour(dt::DateTime) -> Hour
```

The hour part of a DateTime as a `Hour`.

source

Base.Dates.Minute – Method.

```
Minute(dt::DateTime) -> Minute
```

The minute part of a DateTime as a `Minute`.

source

Base.Dates.Second – Method.

```
Second(dt::DateTime) -> Second
```

The second part of a DateTime as a `Second`.

source

Base.Dates.Millisecond – Method.

```
Millisecond(dt::DateTime) -> Millisecond
```

The millisecond part of a DateTime as a `Millisecond`.

source

Base.Dates.Microsecond – Method.

```
Microsecond(dt::Time) -> Microsecond
```

The microsecond part of a Time as a `Microsecond`.

source

Base.Dates.Nanosecond – Method.

```
Nanosecond(dt::Time) -> Nanosecond
```

The nanosecond part of a Time as a `Nanosecond`.

source

Base.Dates.yearmonth – Function.

```
yearmonth(dt::TimeType) -> (Int64, Int64)
```

Simultaneously return the year and month parts of a `Date` or `DateTime`.

source

Base.Dates.monthday – Function.

```
monthday(dt::TimeType) -> (Int64, Int64)
```

Simultaneously return the month and day parts of a `Date` or `DateTime`.

source

Base.Dates.yearmonthday – Function.

```
yearmonthday(dt::TimeType) -> (Int64, Int64, Int64)
```

Simultaneously return the year, month and day parts of a `Date` or `DateTime`.

source

Query Functions

Base.Dates.dayname – Function.

```
dayname(dt::TimeType; locale="english") -> AbstractString
```

Return the full day name corresponding to the day of the week of the Date or DateTime in the given locale.

source

Base.Dates.dayabbr – Function.

```
dayabbr(dt::TimeType; locale="english") -> AbstractString
```

Return the abbreviated name corresponding to the day of the week of the Date or DateTime in the given locale.

source

Base.Dates.dayofweek – Function.

```
dayofweek(dt::TimeType) -> Int64
```

Returns the day of the week as an Int64 with 1 = Monday, 2 = Tuesday, etc..

source

Base.Dates.dayofmonth – Function.

```
dayofmonth(dt::TimeType) -> Int64
```

The day of month of a Date or DateTime as an Int64.

source

Base.Dates.dayofweekofmonth – Function.

```
dayofweekofmonth(dt::TimeType) -> Int
```

For the day of week of dt, returns which number it is in dt's month. So if the day of the week of dt is Monday, then 1 = First Monday of the month, 2 = Second Monday of the month, etc. In the range 1:5.

source

Base.Dates.daysofweekinmonth – Function.

```
daysofweekinmonth(dt::TimeType) -> Int
```

For the day of week of dt, returns the total number of that day of the week in dt's month. Returns 4 or 5. Useful in temporal expressions for specifying the last day of a week in a month by including dayofweekofmonth(dt) == daysofweekinmonth(dt) in the adjuster function.

source

Base.Dates.monthname – Function.

```
monthname(dt::TimeType; locale="english") -> AbstractString
```

Return the full name of the month of the Date or DateTime in the given locale.

source

Base.Dates.monthabbr – Function.

```
monthabbr(dt::TimeType; locale="english") -> AbstractString
```

Return the abbreviated month name of the Date or DateTime in the given locale.

source

Base.Dates.daysinmonth – Function.

```
daysinmonth(dt::TimeType) -> Int
```

Returns the number of days in the month of dt. Value will be 28, 29, 30, or 31.

source

Base.Dates.isleapyear – Function.

```
isleapyear(dt::TimeType) -> Bool
```

Returns true if the year of dt is a leap year.

source

Base.Dates.dayofyear – Function.

```
dayofyear(dt::TimeType) -> Int
```

Returns the day of the year for dt with January 1st being day 1.

source

Base.Dates.daysinyear – Function.

```
daysinyear(dt::TimeType) -> Int
```

Returns 366 if the year of dt is a leap year, otherwise returns 365.

source

Base.Dates.quarterofyear – Function.

```
quarterofyear(dt::TimeType) -> Int
```

Returns the quarter that dt resides in. Range of value is 1:4.

source

Base.Dates.dayofquarter – Function.

```
dayofquarter(dt::TimeType) -> Int
```

Returns the day of the current quarter of dt. Range of value is 1:92.

source

Adjuster Functions

Base.trunc – Method.

```
trunc(dt::TimeType, ::Type{Period}) -> TimeType
```

Truncates the value of dt according to the provided Period type. E.g. if dt is 1996-01-01T12:30:00, then trunc(dt,Day) == 1996-01-01T00:00:00.

> source

Base.Dates.firstdayofweek – Function.

```
firstdayofweek(dt::TimeType) -> TimeType
```

Adjusts dt to the Monday of its week.

> source

Base.Dates.lastdayofweek – Function.

```
lastdayofweek(dt::TimeType) -> TimeType
```

Adjusts dt to the Sunday of its week.

> source

Base.Dates.firstdayofmonth – Function.

```
firstdayofmonth(dt::TimeType) -> TimeType
```

Adjusts dt to the first day of its month.

> source

Base.Dates.lastdayofmonth – Function.

```
lastdayofmonth(dt::TimeType) -> TimeType
```

Adjusts dt to the last day of its month.

> source

Base.Dates.firstdayofyear – Function.

```
firstdayofyear(dt::TimeType) -> TimeType
```

Adjusts dt to the first day of its year.

> source

Base.Dates.lastdayofyear – Function.

```
lastdayofyear(dt::TimeType) -> TimeType
```

Adjusts dt to the last day of its year.

> source

Base.Dates.firstdayofquarter – Function.

```
firstdayofquarter(dt::TimeType) -> TimeType
```

Adjusts dt to the first day of its quarter.

source

Base.Dates.lastdayofquarter – Function.

```
lastdayofquarter(dt::TimeType) -> TimeType
```

Adjusts dt to the last day of its quarter.

source

Base.Dates.tonext – Method.

```
tonext(dt::TimeType, dow::Int; same::Bool=false) -> TimeType
```

Adjusts dt to the next day of week corresponding to dow with 1 = Monday, 2 = Tuesday, etc. Setting same=true allows the current dt to be considered as the next dow, allowing for no adjustment to occur.

source

Base.Dates.toprev – Method.

```
toprev(dt::TimeType, dow::Int; same::Bool=false) -> TimeType
```

Adjusts dt to the previous day of week corresponding to dow with 1 = Monday, 2 = Tuesday, etc. Setting same=true allows the current dt to be considered as the previous dow, allowing for no adjustment to occur.

source

Base.Dates.tofirst – Function.

```
tofirst(dt::TimeType, dow::Int; of=Month) -> TimeType
```

Adjusts dt to the first dow of its month. Alternatively, of=Year will adjust to the first dow of the year.

source

Base.Dates.tolast – Function.

```
tolast(dt::TimeType, dow::Int; of=Month) -> TimeType
```

Adjusts dt to the last dow of its month. Alternatively, of=Year will adjust to the last dow of the year.

source

Base.Dates.tonext – Method.

```
tonext(func::Function, dt::TimeType; step=Day(1), limit=10000, same=false) -> TimeType
```

Adjusts dt by iterating at most limit iterations by step increments until func returns true. func must take a single TimeType argument and return a Bool. same allows dt to be considered in satisfying func.

source

Base.Dates.toprev – Method.

```
toprev(func::Function, dt::TimeType; step=Day(-1), limit=10000, same=false) -> TimeType
```

Adjusts dt by iterating at most limit iterations by step increments until func returns true. func must take a single TimeType argument and return a Bool. same allows dt to be considered in satisfying func.

source

Periods

`Base.Dates.Period` – Method.

```
Year(v)
Month(v)
Week(v)
Day(v)
Hour(v)
Minute(v)
Second(v)
Millisecond(v)
Microsecond(v)
Nanosecond(v)
```

Construct a `Period` type with the given v value. Input must be losslessly convertible to an `Int64`.

source

`Base.Dates.CompoundPeriod` – Method.

```
CompoundPeriod(periods) -> CompoundPeriod
```

Construct a `CompoundPeriod` from a `Vector` of `Periods`. All `Periods` of the same type will be added together.

Examples

```
julia> Dates.CompoundPeriod(Dates.Hour(12), Dates.Hour(13))
25 hours

julia> Dates.CompoundPeriod(Dates.Hour(-1), Dates.Minute(1))
-1 hour, 1 minute

julia> Dates.CompoundPeriod(Dates.Month(1), Dates.Week(-2))
1 month, -2 weeks

julia> Dates.CompoundPeriod(Dates.Minute(50000))
50000 minutes
```

source

`Base.Dates.default` – Function.

```
default(p::Period) -> Period
```

Returns a sensible "default" value for the input Period by returning `T(1)` for Year, Month, and Day, and `T(0)` for Hour, Minute, Second, and Millisecond.

source

Rounding Functions

`Date` and `DateTime` values can be rounded to a specified resolution (e.g., 1 month or 15 minutes) with `floor`, `ceil`, or `round`.

`Base.floor` – Method.

```
floor(dt::TimeType, p::Period) -> TimeType
```

Returns the nearest `Date` or `DateTime` less than or equal to `dt` at resolution `p`.

For convenience, p may be a type instead of a value: `floor(dt, Dates.Hour)` is a shortcut for `floor(dt, Dates.Hour(1))`.

```julia
julia> floor(Date(1985, 8, 16), Dates.Month)
1985-08-01

julia> floor(DateTime(2013, 2, 13, 0, 31, 20), Dates.Minute(15))
2013-02-13T00:30:00

julia> floor(DateTime(2016, 8, 6, 12, 0, 0), Dates.Day)
2016-08-06T00:00:00
```

source

`Base.ceil` – Method.

```
ceil(dt::TimeType, p::Period) -> TimeType
```

Returns the nearest `Date` or `DateTime` greater than or equal to `dt` at resolution `p`.

For convenience, p may be a type instead of a value: `ceil(dt, Dates.Hour)` is a shortcut for `ceil(dt, Dates.Hour(1))`.

```julia
julia> ceil(Date(1985, 8, 16), Dates.Month)
1985-09-01

julia> ceil(DateTime(2013, 2, 13, 0, 31, 20), Dates.Minute(15))
2013-02-13T00:45:00

julia> ceil(DateTime(2016, 8, 6, 12, 0, 0), Dates.Day)
2016-08-07T00:00:00
```

source

`Base.round` – Method.

```
round(dt::TimeType, p::Period, [r::RoundingMode]) -> TimeType
```

Returns the `Date` or `DateTime` nearest to `dt` at resolution `p`. By default (RoundNearestTiesUp), ties (e.g., rounding 9:30 to the nearest hour) will be rounded up.

For convenience, p may be a type instead of a value: `round(dt, Dates.Hour)` is a shortcut for `round(dt, Dates.Hour(1))`.

```julia
julia> round(Date(1985, 8, 16), Dates.Month)
1985-08-01

julia> round(DateTime(2013, 2, 13, 0, 31, 20), Dates.Minute(15))
2013-02-13T00:30:00

julia> round(DateTime(2016, 8, 6, 12, 0, 0), Dates.Day)
2016-08-07T00:00:00
```

Valid rounding modes for `round(::TimeType, ::Period, ::RoundingMode)` are RoundNearestTiesUp (default), RoundDown (`floor`), and RoundUp (`ceil`).

source

The following functions are not exported:

`Base.Dates.floorceil` – Function.

```
floorceil(dt::TimeType, p::Period) -> (TimeType, TimeType)
```

Simultaneously return the `floor` and `ceil` of a `Date` or `DateTime` at resolution p. More efficient than calling both `floor` and `ceil` individually.

source

`Base.Dates.epochdays2date` – Function.

```
epochdays2date(days) -> Date
```

Takes the number of days since the rounding epoch (`0000-01-01T00:00:00`) and returns the corresponding `Date`.

source

`Base.Dates.epochms2datetime` – Function.

```
epochms2datetime(milliseconds) -> DateTime
```

Takes the number of milliseconds since the rounding epoch (`0000-01-01T00:00:00`) and returns the corresponding `DateTime`.

source

`Base.Dates.date2epochdays` – Function.

```
date2epochdays(dt::Date) -> Int64
```

Takes the given `Date` and returns the number of days since the rounding epoch (`0000-01-01T00:00:00`) as an `Int64`.

source

`Base.Dates.datetime2epochms` – Function.

```
datetime2epochms(dt::DateTime) -> Int64
```

Takes the given `DateTime` and returns the number of milliseconds since the rounding epoch (`0000-01-01T00:00:00`) as an `Int64`.

source

Conversion Functions

`Base.Dates.today` – Function.

```
today() -> Date
```

Returns the date portion of `now()`.

source

`Base.Dates.unix2datetime` – Function.

```
unix2datetime(x) -> DateTime
```

Takes the number of seconds since unix epoch `1970-01-01T00:00:00` and converts to the corresponding `Date-Time`.

source

`Base.Dates.datetime2unix` – Function.

```
datetime2unix(dt::DateTime) -> Float64
```

Takes the given `DateTime` and returns the number of seconds since the unix epoch `1970-01-01T00:00:00` as a `Float64`.

source

`Base.Dates.julian2datetime` – Function.

```
julian2datetime(julian_days) -> DateTime
```

Takes the number of Julian calendar days since epoch `-4713-11-24T12:00:00` and returns the corresponding `DateTime`.

source

`Base.Dates.datetime2julian` – Function.

```
datetime2julian(dt::DateTime) -> Float64
```

Takes the given `DateTime` and returns the number of Julian calendar days since the julian epoch `-4713-11-24T12:00:00` as a `Float64`.

source

`Base.Dates.rata2datetime` – Function.

```
rata2datetime(days) -> DateTime
```

Takes the number of Rata Die days since epoch `0000-12-31T00:00:00` and returns the corresponding `Date-Time`.

source

`Base.Dates.datetime2rata` – Function.

```
datetime2rata(dt::TimeType) -> Int64
```

Returns the number of Rata Die days since epoch from the given `Date` or `DateTime`.

source

Constants

Days of the Week:

Months of the Year:

Variable	Abbr.	Value (Int)
Monday	Mon	1
Tuesday	Tue	2
Wednesday	Wed	3
Thursday	Thu	4
Friday	Fri	5
Saturday	Sat	6
Sunday	Sun	7

Variable	Abbr.	Value (Int)
January	Jan	1
February	Feb	2
March	Mar	3
April	Apr	4
May	May	5
June	Jun	6
July	Jul	7
August	Aug	8
September	Sep	9
October	Oct	10
November	Nov	11
December	Dec	12

Chapter 60

Iteration utilities

– Function.

```
zip(iters...)
```

For a set of iterable objects, returns an iterable of tuples, where the `i`th tuple contains the `i`th component of each input iterable.

Note that `zip` is its own inverse: `collect(zip(zip(a...)...)) == collect(a)`.

```
julia> a = 1:5
1:5

julia> b = ["e","d","b","c","a"]
5-element Array{String,1}:
 "e"
 "d"
 "b"
 "c"
 "a"

julia> c = zip(a,b)
Base.Iterators.Zip2{UnitRange{Int64},Array{String,1}}(1:5, String["e", "d", "b", "c", "a"])

julia> length(c)
5

julia> first(c)
(1, "e")
```

source

– Function.

```
enumerate(iter)
```

An iterator that yields `(i, x)` where `i` is a counter starting at 1, and `x` is the `i`th value from the given iterator. It's useful when you need not only the values x over which you are iterating, but also the number of iterations so far. Note that `i` may not be valid for indexing `iter`; it's also possible that `x != iter[i]`, if `iter` has indices that do not start at 1. See the `enumerate(IndexLinear(), iter)` method if you want to ensure that `i` is an index.

```julia
julia> a = ["a", "b", "c"];

julia> for (index, value) in enumerate(a)
           println("$index $value")
       end
1 a
2 b
3 c
```

```
enumerate(IndexLinear(), A)
enumerate(IndexCartesian(), A)
enumerate(IndexStyle(A), A)
```

An iterator that accesses each element of the array A, returning (i, x), where i is the index for the element and x = A[i]. This is similar to enumerate(A), except i will always be a valid index for A.

Specifying IndexLinear() ensures that i will be an integer; specifying IndexCartesian() ensures that i will be a CartesianIndex; specifying IndexStyle(A) chooses whichever has been defined as the native indexing style for array A.

```julia
julia> A = ["a" "d"; "b" "e"; "c" "f"];

julia> for (index, value) in enumerate(IndexStyle(A), A)
           println("$index $value")
       end
1 a
2 b
3 c
4 d
5 e
6 f

julia> S = view(A, 1:2, :);

julia> for (index, value) in enumerate(IndexStyle(S), S)
           println("$index $value")
       end
CartesianIndex{2}((1, 1)) a
CartesianIndex{2}((2, 1)) b
CartesianIndex{2}((1, 2)) d
CartesianIndex{2}((2, 2)) e
```

Note that enumerate(A) returns i as a *counter* (always starting at 1), whereas enumerate(IndexLinear(), A) returns i as an *index* (starting at the first linear index of A, which may or may not be 1).

See also: IndexStyle, indices.

Base.Iterators.rest – Function.

```
rest(iter, state)
```

An iterator that yields the same elements as iter, but starting at the given state.

Base.Iterators.countfrom – Function.

```
countfrom(start=1, step=1)
```

An iterator that counts forever, starting at `start` and incrementing by `step`.

source

Base.Iterators.take – Function.

```
take(iter, n)
```

An iterator that generates at most the first n elements of `iter`.

```julia
julia> a = 1:2:11
1:2:11

julia> collect(a)
6-element Array{Int64,1}:
  1
  3
  5
  7
  9
 11

julia> collect(Iterators.take(a,3))
3-element Array{Int64,1}:
 1
 3
 5
```

source

Base.Iterators.drop – Function.

```
drop(iter, n)
```

An iterator that generates all but the first n elements of `iter`.

```julia
julia> a = 1:2:11
1:2:11

julia> collect(a)
6-element Array{Int64,1}:
  1
  3
  5
  7
  9
 11

julia> collect(Iterators.drop(a,4))
2-element Array{Int64,1}:
  9
 11
```

source

Base.Iterators.cycle – Function.

```
cycle(iter)
```

An iterator that cycles through `iter` forever.

source

Base.Iterators.repeated – Function.

```
repeated(x[, n::Int])
```

An iterator that generates the value x forever. If n is specified, generates x that many times (equivalent to `take(repeated(x), n)`).

```
julia> a = Iterators.repeated([1 2], 4);

julia> collect(a)
4-element Array{Array{Int64,2},1}:
 [1 2]
 [1 2]
 [1 2]
 [1 2]
```

source

Base.Iterators.product – Function.

```
product(iters...)
```

Returns an iterator over the product of several iterators. Each generated element is a tuple whose `ith` element comes from the `ith` argument iterator. The first iterator changes the fastest. Example:

```
julia> collect(Iterators.product(1:2,3:5))
2×3 Array{Tuple{Int64,Int64},2}:
 (1, 3)  (1, 4)  (1, 5)
 (2, 3)  (2, 4)  (2, 5)
```

source

Base.Iterators.flatten – Function.

```
flatten(iter)
```

Given an iterator that yields iterators, return an iterator that yields the elements of those iterators. Put differently, the elements of the argument iterator are concatenated. Example:

```
julia> collect(Iterators.flatten((1:2, 8:9)))
4-element Array{Int64,1}:
 1
 2
 8
 9
```

source

Base.Iterators.partition – Function.

```
partition(collection, n)
```

Iterate over a collection n elements at a time.

```
julia> collect(Iterators.partition([1,2,3,4,5], 2))
3-element Array{Array{Int64,1},1}:
 [1, 2]
 [3, 4]
 [5]
```

source

Chapter 61

Unit Testing

61.1   Testing Base Julia

Julia is under rapid development and has an extensive test suite to verify functionality across multiple platforms. If you build Julia from source, you can run this test suite with `make  test`. In a binary install, you can run the test suite using `Base.runtests()`.

`Base.runtests` – Function.

```
runtests([tests=["all"] [, numcores=ceil(Int, Sys.CPU_CORES / 2) ]])
```

Run the Julia unit tests listed in `tests`, which can be either a string or an array of strings, using `numcores` processors. (not exported)

source

61.2   Basic Unit Tests

The `Base.Test` module provides simple *unit testing* functionality. Unit testing is a way to see if your code is correct by checking that the results are what you expect. It can be helpful to ensure your code still works after you make changes, and can be used when developing as a way of specifying the behaviors your code should have when complete.

Simple unit testing can be performed with the `@test()` and `@test_throws()` macros:

`Base.Test.@test` – Macro.

```
@test ex
@test f(args...) key=val ...
```

Tests that the expression ex evaluates to `true`. Returns a `Pass Result` if it does, a `Fail Result` if it is `false`, and an `Error Result` if it could not be evaluated.

The `@test f(args...) key=val...` form is equivalent to writing `@test f(args..., key=val...)` which can be useful when the expression is a call using infix syntax such as approximate comparisons:

```
@test a ≈ b atolε=
```

This is equivalent to the uglier test `@test ≈(a, b, atol=ε)`. It is an error to supply more than one expression unless the first is a call expression and the rest are assignments (k=v).

source

`Base.Test.@test_throws` – Macro.

```
@test_throws extype ex
```

Tests that the expression ex throws an exception of type `extype`. Note that `@test_throws` does not support a trailing keyword form.

source

For example, suppose we want to check our new function `foo(x)` works as expected:

```
julia> using Base.Test

julia> foo(x) = length(x)^2
foo (generic function with 1 method)
```

If the condition is true, a `Pass` is returned:

```
julia> @test foo("bar") == 9
Test Passed
  Expression: foo("bar") == 9
   Evaluated: 9 == 9

julia> @test foo("fizz") >= 10
Test Passed
  Expression: foo("fizz") >= 10
   Evaluated: 16 >= 10
```

If the condition is false, then a `Fail` is returned and an exception is thrown:

```
julia> @test foo("f") == 20
Test Failed
  Expression: foo("f") == 20
   Evaluated: 1 == 20
ERROR: There was an error during testing
 in record at test.jl:268
 in do_test at test.jl:191
```

If the condition could not be evaluated because an exception was thrown, which occurs in this case because `length()` is not defined for symbols, an `Error` object is returned and an exception is thrown:

```
julia> @test foo(:cat) == 1
Error During Test
  Test threw an exception of type MethodError
  Expression: foo(:cat) == 1
  MethodError: `length` has no method matching length(::Symbol)
   in foo at none:1
   in anonymous at test.jl:159
   in do_test at test.jl:180
ERROR: There was an error during testing
 in record at test.jl:268
 in do_test at test.jl:191
```

If we expect that evaluating an expression *should* throw an exception, then we can use `@test_throws()` to check that this occurs:

```
julia> @test_throws MethodError foo(:cat)
Test Passed
  Expression: foo(:cat)
   Evaluated: MethodError
```

61.3 Working with Test Sets

Typically a large number of tests are used to make sure functions work correctly over a range of inputs. In the event a test fails, the default behavior is to throw an exception immediately. However, it is normally preferable to run the rest of the tests first to get a better picture of how many errors there are in the code being tested.

The `@testset()` macro can be used to group tests into *sets*. All the tests in a test set will be run, and at the end of the test set a summary will be printed. If any of the tests failed, or could not be evaluated due to an error, the test set will then throw a `TestSetException`.

`Base.Test.@testset` – Macro.

```
@testset [CustomTestSet] [option=val  ...] ["description"] begin ... end
@testset [CustomTestSet] [option=val  ...] ["description $v"] for v in (...) ... end
@testset [CustomTestSet] [option=val  ...] ["description $v, $w"] for v in (...), w in (...)
      ... end
```

Starts a new test set, or multiple test sets if a `for` loop is provided.

If no custom testset type is given it defaults to creating a `DefaultTestSet`. `DefaultTestSet` records all the results and, if there are any `Fails` or `Errors`, throws an exception at the end of the top-level (non-nested) test set, along with a summary of the test results.

Any custom testset type (subtype of `AbstractTestSet`) can be given and it will also be used for any nested `@testset` invocations. The given options are only applied to the test set where they are given. The default test set type does not take any options.

The description string accepts interpolation from the loop indices. If no description is provided, one is constructed based on the variables.

By default the `@testset` macro will return the testset object itself, though this behavior can be customized in other testset types. If a `for` loop is used then the macro collects and returns a list of the return values of the `finish` method, which by default will return a list of the testset objects used in each iteration.

source

We can put our tests for the `foo(x)` function in a test set:

```
julia> @testset "Foo Tests" begin
           @test foo("a")   == 1
           @test foo("ab")  == 4
           @test foo("abc") == 9
       end
Test Summary: | Pass  Total
Foo Tests     |    3      3
```

Test sets can also be nested:

```
julia> @testset "Foo Tests" begin
           @testset "Animals" begin
               @test foo("cat") == 9
```

```
                @test foo("dog") == foo("cat")
            end
            @testset "Arrays $i" for i in 1:3
                @test foo(zeros(i)) == i^2
                @test foo(ones(i)) == i^2
            end
        end
Test Summary: | Pass  Total
Foo Tests     |   8      8
```

In the event that a nested test set has no failures, as happened here, it will be hidden in the summary. If we do have a test failure, only the details for the failed test sets will be shown:

```
julia> @testset "Foo Tests" begin
            @testset "Animals" begin
                @testset "Felines" begin
                    @test foo("cat") == 9
                end
                @testset "Canines" begin
                    @test foo("dog") == 9
                end
            end
            @testset "Arrays" begin
                @test foo(zeros(2)) == 4
                @test foo(ones(4)) == 15
            end
        end

Arrays: Test Failed
  Expression: foo(ones(4)) == 15
   Evaluated: 16 == 15
 in record at test.jl:297
 in do_test at test.jl:191
Test Summary: | Pass  Fail  Total
Foo Tests     |   3     1      4
  Animals     |   2            2
  Arrays      |   1     1      2
ERROR: Some tests did not pass: 3 passed, 1 failed, 0 errored, 0 broken.
 in finish at test.jl:362
```

61.4   Other Test Macros

As calculations on floating-point values can be imprecise, you can perform approximate equality checks using either @test a ≈ b (where ≈, typed via tab completion of \approx, is the isapprox() function) or use isapprox() directly.

```
julia> @test 1 ≈ 0.999999999

julia> @test 1 ≈ 0.999999
ERROR: test failed: 1 isapprox 0.999999
 in expression: 1 ≈ 0.999999
 in error at error.jl:21
 in default_handler at test.jl:30
 in do_test at test.jl:53
```

```
@inferred f(x)
```

Tests that the call expression `f(x)` returns a value of the same type inferred by the compiler. It is useful to check for type stability.

`f(x)` can be any call expression. Returns the result of `f(x)` if the types match, and an `Error Result` if it finds different types.

```
julia> using Base.Test

julia> f(a,b,c) = b > 1 ? 1 : 1.0
f (generic function with 1 method)

julia> typeof(f(1,2,3))
Int64

julia> @code_warntype f(1,2,3)
Variables:
  #self#::#f
  a::Int64
  b::Int64
  c::Int64

Body:
  begin
      unless (Base.slt_int)(1, b::Int64)::Bool goto 3
      return 1
      3:
      return 1.0
  end::UNION{FLOAT64, INT64}

julia> @inferred f(1,2,3)
ERROR: return type Int64 does not match inferred return type Union{Float64, Int64}
Stacktrace:
 [1] error(::String) at ./error.jl:21

julia> @inferred max(1,2)
2
```

source

```
@test_warn msg expr
```

Test whether evaluating `expr` results in `STDERR` output that contains the `msg` string or matches the `msg` regular expression. If `msg` is a boolean function, tests whether `msg(output)` returns `true`. If `msg` is a tuple or array, checks that the error output contains/matches each item in `msg`. Returns the result of evaluating `expr`.

See also `@test_nowarn` to check for the absence of error output.

source

```
@test_nowarn expr
```

Test whether evaluating `expr` results in empty STDERR output (no warnings or other messages). Returns the result of evaluating `expr`.

source

61.5   Broken Tests

If a test fails consistently it can be changed to use the `@test_broken()` macro. This will denote the test as `Broken` if the test continues to fail and alerts the user via an `Error` if the test succeeds.

Base.Test.@test_broken – Macro.

```
@test_broken ex
@test_broken f(args...) key=val ...
```

Indicates a test that should pass but currently consistently fails. Tests that the expression `ex` evaluates to `false` or causes an exception. Returns a `Broken Result` if it does, or an `Error Result` if the expression evaluates to `true`.

The `@test_broken f(args...) key=val...` form works as for the `@test` macro.

source

`@test_skip()` is also available to skip a test without evaluation, but counting the skipped test in the test set reporting. The test will not run but gives a `Broken Result`.

Base.Test.@test_skip – Macro.

```
@test_skip ex
@test_skip f(args...) key=val ...
```

Marks a test that should not be executed but should be included in test summary reporting as `Broken`. This can be useful for tests that intermittently fail, or tests of not-yet-implemented functionality.

The `@test_skip f(args...) key=val...` form works as for the `@test` macro.

source

61.6   Creating Custom `AbstractTestSet` Types

Packages can create their own `AbstractTestSet` subtypes by implementing the `record` and `finish` methods. The subtype should have a one-argument constructor taking a description string, with any options passed in as keyword arguments.

Base.Test.record – Function.

