Efficient Scheduling for Multiple Database Systems on Shared Hardware

Yi Jiang, Hamish Nicholson, Anastasia Ailamaki
firstname.lastname@epfl.ch

1. Cloud economics demands multi-tenancy

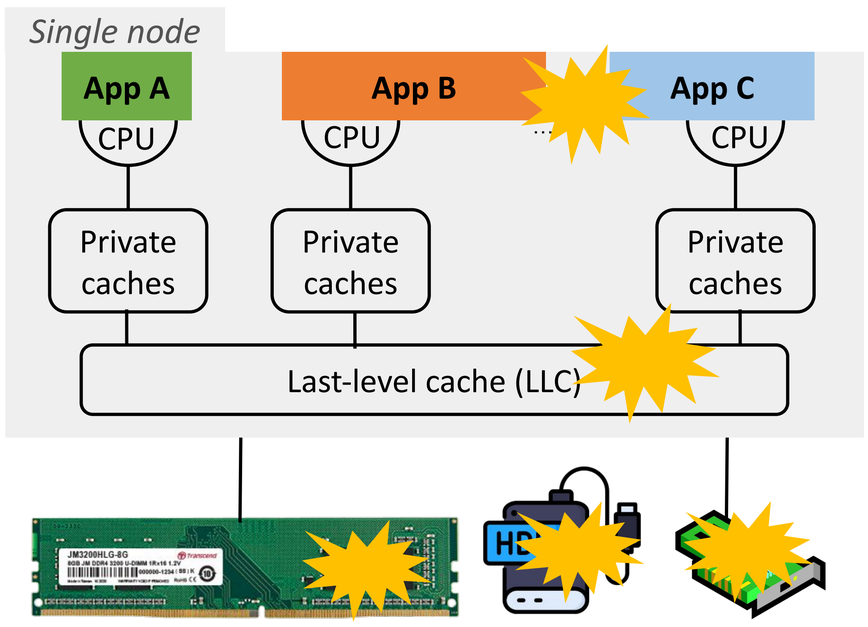

Performance

Tradeoff

Resource Efficiency

Goal:
- Improve resource efficiency
- Meet SLAs

Under co-location

2. Database workloads are diverse

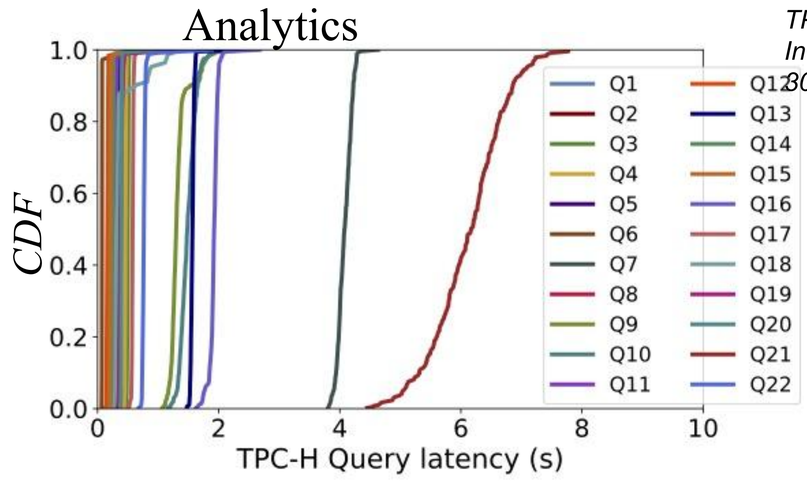

TPC-H in DuckDB, sf=100
Intel Xeon 6952P, 2x96 cores
300MB LLC, 512GB memory

Analytical queries
- **.1s - hours**
- Highly variable
- ms-scale scheduling interval

