

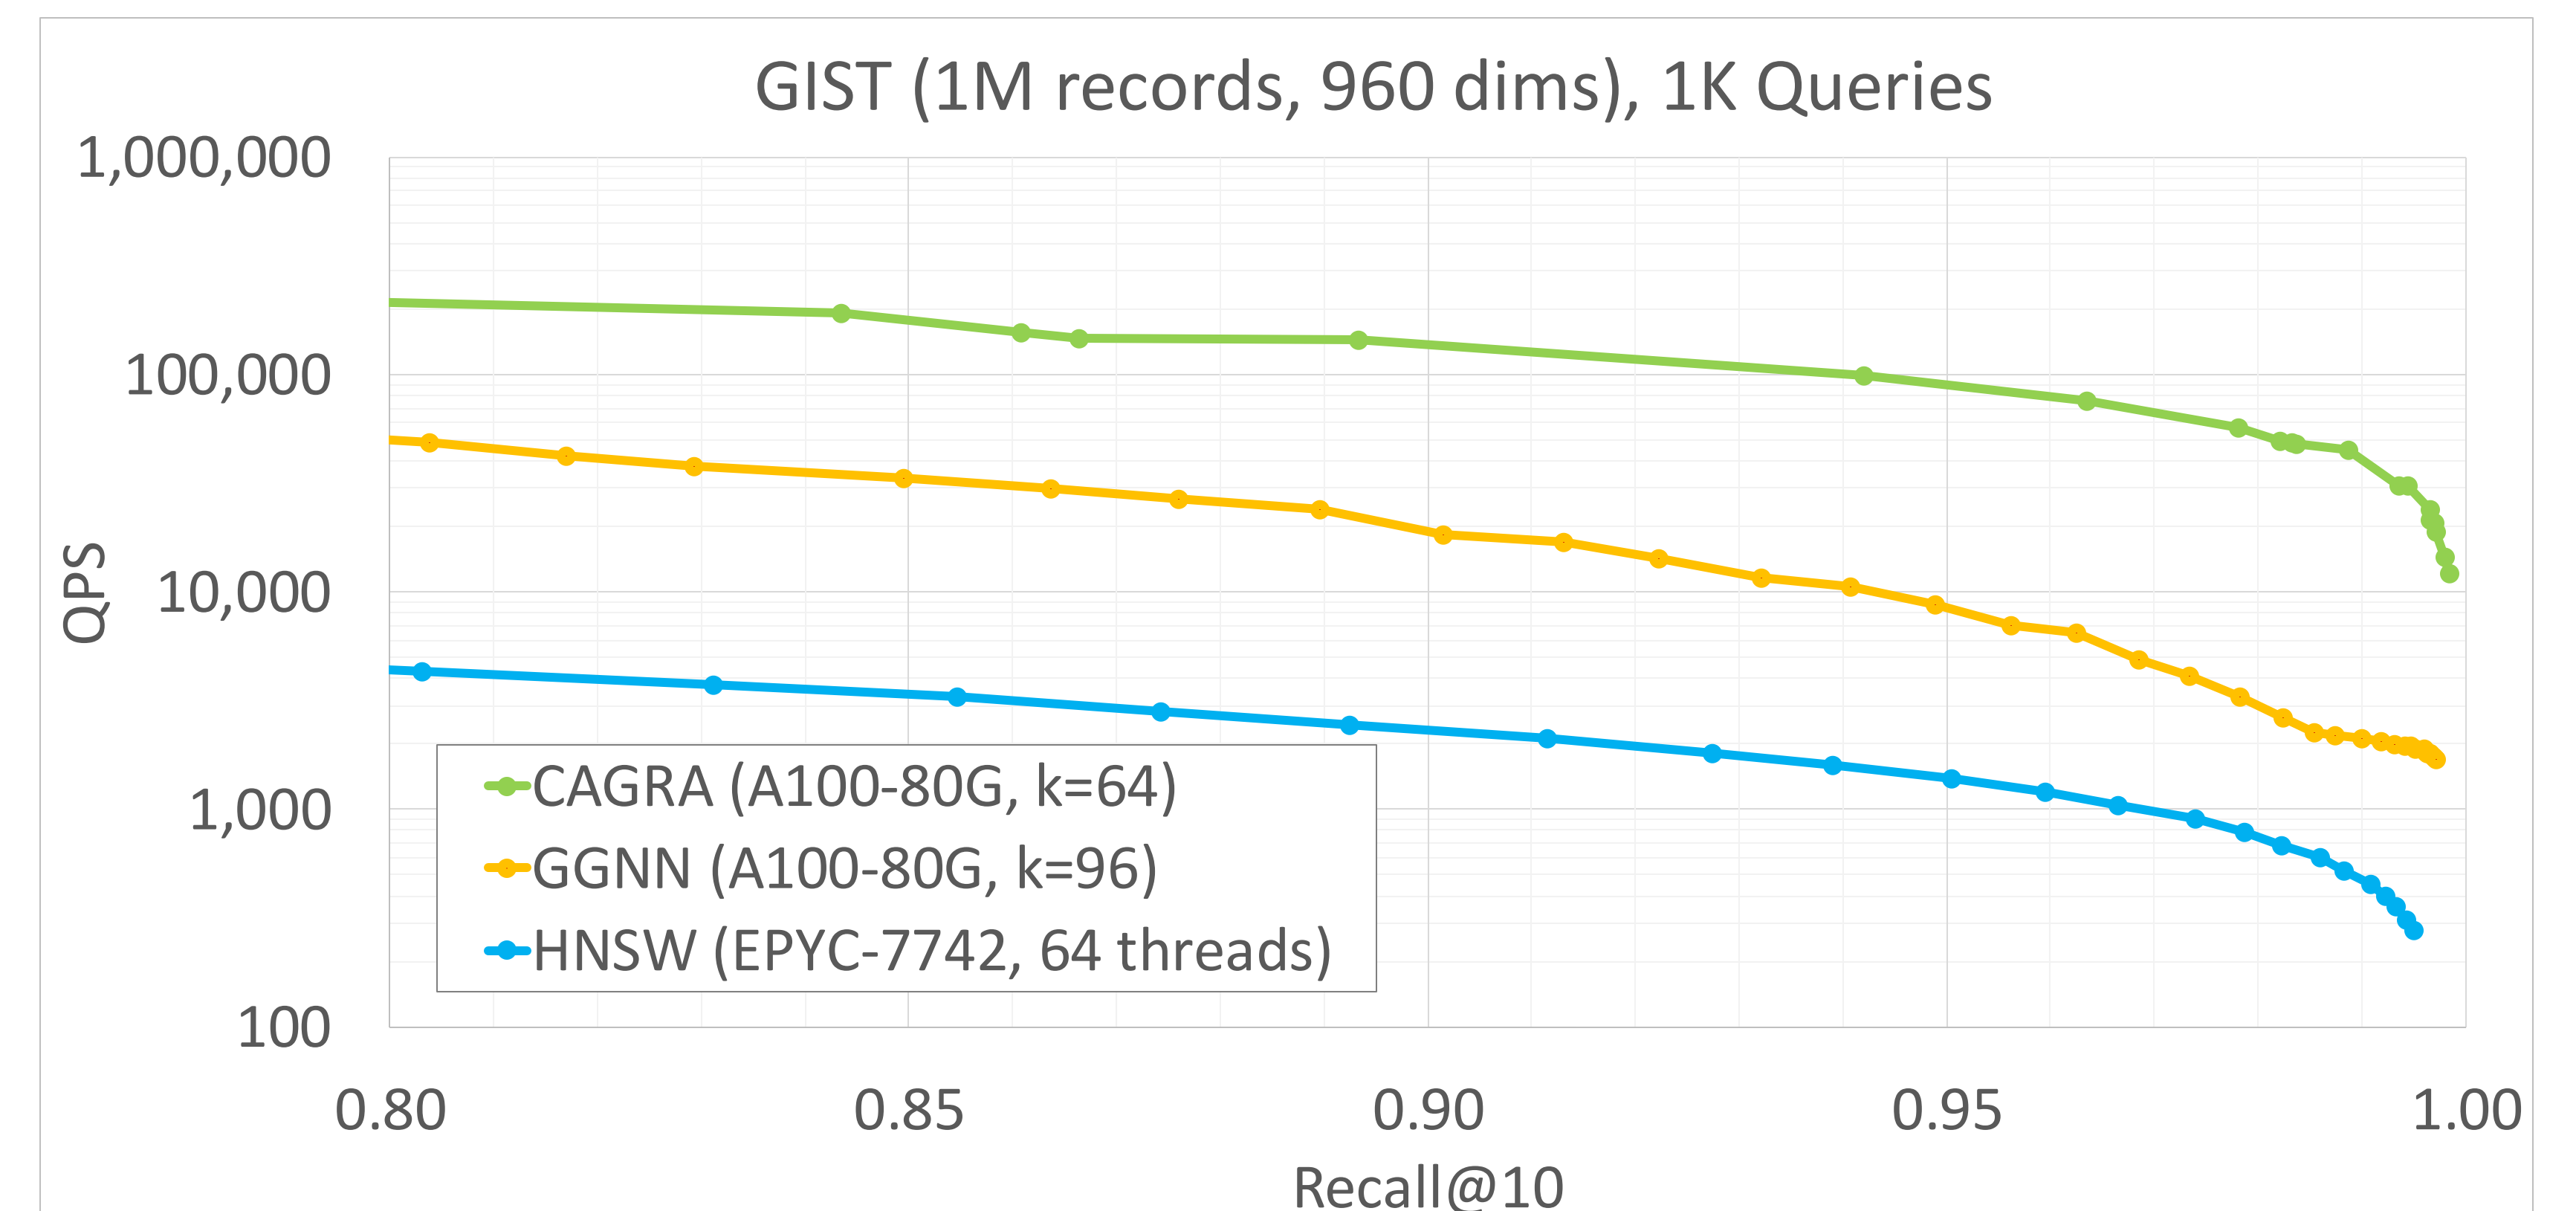
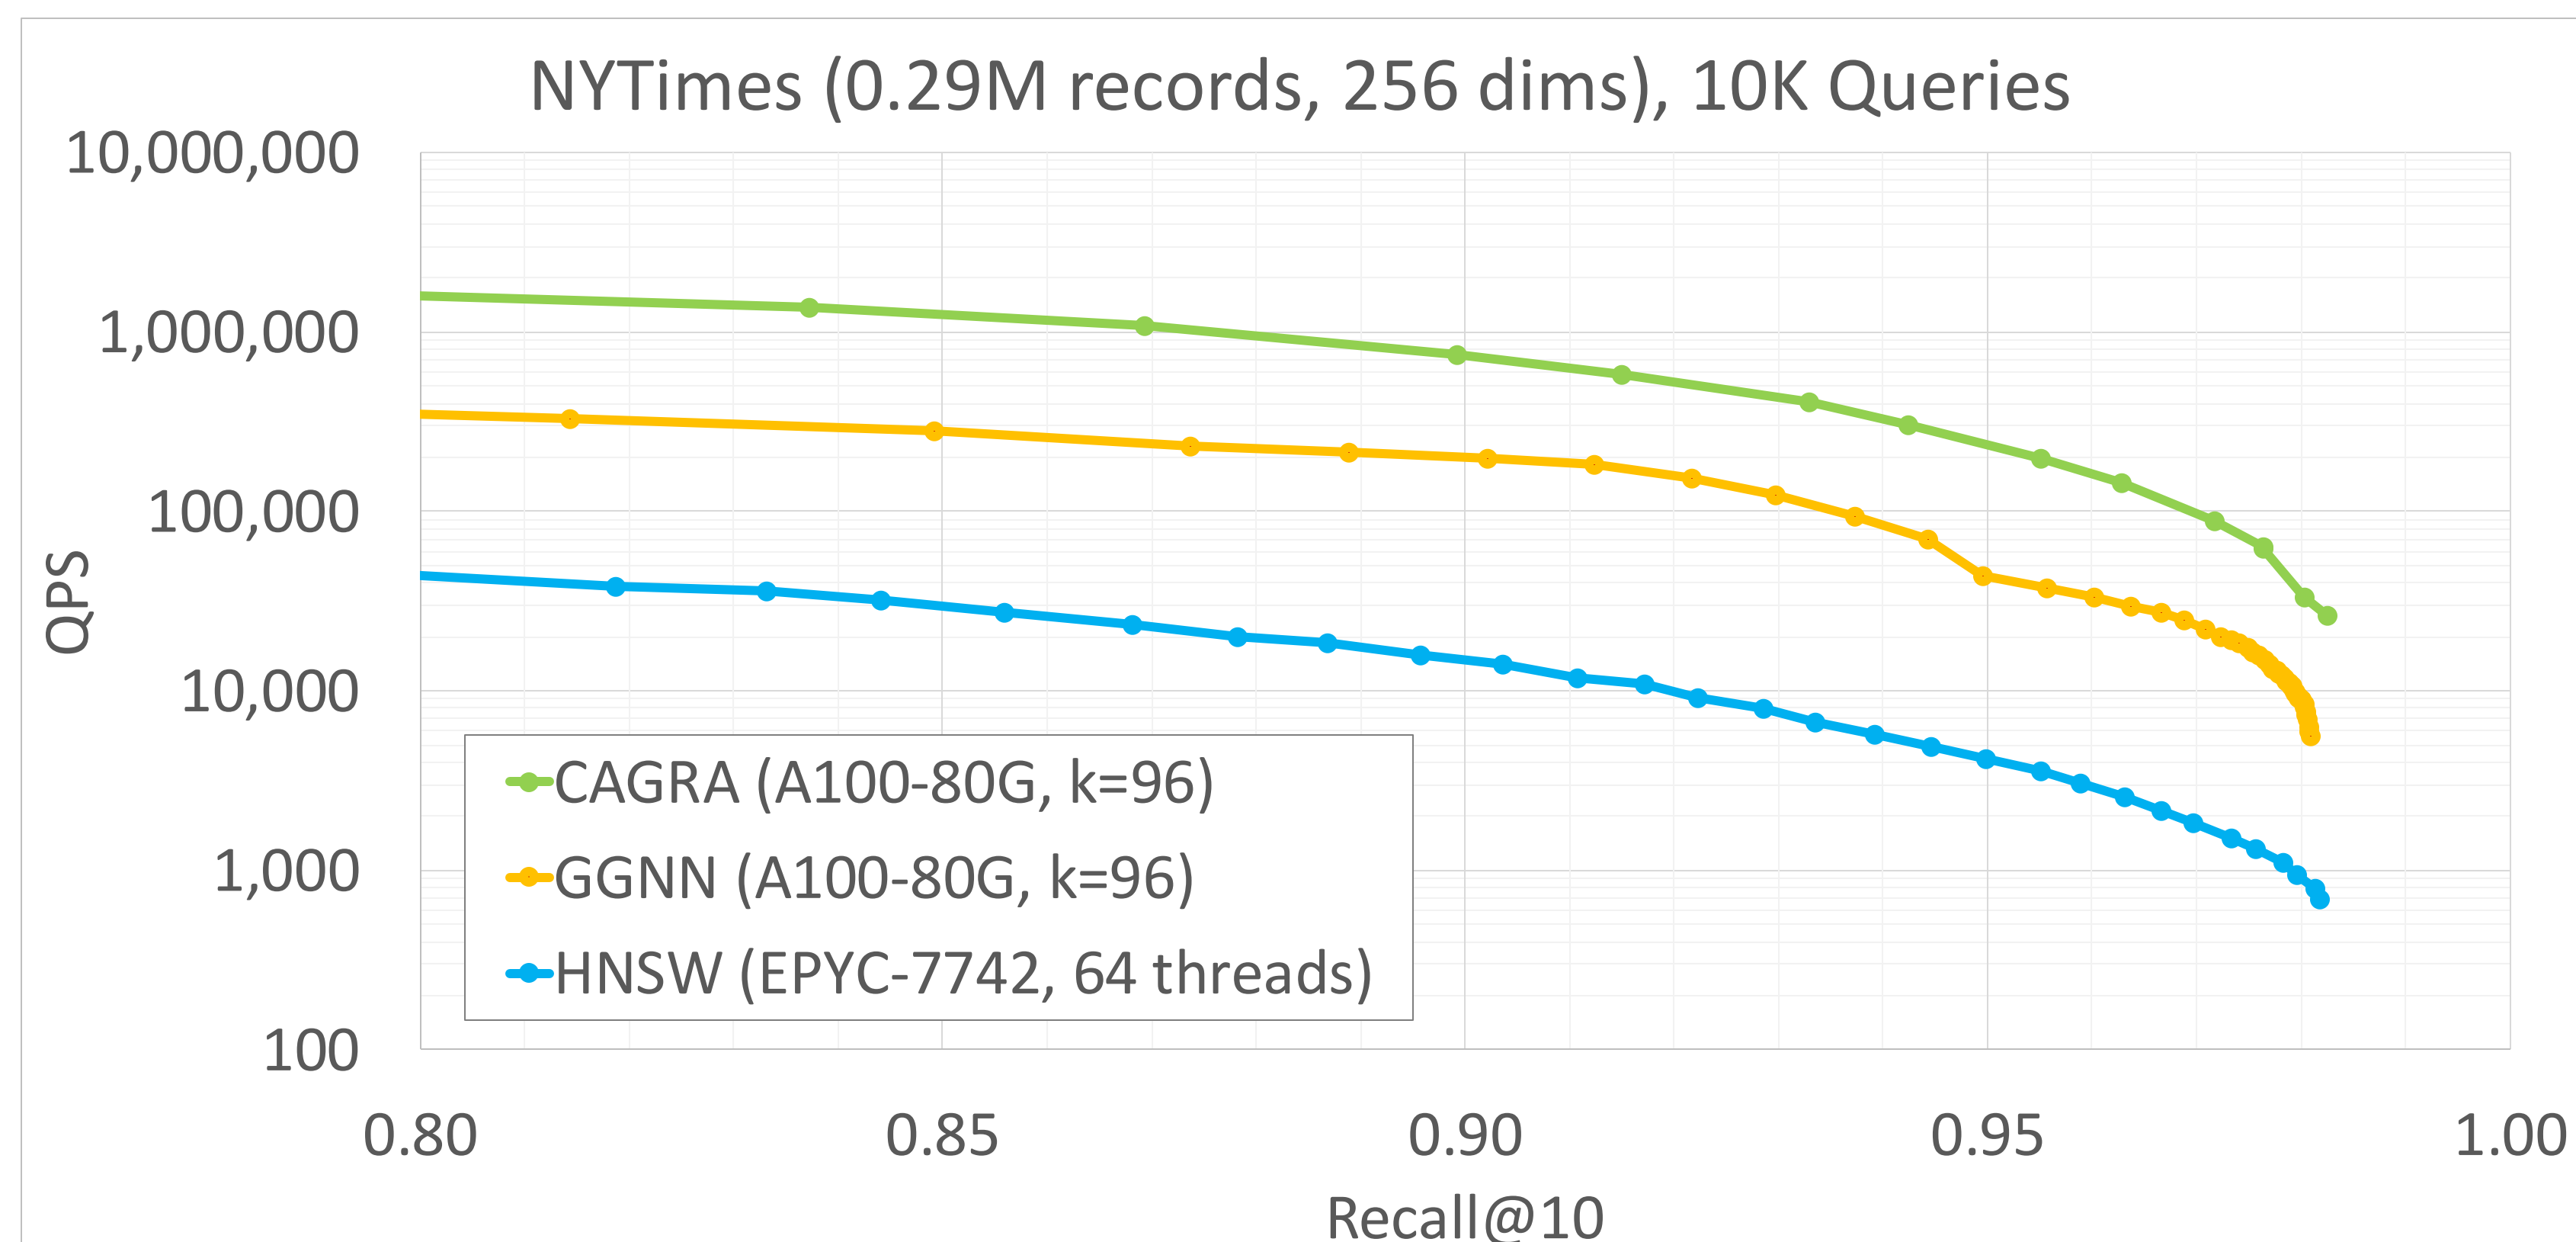