```
record(ts::AbstractTestSet, res::Result)
```

Record a result to a testset. This function is called by the `@testset` infrastructure each time a contained `@test` macro completes, and is given the test result (which could be an `Error`). This will also be called with an `Error` if an exception is thrown inside the test block but outside of a `@test` context.

source

Base.Test.finish – Function.

```
finish(ts::AbstractTestSet)
```

Do any final processing necessary for the given testset. This is called by the `@testset` infrastructure after a test block executes. One common use for this function is to record the testset to the parent's results list, using `get_testset`.

source

`Base.Test` takes responsibility for maintaining a stack of nested testsets as they are executed, but any result accumulation is the responsibility of the `AbstractTestSet` subtype. You can access this stack with the `get_testset` and `get_testset_depth` methods. Note that these functions are not exported.

`Base.Test.get_testset` – Function.

```
get_testset()
```

Retrieve the active test set from the task's local storage. If no test set is active, use the fallback default test set.

source

`Base.Test.get_testset_depth` – Function.

```
get_testset_depth()
```

Returns the number of active test sets, not including the defaut test set

source

`Base.Test` also makes sure that nested `@testset` invocations use the same `AbstractTestSet` subtype as their parent unless it is set explicitly. It does not propagate any properties of the testset. Option inheritance behavior can be implemented by packages using the stack infrastructure that `Base.Test` provides.

Defining a basic `AbstractTestSet` subtype might look like:

```julia
import Base.Test: record, finish
using Base.Test: AbstractTestSet, Result, Pass, Fail, Error
using Base.Test: get_testset_depth, get_testset
struct CustomTestSet <: Base.Test.AbstractTestSet
    description::AbstractString
    foo::Int
    results::Vector
    # constructor takes a description string and options keyword arguments
    CustomTestSet(desc; foo=1) = new(desc, foo, [])
end

record(ts::CustomTestSet, child::AbstractTestSet) = push!(ts.results, child)
record(ts::CustomTestSet, res::Result) = push!(ts.results, res)
function finish(ts::CustomTestSet)
    # just record if we're not the top-level parent
    if get_testset_depth() > 0
        record(get_testset(), ts)
    end
    ts
end
```

And using that testset looks like:

```julia
@testset CustomTestSet foo=4 "custom testset inner 2" begin
    # this testset should inherit the type, but not the argument.
```

```
    @testset "custom testset inner" begin
        @test true
    end
end
```

Chapter 62

C Interface

ccall – Keyword.

```
ccall((symbol, library) or function_pointer, ReturnType, (ArgumentType1, ...), ArgumentValue1
    , ...)
```

Call function in C-exported shared library, specified by (`function name`, `library`) tuple, where each component is a string or symbol.

Note that the argument type tuple must be a literal tuple, and not a tuple-valued variable or expression. Alternatively, `ccall` may also be used to call a function pointer, such as one returned by `dlsym`.

Each `ArgumentValue` to the `ccall` will be converted to the corresponding `ArgumentType`, by automatic insertion of calls to `unsafe_convert(ArgumentType, cconvert(ArgumentType, ArgumentValue))`. (See also the documentation for each of these functions for further details.) In most cases, this simply results in a call to `convert(ArgumentType, ArgumentValue)`.

source

Core.Intrinsics.cglobal – Function.

```
cglobal((symbol, library) [, type=Void])
```

Obtain a pointer to a global variable in a C-exported shared library, specified exactly as in ccall. Returns a `Ptr{Type}`, defaulting to `Ptr{Void}` if no Type argument is supplied. The values can be read or written by unsafe_load or unsafe_store!, respectively.

source

Base.cfunction – Function.

```
cfunction(function::Function, ReturnType::Type, (ArgumentTypes...))
```

Generate C-callable function pointer from Julia function. Type annotation of the return value in the callback function is a must for situations where Julia cannot infer the return type automatically.

For example:

```
function foo()
    # body

    retval::Float64
end

bar = cfunction(foo, Float64, ())
```

source

Base.unsafe_convert – Function.

```
unsafe_convert(T,x)
```

Convert x to a value of type T

In cases where convert would need to take a Julia object and turn it into a Ptr, this function should be used to define and perform that conversion.

Be careful to ensure that a Julia reference to x exists as long as the result of this function will be used. Accordingly, the argument x to this function should never be an expression, only a variable name or field reference. For example, x=a.b.c is acceptable, but x=[a,b,c] is not.

The unsafe prefix on this function indicates that using the result of this function after the x argument to this function is no longer accessible to the program may cause undefined behavior, including program corruption or segfaults, at any later time.

source

Base.cconvert – Function.

```
cconvert(T,x)
```

Convert x to a value of type T, typically by calling convert(T,x)

In cases where x cannot be safely converted to T, unlike convert, cconvert may return an object of a type different from T, which however is suitable for unsafe_convert to handle.

Neither convert nor cconvert should take a Julia object and turn it into a Ptr.

source

Base.unsafe_load – Function.

```
unsafe_load(p::Ptr{T}, i::Integer=1)
```

Load a value of type T from the address of the ith element (1-indexed) starting at p. This is equivalent to the C expression p[i-1].

The unsafe prefix on this function indicates that no validation is performed on the pointer p to ensure that it is valid. Incorrect usage may segfault your program or return garbage answers, in the same manner as C.

source

Base.unsafe_store! – Function.

```
unsafe_store!(p::Ptr{T}, x, i::Integer=1)
```

Store a value of type T to the address of the ith element (1-indexed) starting at p. This is equivalent to the C expression p[i-1] = x.

The unsafe prefix on this function indicates that no validation is performed on the pointer p to ensure that it is valid. Incorrect usage may corrupt or segfault your program, in the same manner as C.

source

Base.unsafe_copy! – Method.

```
unsafe_copy!(dest::Ptr{T}, src::Ptr{T}, N)
```

Copy N elements from a source pointer to a destination, with no checking. The size of an element is determined by the type of the pointers.

The `unsafe` prefix on this function indicates that no validation is performed on the pointers `dest` and `src` to ensure that they are valid. Incorrect usage may corrupt or segfault your program, in the same manner as C.

> source

Base.unsafe_copy! – Method.

```
unsafe_copy!(dest::Array, do, src::Array, so, N)
```

Copy N elements from a source array to a destination, starting at offset so in the source and do in the destination (1-indexed).

The `unsafe` prefix on this function indicates that no validation is performed to ensure that N is inbounds on either array. Incorrect usage may corrupt or segfault your program, in the same manner as C.

> source

Base.copy! – Method.

```
copy!(dest, src) -> dest
```

Copy all elements from collection `src` to array `dest`.

> source

Base.copy! – Method.

```
copy!(dest, do, src, so, N)
```

Copy N elements from collection `src` starting at offset so, to array `dest` starting at offset do. Returns `dest`.

> source

Base.pointer – Function.

```
pointer(array [, index])
```

Get the native address of an array or string element. Be careful to ensure that a Julia reference to a exists as long as this pointer will be used. This function is "unsafe" like `unsafe_convert`.

Calling `Ref(array[, index])` is generally preferable to this function.

> source

Base.unsafe_wrap – Method.

```
unsafe_wrap(Array, pointer::Ptr{T}, dims, own=false)
```

Wrap a Julia `Array` object around the data at the address given by `pointer`, without making a copy. The pointer element type T determines the array element type. `dims` is either an integer (for a 1d array) or a tuple of the array dimensions. `own` optionally specifies whether Julia should take ownership of the memory, calling `free` on the pointer when the array is no longer referenced.

This function is labelled "unsafe" because it will crash if `pointer` is not a valid memory address to data of the requested length.

> source

Base.pointer_from_objref – Function.

```
pointer_from_objref(x)
```

Get the memory address of a Julia object as a Ptr. The existence of the resulting Ptr will not protect the object from garbage collection, so you must ensure that the object remains referenced for the whole time that the Ptr will be used.

source

Base.unsafe_pointer_to_objref – Function.

```
unsafe_pointer_to_objref(p::Ptr)
```

Convert a Ptr to an object reference. Assumes the pointer refers to a valid heap-allocated Julia object. If this is not the case, undefined behavior results, hence this function is considered "unsafe" and should be used with care.

source

Base.disable_sigint – Function.

```
disable_sigint(f::Function)
```

Disable Ctrl-C handler during execution of a function on the current task, for calling external code that may call julia code that is not interrupt safe. Intended to be called using do block syntax as follows:

```
disable_sigint() do
    # interrupt-unsafe code
    ...
end
```

This is not needed on worker threads (Threads.threadid() != 1) since the InterruptException will only be delivered to the master thread. External functions that do not call julia code or julia runtime automatically disable sigint during their execution.

source

Base.reenable_sigint – Function.

```
reenable_sigint(f::Function)
```

Re-enable Ctrl-C handler during execution of a function. Temporarily reverses the effect of disable_sigint.

source

Base.systemerror – Function.

```
systemerror(sysfunc, iftrue)
```

Raises a SystemError for errno with the descriptive string sysfunc if iftrue is true

source

Core.Ptr – Type.

```
Ptr{T}
```

A memory address referring to data of type T. However, there is no guarantee that the memory is actually valid, or that it actually represents data of the specified type.

source

`Core.Ref` – Type.

> Ref{T}

An object that safely references data of type T. This type is guaranteed to point to valid, Julia-allocated memory of the correct type. The underlying data is protected from freeing by the garbage collector as long as the Ref itself is referenced.

When passed as a ccall argument (either as a Ptr or Ref type), a Ref object will be converted to a native pointer to the data it references.

There is no invalid (NULL) Ref.

source

`Base.Cchar` – Type.

> Cchar

Equivalent to the native char c-type.

source

`Base.Cuchar` – Type.

> Cuchar

Equivalent to the native unsigned char c-type (UInt8).

source

`Base.Cshort` – Type.

> Cshort

Equivalent to the native signed short c-type (Int16).

source

`Base.Cushort` – Type.

> Cushort

Equivalent to the native unsigned short c-type (UInt16).

source

`Base.Cint` – Type.

> Cint

Equivalent to the native signed int c-type (Int32).

source

`Base.Cuint` – Type.

> Cuint

Equivalent to the native unsigned int c-type (UInt32).

source

Base.Clong – Type.

> | Clong

Equivalent to the native signed long c-type.

> source

Base.Culong – Type.

> | Culong

Equivalent to the native unsigned long c-type.

> source

Base.Clonglong – Type.

> | Clonglong

Equivalent to the native signed long long c-type (Int64).

> source

Base.Culonglong – Type.

> | Culonglong

Equivalent to the native unsigned long long c-type (UInt64).

> source

Base.Cintmax_t – Type.

> | Cintmax_t

Equivalent to the native intmax_t c-type (Int64).

> source

Base.Cuintmax_t – Type.

> | Cuintmax_t

Equivalent to the native uintmax_t c-type (UInt64).

> source

Base.Csize_t – Type.

> | Csize_t

Equivalent to the native size_t c-type (UInt).

> source

Base.Cssize_t – Type.

> | Cssize_t

Equivalent to the native ssize_t c-type.

> source

`Base.Cptrdiff_t` – Type.

> `Cptrdiff_t`

Equivalent to the native `ptrdiff_t` c-type (`Int`).

source

`Base.Cwchar_t` – Type.

> `Cwchar_t`

Equivalent to the native `wchar_t` c-type (`Int32`).

source

`Base.Cfloat` – Type.

> `Cfloat`

Equivalent to the native `float` c-type (`Float32`).

source

`Base.Cdouble` – Type.

> `Cdouble`

Equivalent to the native `double` c-type (`Float64`).

source

Chapter 63

LLVM Interface

Core.Intrinsics.llvmcall – Function.

```
llvmcall(IR::String, ReturnType, (ArgumentType1, ...), ArgumentValue1, ...)
llvmcall((declarations::String, IR::String), ReturnType, (ArgumentType1, ...), ArgumentValue1
    , ...)
```

Call LLVM IR string in the first argument. Similar to an LLVM function `define` block, arguments are available as consecutive unnamed SSA variables (%0, %1, etc.).

The optional declarations string contains external functions declarations that are necessary for llvm to compile the IR string. Multiple declarations can be passed in by separating them with line breaks.

Note that the argument type tuple must be a literal tuple, and not a tuple-valued variable or expression.

Each `ArgumentValue` to `llvmcall` will be converted to the corresponding `ArgumentType`, by automatic insertion of calls to `unsafe_convert(ArgumentType, cconvert(ArgumentType, ArgumentValue))`. (see also the documentation for each of these functions for further details). In most cases, this simply results in a call to `convert(ArgumentType, ArgumentValue)`.

See `test/llvmcall.jl` for usage examples.

source

Chapter 64

C Standard Library

– Function.

```
malloc(size::Integer) -> Ptr{Void}
```

Call `malloc` from the C standard library.

source

– Function.

```
calloc(num::Integer, size::Integer) -> Ptr{Void}
```

Call `calloc` from the C standard library.

source

– Function.

```
realloc(addr::Ptr, size::Integer) -> Ptr{Void}
```

Call `realloc` from the C standard library.

See warning in the documentation for `free` regarding only using this on memory originally obtained from `malloc`.

source

– Function.

```
free(addr::Ptr)
```

Call `free` from the C standard library. Only use this on memory obtained from `malloc`, not on pointers retrieved from other C libraries. `Ptr` objects obtained from C libraries should be freed by the free functions defined in that library, to avoid assertion failures if multiple `libc` libraries exist on the system.

source

– Function.

```
errno([code])
```

Get the value of the C library's `errno`. If an argument is specified, it is used to set the value of `errno`.

The value of `errno` is only valid immediately after a `ccall` to a C library routine that sets it. Specifically, you cannot call `errno` at the next prompt in a REPL, because lots of code is executed between prompts.

source

Base.Libc.strerror – Function.

```
strerror(n=errno())
```

Convert a system call error code to a descriptive string

source

Base.Libc.GetLastError – Function.

```
GetLastError()
```

Call the Win32 GetLastError function [only available on Windows].

source

Base.Libc.FormatMessage – Function.

```
FormatMessage(n=GetLastError())
```

Convert a Win32 system call error code to a descriptive string [only available on Windows].

source

Base.Libc.time – Method.

```
time(t::TmStruct)
```

Converts a TmStruct struct to a number of seconds since the epoch.

source

Base.Libc.strftime – Function.

```
strftime([format], time)
```

Convert time, given as a number of seconds since the epoch or a TmStruct, to a formatted string using the given format. Supported formats are the same as those in the standard C library.

source

Base.Libc.strptime – Function.

```
strptime([format], timestr)
```

Parse a formatted time string into a TmStruct giving the seconds, minute, hour, date, etc. Supported formats are the same as those in the standard C library. On some platforms, timezones will not be parsed correctly. If the result of this function will be passed to time to convert it to seconds since the epoch, the isdst field should be filled in manually. Setting it to -1 will tell the C library to use the current system settings to determine the timezone.

source

Base.Libc.TmStruct – Type.

```
TmStruct([seconds])
```

Convert a number of seconds since the epoch to broken-down format, with fields sec, min, hour, mday, month, year, wday, yday, and isdst.

source

Base.Libc.flush_cstdio – Function.

```
flush_cstdio()
```

Flushes the C `stdout` and `stderr` streams (which may have been written to by external C code).

source

Chapter 65

Dynamic Linker

The names in `Base.Libdl` are not exported and need to be called e.g. as `Libdl.dlopen()`.

`Base.Libdl.dlopen` – Function.

```
dlopen(libfile::AbstractString [, flags::Integer])
```

Load a shared library, returning an opaque handle.

The extension given by the constant `dlext` (`.so`, `.dll`, or `.dylib`) can be omitted from the `libfile` string, as it is automatically appended if needed. If `libfile` is not an absolute path name, then the paths in the array `DL_LOAD_PATH` are searched for `libfile`, followed by the system load path.

The optional flags argument is a bitwise-or of zero or more of `RTLD_LOCAL`, `RTLD_GLOBAL`, `RTLD_LAZY`, `RTLD_NOW`, `RTLD_NODELETE`, `RTLD_NOLOAD`, `RTLD_DEEPBIND`, and `RTLD_FIRST`. These are converted to the corresponding flags of the POSIX (and/or GNU libc and/or MacOS) dlopen command, if possible, or are ignored if the specified functionality is not available on the current platform. The default flags are platform specific. On MacOS the default `dlopen` flags are `RTLD_LAZY|RTLD_DEEPBIND|RTLD_GLOBAL` while on other platforms the defaults are `RTLD_LAZY|RTLD_DEEPBIND|RTLD_LOCAL`. An important usage of these flags is to specify non default behavior for when the dynamic library loader binds library references to exported symbols and if the bound references are put into process local or global scope. For instance `RTLD_LAZY|RTLD_DEEPBIND|RTLD_GLOBAL` allows the library's symbols to be available for usage in other shared libraries, addressing situations where there are dependencies between shared libraries.

> source

`Base.Libdl.dlopen_e` – Function.

```
dlopen_e(libfile::AbstractString [, flags::Integer])
```

Similar to dlopen, except returns a NULL pointer instead of raising errors.

> source

`Base.Libdl.RTLD_NOW` – Constant.

```
RTLD_DEEPBIND
RTLD_FIRST
RTLD_GLOBAL
RTLD_LAZY
RTLD_LOCAL
RTLD_NODELETE
RTLD_NOLOAD
RTLD_NOW
```

Enum constant for dlopen. See your platform man page for details, if applicable.

> source

`Base.Libdl.dlsym` – Function.

> | dlsym(handle, sym)

Look up a symbol from a shared library handle, return callable function pointer on success.

> source

`Base.Libdl.dlsym_e` – Function.

> | dlsym_e(handle, sym)

Look up a symbol from a shared library handle, silently return NULL pointer on lookup failure.

> source

`Base.Libdl.dlclose` – Function.

> | dlclose(handle)

Close shared library referenced by handle.

> source

`Base.Libdl.dlext` – Constant.

> | dlext

File extension for dynamic libraries (e.g. dll, dylib, so) on the current platform.

> source

`Base.Libdl.find_library` – Function.

> | find_library(names, locations)

Searches for the first library in names in the paths in the `locations` list, DL_LOAD_PATH, or system library paths (in that order) which can successfully be dlopen'd. On success, the return value will be one of the names (potentially prefixed by one of the paths in locations). This string can be assigned to a `global const` and used as the library name in future `ccall`'s. On failure, it returns the empty string.

> source

`Base.Libdl.DL_LOAD_PATH` – Constant.

> | DL_LOAD_PATH

When calling dlopen, the paths in this list will be searched first, in order, before searching the system locations for a valid library handle.

> source

Chapter 66

Profiling

Base.Profile.@profile – Macro.

```
@profile
```

`@profile <expression>` runs your expression while taking periodic backtraces. These are appended to an internal buffer of backtraces.

source

The methods in `Base.Profile` are not exported and need to be called e.g. as `Profile.print()`.

Base.Profile.clear – Function.

```
clear()
```

Clear any existing backtraces from the internal buffer.

source

Base.Profile.print – Function.

```
print([io::IO = STDOUT,] [data::Vector]; kwargs...)
```

Prints profiling results to `io` (by default, STDOUT). If you do not supply a `data` vector, the internal buffer of accumulated backtraces will be used.

The keyword arguments can be any combination of:

- `format` – Determines whether backtraces are printed with (default, `:tree`) or without (`:flat`) indentation indicating tree structure.
- `C` – If `true`, backtraces from C and Fortran code are shown (normally they are excluded).
- `combine` – If `true` (default), instruction pointers are merged that correspond to the same line of code.
- `maxdepth` – Limits the depth higher than `maxdepth` in the `:tree` format.
- `sortedby` – Controls the order in `:flat` format. `:filefuncline` (default) sorts by the source line, whereas `:count` sorts in order of number of collected samples.
- `noisefloor` – Limits frames that exceed the heuristic noise floor of the sample (only applies to format `:tree`). A suggested value to try for this is 2.0 (the default is 0). This parameter hides samples for which n <= noisefloor * √N, where n is the number of samples on this line, and N is the number of samples for the callee.

- mincount – Limits the printout to only those lines with at least `mincount` occurrences.

source

```
print([io::IO = STDOUT,] data::Vector, lidict::LineInfoDict; kwargs...)
```

Prints profiling results to `io`. This variant is used to examine results exported by a previous call to `retrieve`. Supply the vector `data` of backtraces and a dictionary `lidict` of line information.

See `Profile.print([io], data)` for an explanation of the valid keyword arguments.

source

`Base.Profile.init` – Function.

```
init(; n::Integer, delay::Float64)
```

Configure the `delay` between backtraces (measured in seconds), and the number `n` of instruction pointers that may be stored. Each instruction pointer corresponds to a single line of code; backtraces generally consist of a long list of instruction pointers. Default settings can be obtained by calling this function with no arguments, and each can be set independently using keywords or in the order `(n, delay)`.

source

`Base.Profile.fetch` – Function.

```
fetch() -> data
```

Returns a reference to the internal buffer of backtraces. Note that subsequent operations, like `clear`, can affect `data` unless you first make a copy. Note that the values in `data` have meaning only on this machine in the current session, because it depends on the exact memory addresses used in JIT-compiling. This function is primarily for internal use; `retrieve` may be a better choice for most users.

source

`Base.Profile.retrieve` – Function.

```
retrieve() -> data, lidict
```

"Exports" profiling results in a portable format, returning the set of all backtraces (`data`) and a dictionary that maps the (session-specific) instruction pointers in `data` to `LineInfo` values that store the file name, function name, and line number. This function allows you to save profiling results for future analysis.

source

`Base.Profile.callers` – Function.

```
callers(funcname, [data, lidict], [filename=<filename>], [linerange=<start:stop>]) -> Vector{
    Tuple{count, lineinfo}}
```

Given a previous profiling run, determine who called a particular function. Supplying the filename (and optionally, range of line numbers over which the function is defined) allows you to disambiguate an overloaded method. The returned value is a vector containing a count of the number of calls and line information about the caller. One can optionally supply backtrace `data` obtained from `retrieve`; otherwise, the current internal profile buffer is used.

source

`Base.Profile.clear_malloc_data` – Function.

```
clear_malloc_data()
```

Clears any stored memory allocation data when running julia with `--track-allocation`. Execute the command(s) you want to test (to force JIT-compilation), then call `clear_malloc_data`. Then execute your command(s) again, quit Julia, and examine the resulting `*.mem` files.

source

Chapter 67

StackTraces

`Base.StackTraces.StackFrame` – Type.

> `StackFrame`

Stack information representing execution context, with the following fields:

- `func::Symbol`

  The name of the function containing the execution context.

- `linfo::Nullable{Core.MethodInstance}`

  The MethodInstance containing the execution context (if it could be found).

- `file::Symbol`

  The path to the file containing the execution context.

- `line::Int`

  The line number in the file containing the execution context.

- `from_c::Bool`

  True if the code is from C.

- `inlined::Bool`

  True if the code is from an inlined frame.

- `pointer::UInt64`

  Representation of the pointer to the execution context as returned by `backtrace`.

source

`Base.StackTraces.StackTrace` – Type.

> `StackTrace`

An alias for `Vector{StackFrame}` provided for convenience; returned by calls to `stacktrace` and `catch_stack-trace`.

source

`Base.StackTraces.stacktrace` – Function.

> `stacktrace([trace::Vector{Ptr{Void}},] [c_funcs::Bool=false]) -> StackTrace`

Returns a stack trace in the form of a vector of `StackFrames`. (By default stacktrace doesn't return C functions, but this can be enabled.) When called without specifying a trace, `stacktrace` first calls `backtrace`.

> source

Base.StackTraces.catch_stacktrace – Function.

```
catch_stacktrace([c_funcs::Bool=false]) -> StackTrace
```

Returns the stack trace for the most recent error thrown, rather than the current execution context.

> source

The following methods and types in `Base.StackTraces` are not exported and need to be called e.g. as `Stack-Traces.lookup(ptr)`.

Base.StackTraces.lookup – Function.

```
lookup(pointer::Union{Ptr{Void}, UInt}) -> Vector{StackFrame}
```

Given a pointer to an execution context (usually generated by a call to `backtrace`), looks up stack frame context information. Returns an array of frame information for all functions inlined at that point, innermost function first.

> source

Base.StackTraces.remove_frames! – Function.

```
remove_frames!(stack::StackTrace, name::Symbol)
```

Takes a `StackTrace` (a vector of `StackFrames`) and a function name (a `Symbol`) and removes the `StackFrame` specified by the function name from the `StackTrace` (also removing all frames above the specified function). Primarily used to remove `StackTraces` functions from the `StackTrace` prior to returning it.

> source

```
remove_frames!(stack::StackTrace, m::Module)
```

Returns the `StackTrace` with all `StackFrames` from the provided `Module` removed.

> source

Chapter 68

SIMD Support

Type `VecElement{T}` is intended for building libraries of SIMD operations. Practical use of it requires using `llvm-call`. The type is defined as:

```
struct VecElement{T}
    value::T
end
```

It has a special compilation rule: a homogeneous tuple of `VecElement{T}` maps to an LLVM `vector` type when T is a primitive bits type and the tuple length is in the set {2-6,8-10,16}.

At `-O3`, the compiler *might* automatically vectorize operations on such tuples. For example, the following program, when compiled with `julia -O3` generates two SIMD addition instructions (addps) on x86 systems:

```
const m128 = NTuple{4,VecElement{Float32}}

function add(a::m128, b::m128)
    (VecElement(a[1].value+b[1].value),
     VecElement(a[2].value+b[2].value),
     VecElement(a[3].value+b[3].value),
     VecElement(a[4].value+b[4].value))
end

triple(c::m128) = add(add(c,c),c)

code_native(triple,(m128,))
```

However, since the automatic vectorization cannot be relied upon, future use will mostly be via libraries that use `llvmcall`.

Part V

Developer Documentation

Chapter 69

Reflection and introspection

Julia provides a variety of runtime reflection capabilities.

69.1 Module bindings

The exported names for a `Module` are available using `names(m::Module)`, which will return an array of `Symbol` elements representing the exported bindings. `names(m::Module, true)` returns symbols for all bindings in `m`, regardless of export status.

69.2 DataType fields

The names of `DataType` fields may be interrogated using `fieldnames()`. For example, given the following type, `fieldnames(Point)` returns an arrays of `Symbol` elements representing the field names:

```julia
julia> struct Point
           x::Int
           y
       end

julia> fieldnames(Point)
2-element Array{Symbol,1}:
 :x
 :y
```

The type of each field in a `Point` object is stored in the `types` field of the `Point` variable itself:

```julia
julia> Point.types
svec(Int64, Any)
```

While x is annotated as an `Int`, y was unannotated in the type definition, therefore y defaults to the Any type.

Types are themselves represented as a structure called `DataType`:

```julia
julia> typeof(Point)
DataType
```

Note that `fieldnames(DataType)` gives the names for each field of `DataType` itself, and one of these fields is the `types` field observed in the example above.

69.3   Subtypes

The *direct* subtypes of any `DataType` may be listed using `subtypes()`. For example, the abstract `DataType` `AbstractFloat` has four (concrete) subtypes:

```julia
julia> subtypes(AbstractFloat)
4-element Array{Union{DataType, UnionAll},1}:
 BigFloat
 Float16
 Float32
 Float64
```

Any abstract subtype will also be included in this list, but further subtypes thereof will not; recursive application of `subtypes()` may be used to inspect the full type tree.

69.4   DataType layout

The internal representation of a `DataType` is critically important when interfacing with C code and several functions are available to inspect these details. `isbits(T::DataType)` returns true if T is stored with C-compatible alignment. `fieldoffset(T::DataType, i::Integer)` returns the (byte) offset for field *i* relative to the start of the type.

69.5   Function methods

The methods of any generic function may be listed using `methods()`. The method dispatch table may be searched for methods accepting a given type using `methodswith()`.

69.6   Expansion and lowering

As discussed in the Metaprogramming section, the `macroexpand()` function gives the unquoted and interpolated expression (Expr) form for a given macro. To use `macroexpand`, `quote` the expression block itself (otherwise, the macro will be evaluated and the result will be passed instead!). For example:

```julia
julia> macroexpand( :(@edit println("")) )
:((Base.edit)(println, (Base.typesof)("")))
```

The functions `Base.Meta.show_sexpr()` and `dump()` are used to display S-expr style views and depth-nested detail views for any expression.

Finally, the `expand()` function gives the `lowered` form of any expression and is of particular interest for understanding both macros and top-level statements such as function declarations and variable assignments:

```julia
julia> expand( :(f() = 1) )
:(begin
        $(Expr(:method, :f))
        $(Expr(:method, :f, :((Core.svec)((Core.svec)((Core.Typeof)(f)), (Core.svec)()))),
↪   CodeInfo(:(begin  # none, line 1:
        return 1
    end)), false))
        return f
    end)
```

69.7 Intermediate and compiled representations

Inspecting the lowered form for functions requires selection of the specific method to display, because generic functions may have many methods with different type signatures. For this purpose, method-specific code-lowering is available using `code_lowered(f::Function, (Argtypes...))`, and the type-inferred form is available using `code_typed(f::Function, (Argtypes...))`. `code_warntype(f::Function, (Argtypes...))` adds highlighting to the output of `code_typed()` (see `@code_warntype`).

Closer to the machine, the LLVM intermediate representation of a function may be printed using by `code_llvm(f::Function, (Argtypes...))`, and finally the compiled machine code is available using `code_native(f::Function, (Argtypes...))` (this will trigger JIT compilation/code generation for any function which has not previously been called).

For convenience, there are macro versions of the above functions which take standard function calls and expand argument types automatically:

```
julia> @code_llvm +(1,1)

; Function Attrs: sspreq
define i64 @"julia_+_130862"(i64, i64) #0 {
top:
    %2 = add i64 %1, %0, !dbg !8
    ret i64 %2, !dbg !8
}
```

(likewise @code_typed, @code_warntype, @code_lowered, and @code_native)

Chapter 70

Documentation of Julia's Internals

70.1   Initialization of the Julia runtime

How does the Julia runtime execute `julia -e 'println("Hello World!")'` ?

main()

Execution starts at `main()` in `ui/repl.c`.

`main()` calls `libsupport_init()` to set the C library locale and to initialize the "ios" library (see `ios_init_std-streams()` and Legacy `ios.c` library).

Next `parse_opts()` is called to process command line options. Note that `parse_opts()` only deals with options that affect code generation or early initialization. Other options are handled later by `process_options()` in `base/client.jl`.

`parse_opts()` stores command line options in the global `jl_options` struct.

julia_init()

`julia_init()` in `task.c` is called by `main()` and calls `_julia_init()` in `init.c`.

`_julia_init()` begins by calling `libsupport_init()` again (it does nothing the second time).

`restore_signals()` is called to zero the signal handler mask.

`jl_resolve_sysimg_location()` searches configured paths for the base system image. See Building the Julia system image.

`jl_gc_init()` sets up allocation pools and lists for weak refs, preserved values and finalization.

`jl_init_frontend()` loads and initializes a pre-compiled femtolisp image containing the scanner/parser.

`jl_init_types()` creates `jl_datatype_t` type description objects for the built-in types defined in `julia.h`. e.g.