Transactions $17\ \mu s$
- **10 - 100's us**
- us-scale scheduling interval

3. Co-located database systems cannot fully utilize the underlying hardware

Existing resource schedulers
- Avoid application co-location at potential interference
  - *Bolt, Quasar, Borg, Heracles*
- Partition shared resources at runtime to reduce interference
  - *Ubik, Rubik, PARTIES, Caladan*

However
- Fixed decision interval *(PARTIES, CLITE, Aurora Serverless)*
- Not considering the relative importance:
  - Task latency and resource partitioning overhead
- Not adjusting the full resource spectrum *(Caladan)*

Persistent temporal mismatch between:

The latency to (re)configure resource partitions

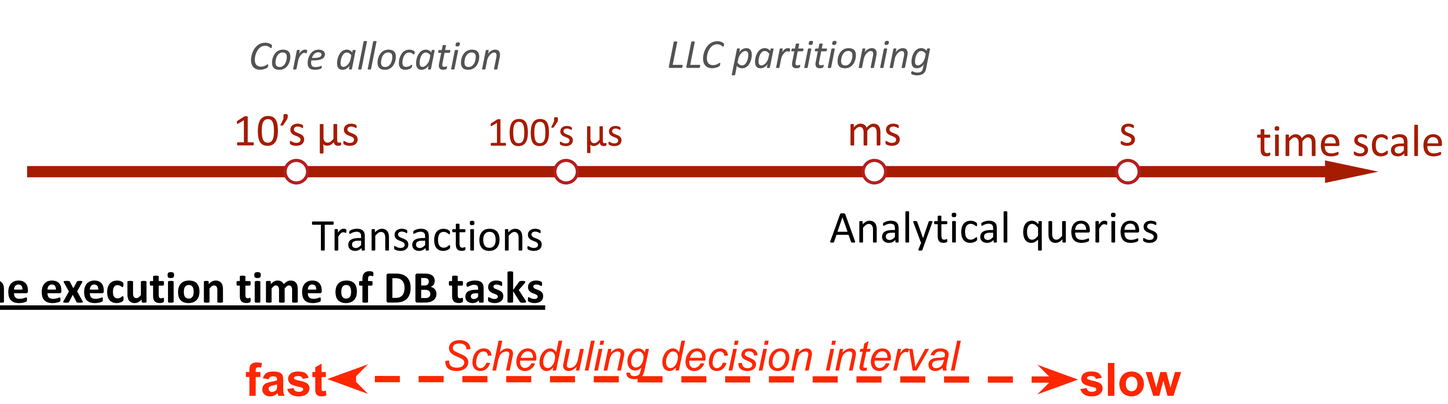

The execution time of DB tasks

fast ← *Scheduling decision interval* → **slow**