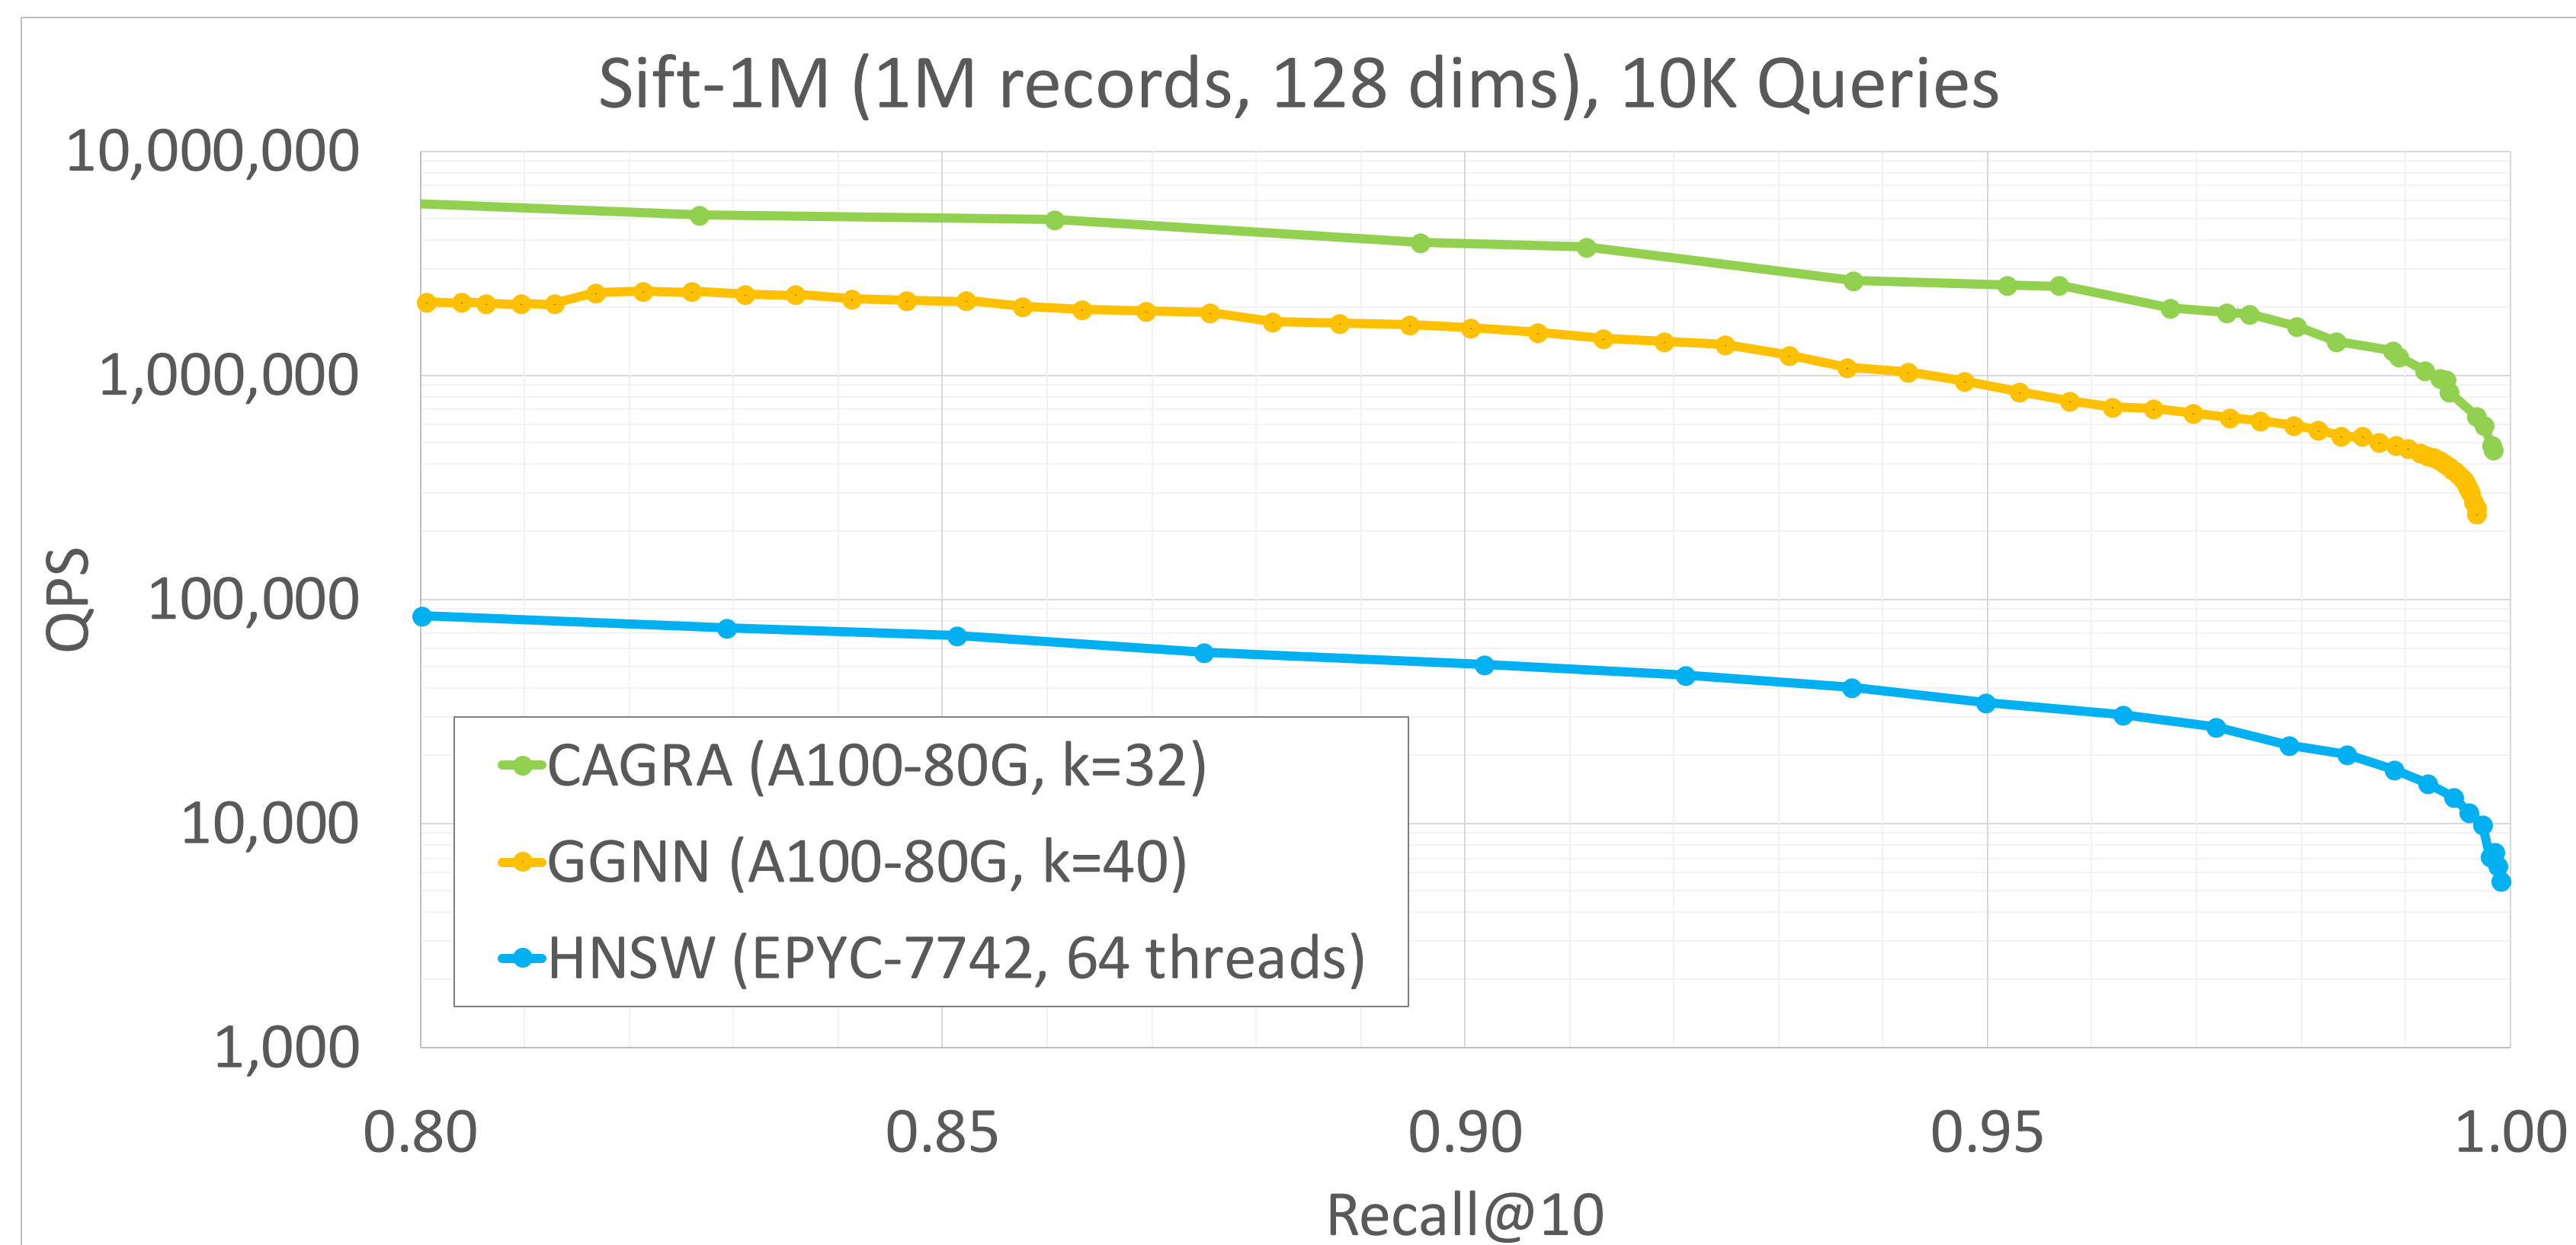


Performance comparison