```
jl_any_type = jl_new_abstracttype(jl_symbol("Any"), NULL, jl_null);
jl_any_type->super = jl_any_type;

jl_type_type = jl_new_abstracttype(jl_symbol("Type"), jl_any_type, jl_null);

jl_int32_type = jl_new_bitstype(jl_symbol("Int32"),
                                jl_any_type, jl_null, 32);
```

`jl_init_tasks()` creates the `jl_datatype_t* jl_task_type` object; initializes the global `jl_root_task` struct; and sets `jl_current_task` to the root task.

`jl_init_codegen()` initializes the LLVM library.

`jl_init_serializer()` initializes 8-bit serialization tags for 256 frequently used `jl_value_t` values. The serialization mechanism uses these tags as shorthand (in lieu of storing whole objects) to save storage space.

If there is no sysimg file (`!jl_options.image_file`) then the `Core` and `Main` modules are created and `boot.jl` is evaluated:

`jl_core_module = jl_new_module(jl_symbol("Core"))` creates the Julia `Core` module.

`jl_init_intrinsic_functions()` creates a new Julia module `Intrinsics` containing constant `jl_intrinsic_type` symbols. These define an integer code for each intrinsic function. `emit_intrinsic()` translates these symbols into LLVM instructions during code generation.

`jl_init_primitives()` hooks C functions up to Julia function symbols. e.g. the symbol `Base.is()` is bound to C function pointer `jl_f_is()` by calling `add_builtin_func("eval", jl_f_top_eval)`.

`jl_new_main_module()` creates the global "Main" module and sets `jl_current_task->current_module = jl_main_module`.

Note: `_julia_init()` then sets `jl_root_task->current_module = jl_core_module`. `jl_root_task` is an alias of `jl_current_task` at this point, so the `current_module` set by `jl_new_main_module()` above is overwritten.

`jl_load("boot.jl", sizeof("boot.jl"))` calls `jl_parse_eval_all` which repeatedly calls `jl_toplevel_eval_flex()` to execute `boot.jl`. <!– TODO – drill down into eval? –>

`jl_get_builtin_hooks()` initializes global C pointers to Julia globals defined in `boot.jl`.

`jl_init_box_caches()` pre-allocates global boxed integer value objects for values up to 1024. This speeds up allocation of boxed ints later on. e.g.:

```
jl_value_t *jl_box_uint8(uint32_t x)
{
    return boxed_uint8_cache[(uint8_t)x];
}
```

`_julia_init()` iterates over the `jl_core_module->bindings.table` looking for `jl_datatype_t` values and sets the type name's module prefix to `jl_core_module`.

`jl_add_standard_imports(jl_main_module)` does "using Base" in the "Main" module.

Note: `_julia_init()` now reverts to `jl_root_task->current_module = jl_main_module` as it was before being set to `jl_core_module` above.

Platform specific signal handlers are initialized for SIGSEGV (OSX, Linux), and SIGFPE (Windows).

Other signals (`SIGINFO, SIGBUS, SIGILL, SIGTERM, SIGABRT, SIGQUIT, SIGSYS` and `SIGPIPE`) are hooked up to `sigdie_handler()` which prints a backtrace.

`jl_init_restored_modules()` calls `jl_module_run_initializer()` for each deserialized module to run the `__init__()` function.

Finally `sigint_handler()` is hooked up to SIGINT and calls `jl_throw(jl_interrupt_exception)`.

`_julia_init()` then returns back to `main()` in `ui/repl.c` and `main()` calls `true_main(argc, (char**)argv)`.

sysimg

If there is a sysimg file, it contains a pre-cooked image of the `Core` and `Main` modules (and whatever else is created by `boot.jl`). See Building the Julia system image.

`jl_restore_system_image()` deserializes the saved sysimg into the current Julia runtime environment and initialization continues after `jl_init_box_caches()` below...

Note: `jl_restore_system_image()` (and dump.c in general) uses the Legacy ios.c library.

true_main()

`true_main()` loads the contents of `argv[]` into Base.ARGS.

If a `.jl` "program" file was supplied on the command line, then `exec_program()` calls `jl_load(program,len)` which calls `jl_parse_eval_all` which repeatedly calls `jl_toplevel_eval_flex()` to execute the program.

However, in our example (julia -e 'println("Hello World!")'), `jl_get_global(jl_base_module, jl_symbol("_start"))` looks up Base._start and `jl_apply()` executes it.

Base._start

Base._start calls Base.process_options which calls `jl_parse_input_line("println("Hello World!")")` to create an expression object and Base.eval() to execute it.

Base.eval

Base.eval() was mapped to jl_f_top_eval by jl_init_primitives().

`jl_f_top_eval()` calls jl_toplevel_eval_in(jl_main_module, ex), where ex is the parsed expression println("Hello World!").

`jl_toplevel_eval_in()` calls `jl_toplevel_eval_flex()` which calls eval() in interpreter.c.

The stack dump below shows how the interpreter works its way through various methods of Base.println() and Base.print() before arriving at write(s::IO, a::Array{T}) where T which does ccall(jl_uv_write()).

`jl_uv_write()` calls uv_write() to write "Hello World!" to JL_STDOUT. See Libuv wrappers for stdio.:

```
Hello World!
```

Since our example has just one function call, which has done its job of printing "Hello World!", the stack now rapidly unwinds back to main().

jl_atexit_hook()

main() calls jl_atexit_hook(). This calls _atexit for each module, then calls jl_gc_run_all_finalizers() and cleans up libuv handles.

julia_save()

Finally, main() calls julia_save(), which if requested on the command line, saves the runtime state to a new system image. See jl_compile_all() and jl_save_system_image().

70.2   Julia ASTs

Julia has two representations of code. First there is a surface syntax AST returned by the parser (e.g. the parse() function), and manipulated by macros. It is a structured representation of code as it is written, constructed by julia-parser.scm from a character stream. Next there is a lowered form, or IR (intermediate representation), which is used by type inference and code generation. In the lowered form there are fewer types of nodes, all macros are expanded, and all control flow is converted to explicit branches and sequences of statements. The lowered form is constructed by julia-syntax.scm.

Stack frame	Source code	Notes
`jl_uv_write()`	`jl_uv.c`	called though `ccall`
`julia_write_282942`	`stream.jl`	function `write!(s::IO, a::Array{T}) where T`
`julia_print_284639`	`ascii.jl`	`print(io::IO, s::String) = (write(io, s); nothing)`
`jlcall_print_284639`		
`jl_apply()`	`julia.h`	
`jl_trampoline()`	`builtins.c`	
`jl_apply()`	`julia.h`	
`jl_apply_generic()`	`gf.c`	`Base.print(Base.TTY, String)`
`jl_apply()`	`julia.h`	
`jl_trampoline()`	`builtins.c`	
`jl_apply()`	`julia.h`	
`jl_apply_generic()`	`gf.c`	`Base.print(Base.TTY, String, Char, Char...)`
`jl_apply()`	`julia.h`	
`jl_f_apply()`	`builtins.c`	
`jl_apply()`	`julia.h`	
`jl_trampoline()`	`builtins.c`	
`jl_apply()`	`julia.h`	
`jl_apply_generic()`	`gf.c`	`Base.println(Base.TTY, String, String...)`
`jl_apply()`	`julia.h`	
`jl_trampoline()`	`builtins.c`	
`jl_apply()`	`julia.h`	
`jl_apply_generic()`	`gf.c`	`Base.println(String,)`
`jl_apply()`	`julia.h`	
`do_call()`	`inter-preter.c`	
`eval()`	`inter-preter.c`	
`jl_inter-pret_toplevel_expr()`	`inter-preter.c`	
`jl_toplevel_eval_flex()`	`toplevel.c`	
`jl_toplevel_eval()`	`toplevel.c`	
`jl_toplevel_eval_in()`	`builtins.c`	
`jl_f_top_eval()`	`builtins.c`	

First we will focus on the lowered form, since it is more important to the compiler. It is also less obvious to the human, since it results from a significant rearrangement of the input syntax.

Lowered form

The following data types exist in lowered form:

- `Expr`

  Has a node type indicated by the `head` field, and an `args` field which is a `Vector{Any}` of subexpressions.

- `Slot`

  Identifies arguments and local variables by consecutive numbering. `Slot` is an abstract type with subtypes `SlotNumber` and `TypedSlot`. Both types have an integer-valued `id` field giving the slot index. Most slots have the same type at all uses, and so are represented with `SlotNumber`. The types of these slots are found in the

slottypes field of their MethodInstance object. Slots that require per-use type annotations are represented with TypedSlot, which has a typ field.

- CodeInfo

  Wraps the IR of a method.

- LineNumberNode

  Contains a single number, specifying the line number the next statement came from.

- LabelNode

  Branch target, a consecutively-numbered integer starting at 0.

- GotoNode

  Unconditional branch.

- QuoteNode

  Wraps an arbitrary value to reference as data. For example, the function `f() = :a` contains a QuoteNode whose `value` field is the symbol a, in order to return the symbol itself instead of evaluating it.

- GlobalRef

  Refers to global variable name in module mod.

- SSAValue

  Refers to a consecutively-numbered (starting at 0) static single assignment (SSA) variable inserted by the compiler.

- NewvarNode

  Marks a point where a variable is created. This has the effect of resetting a variable to undefined.

Expr types

These symbols appear in the head field of Exprs in lowered form.

- call

  Function call (dynamic dispatch). `args[1]` is the function to call, `args[2:end]` are the arguments.

- invoke

  Function call (static dispatch). `args[1]` is the MethodInstance to call, `args[2:end]` are the arguments (including the function that is being called, at `args[2]`).

- static_parameter

  Reference a static parameter by index.

- line

  Line number and file name metadata. Unlike a LineNumberNode, can also contain a file name.

- gotoifnot

  Conditional branch. If `args[1]` is false, goes to label identified in `args[2]`.

- =

  Assignment.

- `method`

  Adds a method to a generic function and assigns the result if necessary.

  Has a 1-argument form and a 4-argument form. The 1-argument form arises from the syntax `function foo end`. In the 1-argument form, the argument is a symbol. If this symbol already names a function in the current scope, nothing happens. If the symbol is undefined, a new function is created and assigned to the identifier specified by the symbol. If the symbol is defined but names a non-function, an error is raised. The definition of "names a function" is that the binding is constant, and refers to an object of singleton type. The rationale for this is that an instance of a singleton type uniquely identifies the type to add the method to. When the type has fields, it wouldn't be clear whether the method was being added to the instance or its type.

  The 4-argument form has the following arguments:

  - `args[1]`

    A function name, or `false` if unknown. If a symbol, then the expression first behaves like the 1-argument form above. This argument is ignored from then on. When this is `false`, it means a method is being added strictly by type, `(::T)(x) = x`.

  - `args[2]`

    A `SimpleVector` of argument type data. `args[2][1]` is a `SimpleVector` of the argument types, and `args[2][2]` is a `SimpleVector` of type variables corresponding to the method's static parameters.

  - `args[3]`

    A `CodeInfo` of the method itself. For "out of scope" method definitions (adding a method to a function that also has methods defined in different scopes) this is an expression that evaluates to a `:lambda` expression.

  - `args[4]`

    `true` or `false`, identifying whether the method is staged (`@generated function`).

- `const`

  Declares a (global) variable as constant.

- `null`

  Has no arguments; simply yields the value `nothing`.

- `new`

  Allocates a new struct-like object. First argument is the type. The new pseudo-function is lowered to this, and the type is always inserted by the compiler. This is very much an internal-only feature, and does no checking. Evaluating arbitrary new expressions can easily segfault.

- `return`

  Returns its argument as the value of the enclosing function.

- `the_exception`

  Yields the caught exception inside a `catch` block. This is the value of the run time system variable `jl_exception_in_transit`.

- `enter`

  Enters an exception handler (`setjmp`). `args[1]` is the label of the catch block to jump to on error.

- `leave`

  Pop exception handlers. `args[1]` is the number of handlers to pop.

- `inbounds`

  Controls turning bounds checks on or off. A stack is maintained; if the first argument of this expression is true or false (`true` means bounds checks are disabled), it is pushed onto the stack. If the first argument is `:pop`, the stack is popped.

- `boundscheck`

  Indicates the beginning or end of a section of code that performs a bounds check. Like `inbounds`, a stack is maintained, and the second argument can be one of: `true`, `false`, or `:pop`.

- `copyast`

  Part of the implementation of quasi-quote. The argument is a surface syntax AST that is simply copied recursively and returned at run time.

- `meta`

  Metadata. `args[1]` is typically a symbol specifying the kind of metadata, and the rest of the arguments are free-form. The following kinds of metadata are commonly used:

  - `:inline` and `:noinline`: Inlining hints.
  - `:push_loc`: enters a sequence of statements from a specified source location.
    * `args[2]` specifies a filename, as a symbol.
    * `args[3]` optionally specifies the name of an (inlined) function that originally contained the code.
  - `:pop_loc`: returns to the source location before the matching `:push_loc`.

Method

A unique'd container describing the shared metadata for a single method.

- `name`, `module`, `file`, `line`, `sig`

  Metadata to uniquely identify the method for the computer and the human.

- `ambig`

  Cache of other methods that may be ambiguous with this one.

- `specializations`

  Cache of all MethodInstance ever created for this Method, used to ensure uniqueness. Uniqueness is required for efficiency, especially for incremental precompile and tracking of method invalidation.

- `source`

  The original source code (usually compressed).

- `roots`

  Pointers to non-AST things that have been interpolated into the AST, required by compression of the AST, type-inference, or the generation of native code.

- `nargs`, `isva`, `called`, `isstaged`, `pure`

  Descriptive bit-fields for the source code of this Method.

- `min_world` / `max_world`

  The range of world ages for which this method is visible to dispatch.

MethodInstance

A unique'd container describing a single callable signature for a Method. See especially Proper maintenance and care of multi-threading locks for important details on how to modify these fields safely.

- `specTypes`

  The primary key for this MethodInstance. Uniqueness is guaranteed through a `def.specializations` lookup.

- `def`

  The `Method` that this function describes a specialization of. Or #undef, if this is a top-level Lambda that is not part of a Method.

- `sparam_vals`

  The values of the static parameters in `specTypes` indexed by `def.sparam_syms`. For the `MethodInstance` at `Method.unspecialized`, this is the empty `SimpleVector`. But for a runtime `MethodInstance` from the `MethodTable` cache, this will always be defined and indexable.

- `rettype`

  The inferred return type for the `specFunctionObject` field, which (in most cases) is also the computed return type for the function in general.

- `inferred`

  May contain a cache of the inferred source for this function, or other information about the inference result such as a constant return value may be put here (if `jlcall_api == 2`), or it could be set to `nothing` to just indicate `rettype` is inferred.

- `ftpr`

  The generic jlcall entry point.

- `jlcall_api`

  The ABI to use when calling `fptr`. Some significant ones include:

  - 0 - Not compiled yet

  - 1 - JL_CALLABLE `jl_value_t *(*)(jl_function_t *f, jl_value_t *args[nargs], uint32_t nargs)`

  - 2 - Constant (value stored in `inferred`)

  - 3 - With Static-parameters forwarded `jl_value_t *(*)(jl_svec_t *sparams, jl_function_t *f, jl_value_t *args[nargs], uint32_t nargs)`

  - 4 - Run in interpreter `jl_value_t *(*)(jl_method_instance_t *meth, jl_function_t *f, jl_value_t *args[nargs], uint32_t nargs)`

- `min_world` / `max_world`

  The range of world ages for which this method instance is valid to be called.

CodeInfo

A temporary container for holding lowered source code.

- `code`

  An Any array of statements

- `slotnames`

  An array of symbols giving the name of each slot (argument or local variable).

- `slottypes`

  An array of types for the slots.

- `slotflags`

  A `UInt8` array of slot properties, represented as bit flags:

    - 2 - assigned (only false if there are *no* assignment statements with this var on the left)
    - 8 - const (currently unused for local variables)
    - 16 - statically assigned once
    - 32 - might be used before assigned. This flag is only valid after type inference.

- `ssavaluetypes`

  Either an array or an `Int`.

  If an `Int`, it gives the number of compiler-inserted temporary locations in the function. If an array, specifies a type for each location.

Boolean properties:

- `inferred`

  Whether this has been produced by type inference.

- `inlineable`

  Whether this should be inlined.

- `propagate_inbounds`

  Whether this should should propagate @inbounds when inlined for the purpose of eliding @boundscheck blocks.

- `pure`

  Whether this is known to be a pure function of its arguments, without respect to the state of the method caches or other mutable global state.

Surface syntax AST

Front end ASTs consist entirely of `Expr`s and atoms (e.g. symbols, numbers). There is generally a different expression head for each visually distinct syntactic form. Examples will be given in s-expression syntax. Each parenthesized list corresponds to an Expr, where the first element is the head. For example `(call f x)` corresponds to `Expr(:call, :f, :x)` in Julia.

Input	AST
f(x)	(call f x)
f(x, y=1, z=2)	(call f x (kw y 1) (kw z 2))
f(x; y=1)	(call f (parameters (kw y 1)) x)
f(x...)	(call f (... x))

Calls

do syntax:

```
f(x) do a,b
    body
end
```

parses as `(call f (-> (tuple a b) (block body)) x)`.

Operators

Most uses of operators are just function calls, so they are parsed with the head `call`. However some operators are special forms (not necessarily function calls), and in those cases the operator itself is the expression head. In julia-parser.scm these are referred to as "syntactic operators". Some operators (+ and *) use N-ary parsing; chained calls are parsed as a single N-argument call. Finally, chains of comparisons have their own special expression structure.

Input	AST
x+y	(call + x y)
a+b+c+d	(call + a b c d)
2x	(call * 2 x)
a&&b	(&& a b)
x += 1	(+= x 1)
a ? 1 : 2	(if a 1 2)
a:b	(: a b)
a:b:c	(: a b c)
a,b	(tuple a b)
a==b	(call == a b)
1<i<=n	(comparison 1 < i <= n)
a.b	(. a (quote b))
a.(b)	(. a b)

Bracketed forms

Macros

Strings

Doc string syntax:

```
"some docs"
f(x) = x
```

parses as `(macrocall (|.| Core '@doc) "some docs" (= (call f x) (block x)))`.

Input	AST
`a[i]`	`(ref a i)`
`t[i;j]`	`(typed_vcat t i j)`
`t[i j]`	`(typed_hcat t i j)`
`t[a b; c d]`	`(typed_vcat t (row a b) (row c d))`
`a{b}`	`(curly a b)`
`a{b;c}`	`(curly a (parameters c) b)`
`[x]`	`(vect x)`
`[x,y]`	`(vect x y)`
`[x;y]`	`(vcat x y)`
`[x y]`	`(hcat x y)`
`[x y; z t]`	`(vcat (row x y) (row z t))`
`[x for y in z, a in b]`	`(comprehension x (= y z) (= a b))`
`T[x for y in z]`	`(typed_comprehension T x (= y z))`
`(a, b, c)`	`(tuple a b c)`
`(a; b; c)`	`(block a (block b c))`

Input	AST
`@m x y`	`(macrocall @m x y)`
`Base.@m x y`	`(macrocall (. Base (quote @m)) x y)`
`@Base.m x y`	`(macrocall (. Base (quote @m)) x y)`

Input	AST
`"a"`	`"a"`
`x"y"`	`(macrocall @x_str "y")`
`x"y"z`	`(macrocall @x_str "y" "z")`
`"x = $x"`	`(string "x = " x)`
`` `a b c` ``	`(macrocall @cmd "a b c")`

Imports and such

Numbers

Julia supports more number types than many scheme implementations, so not all numbers are represented directly as scheme numbers in the AST.

Block forms

A block of statements is parsed as (`block stmt1 stmt2 ...`).

If statement:

```
if a
    b
elseif c
    d
else e
    f
end
```

parses as:

Input	AST
`import a`	`(import a)`
`import a.b.c`	`(import a b c)`
`import ...a`	`(import . . . a)`
`import a.b, c.d`	`(toplevel (import a b) (import c d))`
`import Base: x`	`(import Base x)`
`import Base: x, y`	`(toplevel (import Base x) (import Base y))`
`export a, b`	`(export a b)`

Input	AST
`11111111111111111111`	`(macrocall @int128_str "11111111111111111111")`
`0xfffffffffffffffff`	`(macrocall @uint128_str "0xfffffffffffffffff")`
`1111...many digits...`	`(macrocall @big_str "1111....")`

```
(if a (block (line 2) b)
    (block (line 3) (if c (block (line 4) d)
                        (block (line 5) e (line 6) f))))
```

A `while` loop parses as `(while condition body)`.

A `for` loop parses as `(for (= var iter) body)`. If there is more than one iteration specification, they are parsed as a block: `(for (block (= v1 iter1) (= v2 iter2)) body)`.

`break` and `continue` are parsed as 0-argument expressions `(break)` and `(continue)`.

`let` is parsed as `(let body (= var1 val1) (= var2 val2) ...)`.

A basic function definition is parsed as `(function (call f x) body)`. A more complex example:

```
function f{T}(x::T; k = 1)
    return x+1
end
```

parses as:

```
(function (call (curly f T) (parameters (kw k 1))
                (:: x T))
          (block (line 2 file.jl) (return (call + x 1))))
```

Type definition:

```
mutable struct Foo{T<:S}
    x::T
end
```

parses as:

```
(type #t (curly Foo (<: T S))
      (block (line 2 none) (:: x T)))
```

The first argument is a boolean telling whether the type is mutable.

`try` blocks parse as `(try try_block var catch_block finally_block)`. If no variable is present after `catch`, `var` is `#f`. If there is no `finally` clause, then the last argument is not present.

70.3 More about types

If you've used Julia for a while, you understand the fundamental role that types play. Here we try to get under the hood, focusing particularly on Parametric Types.

Types and sets (and `Any` and `Union{}`/`Bottom`)

It's perhaps easiest to conceive of Julia's type system in terms of sets. While programs manipulate individual values, a type refers to a set of values. This is not the same thing as a collection; for example a `Set` of values is itself a single `Set` value. Rather, a type describes a set of *possible* values, expressing uncertainty about which value we have.

A *concrete* type T describes the set of values whose direct tag, as returned by the `typeof` function, is T. An *abstract* type describes some possibly-larger set of values.

Any describes the entire universe of possible values. `Integer` is a subset of Any that includes Int, `Int8`, and other concrete types. Internally, Julia also makes heavy use of another type known as `Bottom`, which can also be written as `Union{}`. This corresponds to the empty set.

Julia's types support the standard operations of set theory: you can ask whether T1 is a "subset" (subtype) of T2 with `T1 <: T2`. Likewise, you intersect two types using `typeintersect`, take their union with `Union`, and compute a type that contains their union with `typejoin`:

```
julia> typeintersect(Int, Float64)
Union{}

julia> Union{Int, Float64}
Union{Float64, Int64}

julia> typejoin(Int, Float64)
Real

julia> typeintersect(Signed, Union{UInt8, Int8})
Int8

julia> Union{Signed, Union{UInt8, Int8}}
Union{Signed, UInt8}

julia> typejoin(Signed, Union{UInt8, Int8})
Integer

julia> typeintersect(Tuple{Integer,Float64}, Tuple{Int,Real})
Tuple{Int64,Float64}

julia> Union{Tuple{Integer,Float64}, Tuple{Int,Real}}
Union{Tuple{Int64,Real}, Tuple{Integer,Float64}}

julia> typejoin(Tuple{Integer,Float64}, Tuple{Int,Real})
Tuple{Integer,Real}
```

While these operations may seem abstract, they lie at the heart of Julia. For example, method dispatch is implemented by stepping through the items in a method list until reaching one for which the type of the argument tuple is a subtype of the method signature. For this algorithm to work, it's important that methods be sorted by their specificity, and that the search begins with the most specific methods. Consequently, Julia also implements a partial order on types; this is achieved by functionality that is similar to `<:`, but with differences that will be discussed below.

UnionAll types

Julia's type system can also express an *iterated union* of types: a union of types over all values of some variable. This is needed to describe parametric types where the values of some parameters are not known.

For example, :obj:`Array` has two parameters as in `Array{Int,2}`. If we did not know the element type, we could write `Array{T,2} where T`, which is the union of `Array{T,2}` for all values of T: `Union{Array{Int8,2}, Ar-ray{Int16,2}, ...}`.

Such a type is represented by a `UnionAll` object, which contains a variable (T in this example, of type `TypeVar`), and a wrapped type (`Array{T,2}` in this example).

Consider the following methods:

```
f1(A::Array) = 1
f2(A::Array{Int}) = 2
f3(A::Array{T}) where {T<:Any} = 3
f4(A::Array{Any}) = 4
```

The signature of `f3` is a `UnionAll` type wrapping a tuple type. All but `f4` can be called with a = `[1,2]`; all but `f2` can be called with b = `Any[1,2]`.

Let's look at these types a little more closely:

```
julia> dump(Array)
UnionAll
  var: TypeVar
    name: Symbol T
    lb: Core.TypeofBottom Union{}
    ub: Any
  body: UnionAll
    var: TypeVar
      name: Symbol N
      lb: Core.TypeofBottom Union{}
      ub: Any
    body: Array{T,N} <: DenseArray{T,N}
```

This indicates that `Array` actually names a `UnionAll` type. There is one `UnionAll` type for each parameter, nested. The syntax `Array{Int,2}` is equivalent to `Array{Int}{2}`; internally each `UnionAll` is instantiated with a particular variable value, one at a time, outermost-first. This gives a natural meaning to the omission of trailing type parameters; `Array{Int}` gives a type equivalent to `Array{Int,N} where N`.

A `TypeVar` is not itself a type, but rather should be considered part of the structure of a `UnionAll` type. Type variables have lower and upper bounds on their values (in the fields `lb` and ub). The symbol name is purely cosmetic. Internally, `TypeVar`s are compared by address, so they are defined as mutable types to ensure that "different" type variables can be distinguished. However, by convention they should not be mutated.

One can construct `TypeVar`s manually:

```
julia> TypeVar(:V, Signed, Real)
Signed<:V<:Real
```

There are convenience versions that allow you to omit any of these arguments except the name symbol.

The syntax `Array{T} where T<:Integer` is lowered to

```
let T = TypeVar(:T,Integer)
    UnionAll(T, Array{T})
end
```

so it is seldom necessary to construct a `TypeVar` manually (indeed, this is to be avoided).

Free variables

The concept of a *free* type variable is extremely important in the type system. We say that a variable V is free in type T if T does not contain the `UnionAll` that introduces variable V. For example, the type `Array{Array{V} where V<:Integer}` has no free variables, but the `Array{V}` part inside of it does have a free variable, V.

A type with free variables is, in some sense, not really a type at all. Consider the type `Array{Array{T}} where T`, which refers to all homogeneous arrays of arrays. The inner type `Array{T}`, seen by itself, might seem to refer to any kind of array. However, every element of the outer array must have the *same* array type, so `Array{T}` cannot refer to just any old array. One could say that `Array{T}` effectively "occurs" multiple times, and T must have the same value each "time".

For this reason, the function `jl_has_free_typevars` in the C API is very important. Types for which it returns true will not give meaningful answers in subtyping and other type functions.

TypeNames

The following two `Array` types are functionally equivalent, yet print differently:

```julia
julia> TV, NV = TypeVar(:T), TypeVar(:N)
(T, N)

julia> Array
Array

julia> Array{TV,NV}
Array{T,N}
```

These can be distinguished by examining the name field of the type, which is an object of type TypeName:

```julia
julia> dump(Array{Int,1}.name)
TypeName
  name: Symbol Array
  module: Module Core
  names: empty SimpleVector
  wrapper: UnionAll
    var: TypeVar
      name: Symbol T
      lb: Core.TypeofBottom Union{}
      ub: Any
    body: UnionAll
      var: TypeVar
        name: Symbol N
        lb: Core.TypeofBottom Union{}
        ub: Any
      body: Array{T,N} <: DenseArray{T,N}
  cache: SimpleVector
    ...
```

```
    linearcache: SimpleVector
      ...

    hash: Int64 -7900426068641098781
    mt: MethodTable
      name: Symbol Array
      defs: Void nothing
      cache: Void nothing
      max_args: Int64 0
      kwsorter: #undef
      module: Module Core
      : Int64 0
      : Int64 0
```

In this case, the relevant field is `wrapper`, which holds a reference to the top-level type used to make new `Array` types.

