

Microarchitectural Analysis of Graph BI Queries on RDBMS

Rathijit Sen

Yuanyuan Tian

Microsoft Gray Systems Lab

16th LDBC TUC Meeting, 2023

Motivation

- Graph databases on RDBMS

- Popular alternative to native graph dbs
- Interest in performance evaluations
- Less focus on hardware utilization

- How well is current hardware being utilized?
Focus: CPU server

- Hardware trends

- Increasing # CPU cores/socket
- HBM for CPUs
- Heterogeneity
- Specialization
- ...


- What capabilities are desirable for new backends?

LDDBC SNB BI on RDBMS (DuckDB)

- People interact by posting/reacting to messages in forums
- In-memory DuckDB database, Pyro5 server process
- Directed and undirected edges
- 20 BI read queries (+variants for parameter properties)
- Not fully implemented
 - Limits on shortest path computations (with weighted edges)
 - Q15, Q19, Q20
- Dual-socket CPU server, 64 logical cores (HT), 750 GB mem
 - SF=100, peak mem RSS: 248+ GB

Core Utilization & IPC


- How well are the many cores used?
- How efficiently are instructions executed?



- Inefficiencies => performance bottlenecks, potential room for improvement

Pipeline slots utilization


- How well are pipeline slots utilized?



- Alternative platforms/backends?

Memory subsystem bottlenecks


- How well is the memory subsystem performing?



- Offchip, but also onchip cache hierarchy bottlenecks
- DRAM: more stall cycles due to latency rather than due to bandwidth limits
- NUMA overheads

TLB/virtual memory bottlenecks

- Can huge pages improve performance?
- Transparent Huge Pages (THP)



- Workload speedup: 23%

Summary

- Microarchitectural Analysis of Graph BI workload on an RDBMS
 - Under-utilized cores and pipeline slots, low IPC
 - Bottlenecks in memory subsystem, onchip & offchip
 - TLB overheads
 - Additional results in DaMoN '23 paper
- New backends
 - Efficient instruction delivery & data access, memory capacity
 - Maybe weaker cores
- Microarchitectural analysis as complementary technique to algorithmic and software analyses



Thank You