between CAGRA, GGNN and HNSW

Akira Naruse, Senior developer technology engineer | 2023/2/20

Large Batch Performance

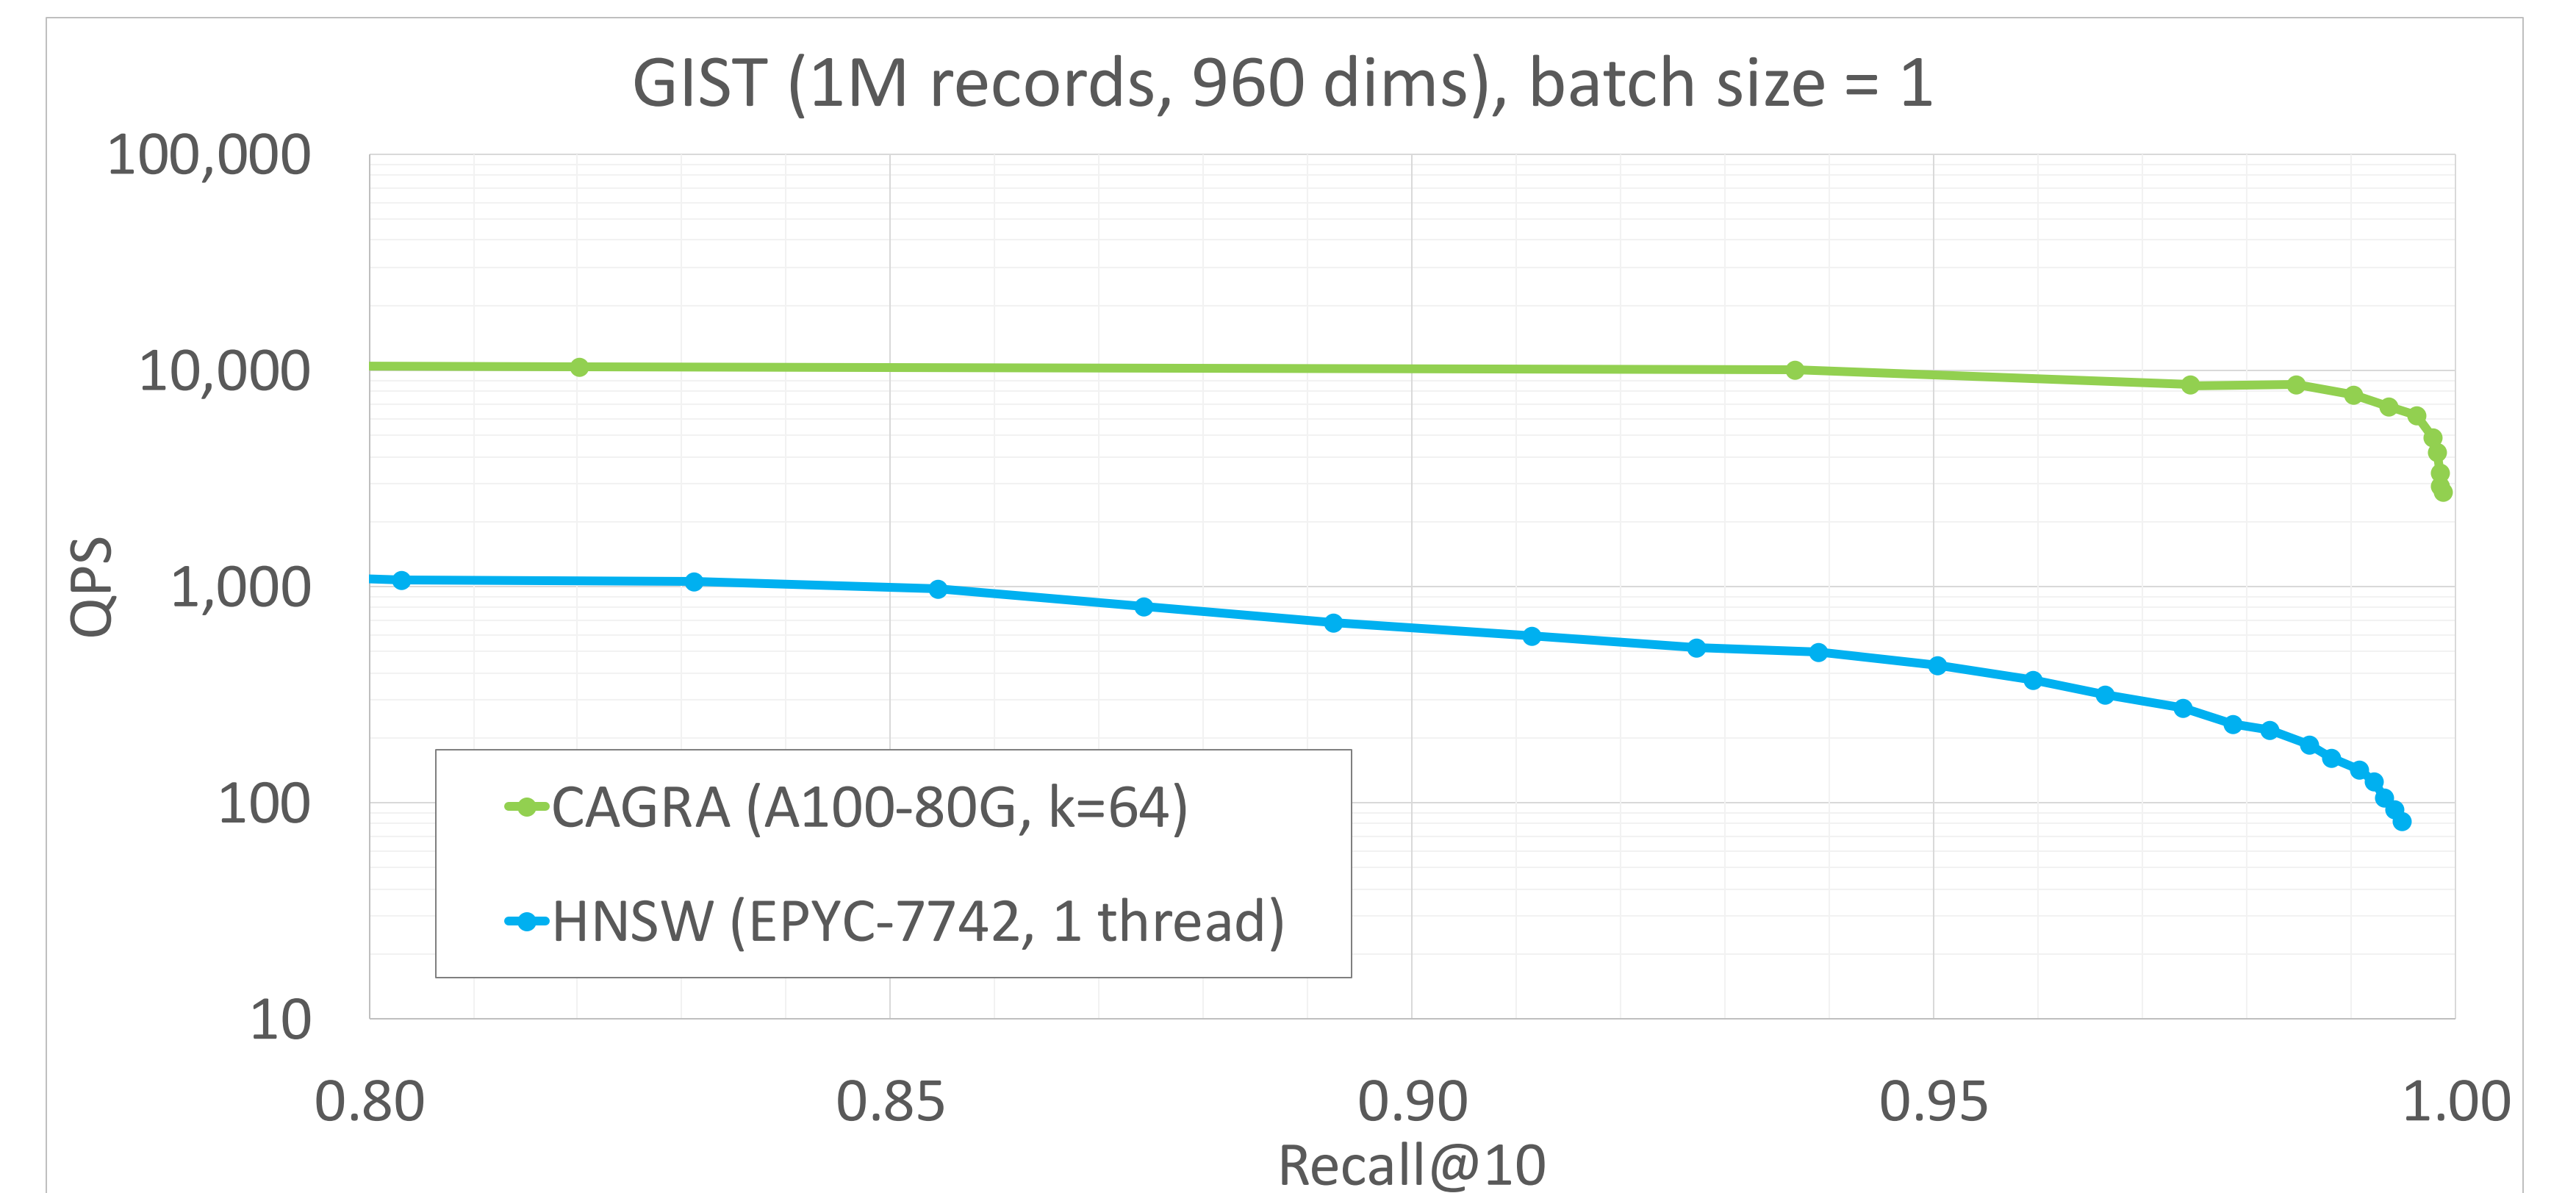