```julia
julia> pointer_from_objref(Array)
Ptr{Void} @0x00007fcc7de64850

julia> pointer_from_objref(Array.body.body.name.wrapper)
Ptr{Void} @0x00007fcc7de64850

julia> pointer_from_objref(Array{TV,NV})
Ptr{Void} @0x00007fcc80c4d930

julia> pointer_from_objref(Array{TV,NV}.name.wrapper)
Ptr{Void} @0x00007fcc7de64850
```

The `wrapper` field of `Array` points to itself, but for `Array{TV,NV}` it points back to the original definition of the type.

What about the other fields? `hash` assigns an integer to each type. To examine the `cache` field, it's helpful to pick a type that is less heavily used than Array. Let's first create our own type:

```julia
julia> struct MyType{T,N} end

julia> MyType{Int,2}
MyType{Int64,2}

julia> MyType{Float32, 5}
MyType{Float32,5}

julia> MyType.body.body.name.cache
svec(MyType{Float32,5}, MyType{Int64,2}, #undef, #undef, #undef, #undef, #undef, #undef)
```

(The cache is pre-allocated to have length 8, but only the first two entries are populated.) Consequently, when you instantiate a parametric type, each concrete type gets saved in a type cache. However, instances containing free type variables are not cached.

Tuple types

Tuple types constitute an interesting special case. For dispatch to work on declarations like `x::Tuple`, the type has to be able to accommodate any tuple. Let's check the parameters:

```
julia> Tuple
Tuple

julia> Tuple.parameters
svec(Vararg{Any,N} where N)
```

Unlike other types, tuple types are covariant in their parameters, so this definition permits `Tuple` to match any type of tuple:

```
julia> typeintersect(Tuple, Tuple{Int,Float64})
Tuple{Int64,Float64}

julia> typeintersect(Tuple{Vararg{Any}}, Tuple{Int,Float64})
Tuple{Int64,Float64}
```

However, if a variadic (`Vararg`) tuple type has free variables it can describe different kinds of tuples:

```
julia> typeintersect(Tuple{Vararg{T} where T}, Tuple{Int,Float64})
Tuple{Int64,Float64}

julia> typeintersect(Tuple{Vararg{T}} where T, Tuple{Int,Float64})
Union{}
```

Notice that when T is free with respect to the `Tuple` type (i.e. its binding `UnionAll` type is outside the `Tuple` type), only one T value must work over the whole type. Therefore a heterogeneous tuple does not match.

Finally, it's worth noting that `Tuple{}` is distinct:

```
julia> Tuple{}
Tuple{}

julia> Tuple{}.parameters
svec()

julia> typeintersect(Tuple{}, Tuple{Int})
Union{}
```

What is the "primary" tuple-type?

```
julia> pointer_from_objref(Tuple)
Ptr{Void} @0x00007f5998a04370

julia> pointer_from_objref(Tuple{})
Ptr{Void} @0x00007f5998a570d0

julia> pointer_from_objref(Tuple.name.wrapper)
Ptr{Void} @0x00007f5998a04370

julia> pointer_from_objref(Tuple{}.name.wrapper)
Ptr{Void} @0x00007f5998a04370
```

so `Tuple == Tuple{Vararg{Any}}` is indeed the primary type.

Diagonal types

Consider the type `Tuple{T,T} where T`. A method with this signature would look like:

```
f(x::T, y::T) where {T} = ...
```

According to the usual interpretation of a `UnionAll` type, this T ranges over all types, including Any, so this type should be equivalent to `Tuple{Any,Any}`. However, this interpretation causes some practical problems.

First, a value of T needs to be available inside the method definition. For a call like `f(1, 1.0)`, it's not clear what T should be. It could be `Union{Int,Float64}`, or perhaps Real. Intuitively, we expect the declaration `x::T` to mean T `=== typeof(x)`. To make sure that invariant holds, we need `typeof(x) === typeof(y) === T` in this method. That implies the method should only be called for arguments of the exact same type.

It turns out that being able to dispatch on whether two values have the same type is very useful (this is used by the promotion system for example), so we have multiple reasons to want a different interpretation of `Tuple{T,T} where T`. To make this work we add the following rule to subtyping: if a variable occurs more than once in covariant position, it is restricted to ranging over only concrete types. ("Covariant position" means that only `Tuple` and `Union` types occur between an occurrence of a variable and the `UnionAll` type that introduces it.) Such variables are called "diagonal variables" or "concrete variables".

So for example, `Tuple{T,T} where T` can be seen as `Union{Tuple{Int8,Int8}, Tuple{Int16,Int16}, ...}`, where T ranges over all concrete types. This gives rise to some interesting subtyping results. For example `Tuple{Real,Real}` is not a subtype of `Tuple{T,T} where T`, because it includes some types like `Tuple{Int8,Int16}` where the two elements have different types. `Tuple{Real,Real}` and `Tuple{T,T} where T` have the non-trivial intersection `Tuple{T,T} where T<:Real`. However, `Tuple{Real}` *is* a subtype of `Tuple{T} where T`, because in that case T occurs only once and so is not diagonal.

Next consider a signature like the following:

```
f(a::Array{T}, x::T, y::T) where {T} = ...
```

In this case, T occurs in invariant position inside `Array{T}`. That means whatever type of array is passed unambiguously determines the value of T -- we say T has an *equality constraint* on it. Therefore in this case the diagonal rule is not really necessary, since the array determines T and we can then allow x and y to be of any subtypes of T. So variables that occur in invariant position are never considered diagonal. This choice of behavior is slightly controversial -- some feel this definition should be written as

```
f(a::Array{T}, x::S, y::S) where {T, S<:T} = ...
```

to clarify whether x and y need to have the same type. In this version of the signature they would, or we could introduce a third variable for the type of y if x and y can have different types.

The next complication is the interaction of unions and diagonal variables, e.g.

```
f(x::Union{Void,T}, y::T) where {T} = ...
```

Consider what this declaration means. y has type T. x then can have either the same type T, or else be of type Void. So all of the following calls should match:

```
f(1, 1)
f("", "")
f(2.0, 2.0)
f(nothing, 1)
f(nothing, "")
f(nothing, 2.0)
```

These examples are telling us something: when x is `nothing::Void`, there are no extra constraints on y. It is as if the method signature had `y::Any`. This means that whether a variable is diagonal is not a static property based on where it appears in a type. Rather, it depends on where a variable appears when the subtyping algorithm *uses* it. When x has type `Void`, we don't need to use the T in `Union{Void,T}`, so T does not "occur". Indeed, we have the following type equivalence:

```
(Tuple{Union{Void,T},T} where T) == Union{Tuple{Void,Any}, Tuple{T,T} where T}
```

Subtyping diagonal variables

The subtyping algorithm for diagonal variables has two components: (1) identifying variable occurrences, and (2) ensuring that diagonal variables range over concrete types only.

The first task is accomplished by keeping counters `occurs_inv` and `occurs_cov` (in `src/subtype.c`) for each variable in the environment, tracking the number of invariant and covariant occurrences, respectively. A variable is diagonal when `occurs_inv == 0 && occurs_cov > 1`.

The second task is accomplished by imposing a condition on a variable's lower bound. As the subtyping algorithm runs, it narrows the bounds of each variable (raising lower bounds and lowering upper bounds) to keep track of the range of variable values for which the subtype relation would hold. When we are done evaluating the body of a `UnionAll` type whose variable is diagonal, we look at the final values of the bounds. Since the variable must be concrete, a contradiction occurs if its lower bound could not be a subtype of a concrete type. For example, an abstract type like `AbstractArray` cannot be a subtype of a concrete type, but a concrete type like `Int` can be, and the empty type `Bottom` can be as well. If a lower bound fails this test the algorithm stops with the answer `false`.

For example, in the problem `Tuple{Int,String} <: Tuple{T,T} where T`, we derive that this would be true if T were a supertype of `Union{Int,String}`. However, `Union{Int,String}` is an abstract type, so the relation does not hold.

This concreteness test is done by the function `is_leaf_bound`. Note that this test is slightly different from `jl_is_leaf_type`, since it also returns `true` for `Bottom`. Currently this function is heuristic, and does not catch all possible concrete types. The difficulty is that whether a lower bound is concrete might depend on the values of other type variable bounds. For example, `Vector{T}` is equivalent to the concrete type `Vector{Int}` only if both the upper and lower bounds of T equal `Int`. We have not yet worked out a complete algorithm for this.

Introduction to the internal machinery

Most operations for dealing with types are found in the files `jltypes.c` and `subtype.c`. A good way to start is to watch subtyping in action. Build Julia with `make debug` and fire up Julia within a debugger. gdb debugging tips has some tips which may be useful.

Because the subtyping code is used heavily in the REPL itself–and hence breakpoints in this code get triggered often–it will be easiest if you make the following definition:

```julia
julia> function mysubtype(a,b)
           ccall(:jl_breakpoint, Void, (Any,), nothing)
           issubtype(a, b)
       end
```

and then set a breakpoint in `jl_breakpoint`. Once this breakpoint gets triggered, you can set breakpoints in other functions.

As a warm-up, try the following:

```
mysubtype(Tuple{Int,Float64}, Tuple{Integer,Real})
```

We can make it more interesting by trying a more complex case:

```
mysubtype(Tuple{Array{Int,2}, Int8}, Tuple{Array{T}, T} where T)
```

Subtyping and method sorting

The `type_morespecific` functions are used for imposing a partial order on functions in method tables (from most-to-least specific). Specificity is strict; if a is more specific than b, then a does not equal b and b is not more specific than a.

If a is a strict subtype of b, then it is automatically considered more specific. From there, `type_morespecific` employs some less formal rules. For example, `subtype` is sensitive to the number of arguments, but `type_more-specific` may not be. In particular, `Tuple{Int,AbstractFloat}` is more specific than `Tuple{Integer}`, even though it is not a subtype. (Of `Tuple{Int,AbstractFloat}` and `Tuple{Integer,Float64}`, neither is more specific than the other.) Likewise, `Tuple{Int,Vararg{Int}}` is not a subtype of `Tuple{Integer}`, but it is considered more specific. However, `morespecific` does get a bonus for length: in particular, `Tuple{Int,Int}` is more specific than `Tuple{Int,Vararg{Int}}`.

If you're debugging how methods get sorted, it can be convenient to define the function:

```
type_morespecific(a, b) = ccall(:jl_type_morespecific, Cint, (Any,Any), a, b)
```

which allows you to test whether tuple type a is more specific than tuple type b.

70.4   Memory layout of Julia Objects

Object layout (jl_value_t)

The `jl_value_t` struct is the name for a block of memory owned by the Julia Garbage Collector, representing the data associated with a Julia object in memory. Absent any type information, it is simply an opaque pointer:

```
typedef struct jl_value_t* jl_pvalue_t;
```

Each `jl_value_t` struct is contained in a `jl_typetag_t` struct that contains metadata information about the Julia object, such as its type and garbage collector (gc) reachability:

```
typedef struct {
    opaque metadata;
    jl_value_t value;
} jl_typetag_t;
```

The type of any Julia object is an instance of a leaf `jl_datatype_t` object. The `jl_typeof()` function can be used to query for it:

```
jl_value_t *jl_typeof(jl_value_t *v);
```

The layout of the object depends on its type. Reflection methods can be used to inspect that layout. A field can be accessed by calling one of the get-field methods:

```
jl_value_t *jl_get_nth_field_checked(jl_value_t *v, size_t i);
jl_value_t *jl_get_field(jl_value_t *o, char *fld);
```

If the field types are known, a priori, to be all pointers, the values can also be extracted directly as an array access:

```
jl_value_t *v = value->fieldptr[n];
```

As an example, a "boxed" `uint16_t` is stored as follows:

```
struct {
    opaque metadata;
    struct {
        uint16_t data;          // -- 2 bytes
    } jl_value_t;
};
```

This object is created by `jl_box_uint16()`. Note that the `jl_value_t` pointer references the data portion, not the metadata at the top of the struct.

A value may be stored "unboxed" in many circumstances (just the data, without the metadata, and possibly not even stored but just kept in registers), so it is unsafe to assume that the address of a box is a unique identifier. The "egal" test (corresponding to the `===` function in Julia), should instead be used to compare two unknown objects for equivalence:

```
int jl_egal(jl_value_t *a, jl_value_t *b);
```

This optimization should be relatively transparent to the API, since the object will be "boxed" on-demand, whenever a `jl_value_t` pointer is needed.

Note that modification of a `jl_value_t` pointer in memory is permitted only if the object is mutable. Otherwise, modification of the value may corrupt the program and the result will be undefined. The mutability property of a value can be queried for with:

```
int jl_is_mutable(jl_value_t *v);
```

If the object being stored is a `jl_value_t`, the Julia garbage collector must be notified also:

```
void jl_gc_wb(jl_value_t *parent, jl_value_t *ptr);
```

However, the Embedding Julia section of the manual is also required reading at this point, for covering other details of boxing and unboxing various types, and understanding the gc interactions.

Mirror structs for some of the built-in types are defined in `julia.h`. The corresponding global `jl_datatype_t` objects are created by `jl_init_types` in `jltypes.c`.

Garbage collector mark bits

The garbage collector uses several bits from the metadata portion of the `jl_typetag_t` to track each object in the system. Further details about this algorithm can be found in the comments of the garbage collector implementation in `gc.c`.

Object allocation

Most new objects are allocated by `jl_new_structv()`:

```
jl_value_t *jl_new_struct(jl_datatype_t *type, ...);
jl_value_t *jl_new_structv(jl_datatype_t *type, jl_value_t **args, uint32_t na);
```

Although, `isbits` objects can be also constructed directly from memory:

```
jl_value_t *jl_new_bits(jl_value_t *bt, void *data)
```

And some objects have special constructors that must be used instead of the above functions:

Types:

```
jl_datatype_t *jl_apply_type(jl_datatype_t *tc, jl_tuple_t *params);
jl_datatype_t *jl_apply_array_type(jl_datatype_t *type, size_t dim);
jl_uniontype_t *jl_new_uniontype(jl_tuple_t *types);
```

While these are the most commonly used options, there are more low-level constructors too, which you can find declared in `julia.h`. These are used in `jl_init_types()` to create the initial types needed to bootstrap the creation of the Julia system image.

Tuples:

```
jl_tuple_t *jl_tuple(size_t n, ...);
jl_tuple_t *jl_tuplev(size_t n, jl_value_t **v);
jl_tuple_t *jl_alloc_tuple(size_t n);
```

The representation of tuples is highly unique in the Julia object representation ecosystem. In some cases, a `Base.tuple()` object may be an array of pointers to the objects contained by the tuple equivalent to:

```
typedef struct {
    size_t length;
    jl_value_t *data[length];
} jl_tuple_t;
```

However, in other cases, the tuple may be converted to an anonymous `isbits` type and stored unboxed, or it may not stored at all (if it is not being used in a generic context as a jl_value_t*).

Symbols:

```
jl_sym_t *jl_symbol(const char *str);
```

Functions and MethodInstance:

```
jl_function_t *jl_new_generic_function(jl_sym_t *name);
jl_method_instance_t *jl_new_method_instance(jl_value_t *ast, jl_tuple_t *sparams);
```

Arrays:

```
jl_array_t *jl_new_array(jl_value_t *atype, jl_tuple_t *dims);
jl_array_t *jl_new_arrayv(jl_value_t *atype, ...);
jl_array_t *jl_alloc_array_1d(jl_value_t *atype, size_t nr);
jl_array_t *jl_alloc_array_2d(jl_value_t *atype, size_t nr, size_t nc);
jl_array_t *jl_alloc_array_3d(jl_value_t *atype, size_t nr, size_t nc, size_t z);
jl_array_t *jl_alloc_vec_any(size_t n);
```

Note that many of these have alternative allocation functions for various special-purposes. The list here reflects the more common usages, but a more complete list can be found by reading the `julia.h` header file.

Internal to Julia, storage is typically allocated by `newstruct()` (or `newobj()` for the special types):

```
jl_value_t *newstruct(jl_value_t *type);
jl_value_t *newobj(jl_value_t *type, size_t nfields);
```

And at the lowest level, memory is getting allocated by a call to the garbage collector (in `gc.c`), then tagged with its type:

```
jl_value_t *jl_gc_allocobj(size_t nbytes);
void jl_set_typeof(jl_value_t *v, jl_datatype_t *type);
```

Note that all objects are allocated in multiples of 4 bytes and aligned to the platform pointer size. Memory is allocated from a pool for smaller objects, or directly with `malloc()` for large objects.

Singleton Types

Singleton types have only one instance and no data fields. Singleton instances have a size of 0 bytes, and consist only of their metadata. e.g. `nothing::Void`.

See Singleton Types and Nothingness and missing values

70.5 Eval of Julia code

One of the hardest parts about learning how the Julia Language runs code is learning how all of the pieces work together to execute a block of code.

Each chunk of code typically makes a trip through many steps with potentially unfamiliar names, such as (in no particular order): flisp, AST, C++, LLVM, `eval`, `typeinf`, `macroexpand`, sysimg (or system image), bootstrapping, compile, parse, execute, JIT, interpret, box, unbox, intrinsic function, and primitive function, before turning into the desired result (hopefully).

Definitions

- REPL

  REPL stands for Read-Eval-Print Loop. It's just what we call the command line environment for short.

- AST

  Abstract Syntax Tree The AST is the digital representation of the code structure. In this form the code has been tokenized for meaning so that it is more suitable for manipulation and execution.

Julia Execution

The 10,000 foot view of the whole process is as follows:

1. The user starts `julia`.

2. The C function `main()` from `ui/repl.c` gets called. This function processes the command line arguments, filling in the `jl_options` struct and setting the variable ARGS. It then initializes Julia (by calling `julia_init in task.c`, which may load a previously compiled sysimg). Finally, it passes off control to Julia by calling `Base._start()`.

3. When `_start()` takes over control, the subsequent sequence of commands depends on the command line arguments given. For example, if a filename was supplied, it will proceed to execute that file. Otherwise, it will start an interactive REPL.

4. Skipping the details about how the REPL interacts with the user, let's just say the program ends up with a block of code that it wants to run.

5. If the block of code to run is in a file, `jl_load(char *filename)` gets invoked to load the file and parse it. Each fragment of code is then passed to `eval` to execute.

6. Each fragment of code (or AST), is handed off to `eval()` to turn into results.

7. `eval()` takes each code fragment and tries to run it in `jl_toplevel_eval_flex()`.

8. `jl_toplevel_eval_flex()` decides whether the code is a "toplevel" action (such as `using` or `module`), which would be invalid inside a function. If so, it passes off the code to the toplevel interpreter.

9. `jl_toplevel_eval_flex()` then expands the code to eliminate any macros and to "lower" the AST to make it simpler to execute.

10. `jl_toplevel_eval_flex()` then uses some simple heuristics to decide whether to JIT compiler the AST or to interpret it directly.

11. The bulk of the work to interpret code is handled by `eval` in `interpreter.c`.

12. If instead, the code is compiled, the bulk of the work is handled by `codegen.cpp`. Whenever a Julia function is called for the first time with a given set of argument types, type inference will be run on that function. This information is used by the codegen step to generate faster code.

13. Eventually, the user quits the REPL, or the end of the program is reached, and the `_start()` method returns.

14. Just before exiting, `main()` calls `jl_atexit_hook(exit_code)`. This calls `Base._atexit()` (which calls any functions registered to `atexit()` inside Julia). Then it calls `jl_gc_run_all_finalizers()`. Finally, it gracefully cleans up all `libuv` handles and waits for them to flush and close.

Parsing

The Julia parser is a small lisp program written in femtolisp, the source-code for which is distributed inside Julia in `src/flisp`.

The interface functions for this are primarily defined in `jlfrontend.scm`. The code in `ast.c` handles this handoff on the Julia side.

The other relevant files at this stage are `julia-parser.scm`, which handles tokenizing Julia code and turning it into an AST, and `julia-syntax.scm`, which handles transforming complex AST representations into simpler, "lowered" AST representations which are more suitable for analysis and execution.

Macro Expansion

When `eval()` encounters a macro, it expands that AST node before attempting to evaluate the expression. Macro expansion involves a handoff from `eval()` (in Julia), to the parser function `jl_macroexpand()` (written in `flisp`) to the Julia macro itself (written in - what else - Julia) via `fl_invoke_julia_macro()`, and back.

Typically, macro expansion is invoked as a first step during a call to `expand()`/`jl_expand()`, although it can also be invoked directly by a call to `macroexpand()`/`jl_macroexpand()`.

Type Inference

Type inference is implemented in Julia by `typeinf() in inference.jl`. Type inference is the process of examining a Julia function and determining bounds for the types of each of its variables, as well as bounds on the type of the return value from the function. This enables many future optimizations, such as unboxing of known immutable values, and compile-time hoisting of various run-time operations such as computing field offsets and function pointers. Type inference may also include other steps such as constant propagation and inlining.

More Definitions

- JIT
  Just-In-Time Compilation The process of generating native-machine code into memory right when it is needed.

- LLVM

  Low-Level Virtual Machine (a compiler) The Julia JIT compiler is a program/library called libLLVM. Codegen in Julia refers both to the process of taking a Julia AST and turning it into LLVM instructions, and the process of LLVM optimizing that and turning it into native assembly instructions.

- C++

  The programming language that LLVM is implemented in, which means that codegen is also implemented in this language. The rest of Julia's library is implemented in C, in part because its smaller feature set makes it more usable as a cross-language interface layer.

- box

  This term is used to describe the process of taking a value and allocating a wrapper around the data that is tracked by the garbage collector (gc) and is tagged with the object's type.

- unbox

  The reverse of boxing a value. This operation enables more efficient manipulation of data when the type of that data is fully known at compile-time (through type inference).

- generic function

  A Julia function composed of multiple "methods" that are selected for dynamic dispatch based on the argument type-signature

- anonymous function or "method"

  A Julia function without a name and without type-dispatch capabilities

- primitive function

  A function implemented in C but exposed in Julia as a named function "method" (albeit without generic function dispatch capabilities, similar to a anonymous function)

- intrinsic function

  A low-level operation exposed as a function in Julia. These pseudo-functions implement operations on raw bits such as add and sign extend that cannot be expressed directly in any other way. Since they operate on bits directly, they must be compiled into a function and surrounded by a call to `Core.Intrinsics.box(T, ...)` to reassign type information to the value.

JIT Code Generation

Codegen is the process of turning a Julia AST into native machine code.

The JIT environment is initialized by an early call to `jl_init_codegen in codegen.cpp`.

On demand, a Julia method is converted into a native function by the function `emit_function(jl_method_in-stance_t*)`. (note, when using the MCJIT (in LLVM v3.4+), each function must be JIT into a new module.) This function recursively calls `emit_expr()` until the entire function has been emitted.

Much of the remaining bulk of this file is devoted to various manual optimizations of specific code patterns. For example, `emit_known_call()` knows how to inline many of the primitive functions (defined in `builtins.c`) for various combinations of argument types.

Other parts of codegen are handled by various helper files:

- `debuginfo.cpp`

  Handles backtraces for JIT functions

- `ccall.cpp`

  Handles the ccall and llvmcall FFI, along with various `abi_*.cpp` files

- `intrinsics.cpp`

  Handles the emission of various low-level intrinsic functions

  Bootstrapping

  The process of creating a new system image is called "bootstrapping".

  The etymology of this word comes from the phrase "pulling oneself up by the bootstraps", and refers to the idea of starting from a very limited set of available functions and definitions and ending with the creation of a full-featured environment.

System Image

The system image is a precompiled archive of a set of Julia files. The `sys.ji` file distributed with Julia is one such system image, generated by executing the file `sysimg.jl`, and serializing the resulting environment (including Types, Functions, Modules, and all other defined values) into a file. Therefore, it contains a frozen version of the `Main`, `Core`, and `Base` modules (and whatever else was in the environment at the end of bootstrapping). This serializer/deserializer is implemented by `jl_save_system_image/jl_restore_system_image` in `dump.c`.

If there is no sysimg file (`jl_options.image_file == NULL`), this also implies that `--build` was given on the command line, so the final result should be a new sysimg file. During Julia initialization, minimal `Core` and `Main` modules are created. Then a file named `boot.jl` is evaluated from the current directory. Julia then evaluates any file given as a command line argument until it reaches the end. Finally, it saves the resulting environment to a "sysimg" file for use as a starting point for a future Julia run.

70.6   Calling Conventions

Julia uses three calling conventions for four distinct purposes:

Name	Prefix	Purpose
Native	`julia_`	Speed via specialized signatures
JL Call	`jlcall_`	Wrapper for generic calls
JL Call	`jl_`	Builtins
C ABI	`jlcapi_`	Wrapper callable from C

Julia Native Calling Convention

The native calling convention is designed for fast non-generic calls. It usually uses a specialized signature.

- LLVM ghosts (zero-length types) are omitted.

- LLVM scalars and vectors are passed by value.

- LLVM aggregates (arrays and structs) are passed by reference.

A small return values is returned as LLVM return values. A large return values is returned via the "structure return" (`sret`) convention, where the caller provides a pointer to a return slot.

An argument or return values thta is a homogeneous tuple is sometimes represented as an LLVM vector instead of an LLVM array.

JL Call Convention

The JL Call convention is for builtins and generic dispatch. Hand-written functions using this convention are declared via the macro `JL_CALLABLE`. The convention uses exactly 3 parameters:

- `F` - Julia representation of function that is being applied

- `args` - pointer to array of pointers to boxes

- `nargs` - length of the array

The return value is a pointer to a box.

C ABI

C ABI wrappers enable calling Julia from C. The wrapper calls a function using the native calling convention.

Tuples are always represented as C arrays.

70.7 High-level Overview of the Native-Code Generation Process

Representation of Pointers

When emitting code to an object file, pointers will be emitted as relocations. The deserialization code will ensure any object that pointed to one of these constants gets recreated and contains the right runtime pointer.

Otherwise, they will be emitted as literal constants.

To emit one of these objects, call `literal_pointer_val`. It'll handle tracking the Julia value and the LLVM global, ensuring they are valid both for the current runtime and after deserialization.

When emitted into the object file, these globals are stored as references in a large `gvals` table. This allows the deserializer to reference them by index, and implement a custom manual mechanism similar to a Global Offset Table (GOT) to restore them.

Function pointers are handled similarly. They are stored as values in a large `fvals` table. Like globals, this allows the deserializer to reference them by index.

Note that `extern` functions are handled separately, with names, via the usual symbol resolution mechanism in the linker.

Note too that `ccall` functions are also handled separately, via a manual GOT and Procedure Linkage Table (PLT).

Representation of Intermediate Values

Values are passed around in a `jl_cgval_t` struct. This represents an R-value, and includes enough information to determine how to assign or pass it somewhere.

They are created via one of the helper constructors, usually: `mark_julia_type` (for immediate values) and `mark_julia_slot` (for pointers to values).

The function `convert_julia_type` can transform between any two types. It returns an R-value with `cgval.typ` set to `typ`. It'll cast the object to the requested representation, making heap boxes, allocating stack copies, and computing tagged unions as needed to change the representation.

By contrast `update_julia_type` will change `cgval.typ` to `typ`, only if it can be done at zero-cost (i.e. without emitting any code).

Union representation

Inferred union types may be stack allocated via a tagged type representation.

The primitive routines that need to be able to handle tagged unions are:

- mark-type

- load-local

- store-local

- isa

- is

- emit_typeof

- emit_sizeof

- boxed

- unbox

- specialized cc-ret

Everything else should be possible to handle in inference by using these primitives to implement union-splitting.

The representation of the tagged-union is as a pair of `< void* union, byte selector >`. The selector is fixed-size as `byte & 0x7f`, and will union-tag the first 126 isbits. It records the one-based depth-first count into the type-union of the isbits objects inside. An index of zero indicates that the `union*` is actually a tagged heap-allocated `jl_value_t*`, and needs to be treated as normal for a boxed object rather than as a tagged union.

The high bit of the selector (`byte & 0x80`) can be tested to determine if the `void*` is actually a heap-allocated (`jl_value_t*`) box, thus avoiding the cost of re-allocating a box, while maintaining the ability to efficiently handle union-splitting based on the low bits.

It is guaranteed that `byte & 0x7f` is an exact test for the type, if the value can be represented by a tag – it will never be marked `byte = 0x80`. It is not necessary to also test the type-tag when testing `isa`.

The `union*` memory region may be allocated at *any* size. The only constraint is that it is big enough to contain the data currently specified by `selector`. It might not be big enough to contain the union of all types that could be stored there according to the associated Union type field. Use appropriate care when copying.

Specialized Calling Convention Signature Representation

A `jl_returninfo_t` object describes the calling convention details of any callable.

If any of the arguments or return type of a method can be represented unboxed, and the method is not varargs, it'll be given an optimized calling convention signature based on its `specTypes` and `rettype` fields.

The general principles are that:

- Primitive types get passed in int/float registers.

- Tuples of VecElement types get passed in vector registers.

- Structs get passed on the stack.

- Return values are handle similarly to arguments, with a size-cutoff at which they will instead be returned via a hidden sret argument.

The total logic for this is implemented by `get_specsig_function` and `deserves_sret`.

Additionally, if the return type is a union, it may be returned as a pair of values (a pointer and a tag). If the union values can be stack-allocated, then sufficient space to store them will also be passed as a hidden first argument. It is up to the callee whether the returned pointer will point to this space, a boxed object, or even other constant memory.

70.8 Julia Functions

This document will explain how functions, method definitions, and method tables work.

Method Tables

Every function in Julia is a generic function. A generic function is conceptually a single function, but consists of many definitions, or methods. The methods of a generic function are stored in a method table. Method tables (type `MethodTable`) are associated with TypeNames. A TypeName describes a family of parameterized types. For example `Complex{Float32}` and `Complex{Float64}` share the same `Complex` type name object.

All objects in Julia are potentially callable, because every object has a type, which in turn has a TypeName.

Function calls

Given the call `f(x,y)`, the following steps are performed: first, the method table to use is accessed as `typeof(f).name.mt`. Second, an argument tuple type is formed, `Tuple{typeof(f), typeof(x), typeof(y)}`. Note that the type of the function itself is the first element. This is because the type might have parameters, and so needs to take part in dispatch. This tuple type is looked up in the method table.

This dispatch process is performed by `jl_apply_generic`, which takes two arguments: a pointer to an array of the values f, x, and y, and the number of values (in this case 3).

Throughout the system, there are two kinds of APIs that handle functions and argument lists: those that accept the function and arguments separately, and those that accept a single argument structure. In the first kind of API, the "arguments" part does *not* contain information about the function, since that is passed separately. In the second kind of API, the function is the first element of the argument structure.

For example, the following function for performing a call accepts just an `args` pointer, so the first element of the args array will be the function to call:

```
jl_value_t *jl_apply(jl_value_t **args, uint32_t nargs)
```

This entry point for the same functionality accepts the function separately, so the `args` array does not contain the function:

```
jl_value_t *jl_call(jl_function_t *f, jl_value_t **args, int32_t nargs);
```

Adding methods

Given the above dispatch process, conceptually all that is needed to add a new method is (1) a tuple type, and (2) code for the body of the method. `jl_method_def` implements this operation. `jl_first_argument_datatype` is called to extract the relevant method table from what would be the type of the first argument. This is much more complicated than the corresponding procedure during dispatch, since the argument tuple type might be abstract. For example, we can define:

```
(::Union{Foo{Int},Foo{Int8}})(x) = 0
```

which works since all possible matching methods would belong to the same method table.

Creating generic functions

Since every object is callable, nothing special is needed to create a generic function. Therefore `jl_new_generic_func-tion` simply creates a new singleton (0 size) subtype of `Function` and returns its instance. A function can have a mnemonic "display name" which is used in debug info and when printing objects. For example the name of `Base.sin` is `sin`. By convention, the name of the created *type* is the same as the function name, with a # prepended. So `typeof(sin)` is `Base.#sin`.

Closures

A closure is simply a callable object with field names corresponding to captured variables. For example, the following code:

```
function adder(x)
    return y->x+y
end
```

is lowered to (roughly):

```
struct ##1{T}
    x::T
end

(_::##1)(y) = _.x + y

function adder(x)
    return ##1(x)
end
```

Constructors

A constructor call is just a call to a type. The type of most types is `DataType`, so the method table for `DataType` contains most constructor definitions. One wrinkle is the fallback definition that makes all types callable via `convert`:

```
(::Type{T}){T}(args...) = convert(T, args...)::T
```

In this definition the function type is abstract, which is not normally supported. To make this work, all subtypes of `Type` (`Type`, `UnionAll`, `Union`, and `DataType`) currently share a method table via special arrangement.

Builtins

The "builtin" functions, defined in the `Core` module, are:

```
=== typeof sizeof issubtype isa typeassert throw tuple getfield setfield! fieldtype
nfields isdefined arrayref arrayset arraysize applicable invoke apply_type _apply
_expr svec
```

These are all singleton objects whose types are subtypes of `Builtin`, which is a subtype of `Function`. Their purpose is to expose entry points in the run time that use the "jlcall" calling convention:

```
jl_value_t *(jl_value_t*, jl_value_t**, uint32_t)
```

The method tables of builtins are empty. Instead, they have a single catch-all method cache entry (`Tuple{Vararg{Any}}`) whose jlcall fptr points to the correct function. This is kind of a hack but works reasonably well.

Keyword arguments

Keyword arguments work by associating a special, hidden function object with each method table that has definitions with keyword arguments. This function is called the "keyword argument sorter" or "keyword sorter", or "kwsorter", and is stored in the `kwsorter` field of `MethodTable` objects. Every definition in the kwsorter function has the same arguments as some definition in the normal method table, except with a single `Array` argument prepended. This array contains alternating symbols and values that represent the passed keyword arguments. The kwsorter's job is to move keyword arguments into their canonical positions based on name, plus evaluate and substite any needed default value expressions. The result is a normal positional argument list, which is then passed to yet another function.

The easiest way to understand the process is to look at how a keyword argument method definition is lowered. The code:

```julia
function circle(center, radius; color = black, fill::Bool = true, options...)
    # draw
end
```

actually produces *three* method definitions. The first is a function that accepts all arguments (including keywords) as positional arguments, and includes the code for the method body. It has an auto-generated name:

```julia
function #circle#1(color, fill::Bool, options, circle, center, radius)
    # draw
end
```

The second method is an ordinary definition for the original `circle` function, which handles the case where no keyword arguments are passed:

```julia
function circle(center, radius)
    #circle#1(black, true, Any[], circle, center, radius)
end
```

This simply dispatches to the first method, passing along default values. Finally there is the kwsorter definition:

```julia
function (::Core.kwftype(typeof(circle)))(kw::Array, circle, center, radius)
    options = Any[]
    color = arg associated with :color, or black if not found
    fill = arg associated with :fill, or true if not found
    # push remaining elements of kw into options array
    #circle#1(color, fill, options, circle, center, radius)
end
```

The front end generates code to loop over the kw array and pick out arguments in the right order, evaluating default expressions when an argument is not found.

The function `Core.kwftype(t)` fetches (and creates, if necessary) the field `t.name.mt.kwsorter`.

This design has the feature that call sites that don't use keyword arguments require no special handling; everything works as if they were not part of the language at all. Call sites that do use keyword arguments are dispatched directly to the called function's kwsorter. For example the call:

```
circle((0,0), 1.0, color = red; other...)
```

is lowered to:

```
kwfunc(circle)(Any[:color,red,other...], circle, (0,0), 1.0)
```

The unpacking procedure represented here as `other...` actually further unpacks each *element* of `other`, expecting each one to contain two values (a symbol and a value). `kwfunc` (also in `Core`) fetches the kwsorter for the called function. Notice that the original `circle` function is passed through, to handle closures.

Compiler efficiency issues

Generating a new type for every function has potentially serious consequences for compiler resource use when combined with Julia's "specialize on all arguments by default" design. Indeed, the initial implementation of this design suffered from much longer build and test times, higher memory use, and a system image nearly 2x larger than the baseline. In a naive implementation, the problem is bad enough to make the system nearly unusable. Several significant optimizations were needed to make the design practical.

The first issue is excessive specialization of functions for different values of function-valued arguments. Many functions simply "pass through" an argument to somewhere else, e.g. to another function or to a storage location. Such functions do not need to be specialized for every closure that might be passed in. Fortunately this case is easy to distinguish by simply considering whether a function *calls* one of its arguments (i.e. the argument appears in "head position" somewhere). Performance-critical higher-order functions like map certainly call their argument function and so will still be specialized as expected. This optimization is implemented by recording which arguments are called during the `analyze-variables` pass in the front end. When `cache_method` sees an argument in the `Function` type hierarchy passed to a slot declared as `Any` or `Function`, it pretends the slot was declared as ANY (the "don't specialize" hint). This heuristic seems to be extremely effective in practice.

The next issue concerns the structure of method cache hash tables. Empirical studies show that the vast majority of dynamically-dispatched calls involve one or two arguments. In turn, many of these cases can be resolved by considering only the first argument. (Aside: proponents of single dispatch would not be surprised by this at all. However, this argument means "multiple dispatch is easy to optimize in practice", and that we should therefore use it, *not* "we should use single dispatch"!) So the method cache uses the type of the first argument as its primary key. Note, however, that this corresponds to the *second* element of the tuple type for a function call (the first element being the type of the function itself). Typically, type variation in head position is extremely low – indeed, the majority of functions belong to singleton types with no parameters. However, this is not the case for constructors, where a single method table holds constructors for every type. Therefore the Type method table is special-cased to use the *first* tuple type element instead of the second.

The front end generates type declarations for all closures. Initially, this was implemented by generating normal type declarations. However, this produced an extremely large number of constructors, all of which were trivial (simply passing all arguments through to new). Since methods are partially ordered, inserting all of these methods is O(n^2), plus there are just too many of them to keep around. This was optimized by generating `composite_type` expressions directly (bypassing default constructor generation), and using new directly to create closure instances. Not the prettiest thing ever, but you do what you gotta do.

The next problem was the `@test` macro, which generated a 0-argument closure for each test case. This is not really necessary, since each test case is simply run once in place. Therefore I modified `@test` to expand to a try-catch block that records the test result (true, false, or exception raised) and calls the test suite handler on it.

However this caused a new problem. When many tests are grouped together in a single function, e.g. a single top level expression, or some other test grouping function, that function could have a very large number of exception handlers. This triggered a kind of dataflow analysis worst case, where type inference spun around for minutes enumerating possible paths through the forest of handlers. This was fixed by simply bailing out of type inference when it encounters

more than some number of handlers (currently 25). Presumably no performance-critical function will have more than 25 exception handlers. If one ever does, I'm willing to raise the limit to 26.

A minor issue occurs during the bootstrap process due to storing all constructors in a single method table. In the second bootstrap step, where `inference.ji` is compiled using `inference0.ji`, constructors for `inference0`'s types remain in the table, so there are still references to the old inference module and `inference.ji` is 2x the size it should be. This was fixed in `dump.c` by filtering definitions from "replaced modules" out of method tables and caches before saving a system image. A "replaced module" is one that satisfies the condition `m != jl_get_global(m->parent, m->name)` – in other words, some newer module has taken its name and place.

Another type inference worst case was triggered by the following code from the QuadGK.jl package, formerly part of Base:

```julia
function do_quadgk(f, s, n, ::Type{Tw}, abstol, reltol, maxevals, nrm) where Tw
    if eltype(s) <: Real # check for infinite or semi-infinite intervals
        s1 = s[1]; s2 = s[end]; inf1 = isinf(s1); inf2 = isinf(s2)
        if inf1 || inf2
            if inf1 && inf2 # x = t/(1-t^2) coordinate transformation
                return do_quadgk(t -> begin t2 = t*t; den = 1 / (1 - t2);
                                            f(t*den) * (1+t2)*den*den; end,
                                 map(x -> isinf(x) ? copysign(one(x), x) : 2x / (1+hypot(1,2x)),
↪   s),
                                 n, Tw, abstol, reltol, maxevals, nrm)
            end
            s0,si = inf1 ? (s2,s1) : (s1,s2)
            if si < 0 # x = s0 - t/(1-t)
                return do_quadgk(t -> begin den = 1 / (1 - t);
                                            f(s0 - t*den) * den*den; end,
                                 reverse!(map(x -> 1 / (1 + 1 / (s0 - x)), s)),
                                 n, Tw, abstol, reltol, maxevals, nrm)
            else # x = s0 + t/(1-t)
                return do_quadgk(t -> begin den = 1 / (1 - t);
                                            f(s0 + t*den) * den*den; end,
                                 map(x -> 1 / (1 + 1 / (x - s0)), s),
                                 n, Tw, abstol, reltol, maxevals, nrm)
            end
        end
    end
```

This code has a 3-way tail recursion, where each call wraps the current function argument `f` in a different new closure. Inference must consider 3^n (where n is the call depth) possible signatures. This blows up way too quickly, so logic was added to `typeinf_uncached` to immediately widen any argument that is a subtype of `Function` and that grows in depth down the stack.

70.9   Base.Cartesian

The (non-exported) Cartesian module provides macros that facilitate writing multidimensional algorithms. It is hoped that Cartesian will not, in the long term, be necessary; however, at present it is one of the few ways to write compact and performant multidimensional code.

Principles of usage

A simple example of usage is:

```
@nloops 3 i A begin
    s += @nref 3 A i
end
```

which generates the following code:

```
for i_3 = 1:size(A,3)
    for i_2 = 1:size(A,2)
        for i_1 = 1:size(A,1)
            s += A[i_1,i_2,i_3]
        end
    end
end
```

In general, Cartesian allows you to write generic code that contains repetitive elements, like the nested loops in this example. Other applications include repeated expressions (e.g., loop unwinding) or creating function calls with variable numbers of arguments without using the "splat" construct (`i...`).

Basic syntax

The (basic) syntax of `@nloops` is as follows:

- The first argument must be an integer (*not* a variable) specifying the number of loops.

- The second argument is the symbol-prefix used for the iterator variable. Here we used `i`, and variables `i_1`, `i_2`, `i_3` were generated.

- The third argument specifies the range for each iterator variable. If you use a variable (symbol) here, it's taken as `1:size(A,dim)`. More flexibly, you can use the anonymous-function expression syntax described below.

- The last argument is the body of the loop. Here, that's what appears between the `begin...end`.

There are some additional features of `@nloops` described in the reference section.

`@nref` follows a similar pattern, generating `A[i_1,i_2,i_3]` from `@nref 3 A i`. The general practice is to read from left to right, which is why `@nloops` is `@nloops 3 i A expr` (as in `for i_2 = 1:size(A,2)`, where `i_2` is to the left and the range is to the right) whereas `@nref` is `@nref 3 A i` (as in `A[i_1,i_2,i_3]`, where the array comes first).

If you're developing code with Cartesian, you may find that debugging is easier when you examine the generated code, using `macroexpand`:

```
julia> macroexpand(:(@nref 2 A i))
:(A[i_1, i_2])
```

Supplying the number of expressions

The first argument to both of these macros is the number of expressions, which must be an integer. When you're writing a function that you intend to work in multiple dimensions, this may not be something you want to hard-code. If you're writing code that you need to work with older Julia versions, currently you should use the `@ngenerate` macro described in an older version of this documentation.

Starting in Julia 0.4-pre, the recommended approach is to use a `@generated function`. Here's an example:

```
@generated function mysum(A::Array{T,N}) where {T,N}
    quote
        s = zero(T)
        @nloops $N i A begin
            s += @nref $N A i
        end
        s
    end
end
```

Naturally, you can also prepare expressions or perform calculations before the quote block.

Anonymous-function expressions as macro arguments

Perhaps the single most powerful feature in Cartesian is the ability to supply anonymous-function expressions that get evaluated at parsing time. Let's consider a simple example:

```
@nexprs 2 j->(i_j = 1)
```

@nexprs generates n expressions that follow a pattern. This code would generate the following statements:

```
i_1 = 1
i_2 = 1
```

In each generated statement, an "isolated" j (the variable of the anonymous function) gets replaced by values in the range 1:2. Generally speaking, Cartesian employs a LaTeX-like syntax. This allows you to do math on the index j. Here's an example computing the strides of an array:

```
s_1 = 1
@nexprs 3 j->(s_{j+1} = s_j * size(A, j))
```

would generate expressions

```
s_1 = 1
s_2 = s_1 * size(A, 1)
s_3 = s_2 * size(A, 2)
s_4 = s_3 * size(A, 3)
```

Anonymous-function expressions have many uses in practice.

Macro reference   Base.Cartesian.@nloops – Macro.

```
@nloops N itersym rangeexpr bodyexpr
@nloops N itersym rangeexpr preexpr bodyexpr
@nloops N itersym rangeexpr preexpr postexpr bodyexpr
```

Generate N nested loops, using itersym as the prefix for the iteration variables. rangeexpr may be an anonymous-function expression, or a simple symbol var in which case the range is indices(var, d) for dimension d.

Optionally, you can provide "pre" and "post" expressions. These get executed first and last, respectively, in the body of each loop. For example:

```
@nloops 2 i A d -> j_d = min(i_d, 5) begin
    s += @nref 2 A j
end
```

would generate:

```
for i_2 = indices(A, 2)
    j_2 = min(i_2, 5)
    for i_1 = indices(A, 1)
        j_1 = min(i_1, 5)
        s += A[j_1, j_2]
    end
end
```

If you want just a post-expression, supply nothing for the pre-expression. Using parentheses and semicolons, you can supply multi-statement expressions.

source

Base.Cartesian.@nref – Macro.

```
@nref N A indexexpr
```

Generate expressions like A[i_1, i_2, ...]. indexexpr can either be an iteration-symbol prefix, or an anonymous-function expression.

```
julia> @macroexpand Base.Cartesian.@nref 3 A i
:(A[i_1, i_2, i_3])
```

source

Base.Cartesian.@nextract – Macro.

```
@nextract N esym isym
```

Generate N variables esym_1, esym_2, ..., esym_N to extract values from isym. isym can be either a Symbol or anonymous-function expression.

@nextract 2 x y would generate

```
x_1 = y[1]
x_2 = y[2]
```

while @nextract 3 x d->y[2d-1] yields

```
x_1 = y[1]
x_2 = y[3]
x_3 = y[5]
```

source

Base.Cartesian.@nexprs – Macro.

```
@nexprs N expr
```

Generate N expressions. expr should be an anonymous-function expression.

```
julia> @macroexpand Base.Cartesian.@nexprs 4 i -> y[i] = A[i+j]
quote
    y[1] = A[1 + j]
    y[2] = A[2 + j]
    y[3] = A[3 + j]
    y[4] = A[4 + j]
end
```

source

Base.Cartesian.@ncall – Macro.

```
@ncall N f sym...
```

Generate a function call expression. `sym` represents any number of function arguments, the last of which may be an anonymous-function expression and is expanded into `N` arguments.

For example `@ncall 3 func a` generates

```
func(a_1, a_2, a_3)
```

while `@ncall 2 func a b i->c[i]` yields

```
func(a, b, c[1], c[2])
```

source

Base.Cartesian.@ntuple – Macro.

```
@ntuple N expr
```

Generates an N-tuple. `@ntuple 2 i` would generate `(i_1, i_2)`, and `@ntuple 2 k->k+1` would generate `(2,3)`.

source

Base.Cartesian.@nall – Macro.

```
@nall N expr
```

Check whether all of the expressions generated by the anonymous-function expression `expr` evaluate to `true`.

`@nall 3 d->(i_d > 1)` would generate the expression `(i_1 > 1 && i_2 > 1 && i_3 > 1)`. This can be convenient for bounds-checking.

source

Base.Cartesian.@nany – Macro.

```
@nany N expr
```

Check whether any of the expressions generated by the anonymous-function expression `expr` evaluate to `true`.

`@nany 3 d->(i_d > 1)` would generate the expression `(i_1 > 1 || i_2 > 1 || i_3 > 1)`.

source

Base.Cartesian.@nif – Macro.

```
@nif N conditionexpr expr
@nif N conditionexpr expr elseexpr
```

Generates a sequence of `if ... elseif ... else ... end` statements. For example:

```
@nif 3 d->(i_d >= size(A,d)) d->(error("Dimension ", d, " too big")) d->println("All OK")
```

would generate:

```
if i_1 > size(A, 1)
    error("Dimension ", 1, " too big")
elseif i_2 > size(A, 2)
    error("Dimension ", 2, " too big")
else
    println("All OK")
end
```

source

70.10   Talking to the compiler (the `:meta` mechanism)

In some circumstances, one might wish to provide hints or instructions that a given block of code has special properties: you might always want to inline it, or you might want to turn on special compiler optimization passes. Starting with version 0.4, Julia has a convention that these instructions can be placed inside a `:meta` expression, which is typically (but not necessarily) the first expression in the body of a function.

`:meta` expressions are created with macros. As an example, consider the implementation of the `@inline` macro:

```
macro inline(ex)
    esc(isa(ex, Expr) ? pushmeta!(ex, :inline) : ex)
end
```

Here, `ex` is expected to be an expression defining a function. A statement like this:

```
@inline function myfunction(x)
    x*(x+3)
end
```

gets turned into an expression like this:

```
quote
    function myfunction(x)
        Expr(:meta, :inline)
        x*(x+3)
    end
end
```

`Base.pushmeta!(ex, :symbol, args...)` appends `:symbol` to the end of the `:meta` expression, creating a new `:meta` expression if necessary. If `args` is specified, a nested expression containing `:symbol` and these arguments is appended instead, which can be used to specify additional information.

To use the metadata, you have to parse these `:meta` expressions. If your implementation can be performed within Julia, `Base.popmeta!` is very handy: `Base.popmeta!(body, :symbol)` will scan a function *body* expression (one without the function signature) for the first `:meta` expression containing `:symbol`, extract any arguments, and return a tuple `(found::Bool, args::Array{Any})`. If the metadata did not have any arguments, or `:symbol` was not found, the `args` array will be empty.

Not yet provided is a convenient infrastructure for parsing `:meta` expressions from C++.

70.11 SubArrays

Julia's SubArray type is a container encoding a "view" of a parent `AbstractArray`. This page documents some of the design principles and implementation of SubArrays.

Indexing: cartesian vs. linear indexing

Broadly speaking, there are two main ways to access data in an array. The first, often called cartesian indexing, uses N indexes for an N -dimensional `AbstractArray`. For example, a matrix A (2-dimensional) can be indexed in cartesian style as A[i,j]. The second indexing method, referred to as linear indexing, uses a single index even for higher-dimensional objects. For example, if A = reshape(1:12, 3, 4), then the expression A[5] returns the value 5. Julia allows you to combine these styles of indexing: for example, a 3d array A3 can be indexed as A3[i,j], in which case i is interpreted as a cartesian index for the first dimension, and j is a linear index over dimensions 2 and 3.

For `Arrays`, linear indexing appeals to the underlying storage format: an array is laid out as a contiguous block of memory, and hence the linear index is just the offset (+1) of the corresponding entry relative to the beginning of the array. However, this is not true for many other `AbstractArray` types: examples include `SparseMatrixCSC`, arrays that require some kind of computation (such as interpolation), and the type under discussion here, `SubArray`. For these types, the underlying information is more naturally described in terms of cartesian indexes.

You can manually convert from a cartesian index to a linear index with `sub2ind`, and vice versa using `ind2sub`. `getindex` and `setindex!` functions for `AbstractArray` types may include similar operations.

While converting from a cartesian index to a linear index is fast (it's just multiplication and addition), converting from a linear index to a cartesian index is very slow: it relies on the `div` operation, which is one of the slowest low-level operations you can perform with a CPU. For this reason, any code that deals with `AbstractArray` types is best designed in terms of cartesian, rather than linear, indexing.

Index replacement

Consider making 2d slices of a 3d array:

```
S1 = view(A, :, 5, 2:6)
S2 = view(A, 5, :, 2:6)
```

`view` drops "singleton" dimensions (ones that are specified by an `Int`), so both S1 and S2 are two-dimensional Sub-Arrays. Consequently, the natural way to index these is with S1[i,j]. To extract the value from the parent array A, the natural approach is to replace S1[i,j] with A[i,5,(2:6)[j]] and S2[i,j] with A[5,i,(2:6)[j]].

The key feature of the design of SubArrays is that this index replacement can be performed without any runtime overhead.

SubArray design

Type parameters and fields

The strategy adopted is first and foremost expressed in the definition of the type:

```
struct SubArray{T,N,P,I,L} <: AbstractArray{T,N}
    parent::P
    indexes::I
    offset1::Int      # for linear indexing and pointer, only valid when L==true
    stride1::Int      # used only for linear indexing
    ...
end
```

SubArray has 5 type parameters. The first two are the standard element type and dimensionality. The next is the type of the parent AbstractArray. The most heavily-used is the fourth parameter, a Tuple of the types of the indices for each dimension. The final one, L, is only provided as a convenience for dispatch; it's a boolean that represents whether the index types support fast linear indexing. More on that later.

If in our example above A is a Array{Float64, 3}, our S1 case above would be a SubArray{Int64,2,Array{Int64,3},Tu-ple{Colon,Int64,UnitRange{Int64}},false}. Note in particular the tuple parameter, which stores the types of the indices used to create S1. Likewise,

```julia
julia> S1.indexes
(Colon(),5,2:6)
```

Storing these values allows index replacement, and having the types encoded as parameters allows one to dispatch to efficient algorithms.

Index translation

Performing index translation requires that you do different things for different concrete SubArray types. For example, for S1, one needs to apply the i, j indices to the first and third dimensions of the parent array, whereas for S2 one needs to apply them to the second and third. The simplest approach to indexing would be to do the type-analysis at runtime:

```julia
parentindexes = Array{Any}(0)
for thisindex in S.indexes
    ...
    if isa(thisindex, Int)
        # Don't consume one of the input indexes
        push!(parentindexes, thisindex)
    elseif isa(thisindex, AbstractVector)
        # Consume an input index
        push!(parentindexes, thisindex[inputindex[j]])
        j += 1
    elseif isa(thisindex, AbstractMatrix)
        # Consume two input indices
        push!(parentindexes, thisindex[inputindex[j], inputindex[j+1]])
        j += 2
    elseif ...
end
S.parent[parentindexes...]
```

Unfortunately, this would be disastrous in terms of performance: each element access would allocate memory, and involves the running of a lot of poorly-typed code.

The better approach is to dispatch to specific methods to handle each type of stored index. That's what reindex does: it dispatches on the type of the first stored index and consumes the appropriate number of input indices, and then it recurses on the remaining indices. In the case of S1, this expands to

```julia
Base.reindex(S1, S1.indexes, (i, j)) == (i, S1.indexes[2], S1.indexes[3][j])
```

for any pair of indices (i,j) (except CartesianIndexs and arrays thereof, see below).

This is the core of a SubArray; indexing methods depend upon reindex to do this index translation. Sometimes, though, we can avoid the indirection and make it even faster.

Linear indexing

Linear indexing can be implemented efficiently when the entire array has a single stride that separates successive elements, starting from some offset. This means that we can pre-compute these values and represent linear indexing simply as an addition and multiplication, avoiding the indirection of `reindex` and (more importantly) the slow computation of the cartesian coordinates entirely.

For `SubArray` types, the availability of efficient linear indexing is based purely on the types of the indices, and does not depend on values like the size of the parent array. You can ask whether a given set of indices supports fast linear indexing with the internal `Base.viewindexing` function:

```julia
julia> Base.viewindexing(S1.indexes)
IndexCartesian()

julia> Base.viewindexing(S2.indexes)
IndexLinear()
```

This is computed during construction of the `SubArray` and stored in the L type parameter as a boolean that encodes fast linear indexing support. While not strictly necessary, it means that we can define dispatch directly on `SubArray{T,N,A,I,true}` without any intermediaries.

Since this computation doesn't depend on runtime values, it can miss some cases in which the stride happens to be uniform:

```julia
julia> A = reshape(1:4*2, 4, 2)
4×2 Base.ReshapedArray{Int64,2,UnitRange{Int64},Tuple{}}:
 1  5
 2  6
 3  7
 4  8

julia> diff(A[2:2:4,:][:])
3-element Array{Int64,1}:
 2
 2
 2
```

A view constructed as `view(A, 2:2:4, :)` happens to have uniform stride, and therefore linear indexing indeed could be performed efficiently. However, success in this case depends on the size of the array: if the first dimension instead were odd,