Fast and adaptive core allocation for Txns

+

Adjust the full resource spectrum

Resource scheduling decisions must operate at different frequencies

4. SCoRE system design

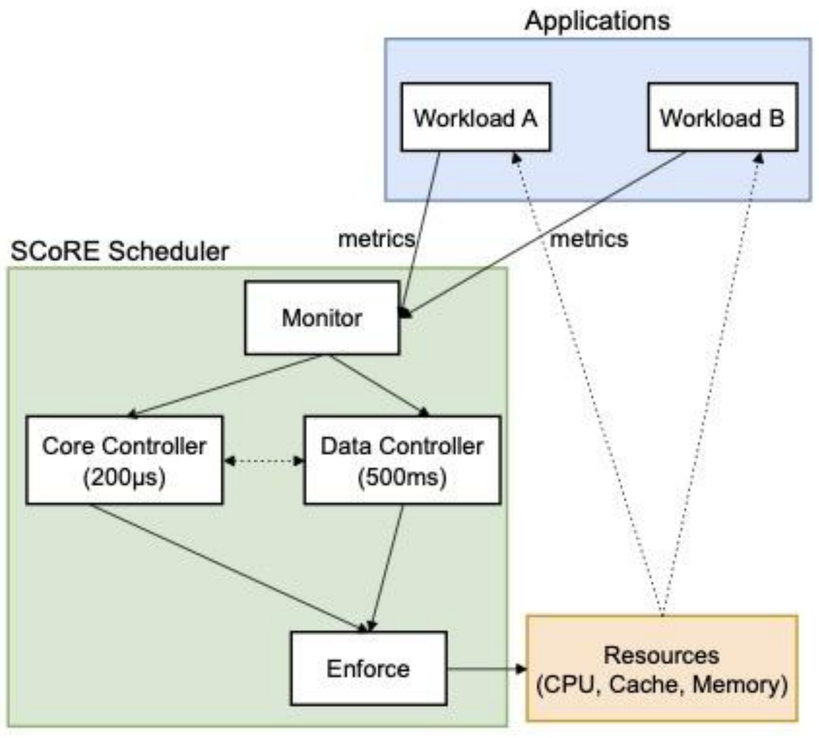

Decouple compute scheduling from data-access resources at fine/coarse granularity

5. Better runtime resource allocation

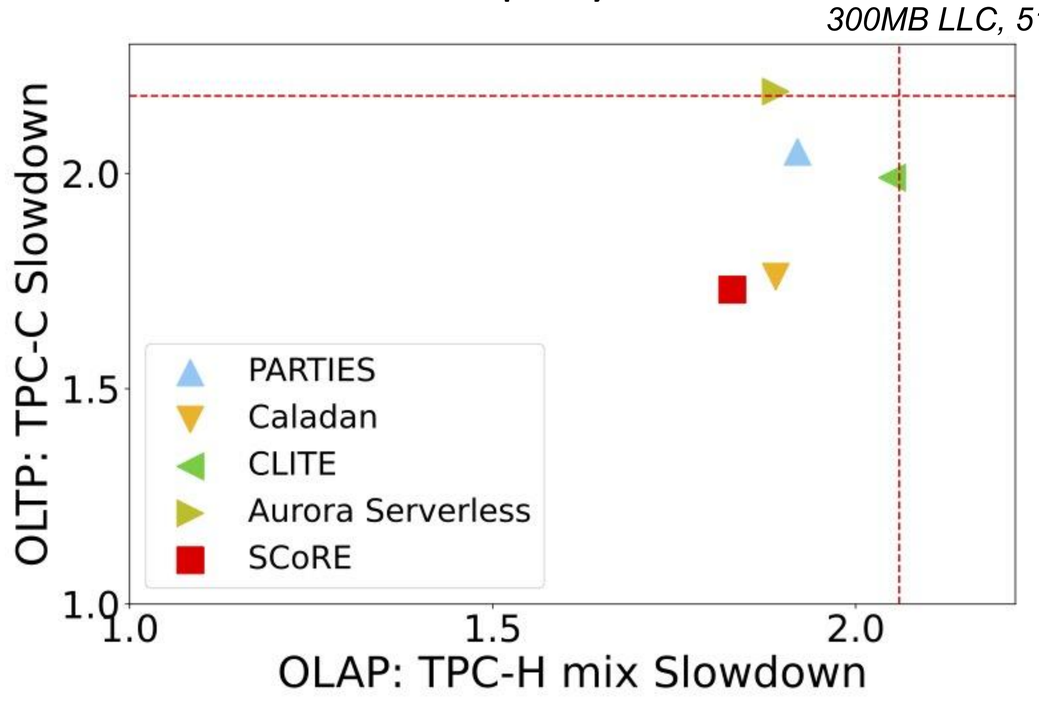

TPC-H in DuckDB, sf=100
TPC-C in Silo, sf=300
Intel Xeon 6952P, 2x96 cores
300MB LLC, 512GB memory

SCoRE's decoupled resource adjustment improves Txns/queries by 3-46%/6-22%

6. Efficient resource scheduling for DB systems under co-location

- Temporal mismatch between task latency & resource partitioning overhead
- Core control decoupling from data-access resource management

DATA-INTENSIVE APPLICATIONS AND SYSTEMS

EPFL