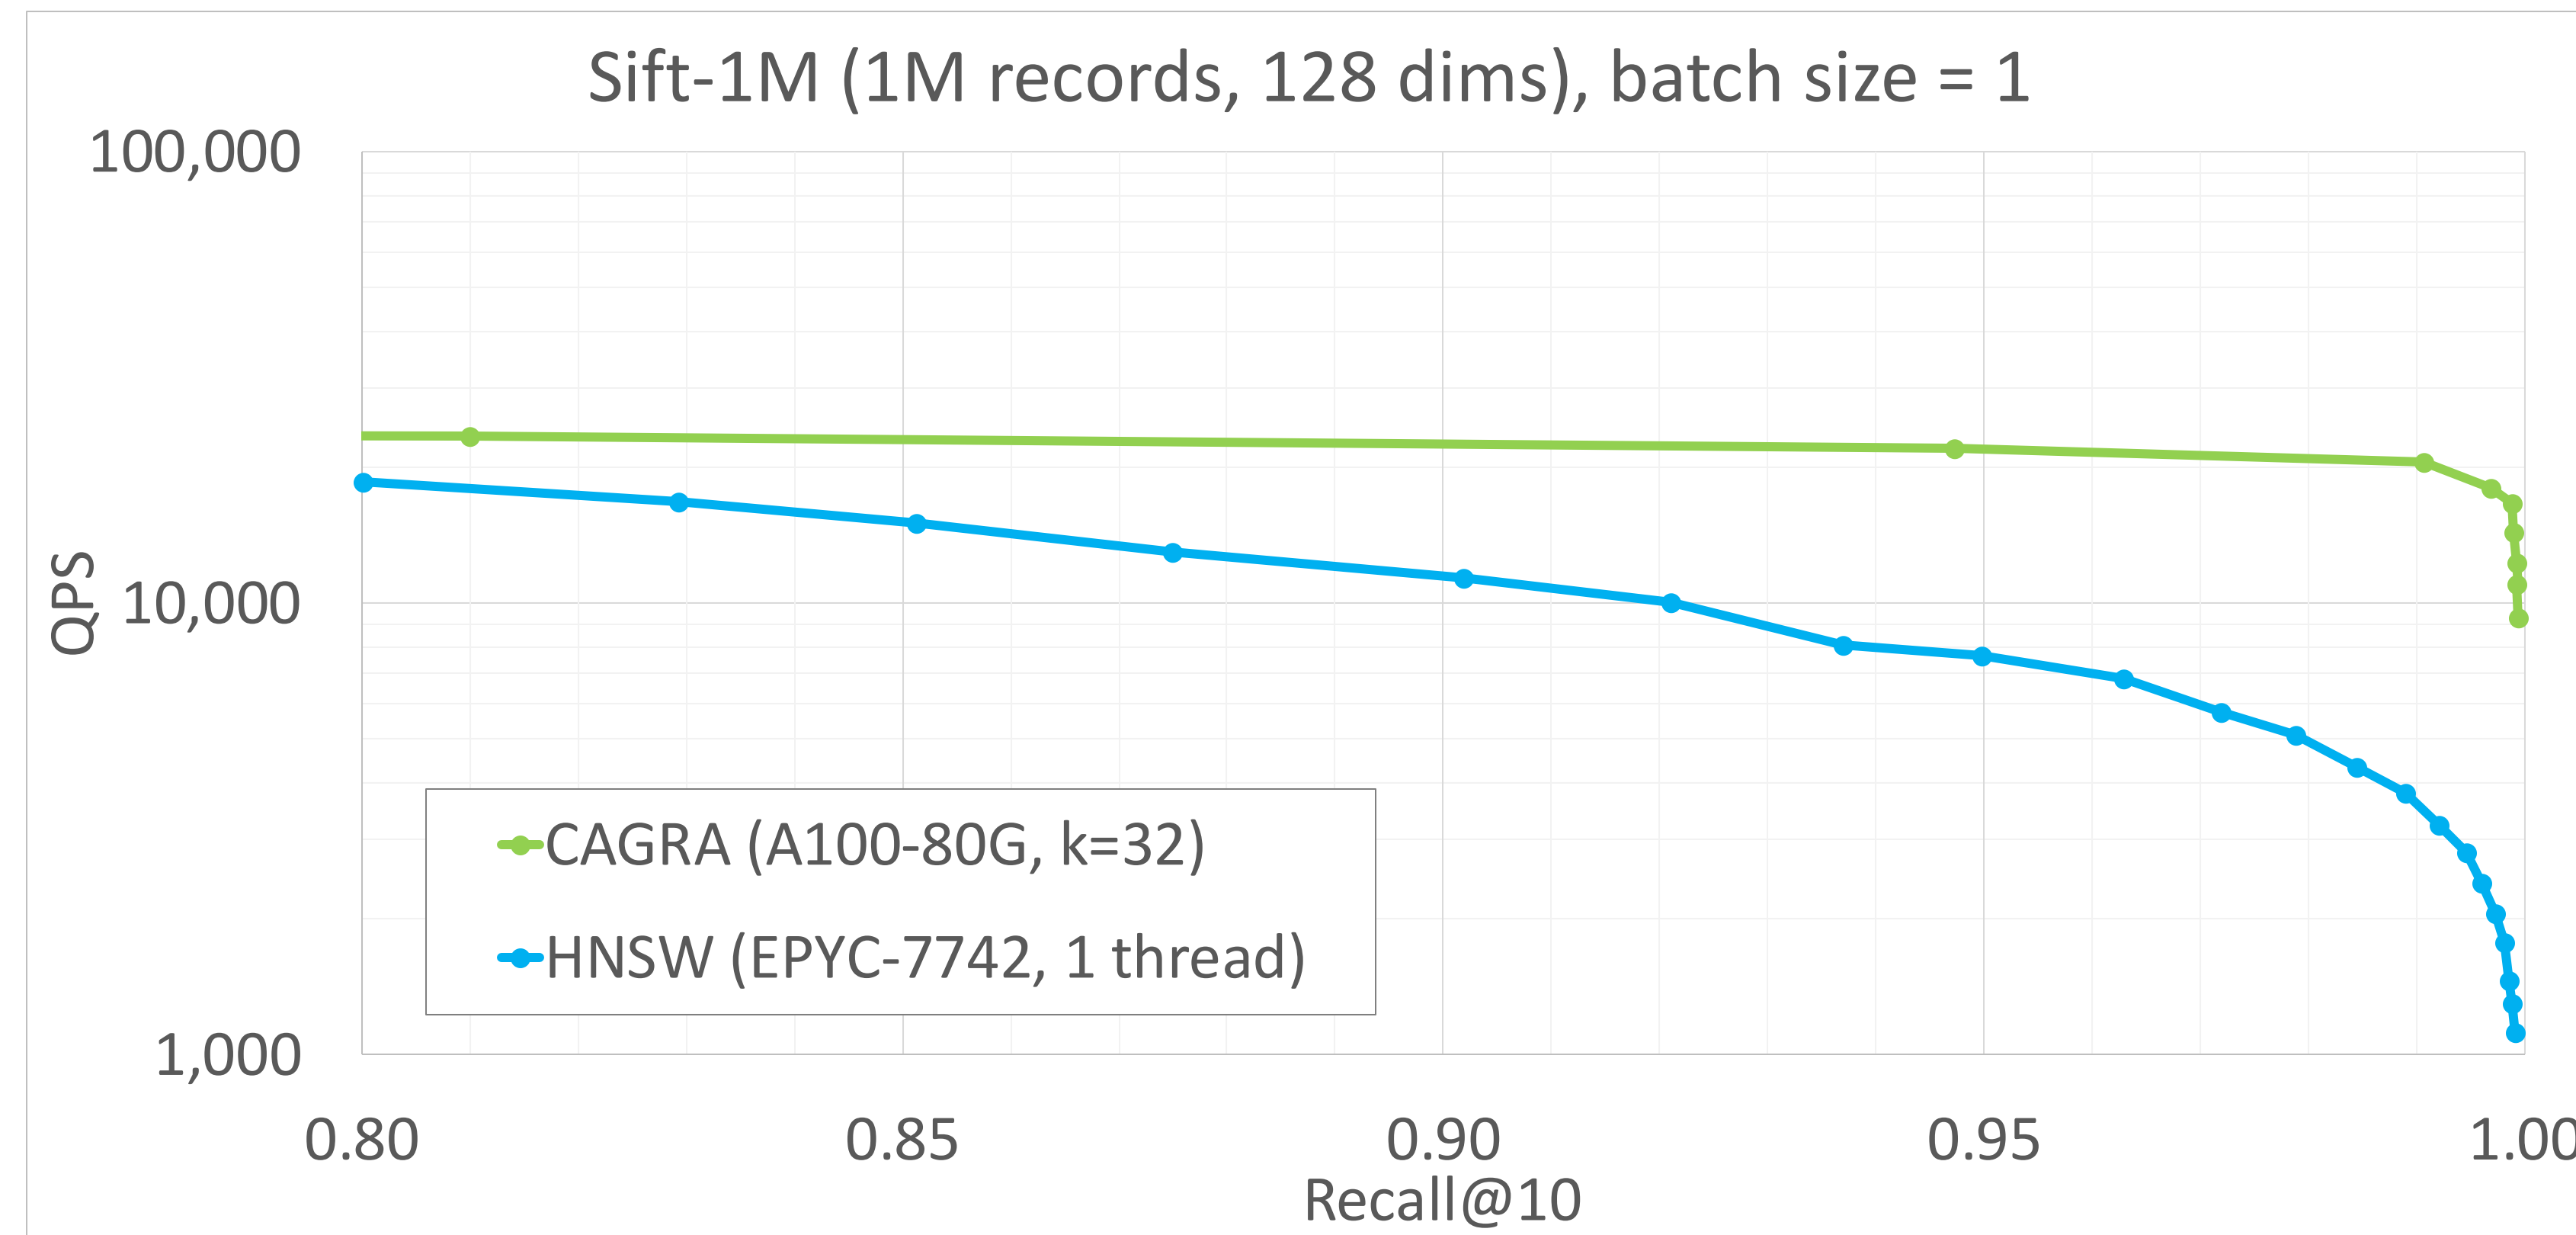
GPU: A100-80G, CPU: EPYC-7742



(*) GGNN: version 0.5, HNSW: hnsplib version 0.7.0, Machine: DGX-A100, CUDA 11.4

Single Query Performance

GPU: A100-80G, CPU: EPYC-7742



(*) HNSW: hnswlib version 0.7.0, Machine: DGX-A100, CUDA 11.4