```julia
julia> A = reshape(1:5*2, 5, 2)
5×2 Base.ReshapedArray{Int64,2,UnitRange{Int64},Tuple{}}:
 1   6
 2   7
 3   8
 4   9
 5  10

julia> diff(A[2:2:4,:][:])
3-element Array{Int64,1}:
 2
 3
 2
```

then `A[2:2:4, :]` does not have uniform stride, so we cannot guarantee efficient linear indexing. Since we have to base this decision based purely on types encoded in the parameters of the SubArray, `S = view(A, 2:2:4, :)` cannot implement efficient linear indexing.

A few details

- Note that the `Base.reindex` function is agnostic to the types of the input indices; it simply determines how and where the stored indices should be reindexed. It not only supports integer indices, but it supports non-scalar indexing, too. This means that views of views don't need two levels of indirection; they can simply re-compute the indices into the original parent array!

- Hopefully by now it's fairly clear that supporting slices means that the dimensionality, given by the parameter N, is not necessarily equal to the dimensionality of the parent array or the length of the `indexes` tuple. Neither do user-supplied indices necessarily line up with entries in the `indexes` tuple (e.g., the second user-supplied index might correspond to the third dimension of the parent array, and the third element in the `indexes` tuple).

  What might be less obvious is that the dimensionality of the stored parent array must be equal to the number of effective indices in the `indexes` tuple. Some examples:

  ```
  A = reshape(1:35, 5, 7) # A 2d parent Array
  S = view(A, 2:7)        # A 1d view created by linear indexing
  S = view(A, :, :, 1:1)  # Appending extra indices is supported
  ```

  Naively, you'd think you could just set `S.parent = A` and `S.indexes = (:,:,1:1)`, but supporting this dramatically complicates the reindexing process, especially for views of views. Not only do you need to dispatch on the types of the stored indices, but you need to examine whether a given index is the final one and "merge" any remaining stored indices together. This is not an easy task, and even worse: it's slow since it implicitly depends upon linear indexing.

  Fortunately, this is precisely the computation that `ReshapedArray` performs, and it does so linearly if possible. Consequently, `view` ensures that the parent array is the appropriate dimensionality for the given indices by reshaping it if needed. The inner `SubArray` constructor ensures that this invariant is satisfied.

- `CartesianIndex` and arrays thereof throw a nasty wrench into the `reindex` scheme. Recall that `reindex` simply dispatches on the type of the stored indices in order to determine how many passed indices should be used and where they should go. But with `CartesianIndex`, there's no longer a one-to-one correspondence between the number of passed arguments and the number of dimensions that they index into. If we return to the above example of `Base.reindex(S1, S1.indexes, (i, j))`, you can see that the expansion is incorrect for `i, j = CartesianIndex(), CartesianIndex(2,1)`. It should *skip* the `CartesianIndex()` entirely and return:

  ```
  (CartesianIndex(2,1)[1], S1.indexes[2], S1.indexes[3][CartesianIndex(2,1)[2]])
  ```

  Instead, though, we get:

  ```
  (CartesianIndex(), S1.indexes[2], S1.indexes[3][CartesianIndex(2,1)])
  ```

  Doing this correctly would require *combined* dispatch on both the stored and passed indices across all combinations of dimensionalities in an intractable manner. As such, `reindex` must never be called with `CartesianIndex` indices. Fortunately, the scalar case is easily handled by first flattening the `CartesianIndex` arguments to plain integers. Arrays of `CartesianIndex`, however, cannot be split apart into orthogonal pieces so easily. Before attempting to use `reindex`, `view` must ensure that there are no arrays of `CartesianIndex` in the argument list. If there are, it can simply "punt" by avoiding the `reindex` calculation entirely, constructing a nested `SubArray` with two levels of indirection instead.

70.12 System Image Building

Building the Julia system image

Julia ships with a preparsed system image containing the contents of the Base module, named `sys.ji`. This file is also precompiled into a shared library called `sys.{so,dll,dylib}` on as many platforms as possible, so as to give vastly improved startup times. On systems that do not ship with a precompiled system image file, one can be generated from the source files shipped in Julia's `DATAROOTDIR/julia/base` folder.

This operation is useful for multiple reasons. A user may:

- Build a precompiled shared library system image on a platform that did not ship with one, thereby improving startup times.

- Modify `Base`, rebuild the system image and use the new `Base` next time Julia is started.

- Include a `userimg.jl` file that includes packages into the system image, thereby creating a system image that has packages embedded into the startup environment.

Julia now ships with a script that automates the tasks of building the system image, wittingly named `build_sysimg.jl` that lives in `DATAROOTDIR/julia/`. That is, to include it into a current Julia session, type:

```
include(joinpath(JULIA_HOME, Base.DATAROOTDIR, "julia", "build_sysimg.jl"))
```

This will include a `build_sysimg()` function:

`BuildSysImg.build_sysimg` – Function.

```
build_sysimg(sysimg_path=default_sysimg_path, cpu_target="native", userimg_path=nothing;
    force=false)
```

Rebuild the system image. Store it in `sysimg_path`, which defaults to a file named `sys.ji` that sits in the same folder as `libjulia.{so,dylib}`, except on Windows where it defaults to `JULIA_HOME/../lib/julia/sys.ji`. Use the cpu instruction set given by `cpu_target`. Valid CPU targets are the same as for the `-C` option to `julia`, or the `-march` option to `gcc`. Defaults to `native`, which means to use all CPU instructions available on the current processor. Include the user image file given by `userimg_path`, which should contain directives such as `using MyPackage` to include that package in the new system image. New system image will not replace an older image unless `force` is set to true.

source

Note that this file can also be run as a script itself, with command line arguments taking the place of arguments passed to the `build_sysimg` function. For example, to build a system image in `/tmp/sys.{so,dll,dylib}`, with the core2 CPU instruction set, a user image of `~/userimg.jl` and `force` set to `true`, one would execute:

```
julia build_sysimg.jl /tmp/sys core2 ~/userimg.jl --force
```

70.13 Working with LLVM

This is not a replacement for the LLVM documentation, but a collection of tips for working on LLVM for Julia.

File	Description
`builtins.c`	Builtin functions
`ccall.cpp`	Lowering `ccall`
`cgutils.cpp`	Lowering utilities, notably for array and tuple accesses
`codegen.cpp`	Top-level of code generation, pass list, lowering builtins
`debuginfo.cpp`	Tracks debug information for JIT code
`disasm.cpp`	Handles native object file and JIT code diassembly
`gf.c`	Generic functions
`intrinsics.cpp`	Lowering intrinsics
`llvm-simdloop.cpp`	Custom LLVM pass for `@simd`
`sys.c`	I/O and operating system utility functions

Overview of Julia to LLVM Interface

Julia statically links in LLVM by default. Build with `USE_LLVM_SHLIB=1` to link dynamically.

The code for lowering Julia AST to LLVM IR or interpreting it directly is in directory `src/`.

Some of the `.cpp` files form a group that compile to a single object.

The difference between an intrinsic and a builtin is that a builtin is a first class function that can be used like any other Julia function. An intrinsic can operate only on unboxed data, and therefore its arguments must be statically typed.

Alias Analysis

Julia currently uses LLVM's Type Based Alias Analysis. To find the comments that document the inclusion relationships, look for `static MDNode*` in `src/codegen.cpp`.

The `-O` option enables LLVM's Basic Alias Analysis.

Building Julia with a different version of LLVM

The default version of LLVM is specified in `deps/Versions.make`. You can override it by creating a file called `Make.user` in the top-level directory and adding a line to it such as:

```
LLVM_VER = 3.5.0
```

Besides the LLVM release numerals, you can also use `LLVM_VER = svn` to bulid against the latest development version of LLVM.

Passing options to LLVM

You can pass options to LLVM using *debug* builds of Julia. To create a debug build, run `make debug`. The resulting executable is `usr/bin/julia-debug`. You can pass LLVM options to this executable via the environment variable `JULIA_LLVM_ARGS`. Here are example settings using `bash` syntax:

- `export JULIA_LLVM_ARGS = -print-after-all` dumps IR after each pass.

- `export JULIA_LLVM_ARGS = -debug-only=loop-vectorize` dumps LLVM `DEBUG(...)` diagnostics for loop vectorizer *if* you built Julia with `LLVM_ASSERTIONS=1`. Otherwise you will get warnings about "Unknown command line argument". Counter-intuitively, building Julia with `LLVM_DEBUG=1` is *not* enough to dump DEBUG diagnostics from a pass.

Improving LLVM optimizations for Julia

Improving LLVM code generation usually involves either changing Julia lowering to be more friendly to LLVM's passes, or improving a pass.

If you are planning to improve a pass, be sure to read the LLVM developer policy. The best strategy is to create a code example in a form where you can use LLVM's `opt` tool to study it and the pass of interest in isolation.

1. Create an example Julia code of interest.

2. Use `JULIA_LLVM_ARGS = -print-after-all` to dump the IR.

3. Pick out the IR at the point just before the pass of interest runs.

4. Strip the debug metadata and fix up the TBAA metadata by hand.

The last step is labor intensive. Suggestions on a better way would be appreciated.

70.14   printf() and stdio in the Julia runtime

Libuv wrappers for stdio

`julia.h` defines libuv wrappers for the `stdio.h` streams:

```
uv_stream_t *JL_STDIN;
uv_stream_t *JL_STDOUT;
uv_stream_t *JL_STDERR;
```

... and corresponding output functions:

```
int jl_printf(uv_stream_t *s, const char *format, ...);
int jl_vprintf(uv_stream_t *s, const char *format, va_list args);
```

These `printf` functions are used by the `.c` files in the `src/` and `ui/` directories wherever stdio is needed to ensure that output buffering is handled in a unified way.

In special cases, like signal handlers, where the full libuv infrastructure is too heavy, `jl_safe_printf()` can be used to write(2) directly to STDERR_FILENO:

```
void jl_safe_printf(const char *str, ...);
```

Interface between JL_STD* and Julia code

`Base.STDIN`, `Base.STDOUT` and `Base.STDERR` are bound to the JL_STD* libuv streams defined in the runtime.

Julia's `__init__()` function (in `base/sysimg.jl`) calls `reinit_stdio()` (in `base/stream.jl`) to create Julia objects for `Base.STDIN`, `Base.STDOUT` and `Base.STDERR`.

`reinit_stdio()` uses ccall to retrieve pointers to JL_STD* and calls `jl_uv_handle_type()` to inspect the type of each stream. It then creates a Julia `Base.IOStream`, `Base.TTY` or `Base.PipeEndpoint` object to represent each stream, e.g.:

```
$ julia -e 'println(typeof((STDIN, STDOUT, STDERR)))'
Tuple{Base.TTY,Base.TTY,Base.TTY}

$ julia -e 'println(typeof((STDIN, STDOUT, STDERR)))' < /dev/null 2>/dev/null
Tuple{IOStream,Base.TTY,IOStream}

$ echo hello | julia -e 'println(typeof((STDIN, STDOUT, STDERR)))' | cat
Tuple{Base.PipeEndpoint,Base.PipeEndpoint,Base.TTY}
```

The `Base.read()` and `Base.write()` methods for these streams use `ccall` to call libuv wrappers in `src/jl_uv.c`, e.g.:

```
stream.jl: function write(s::IO, p::Ptr, nb::Integer)
               -> ccall(:jl_uv_write, ...)
  jl_uv.c:         -> int jl_uv_write(uv_stream_t *stream, ...)
                     -> uv_write(uvw, stream, buf, ...)
```

printf() during initialization

The libuv streams relied upon by `jl_printf()` etc., are not available until midway through initialization of the runtime (see `init.c`, `init_stdio()`). Error messages or warnings that need to be printed before this are routed to the standard C library `fwrite()` function by the following mechanism:

In `sys.c`, the JL_STD* stream pointers are statically initialized to integer constants: STD*_FILENO (0, 1 and 2). In `jl_uv.c` the `jl_uv_puts()` function checks its `uv_stream_t*` stream argument and calls `fwrite()` if stream is set to STDOUT_FILENO or STDERR_FILENO.

This allows for uniform use of `jl_printf()` throughout the runtime regardless of whether or not any particular piece of code is reachable before initialization is complete.

Legacy `ios.c` library

The `src/support/ios.c` library is inherited from femtolisp. It provides cross-platform buffered file IO and in-memory temporary buffers.

`ios.c` is still used by:

- `src/flisp/*.c`

- `src/dump.c` – for serialization file IO and for memory buffers.

- `base/iostream.jl` – for file IO (see `base/fs.jl` for libuv equivalent).

Use of `ios.c` in these modules is mostly self-contained and separated from the libuv I/O system. However, there is one place where femtolisp calls through to `jl_printf()` with a legacy `ios_t` stream.

There is a hack in `ios.h` that makes the `ios_t.bm` field line up with the `uv_stream_t.type` and ensures that the values used for `ios_t.bm` to not overlap with valid UV_HANDLE_TYPE values. This allows `uv_stream_t` pointers to point to `ios_t` streams.

This is needed because `jl_printf()` caller `jl_static_show()` is passed an `ios_t` stream by femtolisp's `fl_print()` function. Julia's `jl_uv_puts()` function has special handling for this:

```
if (stream->type > UV_HANDLE_TYPE_MAX) {
    return ios_write((ios_t*)stream, str, n);
}
```

70.15   Bounds checking

Like many modern programming languages, Julia uses bounds checking to ensure program safety when accessing arrays. In tight inner loops or other performance critical situations, you may wish to skip these bounds checks to improve runtime performance. For instance, in order to emit vectorized (SIMD) instructions, your loop body cannot contain branches, and thus cannot contain bounds checks. Consequently, Julia includes an `@inbounds(...)` macro to tell the compiler to skip such bounds checks within the given block. For the built-in `Array` type, the magic happens inside the `arrayref` and `arrayset` intrinsics. User-defined array types instead use the `@boundscheck(...)` macro to achieve context-sensitive code selection.

Eliding bounds checks

The @boundscheck(...) macro marks blocks of code that perform bounds checking. When such blocks appear inside of an @inbounds(...) block, the compiler removes these blocks. When the @boundscheck(...) is nested inside of a calling function containing an @inbounds(...), the compiler will remove the @boundscheck block *only if it is inlined* into the calling function. For example, you might write the method sum as:

```julia
function sum(A::AbstractArray)
    r = zero(eltype(A))
    for i = 1:length(A)
        @inbounds r += A[i]
    end
    return r
end
```

With a custom array-like type MyArray having:

```julia
@inline getindex(A::MyArray, i::Real) = (@boundscheck checkbounds(A,i); A.data[to_index(i)])
```

Then when getindex is inlined into sum, the call to checkbounds(A,i) will be elided. If your function contains multiple layers of inlining, only @boundscheck blocks at most one level of inlining deeper are eliminated. The rule prevents unintended changes in program behavior from code further up the stack.

Propagating inbounds

There may be certain scenarios where for code-organization reasons you want more than one layer between the @inbounds and @boundscheck declarations. For instance, the default getindex methods have the chain getindex(A::AbstractArray, i::Real) calls getindex(IndexStyle(A), A, i) calls _getindex(::IndexLinear, A, i).

To override the "one layer of inlining" rule, a function may be marked with @propagate_inbounds to propagate an inbounds context (or out of bounds context) through one additional layer of inlining.

The bounds checking call hierarchy

The overall hierarchy is:

- checkbounds(A, I...) which calls
    - checkbounds(Bool, A, I...) which calls
        * checkbounds_indices(Bool, indices(A), I) which recursively calls
            · checkindex for each dimension

Here A is the array, and I contains the "requested" indices. indices(A) returns a tuple of "permitted" indices of A.

checkbounds(A, I...) throws an error if the indices are invalid, whereas checkbounds(Bool, A, I...) returns false in that circumstance. checkbounds_indices discards any information about the array other than its indices tuple, and performs a pure indices-vs-indices comparison: this allows relatively few compiled methods to serve a huge variety of array types. Indices are specified as tuples, and are usually compared in a 1-1 fashion with individual dimensions handled by calling another important function, checkindex: typically,

```julia
checkbounds_indices(Bool, (IA1, IA...), (I1, I...)) = checkindex(Bool, IA1, I1) &
                                                      checkbounds_indices(Bool, IA, I)
```

so `checkindex` checks a single dimension. All of these functions, including the unexported `checkbounds_indices` have docstrings accessible with `?` .

If you have to customize bounds checking for a specific array type, you should specialize `checkbounds(Bool, A, I...)`. However, in most cases you should be able to rely on `checkbounds_indices` as long as you supply useful `indices` for your array type.

If you have novel index types, first consider specializing `checkindex`, which handles a single index for a particular dimension of an array. If you have a custom multidimensional index type (similar to `CartesianIndex`), then you may have to consider specializing `checkbounds_indices`.

Note this hierarchy has been designed to reduce the likelihood of method ambiguities. We try to make `checkbounds` the place to specialize on array type, and try to avoid specializations on index types; conversely, `checkindex` is intended to be specialized only on index type (especially, the last argument).

70.16   Proper maintenance and care of multi-threading locks

The following strategies are used to ensure that the code is dead-lock free (generally by addressing the 4th Coffman condition: circular wait).

1. structure code such that only one lock will need to be acquired at a time

2. always acquire shared locks in the same order, as given by the table below

3. avoid constructs that expect to need unrestricted recursion

Locks

Below are all of the locks that exist in the system and the mechanisms for using them that avoid the potential for deadlocks (no Ostrich algorithm allowed here):

The following are definitely leaf locks (level 1), and must not try to acquire any other lock:

- safepoint

    Note that this lock is acquired implicitly by `JL_LOCK` and `JL_UNLOCK`. use the `_NOGC` variants to avoid that for level 1 locks.
    While holding this lock, the code must not do any allocation or hit any safepoints. Note that there are safepoints when doing allocation, enabling / disabling GC, entering / restoring exception frames, and taking / releasing locks.

- shared_map

- finalizers

- pagealloc

- gc_perm_lock

- flisp

    flisp itself is already threadsafe, this lock only protects the `jl_ast_context_list_t` pool

The following is a leaf lock (level 2), and only acquires level 1 locks (safepoint) internally:

- typecache

The following is a level 3 lock, which can only acquire level 1 or level 2 locks internally:

- Method->writelock

The following is a level 4 lock, which can only recurse to acquire level 1, 2, or 3 locks:

- MethodTable->writelock

No Julia code may be called while holding a lock above this point.

The following is a level 6 lock, which can only recurse to acquire locks at lower levels:

- codegen

The following is an almost root lock (level end-1), meaning only the root look may be held when trying to acquire it:

- typeinf

    this one is perhaps one of the most tricky ones, since type-inference can be invoked from many points
    currently the lock is merged with the codegen lock, since they call each other recursively

The following is the root lock, meaning no other lock shall be held when trying to acquire it:

- toplevel

    this should be held while attempting a top-level action (such as making a new type or defining a new method): trying to obtain this lock inside a staged function will cause a deadlock condition!
    additionally, it's unclear if *any* code can safely run in parallel with an arbitrary toplevel expression, so it may require all threads to get to a safepoint first

Broken Locks

The following locks are broken:

- toplevel

    doesn't exist right now
    fix: create it

Shared Global Data Structures

These data structures each need locks due to being shared mutable global state. It is the inverse list for the above lock priority list. This list does not include level 1 leaf resources due to their simplicity.

MethodTable modifications (def, cache, kwsorter type) : MethodTable->writelock

Type declarations : toplevel lock

Type application : typecache lock

Module serializer : toplevel lock

JIT & type-inference : codegen lock

MethodInstance updates : codegen lock

- These fields are generally lazy initialized, using the test-and-test-and-set pattern.

- These are set at construction and immutable:

  - specTypes
  - sparam_vals
  - def

- These are set by `jl_type_infer` (while holding codegen lock):

  - rettype
  - inferred
  - these can also be reset, see `jl_set_lambda_rettype` for that logic as it needs to keep `functionObjectsDecls` in sync

- `inInference` flag:

  - optimization to quickly avoid recurring into `jl_type_infer` while it is already running
  - actual state (of setting `inferred`, then `fptr`) is protected by codegen lock

- Function pointers (`jlcall_api` and `fptr`, `unspecialized_ducttape`):

  - these transition once, from NULL to a value, while the codegen lock is held

- Code-generator cache (the contents of `functionObjectsDecls`):

  - these can transition multiple times, but only while the codegen lock is held
  - it is valid to use old version of this, or block for new versions of this, so races are benign, as long as the code is careful not to reference other data in the method instance (such as `rettype`) and assume it is coordinated, unless also holding the codegen lock

- `compile_traced` flag:

  - unknown

LLVMContext : codegen lock

Method : Method->writelock

- roots array (serializer and codegen)

- invoke / specializations / tfunc modifications

70.17 Arrays with custom indices

Julia 0.5 adds experimental support for arrays with arbitrary indices. Conventionally, Julia's arrays are indexed starting at 1, whereas some other languages start numbering at 0, and yet others (e.g., Fortran) allow you to specify arbitrary starting indices. While there is much merit in picking a standard (i.e., 1 for Julia), there are some algorithms which simplify considerably if you can index outside the range `1:size(A,d)` (and not just `0:size(A,d)-1`, either). Such array types are expected to be supplied through packages.

The purpose of this page is to address the question, "what do I have to do to support such arrays in my own code?" First, let's address the simplest case: if you know that your code will never need to handle arrays with unconventional indexing, hopefully the answer is "nothing." Old code, on conventional arrays, should function essentially without alteration as long as it was using the exported interfaces of Julia.

Generalizing existing code

As an overview, the steps are:

- replace many uses of `size` with `indices`

- replace `1:length(A)` with `linearindices(A)`, and `length(A)` with `length(linearindices(A))`

- replace explicit allocations like `Array{Int}(size(B))` with `similar(Array{Int}, indices(B))`

These are described in more detail below.

Background

Because unconventional indexing breaks deeply-held assumptions throughout the Julia ecosystem, early adopters running code that has not been updated are likely to experience errors. The most frustrating bugs would be incorrect results or segfaults (total crashes of Julia). For example, consider the following function:

```
function mycopy!(dest::AbstractVector, src::AbstractVector)
    length(dest) == length(src) || throw(DimensionMismatch("vectors must match"))
    # OK, now we're safe to use @inbounds, right? (not anymore!)
    for i = 1:length(src)
        @inbounds dest[i] = src[i]
    end
    dest
end
```

This code implicitly assumes that vectors are indexed from 1. Previously that was a safe assumption, so this code was fine, but (depending on what types the user passes to this function) it may no longer be safe. If this code continued to work when passed a vector with non-1 indices, it would either produce an incorrect answer or it would segfault. (If you do get segfaults, to help locate the cause try running julia with the option `--check-bounds=yes`.)

To ensure that such errors are caught, in Julia 0.5 both `length` and `size`**should** throw an error when passed an array with non-1 indexing. This is designed to force users of such arrays to check the code, and inspect it for whether it needs to be generalized.

Using `indices` for bounds checks and loop iteration

`indices(A)` (reminiscent of `size(A)`) returns a tuple of `AbstractUnitRange` objects, specifying the range of valid indices along each dimension of A. When A has unconventional indexing, the ranges may not start at 1. If you just want the range for a particular dimension d, there is `indices(A, d)`.

Base implements a custom range type, `OneTo`, where `OneTo(n)` means the same thing as `1:n` but in a form that guarantees (via the type system) that the lower index is 1. For any new `AbstractArray` type, this is the default returned by `indices`, and it indicates that this array type uses "conventional" 1-based indexing. Note that if you don't want to be bothered supporting arrays with non-1 indexing, you can add the following line:

```
@assert all(x->isa(x, Base.OneTo), indices(A))
```

at the top of any function.

For bounds checking, note that there are dedicated functions `checkbounds` and `checkindex` which can sometimes simplify such tests.

Linear indexing (`linearindices`)

Some algorithms are most conveniently (or efficiently) written in terms of a single linear index, `A[i]` even if A is multi-dimensional. In "true" linear indexing, the indices always range from `1:length(A)`. However, this raises an ambiguity for one-dimensional arrays (a.k.a., `AbstractVector`): does `v[i]` mean linear indexing, or Cartesian indexing with the array's native indices?

For this reason, if you want to use linear indexing in an algorithm, your best option is to get the index range by calling `linearindices(A)`. This will return `indices(A, 1)` if A is an `AbstractVector`, and the equivalent of `1:length(A)` otherwise.

In a sense, one can say that 1-dimensional arrays always use Cartesian indexing. To help enforce this, it's worth noting that `sub2ind(shape, i...)` and `ind2sub(shape, ind)` will throw an error if shape indicates a 1-dimensional array with unconventional indexing (i.e., is a `Tuple{UnitRange}` rather than a tuple of `OneTo`). For arrays with conventional indexing, these functions continue to work the same as always.

Using `indices` and `linearindices`, here is one way you could rewrite `mycopy!`:

```julia
function mycopy!(dest::AbstractVector, src::AbstractVector)
    indices(dest) == indices(src) || throw(DimensionMismatch("vectors must match"))
    for i in linearindices(src)
        @inbounds dest[i] = src[i]
    end
    dest
end
```

Allocating storage using generalizations of `similar`

Storage is often allocated with `Array{Int}(dims)` or `similar(A, args...)`. When the result needs to match the indices of some other array, this may not always suffice. The generic replacement for such patterns is to use `similar(storagetype, shape)`. `storagetype` indicates the kind of underlying "conventional" behavior you'd like, e.g., `Array{Int}` or `BitArray` or even `dims->zeros(Float32, dims)` (which would allocate an all-zeros array). `shape` is a tuple of `Integer` or `AbstractUnitRange` values, specifying the indices that you want the result to use.

Let's walk through a couple of explicit examples. First, if A has conventional indices, then `similar(Array{Int}, indices(A))` would end up calling `Array{Int}(size(A))`, and thus return an array. If A is an `AbstractArray` type with unconventional indexing, then `similar(Array{Int}, indices(A))` should return something that "behaves like" an `Array{Int}` but with a shape (including indices) that matches A. (The most obvious implementation is to allocate an `Array{Int}(size(A))` and then "wrap" it in a type that shifts the indices.)

Note also that `similar(Array{Int}, (indices(A, 2),))` would allocate an `AbstractVector{Int}` (i.e., 1-dimensional array) that matches the indices of the columns of A.

Deprecations

In generalizing Julia's code base, at least one deprecation was unavoidable: earlier versions of Julia defined `first(::Colon)` = 1, meaning that the first index along a dimension indexed by `:` is 1. This definition can no longer be justified, so it was deprecated. There is no provided replacement, because the proper replacement depends on what you are doing and might need to know more about the array. However, it appears that many uses of `first(::Colon)` are really about computing an index offset; when that is the case, a candidate replacement is:

```julia
indexoffset(r::AbstractVector) = first(r) - 1
indexoffset(::Colon) = 0
```

In other words, while `first(:)` does not itself make sense, in general you can say that the offset associated with a colon-index is zero.

Writing custom array types with non-1 indexing

Most of the methods you'll need to define are standard for any `AbstractArray` type, see Abstract Arrays. This page focuses on the steps needed to define unconventional indexing.

Do not implement `size` or `length`

Perhaps the majority of pre-existing code that uses `size` will not work properly for arrays with non-1 indices. For that reason, it is much better to avoid implementing these methods, and use the resulting `MethodError` to identify code that needs to be audited and perhaps generalized.

Do not annotate bounds checks

Julia 0.5 includes @boundscheck to annotate code that can be removed for callers that exploit @inbounds. Initially, it seems far preferable to run with bounds checking always enabled (i.e., omit the @boundscheck annotation so the check always runs).

Custom `AbstractUnitRange` types

If you're writing a non-1 indexed array type, you will want to specialize `indices` so it returns a `UnitRange`, or (perhaps better) a custom `AbstractUnitRange`. The advantage of a custom type is that it "signals" the allocation type for functions like `similar`. If we're writing an array type for which indexing will start at 0, we likely want to begin by creating a new `AbstractUnitRange`, `ZeroRange`, where `ZeroRange(n)` is equivalent to `0:n-1`.

In general, you should probably *not* export ZeroRange from your package: there may be other packages that implement their own ZeroRange, and having multiple distinct ZeroRange types is (perhaps counterintuitively) an advantage: `ModuleA.ZeroRange` indicates that `similar` should create a `ModuleA.ZeroArray`, whereas `ModuleB.ZeroRange` indicates a `ModuleB.ZeroArray` type. This design allows peaceful coexistence among many different custom array types.

Note that the Julia package CustomUnitRanges.jl can sometimes be used to avoid the need to write your own ZeroRange type.

Specializing `indices`

Once you have your `AbstractUnitRange` type, then specialize `indices`:

```
Base.indices(A::ZeroArray) = map(n->ZeroRange(n), A.size)
```

where here we imagine that `ZeroArray` has a field called `size` (there would be other ways to implement this).

In some cases, the fallback definition for `indices(A, d)`:

```
indices(A::AbstractArray{T,N}, d) where {T,N} = d <= N ? indices(A)[d] : OneTo(1)
```

may not be what you want: you may need to specialize it to return something other than `OneTo(1)` when `d > ndims(A)`. Likewise, in `Base` there is a dedicated function `indices1` which is equivalent to `indices(A, 1)` but which avoids checking (at runtime) whether `ndims(A) > 0`. (This is purely a performance optimization.) It is defined as:

```
indices1(A::AbstractArray{T,0}) where {T} = OneTo(1)
indices1(A::AbstractArray) = indices(A)[1]
```

If the first of these (the zero-dimensional case) is problematic for your custom array type, be sure to specialize it appropriately.

Specializing `similar`

Given your custom `ZeroRange` type, then you should also add the following two specializations for `similar`:

```julia
function Base.similar(A::AbstractArray, T::Type, shape::Tuple{ZeroRange,Vararg{ZeroRange}})
    # body
end

function Base.similar(f::Union{Function,DataType}, shape::Tuple{ZeroRange,Vararg{ZeroRange}})
    # body
end
```

Both of these should allocate your custom array type.

Specializing `reshape`

Optionally, define a method

```julia
Base.reshape(A::AbstractArray, shape::Tuple{ZeroRange,Vararg{ZeroRange}}) = ...
```

and you can `reshape` an array so that the result has custom indices.

Summary

Writing code that doesn't make assumptions about indexing requires a few extra abstractions, but hopefully the necessary changes are relatively straightforward.

As a reminder, this support is still experimental. While much of Julia's base code has been updated to support unconventional indexing, without a doubt there are many omissions that will be discovered only through usage. Moreover, at the time of this writing, most packages do not support unconventional indexing. As a consequence, early adopters should be prepared to identify and/or fix bugs. On the other hand, only through practical usage will it become clear whether this experimental feature should be retained in future versions of Julia; consequently, interested parties are encouraged to accept some ownership for putting it through its paces.

70.18   Base.LibGit2

The LibGit2 module provides bindings to libgit2, a portable C library that implements core functionality for the Git version control system. These bindings are currently used to power Julia's package manager. It is expected that this module will eventually be moved into a separate package.

Functionality

Some of this documentation assumes some prior knowledge of the libgit2 API. For more information on some of the objects and methods referenced here, consult the upstream libgit2 API reference.

`Base.LibGit2.AbstractCredentials` – Type.

> Abstract credentials payload
>
> source

`Base.LibGit2.Buffer` – Type.

> `LibGit2.Buffer`

A data buffer for exporting data from libgit2. Matches the `git_buf` struct.

When fetching data from LibGit2, a typical usage would look like:

```
buf_ref = Ref(Buffer())
@check ccall(..., (Ptr{Buffer},), buf_ref)
# operation on buf_ref
free(buf_ref)
```

In particular, note that `LibGit2.free` should be called afterward on the `Ref` object.

source

Base.LibGit2.CachedCredentials – Type.

Credentials that support caching

source

Base.LibGit2.CheckoutOptions – Type.

```
LibGit2.CheckoutOptions
```

Matches the `git_checkout_options` struct.

source

Base.LibGit2.CloneOptions – Type.

```
LibGit2.CloneOptions
```

Matches the `git_clone_options` struct.

source

Base.LibGit2.DiffDelta – Type.

```
LibGit2.DiffDelta
```

Description of changes to one entry. Matches the `git_diff_delta` struct.

The fields represent:

- `status`: One of `Consts.DELTA_STATUS`, indicating whether the file has been added/modified/deleted.
- `flags`: Flags for the delta and the objects on each side. Determines whether to treat the file(s) as binary/-text, whether they exist on each side of the diff, and whether the object ids are known to be correct.
- `similarity`: Used to indicate if a file has been renamed or copied.
- `nfiles`: The number of files in the delta (for instance, if the delta was run on a submodule commit id, it may contain more than one file).
- `old_file`: A `DiffFile` containing information about the file(s) before the changes.
- `new_file`: A `DiffFile` containing information about the file(s) after the changes.

source

Base.LibGit2.DiffFile – Type.

```
LibGit2.DiffFile
```

Description of one side of a delta. Matches the `git_diff_file` struct.

> source

`Base.LibGit2.DiffOptionsStruct` – Type.

> LibGit2.DiffOptionsStruct

Matches the `git_diff_options` struct.

> source

`Base.LibGit2.FetchHead` – Type.

> LibGit2.FetchHead

Contains the information about HEAD during a fetch, including the name and URL of the branch fetched from, the oid of the HEAD, and whether the fetched HEAD has been merged locally.

> source

`Base.LibGit2.FetchOptions` – Type.

> LibGit2.FetchOptions

Matches the `git_fetch_options` struct.

> source

`Base.LibGit2.GitBlob` – Type.

> ```
> GitBlob(repo::GitRepo, hash::AbstractGitHash)
> GitBlob(repo::GitRepo, spec::AbstractString)
> ```

Return a `GitBlob` object from `repo` specified by hash/spec.

- hash is a full (`GitHash`) or partial (`GitShortHash`) hash.
- spec is a textual specification: see the git docs for a full list.

> source

`Base.LibGit2.GitCommit` – Type.

> ```
> GitCommit(repo::GitRepo, hash::AbstractGitHash)
> GitCommit(repo::GitRepo, spec::AbstractString)
> ```

Return a `GitCommit` object from `repo` specified by hash/spec.

- hash is a full (`GitHash`) or partial (`GitShortHash`) hash.
- spec is a textual specification: see the git docs for a full list.

> source

`Base.LibGit2.GitHash` – Type.

> GitHash

A git object identifier, based on the sha-1 hash. It is a 20 byte string (40 hex digits) used to identify a `GitObject` in a repository.

source

`Base.LibGit2.GitObject` – Type.

```
GitObject(repo::GitRepo, hash::AbstractGitHash)
GitObject(repo::GitRepo, spec::AbstractString)
```

Return the specified object (`GitCommit`, `GitBlob`, `GitTree` or `GitTag`) from repo specified by hash/spec.

- hash is a full (`GitHash`) or partial (`GitShortHash`) hash.
- spec is a textual specification: see the git docs for a full list.

source

`Base.LibGit2.GitRemote` – Type.

```
GitRemote(repo::GitRepo, rmt_name::AbstractString, rmt_url::AbstractString) -> GitRemote
```

Look up a remote git repository using its name and URL. Uses the default fetch refspec.

Example

```
repo = LibGit2.init(repo_path)
remote = LibGit2.GitRemote(repo, "upstream", repo_url)
```

source

```
GitRemote(repo::GitRepo, rmt_name::AbstractString, rmt_url::AbstractString, fetch_spec::
    AbstractString) -> GitRemote
```

Look up a remote git repository using the repository's name and URL, as well as specifications for how to fetch from the remote (e.g. which remote branch to fetch from).

Example

```
repo = LibGit2.init(repo_path)
refspec = "+refs/heads/mybranch:refs/remotes/origin/mybranch"
remote = LibGit2.GitRemote(repo, "upstream", repo_url, refspec)
```

source

`Base.LibGit2.GitRemoteAnon` – Function.

```
GitRemoteAnon(repo::GitRepo, url::AbstractString) -> GitRemote
```

Look up a remote git repository using only its URL, not its name.

Example

```
repo = LibGit2.init(repo_path)
remote = LibGit2.GitRemoteAnon(repo, repo_url)
```

source

`Base.LibGit2.GitRepo` – Type.

```
LibGit2.GitRepo(path::AbstractString)
```

Opens a git repository at `path`.

source

Base.LibGit2.GitRepoExt – Function.

```
LibGit2.GitRepoExt(path::AbstractString, flags::Cuint = Cuint(Consts.REPOSITORY_OPEN_DEFAULT)
    )
```

Opens a git repository at `path` with extended controls (for instance, if the current user must be a member of a special access group to read `path`).

source

Base.LibGit2.GitShortHash – Type.

```
GitShortHash
```

This is a shortened form of `GitHash`, which can be used to identify a git object when it is unique.

Internally it is stored as two fields: a full-size `GitHash` (hash) and a length (`len`). Only the initial `len` hex digits of hash are used.

source

Base.LibGit2.GitSignature – Type.

```
LibGit2.GitSignature
```

This is a Julia wrapper around a pointer to a `git_signature` object.

source

Base.LibGit2.GitStatus – Type.

```
LibGit2.GitStatus(repo::GitRepo; status_opts=StatusOptions())
```

Collect information about the status of each file in the git repository `repo` (e.g. is the file modified, staged, etc.). `status_opts` can be used to set various options, for instance whether or not to look at untracked files or whether to include submodules or not.

source

Base.LibGit2.GitTag – Type.

```
GitTag(repo::GitRepo, hash::AbstractGitHash)
GitTag(repo::GitRepo, spec::AbstractString)
```

Return a `GitTag` object from `repo` specified by `hash`/`spec`.

- `hash` is a full (`GitHash`) or partial (`GitShortHash`) hash.
- `spec` is a textual specification: see the git docs for a full list.

source

Base.LibGit2.GitTree – Type.

```
GitTree(repo::GitRepo, hash::AbstractGitHash)
GitTree(repo::GitRepo, spec::AbstractString)
```

Return a `GitTree` object from `repo` specified by hash/spec.

- hash is a full (`GitHash`) or partial (`GitShortHash`) hash.
- spec is a textual specification: see the git docs for a full list.

source

Base.LibGit2.IndexEntry – Type.

```
LibGit2.IndexEntry
```

In-memory representation of a file entry in the index. Matches the `git_index_entry` struct.

source

Base.LibGit2.IndexTime – Type.

```
LibGit2.IndexTime
```

Matches the `git_index_time` struct.

source

Base.LibGit2.MergeOptions – Type.

```
LibGit2.MergeOptions
```

Matches the `git_merge_options` struct.

source

Base.LibGit2.ProxyOptions – Type.

```
LibGit2.ProxyOptions
```

Options for connecting through a proxy.

Matches the `git_proxy_options` struct.

source

Base.LibGit2.PushOptions – Type.

```
LibGit2.PushOptions
```

Matches the `git_push_options` struct.

source

Base.LibGit2.RebaseOperation – Type.

```
LibGit2.RebaseOperation
```

Describes a single instruction/operation to be performed during the rebase. Matches the `git_rebase_operation` struct.

source

Base.LibGit2.RebaseOptions – Type.

> LibGit2.RebaseOptions

Matches the `git_rebase_options` struct.

> source

Base.LibGit2.RemoteCallbacks – Type.

> LibGit2.RemoteCallbacks

Callback settings. Matches the `git_remote_callbacks` struct.

> source

Base.LibGit2.SSHCredentials – Type.

SSH credentials type

> source

Base.LibGit2.SignatureStruct – Type.

> LibGit2.SignatureStruct

An action signature (e.g. for committers, taggers, etc). Matches the `git_signature` struct.

> source

Base.LibGit2.StatusEntry – Type.

> LibGit2.StatusEntry

Providing the differences between the file as it exists in HEAD and the index, and providing the differences between the index and the working directory. Matches the `git_status_entry` struct.

> source

Base.LibGit2.StatusOptions – Type.

> LibGit2.StatusOptions

Options to control how `git_status_foreach_ext()` will issue callbacks. Matches the `git_status_opt_t` struct.

> source

Base.LibGit2.StrArrayStruct – Type.

> LibGit2.StrArrayStruct

A LibGit2 representation of an array of strings. Matches the `git_strarray` struct.

When fetching data from LibGit2, a typical usage would look like:

```
sa_ref = Ref(StrArrayStruct())
@check ccall(..., (Ptr{StrArrayStruct},), sa_ref)
res = convert(Vector{String}, sa_ref[])
free(sa_ref)
```

In particular, note that `LibGit2.free` should be called afterward on the `Ref` object.

Conversely, when passing a vector of strings to LibGit2, it is generally simplest to rely on implicit conversion:

```
strs = String[...]
@check ccall(..., (Ptr{StrArrayStruct},), strs)
```

Note that no call to `free` is required as the data is allocated by Julia.

source

Base.LibGit2.TimeStruct – Type.

```
LibGit2.TimeStruct
```

Time in a signature. Matches the `git_time` struct.

source

Base.LibGit2.UserPasswordCredentials – Type.

Credentials that support only `user` and `password` parameters

source

Base.LibGit2.add_fetch! – Function.

```
add_fetch!(repo::GitRepo, rmt::GitRemote, fetch_spec::String)
```

Add a *fetch* refspec for the specified `rmt`. This refspec will contain information about which branch(es) to fetch from.

Example

```
julia> LibGit2.add_fetch!(repo, remote, "upstream");

julia> LibGit2.fetch_refspecs(remote)
String["+refs/heads/*:refs/remotes/upstream/*"]
```

source

Base.LibGit2.add_push! – Function.

```
add_push!(repo::GitRepo, rmt::GitRemote, push_spec::String)
```

Add a *push* refspec for the specified `rmt`. This refspec will contain information about which branch(es) to push to.

Example

```
julia> LibGit2.add_push!(repo, remote, "refs/heads/master");

julia> remote = LibGit2.get(LibGit2.GitRemote, repo, branch);

julia> LibGit2.push_refspecs(remote)
String["refs/heads/master"]
```

> **Note**
>
> You may need to close and reopen the `GitRemote` in question after updating its push refspecs in order for the change to take effect and for calls to push to work.

source

Base.LibGit2.addblob! – Function.

```
LibGit2.addblob!(repo::GitRepo, path::AbstractString)
```

Reads the file at `path` and adds it to the object database of `repo` as a loose blob. Returns the `GitHash` of the resulting blob.

Example

```
hash_str = hex(commit_oid)
blob_file = joinpath(repo_path, ".git", "objects", hash_str[1:2], hash_str[3:end])
id = LibGit2.addblob!(repo, blob_file)
```

source

Base.LibGit2.authors – Function.

```
authors(repo::GitRepo) -> Vector{Signature}
```

Returns all authors of commits to the `repo` repository.

Example

```
repo = LibGit2.GitRepo(repo_path)
repo_file = open(joinpath(repo_path, test_file), "a")

println(repo_file, commit_msg)
flush(repo_file)
LibGit2.add!(repo, test_file)
sig = LibGit2.Signature("TEST", "TEST@TEST.COM", round(time(), 0), 0)
commit_oid1 = LibGit2.commit(repo, "commit1"; author=sig, committer=sig)
println(repo_file, randstring(10))
flush(repo_file)
LibGit2.add!(repo, test_file)
commit_oid2 = LibGit2.commit(repo, "commit2"; author=sig, committer=sig)

# will be a Vector of [sig, sig]
auths = LibGit2.authors(repo)
```

source

Base.LibGit2.branch – Function.

```
branch(repo::GitRepo)
```

Equivalent to `git branch`. Create a new branch from the current HEAD.

source

Base.LibGit2.branch! – Function.

```
branch!(repo::GitRepo, branch_name::AbstractString, commit::AbstractString=""; kwargs...)
```

Checkout a new git branch in the `repo` repository. `commit` is the `GitHash`, in string form, which will be the start of the new branch. If `commit` is an empty string, the current HEAD will be used.

The keyword arguments are:

- `track::AbstractString=""`: the name of the remote branch this new branch should track, if any. If empty (the default), no remote branch will be tracked.
- `force::Bool=false`: if `true`, branch creation will be forced.
- `set_head::Bool=true`: if `true`, after the branch creation finishes the branch head will be set as the HEAD of `repo`.

Equivalent to `git checkout [-b|-B] <branch_name> [<commit>] [--track <track>]`.

Example

```
repo = LibGit2.GitRepo(repo_path)
LibGit2.branch!(repo, "new_branch", set_head=false)
```

source

Base.LibGit2.checkout! – Function.

```
checkout!(repo::GitRepo, commit::AbstractString=""; force::Bool=true)
```

Equivalent to `git checkout [-f] --detach <commit>`. Checkout the git commit `commit` (a `GitHash` in string form) in `repo`. If `force` is `true`, force the checkout and discard any current changes. Note that this detaches the current HEAD.

Example

```
repo = LibGit2.init(repo_path)
open(joinpath(LibGit2.path(repo), "file1"), "w") do f
    write(f, "111
")
end
LibGit2.add!(repo, "file1")
commit_oid = LibGit2.commit(repo, "add file1")
open(joinpath(LibGit2.path(repo), "file1"), "w") do f
    write(f, "112
")
end
# would fail without the force=true
# since there are modifications to the file
LibGit2.checkout!(repo, string(commit_oid), force=true)
```

source

Base.LibGit2.checkused! – Function.

Checks if credentials were used

source

Checks if credentials were used or failed authentication, see `LibGit2.credentials_callback`

source

Base.LibGit2.clone – Function.

```
clone(repo_url::AbstractString, repo_path::AbstractString; kwargs...)
```

Clone a remote repository located at `repo_url` to the local filesystem location `repo_path`.

The keyword arguments are:

- `branch::AbstractString=""`: which branch of the remote to clone, if not the default repository branch (usually `master`).

- `isbare::Bool=false`: if `true`, clone the remote as a bare repository, which will make `repo_path` itself the git directory instead of `repo_path/.git`. This means that a working tree cannot be checked out. Plays the role of the git CLI argument `--bare`.

- `remote_cb::Ptr{Void}=C_NULL`: a callback which will be used to create the remote before it is cloned. If `C_NULL` (the default), no attempt will be made to create the remote - it will be assumed to already exist.

- `payload::Nullable{P<:AbstractCredentials}=Nullable{AbstractCredentials}()`: provides credentials if necessary, for instance if the remote is a private repository.

Equivalent to `git clone [-b <branch>] [--bare] <repo_url> <repo_path>`.

Examples

```
repo_url = "https://github.com/JuliaLang/Example.jl"
repo1 = LibGit2.clone(repo_url, "test_path")
repo2 = LibGit2.clone(repo_url, "test_path", isbare=true)
julia_url = "https://github.com/JuliaLang/julia"
julia_repo = LibGit2.clone(julia_url, "julia_path", branch="release-0.6")
```

source

`Base.LibGit2.commit` – Function.

Wrapper around `git_commit_create`

source

Commit changes to repository

source

```
LibGit2.commit(rb::GitRebase, sig::GitSignature)
```

Commits the current patch to the rebase `rb`, using `sig` as the committer. Is silent if the commit has already been applied.

source

`Base.LibGit2.create_branch` – Function.

```
LibGit2.create_branch(repo::GitRepo, bname::AbstractString, commit_obj::GitCommit; force::
    Bool=false)
```

Create a new branch in the repository `repo` with name `bname`, which points to commit `commit_obj` (which has to be part of `repo`). If `force` is `true`, overwrite an existing branch named `bname` if it exists. If `force` is `false` and a branch already exists named `bname`, this function will throw an error.

source

`Base.LibGit2.credentials_callback` – Function.

Credentials callback function

Function provides different credential acquisition functionality w.r.t. a connection protocol. If a payload is provided then `payload_ptr` should contain a `LibGit2.AbstractCredentials` object.

For `LibGit2.Consts.CREDTYPE_USERPASS_PLAINTEXT` type, if the payload contains fields: `user` & `pass`, they are used to create authentication credentials. Empty `user` name and `password` trigger an authentication error.

For `LibGit2.Consts.CREDTYPE_SSH_KEY` type, if the payload contains fields: `user`, `prvkey`, `pubkey` & `pass`, they are used to create authentication credentials. Empty `user` name triggers an authentication error.

Credentials are checked in the following order (if supported):

- ssh key pair (`ssh-agent` if specified in payload's `usesshagent` field)
- plain text

Note: Due to the specifics of the `libgit2` authentication procedure, when authentication fails, this function is called again without any indication whether authentication was successful or not. To avoid an infinite loop from repeatedly using the same faulty credentials, the `checkused!` function can be called. This function returns `true` if the credentials were used. Using credentials triggers a user prompt for (re)entering required information. `UserPasswordCredentials` and `CachedCredentials` are implemented using a call counting strategy that prevents repeated usage of faulty credentials.

source

Base.LibGit2.credentials_cb – Function.

C function pointer for `credentials_callback`

source

Base.LibGit2.default_signature – Function.

Return signature object. Free it after use.

source

Base.LibGit2.delete_branch – Function.

```
LibGit2.delete_branch(branch::GitReference)
```

Delete the branch pointed to by `branch`.

source

Base.LibGit2.diff_files – Function.

```
diff_files(repo::GitRepo, branch1::AbstractString, branch2::AbstractString; kwarg...) ->
    Vector{AbstractString}
```

Show which files have changed in the git repository `repo` between branches `branch1` and `branch2`.

The keyword argument is:

- `filter::Set{Consts.DELTA_STATUS}=Set([Consts.DELTA_ADDED, Consts.DELTA_MODIFIED, Consts.DELTA_DELETED]))`, and it sets options for the diff. The default is to show files added, modified, or deleted.

Returns only the *names* of the files which have changed, *not* their contents.

Example

```julia
LibGit2.branch!(repo, "branch/a")
LibGit2.branch!(repo, "branch/b")
# add a file to repo
open(joinpath(LibGit2.path(repo),"file"),"w") do f
    write(f, "hello repo
")
end
LibGit2.add!(repo, "file")
LibGit2.commit(repo, "add file")
# returns ["file"]
filt = Set([LibGit2.Consts.DELTA_ADDED])
files = LibGit2.diff_files(repo, "branch/a", "branch/b", filter=filt)
# returns [] because existing files weren't modified
filt = Set([LibGit2.Consts.DELTA_MODIFIED])
files = LibGit2.diff_files(repo, "branch/a", "branch/b", filter=filt)
```

Equivalent to `git diff --name-only --diff-filter=<filter> <branch1> <branch2>`.

source

`Base.LibGit2.fetch` – Function.

```julia
fetch(rmt::GitRemote, refspecs; options::FetchOptions=FetchOptions(), msg="")
```

Fetch from the specified `rmt` remote git repository, using `refspecs` to determine which remote branch(es) to fetch. The keyword arguments are:

- `options`: determines the options for the fetch, e.g. whether to prune afterwards.

- `msg`: a message to insert into the reflogs.

source

```julia
fetch(repo::GitRepo; kwargs...)
```

Fetches updates from an upstream of the repository `repo`.

The keyword arguments are:

- `remote::AbstractString="origin"`: which remote, specified by name, of `repo` to fetch from. If this is empty, the URL will be used to construct an anonymous remote.

- `remoteurl::AbstractString=""`: the URL of `remote`. If not specified, will be assumed based on the given name of `remote`.

- `refspecs=AbstractString[]`: determines properties of the fetch.

- `payload=Nullable{AbstractCredentials}()`: provides credentials, if necessary, for instance if re-mote is a private repository.

Equivalent to `git fetch [<remoteurl>|<repo>] [<refspecs>]`.

source

`Base.LibGit2.fetch_refspecs` – Function.

```julia
fetch_refspecs(rmt::GitRemote) -> Vector{String}
```

Get the *fetch* refspecs for the specified `rmt`. These refspecs contain information about which branch(es) to fetch from.

source

`Base.LibGit2.fetchhead_foreach_cb` – Function.

C function pointer for `fetchhead_foreach_callback`

source

`Base.LibGit2.ffmerge!` – Function.

Fastforward merge changes into current head

source

`Base.LibGit2.fullname` – Function.

```
LibGit2.fullname(ref::GitReference)
```

Return the name of the reference pointed to by the symbolic reference `ref`. If `ref` is not a symbolic reference, returns an empty string.

source

`Base.LibGit2.get_creds!` – Function.

Obtain the cached credentials for the given host+protocol (credid), or return and store the default if not found

source

`Base.LibGit2.gitdir` – Function.

```
LibGit2.gitdir(repo::GitRepo)
```

Returns the location of the "git" files of `repo`:

- for normal repositories, this is the location of the `.git` folder.
- for bare repositories, this is the location of the repository itself.

See also `workdir`, `path`.

source

`Base.LibGit2.head` – Function.

```
LibGit2.head(repo::GitRepo) -> GitReference
```

Returns a `GitReference` to the current HEAD of `repo`.

source

```
head(pkg::AbstractString) -> String
```

Return current HEAD `GitHash` of the `pkg` repo as a string.

source

`Base.LibGit2.head!` – Function.

```
LibGit2.head!(repo::GitRepo, ref::GitReference) -> GitReference
```

Set the HEAD of repo to the object pointed to by ref.

source

Base.LibGit2.head_oid – Function.

```
LibGit2.head_oid(repo::GitRepo) -> GitHash
```

Lookup the object id of the current HEAD of git repository repo.

source

Base.LibGit2.headname – Function.

```
LibGit2.headname(repo::GitRepo)
```

Lookup the name of the current HEAD of git repository repo. If repo is currently detached, returns the name of
the HEAD it's detached from.

source

Base.LibGit2.init – Function.

```
LibGit2.init(path::AbstractString, bare::Bool=false) -> GitRepo
```

Opens a new git repository at path. If bare is false, the working tree will be created in path/.git. If bare is
true, no working directory will be created.

source

Base.LibGit2.is_ancestor_of – Function.

```
is_ancestor_of(a::AbstractString, b::AbstractString, repo::GitRepo) -> Bool
```

Returns true if a, a GitHash in string form, is an ancestor of b, a GitHash in string form.

Example

```
julia> repo = LibGit2.GitRepo(repo_path);

julia> LibGit2.add!(repo, test_file1);

julia> commit_oid1 = LibGit2.commit(repo, "commit1");

julia> LibGit2.add!(repo, test_file2);

julia> commit_oid2 = LibGit2.commit(repo, "commit2");

julia> LibGit2.is_ancestor_of(string(commit_oid1), string(commit_oid2), repo)
true
```

source

Base.LibGit2.isbinary – Function.

Use a heuristic to guess if a file is binary: searching for NULL bytes and looking for a reasonable ratio of printable
to non-printable characters among the first 8000 bytes.

source

Base.LibGit2.iscommit – Function.

```
iscommit(id::AbstractString, repo::GitRepo) -> Bool
```

Checks if commit id (which is a GitHash in string form) is in the repository.

Example

```
julia> repo = LibGit2.GitRepo(repo_path);

julia> LibGit2.add!(repo, test_file);

julia> commit_oid = LibGit2.commit(repo, "add test_file");

julia> LibGit2.iscommit(string(commit_oid), repo)
true
```

source

Base.LibGit2.isdiff – Function.

```
LibGit2.isdiff(repo::GitRepo, treeish::AbstractString, pathspecs::AbstractString=""; cached::
    Bool=false)
```

Checks if there are any differences between the tree specified by treeish and the tracked files in the working tree (if cached=false) or the index (if cached=true). pathspecs are the specifications for options for the diff.

Example

```
repo = LibGit2.GitRepo(repo_path)
LibGit2.isdiff(repo, "HEAD") # should be false
open(joinpath(repo_path, new_file), "a") do f
    println(f, "here's my cool new file")
end
LibGit2.isdiff(repo, "HEAD") # now true
```

Equivalent to git diff-index <treeish> [-- <pathspecs>].

source

Base.LibGit2.isdirty – Function.

```
LibGit2.isdirty(repo::GitRepo, pathspecs::AbstractString=""; cached::Bool=false) -> Bool
```

Checks if there have been any changes to tracked files in the working tree (if cached=false) or the index (if cached=true). pathspecs are the specifications for options for the diff.

Example

```
repo = LibGit2.GitRepo(repo_path)
LibGit2.isdirty(repo) # should be false
open(joinpath(repo_path, new_file), "a") do f
    println(f, "here's my cool new file")
end
LibGit2.isdirty(repo) # now true
LibGit2.isdirty(repo, new_file) # now true
```

Equivalent to `git diff-index HEAD [-- <pathspecs>]`.

source

Base.LibGit2.isorphan – Function.

```
LibGit2.isorphan(repo::GitRepo)
```

Checks if the current branch is an "orphan" branch, i.e. has no commits. The first commit to this branch will have no parents.

source

Base.LibGit2.lookup_branch – Function.

```
lookup_branch(repo::GitRepo, branch_name::AbstractString, remote::Bool=false) -> Nullable{
    GitReference}
```

Determine if the branch specified by `branch_name` exists in the repository `repo`. If `remote` is `true`, `repo` is assumed to be a remote git repository. Otherwise, it is part of the local filesystem.

`lookup_branch` returns a `Nullable`, which will be null if the requested branch does not exist yet. If the branch does exist, the `Nullable` contains a `GitReference` to the branch.

source

Base.LibGit2.mirror_callback – Function.

Mirror callback function

Function sets `+refs/*:refs/*` refspecs and `mirror` flag for remote reference.

source

Base.LibGit2.mirror_cb – Function.

C function pointer for `mirror_callback`

source

Base.LibGit2.name – Function.

```
LibGit2.name(ref::GitReference)
```

Return the full name of `ref`.

source

```
name(rmt::GitRemote)
```

Get the name of a remote repository, for instance `"origin"`. If the remote is anonymous (see GitRemoteAnon) the name will be an empty string `""`.

Example

```
julia> repo_url = "https://github.com/JuliaLang/Example.jl";

julia> repo = LibGit2.clone(cache_repo, "test_directory");

julia> remote = LibGit2.GitRemote(repo, "origin", repo_url);

julia> name(remote)
"origin"
```

source

```
LibGit2.name(tag::GitTag)
```

The name of tag (e.g. `"v0.5"`).

source

Base.LibGit2.need_update – Function.

```
need_update(repo::GitRepo)
```

Equivalent to `git update-index`. Returns `true` if repo needs updating.

source

Base.LibGit2.objtype – Function.

```
objtype(obj_type::Consts.OBJECT)
```

Returns the type corresponding to the enum value.

source

Base.LibGit2.path – Function.

```
LibGit2.path(repo::GitRepo)
```

The base file path of the repository `repo`.

- for normal repositories, this will typically be the parent directory of the ".git" directory (note: this may be different than the working directory, see `workdir` for more details).
- for bare repositories, this is the location of the "git" files.

See also `gitdir`, `workdir`.

source

Base.LibGit2.peel – Function.

```
peel([T,] ref::GitReference)
```

Recursively peel `ref` until an object of type T is obtained. If no T is provided, then `ref` will be peeled until an object other than a `GitTag` is obtained.

- A `GitTag` will be peeled to the object it references.
- A `GitCommit` will be peeled to a `GitTree`.

  Note
  Only annotated tags can be peeled to `GitTag` objects. Lightweight tags (the default) are references under `refs/tags/` which point directly to `GitCommit` objects.

source

```
peel([T,] obj::GitObject)
```

Recursively peel `obj` until an object of type T is obtained. If no T is provided, then `obj` will be peeled until the type changes.

- A `GitTag` will be peeled to the object it references.
- A `GitCommit` will be peeled to a `GitTree`.

source

Base.LibGit2.posixpath – Function.

```
LibGit2.posixpath(path)
```

Standardise the path string `path` to use POSIX separators.

source

Base.LibGit2.push – Function.

```
push(rmt::GitRemote, refspecs; force::Bool=false, options::PushOptions=PushOptions())
```

Push to the specified `rmt` remote git repository, using `refspecs` to determine which remote branch(es) to push to. The keyword arguments are:

- `force`: if `true`, a force-push will occur, disregarding conflicts.
- `options`: determines the options for the push, e.g. which proxy headers to use.

    Note

    You can add information about the push refspecs in two other ways: by setting an option in the repository's `GitConfig` (with `push.default` as the key) or by calling add_push!. Otherwise you will need to explicitly specify a push refspec in the call to `push` for it to have any effect, like so: `LibGit2.push(repo, refspecs=["refs/heads/master"])`.

source

```
push(repo::GitRepo; kwargs...)
```

Pushes updates to an upstream of `repo`.

The keyword arguments are:

- `remote::AbstractString="origin"`: the name of the upstream remote to push to.
- `remoteurl::AbstractString=""`: the URL of `remote`.
- `refspecs=AbstractString[]`: determines properties of the push.
- `force::Bool=false`: determines if the push will be a force push, overwriting the remote branch.
- `payload=Nullable{AbstractCredentials}()`: provides credentials, if necessary, for instance if `remote` is a private repository.

Equivalent to `git push [<remoteurl>|<repo>] [<refspecs>]`.

source

Base.LibGit2.push_refspecs – Function.

```
push_refspecs(rmt::GitRemote) -> Vector{String}
```

Get the *push* refspecs for the specified `rmt`. These refspecs contain information about which branch(es) to push to.

source

Base.LibGit2.read_tree! – Function.

```
LibGit2.read_tree!(idx::GitIndex, tree::GitTree)
LibGit2.read_tree!(idx::GitIndex, treehash::AbstractGitHash)
```

Read the tree `tree` (or the tree pointed to by `treehash` in the repository owned by `idx`) into the index `idx`. The current index contents will be replaced.

source

Base.LibGit2.rebase! – Function.

```
LibGit2.rebase!(repo::GitRepo, upstream::AbstractString="", newbase::AbstractString="")
```

Attempt an automatic merge rebase of the current branch, from `upstream` if provided, or otherwise from the upstream tracking branch. `newbase` is the branch to rebase onto. By default this is `upstream`.

If any conflicts arise which cannot be automatically resolved, the rebase will abort, leaving the repository and working tree in its original state, and the function will throw a `GitError`. This is roughly equivalent to the following command line statement:

```
git rebase --merge [<upstream>]
if [ -d ".git/rebase-merge" ]; then
    git rebase --abort
fi
```

source

Base.LibGit2.ref_list – Function.

```
LibGit2.ref_list(repo::GitRepo) -> Vector{String}
```

Get a list of all reference names in the `repo` repository.

source

Base.LibGit2.reftype – Function.

```
LibGit2.reftype(ref::GitReference) -> Cint
```

Returns a `Cint` corresponding to the type of `ref`:

- `0` if the reference is invalid
- `1` if the reference is an object id
- `2` if the reference is symbolic

source

Base.LibGit2.remotes – Function.

```
LibGit2.remotes(repo::GitRepo)
```

Returns a vector of the names of the remotes of `repo`.

source

Base.LibGit2.reset! – Function.

Resets credentials for another use

source

Updates some entries, determined by the `pathspecs`, in the index from the target commit tree.

source

Sets the current head to the specified commit oid and optionally resets the index and working tree to match.

source

git reset [<committish>] [–] <pathspecs>...

source

```
reset!(repo::GitRepo, id::GitHash, mode::Cint = Consts.RESET_MIXED)
```

Reset the repository `repo` to its state at `id`, using one of three modes set by `mode`:

1. `Consts.RESET_SOFT` - move HEAD to `id`.
2. `Consts.RESET_MIXED` - default, move HEAD to `id` and reset the index to `id`.
3. `Consts.RESET_HARD` - move HEAD to `id`, reset the index to `id`, and discard all working changes.

Equivalent to `git reset [--soft | --mixed | --hard] <id>`.

Example

```julia
repo = LibGit2.GitRepo(repo_path)
head_oid = LibGit2.head_oid(repo)
open(joinpath(repo_path, "file1"), "w") do f
    write(f, "111
")
end
LibGit2.add!(repo, "file1")
mode = LibGit2.Consts.RESET_HARD
# will discard the changes to file1
# and unstage it
new_head = LibGit2.reset!(repo, head_oid, mode)
```

source

Base.LibGit2.restore – Function.

```
restore(s::State, repo::GitRepo)
```

Return a repository `repo` to a previous `State s`, for example the HEAD of a branch before a merge attempt. `s` can be generated using the `snapshot` function.

source

Base.LibGit2.revcount – Function.

```
LibGit2.revcount(repo::GitRepo, commit1::AbstractString, commit2::AbstractString)
```

List the number of revisions between `commit1` and `commit2` (committish OIDs in string form). Since `commit1` and `commit2` may be on different branches, `revcount` performs a "left-right" revision list (and count), returning a tuple of `Int`s - the number of left and right commits, respectively. A left (or right) commit refers to which side of a symmetric difference in a tree the commit is reachable from.

Equivalent to `git rev-list --left-right --count <commit1> <commit2>`.

source

Base.LibGit2.set_remote_url – Function.

```
set_remote_url(repo::GitRepo, url::AbstractString; remote::AbstractString="origin")
```

Set the url for remote for the git repository repo. The default name of the remote is "origin".

Examples

```
repo_path = joinpath("test_directory", "Example")
repo = LibGit2.init(repo_path)
url1 = "https://github.com/JuliaLang/Example.jl"
LibGit2.set_remote_url(repo, url1, remote="upstream")
url2 = "https://github.com/JuliaLang/Example2.jl"
LibGit2.set_remote_url(repo_path, url2, remote="upstream2")
```

source

```
set_remote_url(path::AbstractString, url::AbstractString; remote::AbstractString="origin")
```

Set the url for remote for the git repository located at path. The default name of the remote is "origin".

source

Base.LibGit2.shortname – Function.

```
LibGit2.shortname(ref::GitReference)
```

Returns a shortened version of the name of ref that's "human-readable".

```
julia> repo = LibGit2.GitRepo(path_to_repo);

julia> branch_ref = LibGit2.head(repo);

julia> LibGit2.name(branch_ref)
"refs/heads/master"

julia> LibGit2.shortname(branch_ref)
"master"
```

source

Base.LibGit2.snapshot – Function.

```
snapshot(repo::GitRepo) -> State
```

Take a snapshot of the current state of the repository repo, storing the current HEAD, index, and any uncommitted work. The output State can be used later during a call to restore to return the repository to the snapshotted state.

source

Base.LibGit2.status – Function.

```
LibGit2.status(repo::GitRepo, path::String)
```

Lookup the status of the file at path in the git repository repo. For instance, this can be used to check if the file at path has been modified and needs to be staged and committed.

source

Base.LibGit2.tag_create – Function.

```
LibGit2.tag_create(repo::GitRepo, tag::AbstractString, commit; kwargs...)
```

Create a new git tag `tag` (e.g. `"v0.5"`) in the repository `repo`, at the commit `commit`.

The keyword arguments are:

- `msg::AbstractString=""`: the message for the tag.
- `force::Bool=false`: if `true`, existing references will be overwritten.
- `sig::Signature=Signature(repo)`: the tagger's signature.

source

Base.LibGit2.tag_delete – Function.

```
LibGit2.tag_delete(repo::GitRepo, tag::AbstractString)
```

Remove the git tag `tag` from the repository `repo`.

source

Base.LibGit2.tag_list – Function.

```
LibGit2.tag_list(repo::GitRepo) -> Vector{String}
```

Get a list of all tags in the git repository `repo`.

source

Base.LibGit2.target – Function.

```
LibGit2.target(tag::GitTag)
```

The `GitHash` of the target object of `tag`.

source

Base.LibGit2.treewalk – Function.

Traverse the entries in a tree and its subtrees in post or pre order.

Function parameter should have following signature:

```
(Cstring, Ptr{Void}, Ptr{Void}) -> Cint
```

source

Base.LibGit2.upstream – Function.

```
upstream(ref::GitReference) -> Nullable{GitReference}
```

Determine if the branch containing `ref` has a specified upstream branch.

`upstream` returns a `Nullable`, which will be null if the requested branch does not have an upstream counterpart. If the upstream branch does exist, the `Nullable` contains a `GitReference` to the upstream branch.

source

Base.LibGit2.url – Function.

```
url(rmt::GitRemote)
```

Get the fetch URL of a remote git repository.

Example

```julia
julia> repo_url = "https://github.com/JuliaLang/Example.jl";

julia> repo = LibGit2.init(mktempdir());

julia> remote = LibGit2.GitRemote(repo, "origin", repo_url);

julia> LibGit2.url(remote)
"https://github.com/JuliaLang/Example.jl"
```

source

Base.LibGit2.with – Function.

Resource management helper function

source

Base.LibGit2.workdir – Function.

```
LibGit2.workdir(repo::GitRepo)
```

The location of the working directory of `repo`. This will throw an error for bare repositories.

> **Note**
>
> This will typically be the parent directory of `gitdir(repo)`, but can be different in some cases: e.g.
> if either the `core.worktree` configuration variable or the `GIT_WORK_TREE` environment variable is
> set.

See also gitdir, path.

source

70.19 Module loading

`Base.require[@ref]` is responsible for loading modules and it also manages the precompilation cache. It is the implementation of the `import` statement.

Experimental features

The features below are experimental and not part of the stable Julia API. Before building upon them inform yourself about the current thinking and whether they might change soon.

Module loading callbacks

It is possible to listen to the modules loaded by `Base.require`, by registering a callback.

```
loaded_packages = Channel{Symbol}()
callback = (mod::Symbol) -> put!(loaded_packages, mod)
push!(Base.package_callbacks, callback)
```

Please note that the symbol given to the callback is a non-unique identifier and it is the responsibility of the callback provider to walk the module chain to determine the fully qualified name of the loaded binding.

The callback below is an example of how to do that:

```julia
# Get the fully-qualified name of a module.
function module_fqn(name::Symbol)
    fqn = Symbol[name]
    mod = getfield(Main, name)
    parent = Base.module_parent(mod)
    while parent !== Main
        push!(fqn, Base.module_name(parent))
        parent = Base.module_parent(parent)
    end
    fqn = reverse!(fqn)
    return join(fqn, '.')
end
```

Chapter 71

Developing/debugging Julia's C code

71.1   Reporting and analyzing crashes (segfaults)

So you managed to break Julia. Congratulations! Collected here are some general procedures you can undergo for common symptoms encountered when something goes awry. Including the information from these debugging steps can greatly help the maintainers when tracking down a segfault or trying to figure out why your script is running slower than expected.

If you've been directed to this page, find the symptom that best matches what you're experiencing and follow the instructions to generate the debugging information requested. Table of symptoms:

- Segfaults during bootstrap (`sysimg.jl`)

- Segfaults when running a script

- Errors during Julia startup

Version/Environment info

No matter the error, we will always need to know what version of Julia you are running. When Julia first starts up, a header is printed out with a version number and date. If your version is `0.2.0` or higher, please include the output of `versioninfo()` in any report you create:

```
julia> versioninfo()
Julia Version 0.6.1-pre.0
Commit dcf39a1dda (2017-06-19 13:06 UTC)
Platform Info:
  OS: macOS (x86_64-apple-darwin16.7.0)
  CPU: Intel(R) Core(TM) i5-5287U CPU @ 2.90GHz
  WORD_SIZE: 64
  BLAS: libopenblas (USE64BITINT DYNAMIC_ARCH NO_AFFINITY Haswell)
  LAPACK: libopenblas64_
  LIBM: libopenlibm
  LLVM: libLLVM-3.9.1 (ORCJIT, broadwell)
```

Segfaults during bootstrap (`sysimg.jl`)

Segfaults toward the end of the `make` process of building Julia are a common symptom of something going wrong while Julia is preparsing the corpus of code in the `base/` folder. Many factors can contribute toward this process

dying unexpectedly, however it is as often as not due to an error in the C-code portion of Julia, and as such must typically be debugged with a debug build inside of gdb. Explicitly:

Create a debug build of Julia:

```
$ cd <julia_root>
$ make debug
```

Note that this process will likely fail with the same error as a normal make incantation, however this will create a debug executable that will offer gdb the debugging symbols needed to get accurate backtraces. Next, manually run the bootstrap process inside of gdb:

```
$ cd base/
$ gdb -x ../contrib/debug_bootstrap.gdb
```

This will start gdb, attempt to run the bootstrap process using the debug build of Julia, and print out a backtrace if (when) it segfaults. You may need to hit <enter> a few times to get the full backtrace. Create a gist with the backtrace, the version info, and any other pertinent information you can think of and open a new issue on Github with a link to the gist.

Segfaults when running a script

The procedure is very similar to Segfaults during bootstrap (`sysimg.jl`). Create a debug build of Julia, and run your script inside of a debugged Julia process:

```
$ cd <julia_root>
$ make debug
$ gdb --args usr/bin/julia-debug <path_to_your_script>
```

Note that gdb will sit there, waiting for instructions. Type r to run the process, and bt to generate a backtrace once it segfaults:

```
(gdb) r
Starting program: /home/sabae/src/julia/usr/bin/julia-debug ./test.jl
...
(gdb) bt
```

Create a gist with the backtrace, the version info, and any other pertinent information you can think of and open a new issue on Github with a link to the gist.

Errors during Julia startup

Occasionally errors occur during Julia's startup process (especially when using binary distributions, as opposed to compiling from source) such as the following:

```
$ julia
exec: error -5
```

These errors typically indicate something is not getting loaded properly very early on in the bootup phase, and our best bet in determining what's going wrong is to use external tools to audit the disk activity of the `julia` process:

- On Linux, use `strace`:

  ```
  $ strace julia
  ```

- On OSX, use `dtruss`:

```
$ dtruss -f julia
```

Create a gist with the `strace`/ `dtruss` ouput, the version info, and any other pertinent information and open a new issue on Github with a link to the gist.

Glossary

A few terms have been used as shorthand in this guide:

- `<julia_root>` refers to the root directory of the Julia source tree; e.g. it should contain folders such as `base`, `deps`, `src`, `test`, etc.....

71.2 gdb debugging tips

Displaying Julia variables

Within gdb, any `jl_value_t*` object `obj` can be displayed using

```
(gdb) call jl_(obj)
```

The object will be displayed in the `julia` session, not in the gdb session. This is a useful way to discover the types and values of objects being manipulated by Julia's C code.

Similarly, if you're debugging some of Julia's internals (e.g., `inference.jl`), you can print `obj` using

```
ccall(:jl_, Void, (Any,), obj)
```

This is a good way to circumvent problems that arise from the order in which julia's output streams are initialized.

Julia's flisp interpreter uses `value_t` objects; these can be displayed with `call fl_print(fl_ctx, ios_stdout, obj)`.

Useful Julia variables for Inspecting

While the addresses of many variables, like singletons, can be be useful to print for many failures, there are a number of additional variables (see `julia.h` for a complete list) that are even more useful.

- (when in `jl_apply_generic`) `mfunc` and `jl_uncompress_ast(mfunc->def, mfunc->code)` :: for figuring out a bit about the call-stack

- `jl_lineno` and `jl_filename` :: for figuring out what line in a test to go start debugging from (or figure out how far into a file has been parsed)

- `$1` :: not really a variable, but still a useful shorthand for referring to the result of the last gdb command (such as `print`)

- `jl_options` :: sometimes useful, since it lists all of the command line options that were successfully parsed

- `jl_uv_stderr` :: because who doesn't like to be able to interact with stdio

Useful Julia functions for Inspecting those variables

- `jl_gdblookup($rip)` :: For looking up the current function and line. (use `$eip` on i686 platforms)

- `jlbacktrace()` :: For dumping the current Julia backtrace stack to stderr. Only usable after `record_back-trace()` has been called.

- `jl_dump_llvm_value(Value*)` :: For invoking `Value->dump()` in gdb, where it doesn't work natively. For example, `f->linfo->functionObject`, `f->linfo->specFunctionObject`, and `to_function(f->linfo)`.

- `Type->dump()` :: only works in lldb. Note: add something like `;1` to prevent lldb from printing its prompt over the output

- `jl_eval_string("expr")` :: for invoking side-effects to modify the current state or to lookup symbols

- `jl_typeof(jl_value_t*)` :: for extracting the type tag of a Julia value (in gdb, call `macro define jl_typeof jl_typeof` first, or pick something short like `ty` for the first arg to define a shorthand)

Inserting breakpoints for inspection from gdb

In your gdb session, set a breakpoint in `jl_breakpoint` like so:

```
(gdb) break jl_breakpoint
```

Then within your Julia code, insert a call to `jl_breakpoint` by adding

```
ccall(:jl_breakpoint, Void, (Any,), obj)
```

where `obj` can be any variable or tuple you want to be accessible in the breakpoint.

It's particularly helpful to back up to the `jl_apply` frame, from which you can display the arguments to a function using, e.g.,

```
(gdb) call jl_(args[0])
```

Another useful frame is `to_function(jl_method_instance_t *li, bool cstyle)`. The `jl_method_instance_t*` argument is a struct with a reference to the final AST sent into the compiler. However, the AST at this point will usually be compressed; to view the AST, call `jl_uncompress_ast` and then pass the result to `jl_`:

```
#2  0x00007ffff7928bf7 in to_function (li=0x2812060, cstyle=false) at codegen.cpp:584
584             abort();
(gdb) p jl_(jl_uncompress_ast(li, li->ast))
```

Inserting breakpoints upon certain conditions

Loading a particular file

Let's say the file is `sysimg.jl`:

```
(gdb) break jl_load if strcmp(fname, "sysimg.jl")==0
```

Calling a particular method

```
(gdb) break jl_apply_generic if strcmp((char*)(jl_symbol_name)(jl_gf_mtable(F)->name), "
    method_to_break")==0
```

Since this function is used for every call, you will make everything 1000x slower if you do this.

Dealing with signals

Julia requires a few signal to function property. The profiler uses SIGUSR2 for sampling and the garbage collector uses SIGSEGV for threads synchronization. If you are debugging some code that uses the profiler or multiple threads, you may want to let the debugger ignore these signals since they can be triggered very often during normal operations. The command to do this in GDB is (replace SIGSEGV with SIGUSRS or other signals you want to ignore):

```
(gdb) handle SIGSEGV noprint nostop pass
```

The corresponding LLDB command is (after the process is started):

```
(lldb) pro hand -p true -s false -n false SIGSEGV
```

If you are debugging a segfault with threaded code, you can set a breakpoint on `jl_critical_error` (`sigdie_handler` should also work on Linux and BSD) in order to only catch the actual segfault rather than the GC synchronization points.

Debugging during Julia's build process (bootstrap)

Errors that occur during `make` need special handling. Julia is built in two stages, constructing `sys0` and `sys.ji`. To see what commands are running at the time of failure, use `make VERBOSE=1`.

At the time of this writing, you can debug build errors during the `sys0` phase from the `base` directory using:

```
julia/base$ gdb --args ../usr/bin/julia-debug -C native --build ../usr/lib/julia/sys0 sysimg.jl
```

You might need to delete all the files in `usr/lib/julia/` to get this to work.

You can debug the `sys.ji` phase using:

```
julia/base$ gdb --args ../usr/bin/julia-debug -C native --build ../usr/lib/julia/sys -J ../usr/
    lib/julia/sys0.ji sysimg.jl
```

By default, any errors will cause Julia to exit, even under gdb. To catch an error "in the act", set a breakpoint in `jl_error` (there are several other useful spots, for specific kinds of failures, including: `jl_too_few_args`, `jl_too_many_args`, and `jl_throw`).

Once an error is caught, a useful technique is to walk up the stack and examine the function by inspecting the related call to `jl_apply`. To take a real-world example:

```
Breakpoint 1, jl_throw (e=0x7ffdf42de400) at task.c:802
802 {
(gdb) p jl_(e)
ErrorException("auto_unbox: unable to determine argument type")
$2 = void
(gdb) bt 10
#0  jl_throw (e=0x7ffdf42de400) at task.c:802
#1  0x00007ffff65412fe in jl_error (str=0x7ffde56be000 <_j_str267> "auto_unbox:
   unable to determine argument type")
   at builtins.c:39
#2  0x00007ffde56bd01a in julia_convert_16886 ()
#3  0x00007ffff6541154 in jl_apply (f=0x7ffdf367f630, args=0x7fffffffc2b0, nargs=2) at julia.h
    :1281
...
```

The most recent `jl_apply` is at frame #3, so we can go back there and look at the AST for the function `julia_convert_16886`. This is the uniqued name for some method of `convert`. `f` in this frame is a `jl_function_t*`, so we can look at the type signature, if any, from the `specTypes` field:

```
(gdb) f 3
#3  0x00007ffff6541154 in jl_apply (f=0x7ffdf367f630, args=0x7fffffffc2b0, nargs=2) at julia.h
    :1281
1281            return f->fptr((jl_value_t*)f, args, nargs);
(gdb) p f->linfo->specTypes
$4 = (jl_tupletype_t *) 0x7ffdf39b1030
(gdb) p jl_( f->linfo->specTypes )
Tuple{Type{Float32}, Float64}           # <-- type signature for julia_convert_16886
```

Then, we can look at the AST for this function:

```
(gdb) p jl_( jl_uncompress_ast(f->linfo, f->linfo->ast) )
Expr(:lambda, Array{Any, 1}[:#s29, :x], Array{Any, 1}[Array{Any, 1}[], Array{Any, 1}[Array{Any,
    1}[:#s29, :Any, 0], Array{Any, 1}[:x, :Any, 0]], Array{Any, 1}[], 0], Expr(:body,
Expr(:line, 90, :float.jl)::Any,
Expr(:return, Expr(:call, :box, :Float32, Expr(:call, :fptrunc, :Float32, :x)::Any)::Any)::Any)::
    Any)::Any
```

Finally, and perhaps most usefully, we can force the function to be recompiled in order to step through the codegen process. To do this, clear the cached `functionObject` from the `jl_lamdbda_info_t*`:

```
(gdb) p f->linfo->functionObject
$8 = (void *) 0x1289d070
(gdb) set f->linfo->functionObject = NULL
```

Then, set a breakpoint somewhere useful (e.g. `emit_function`, `emit_expr`, `emit_call`, etc.), and run codegen:

```
(gdb) p jl_compile(f)
... # your breakpoint here
```

Debugging precompilation errors

Module precompilation spawns a separate Julia process to precompile each module. Setting a breakpoint or catching failures in a precompile worker requires attaching a debugger to the worker. The easiest approach is to set the debugger watch for new process launches matching a given name. For example:

```
(gdb) attach -w -n julia-debug
```

or:

```
(lldb) process attach -w -n julia-debug
```

Then run a script/command to start precompilation. As described earlier, use conditional breakpoints in the parent process to catch specific file-loading events and narrow the debugging window. (some operating systems may require alternative approaches, such as following each `fork` from the parent process)

Mozilla's Record and Replay Framework (rr)

Julia now works out of the box with rr, the lightweight recording and deterministic debugging framework from Mozilla. This allows you to replay the trace of an execution deterministically. The replayed execution's address spaces, register contents, syscall data etc are exactly the same in every run.

A recent version of rr (3.1.0 or higher) is required.

71.3 Using Valgrind with Julia

Valgrind is a tool for memory debugging, memory leak detection, and profiling. This section describes things to keep in mind when using Valgrind to debug memory issues with Julia.

General considerations

By default, Valgrind assumes that there is no self modifying code in the programs it runs. This assumption works fine in most instances but fails miserably for a just-in-time compiler like `julia`. For this reason it is crucial to pass `--smc-check=all-non-file` to `valgrind`, else code may crash or behave unexpectedly (often in subtle ways).

In some cases, to better detect memory errors using Valgrind it can help to compile `julia` with memory pools disabled. The compile-time flag `MEMDEBUG` disables memory pools in Julia, and `MEMDEBUG2` disables memory pools in FemtoLisp. To build `julia` with both flags, add the following line to `Make.user`:

```
CFLAGS = -DMEMDEBUG -DMEMDEBUG2
```

Another thing to note: if your program uses multiple workers processes, it is likely that you want all such worker processes to run under Valgrind, not just the parent process. To do this, pass `--trace-children=yes` to `valgrind`.

Suppressions

Valgrind will typically display spurious warnings as it runs. To reduce the number of such warnings, it helps to provide a suppressions file to Valgrind. A sample suppressions file is included in the Julia source distribution at `contrib/valgrind-julia.supp`.

The suppressions file can be used from the `julia/` source directory as follows:

```
$ valgrind --smc-check=all-non-file --suppressions=contrib/valgrind-julia.supp ./julia progname.
    jl
```

Any memory errors that are displayed should either be reported as bugs or contributed as additional suppressions. Note that some versions of Valgrind are shipped with insufficient default suppressions, so that may be one thing to consider before submitting any bugs.

Running the Julia test suite under Valgrind

It is possible to run the entire Julia test suite under Valgrind, but it does take quite some time (typically several hours). To do so, run the following command from the `julia/test/` directory:

```
valgrind --smc-check=all-non-file --trace-children=yes --suppressions=$PWD/../contrib/valgrind-
    julia.supp ../julia runtests.jl all
```

If you would like to see a report of "definite" memory leaks, pass the flags `--leak-check=full --show-leak-kinds=definite` to `valgrind` as well.

Caveats

Valgrind currently does not support multiple rounding modes, so code that adjusts the rounding mode will behave differently when run under Valgrind.

In general, if after setting `--smc-check=all-non-file` you find that your program behaves differently when run under Valgrind, it may help to pass `--tool=none` to `valgrind` as you investigate further. This will enable the minimal Valgrind machinery but will also run much faster than when the full memory checker is enabled.

71.4 Sanitizer support

General considerations

Using Clang's sanitizers obviously require you to use Clang (USECLANG=1), but there's another catch: most sanitizers require a run-time library, provided by the host compiler, while the instrumented code generated by Julia's JIT relies on functionality from that library. This implies that the LLVM version of your host compiler matches that of the LLVM library used within Julia.

An easy solution is to have an dedicated build folder for providing a matching toolchain, by building with BUILD_LLVM_CLANG=1 and overriding LLVM_USE_CMAKE=1 (Autotool-based builds are incompatible with ASAN). You can then refer to this toolchain from another build folder by specifying USECLANG=1 while overriding the CC and CXX variables.

Address Sanitizer (ASAN)

For detecting or debugging memory bugs, you can use Clang's address sanitizer (ASAN). By compiling with SANITIZE=1 you enable ASAN for the Julia compiler and its generated code. In addition, you can specify LLVM_SANITIZE=1 to sanitize the LLVM library as well. Note that these options incur a high performance and memory cost. For example, using ASAN for Julia and LLVM makes `testall1` takes 8-10 times as long while using 20 times as much memory (this can be reduced to respectively a factor of 3 and 4 by using the options described below).

By default, Julia sets the `allow_user_segv_handler=1` ASAN flag, which is required for signal delivery to work properly. You can define other options using the ASAN_OPTIONS environment flag, in which case you'll need to repeat the default option mentioned before. For example, memory usage can be reduced by specifying `fast_unwind_on_malloc=0` and `malloc_context_size=2`, at the cost of backtrace accuracy. For now, Julia also sets `detect_leaks=0`, but this should be removed in the future.

Memory Sanitizer (MSAN)

For detecting use of uninitialized memory, you can use Clang's memory sanitizer (MSAN) by compiling with SANITIZE_MEMORY=1.