Adaptive
Data Storage & Placement in Distributed Database Systems

Michael Abebe

mtabebe@uwaterloo.ca

August 2022

UNIVERSITY OF WATERLOO

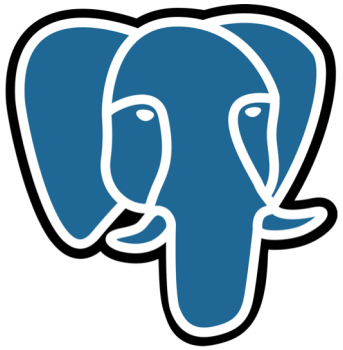

Distributed DBMSs are widely used

Amazon Aurora

Distributed Databases

How and where to store data?

Database Replication

Writers

Readers

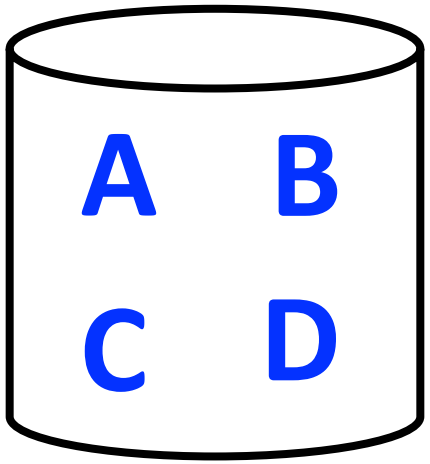

Master

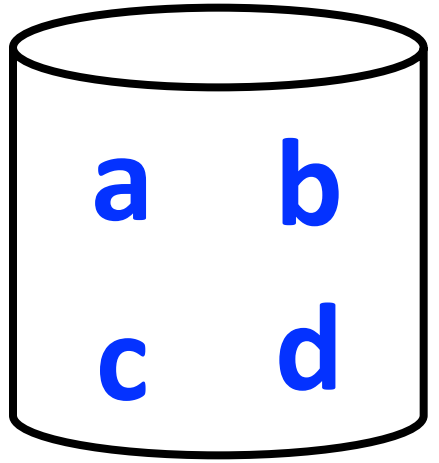

Replica

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

Distribution of Read Load

Master

Replica

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

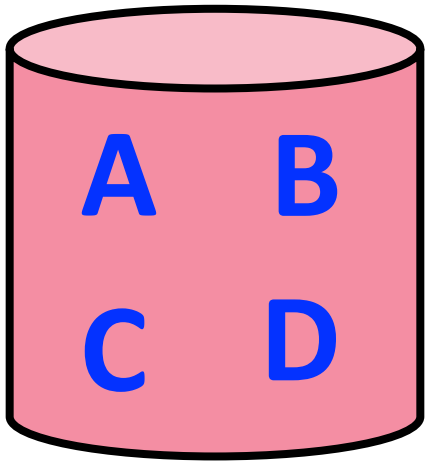

Master

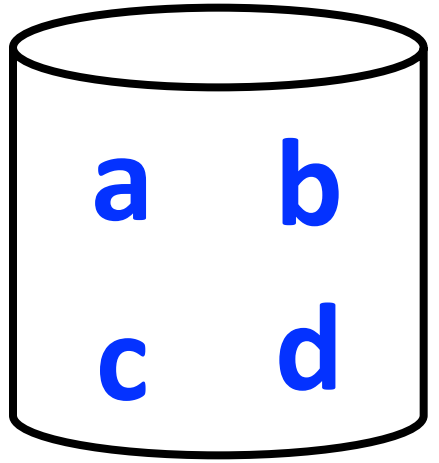

Replica

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

Performance bottleneck

Master

Replica

UNIVERSITY OF WATERLOO

Database Partitioning

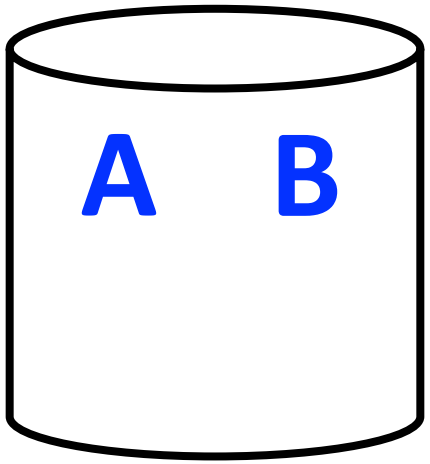

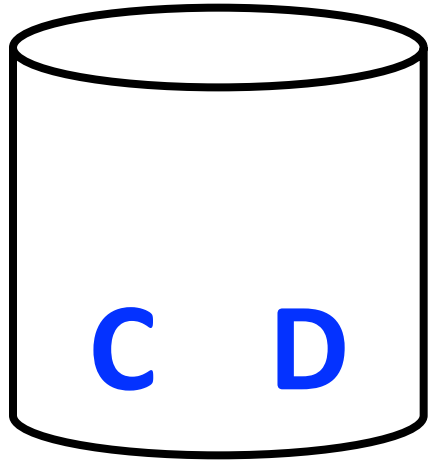

UNIVERSITY OF WATERLOO

Database Partitioning

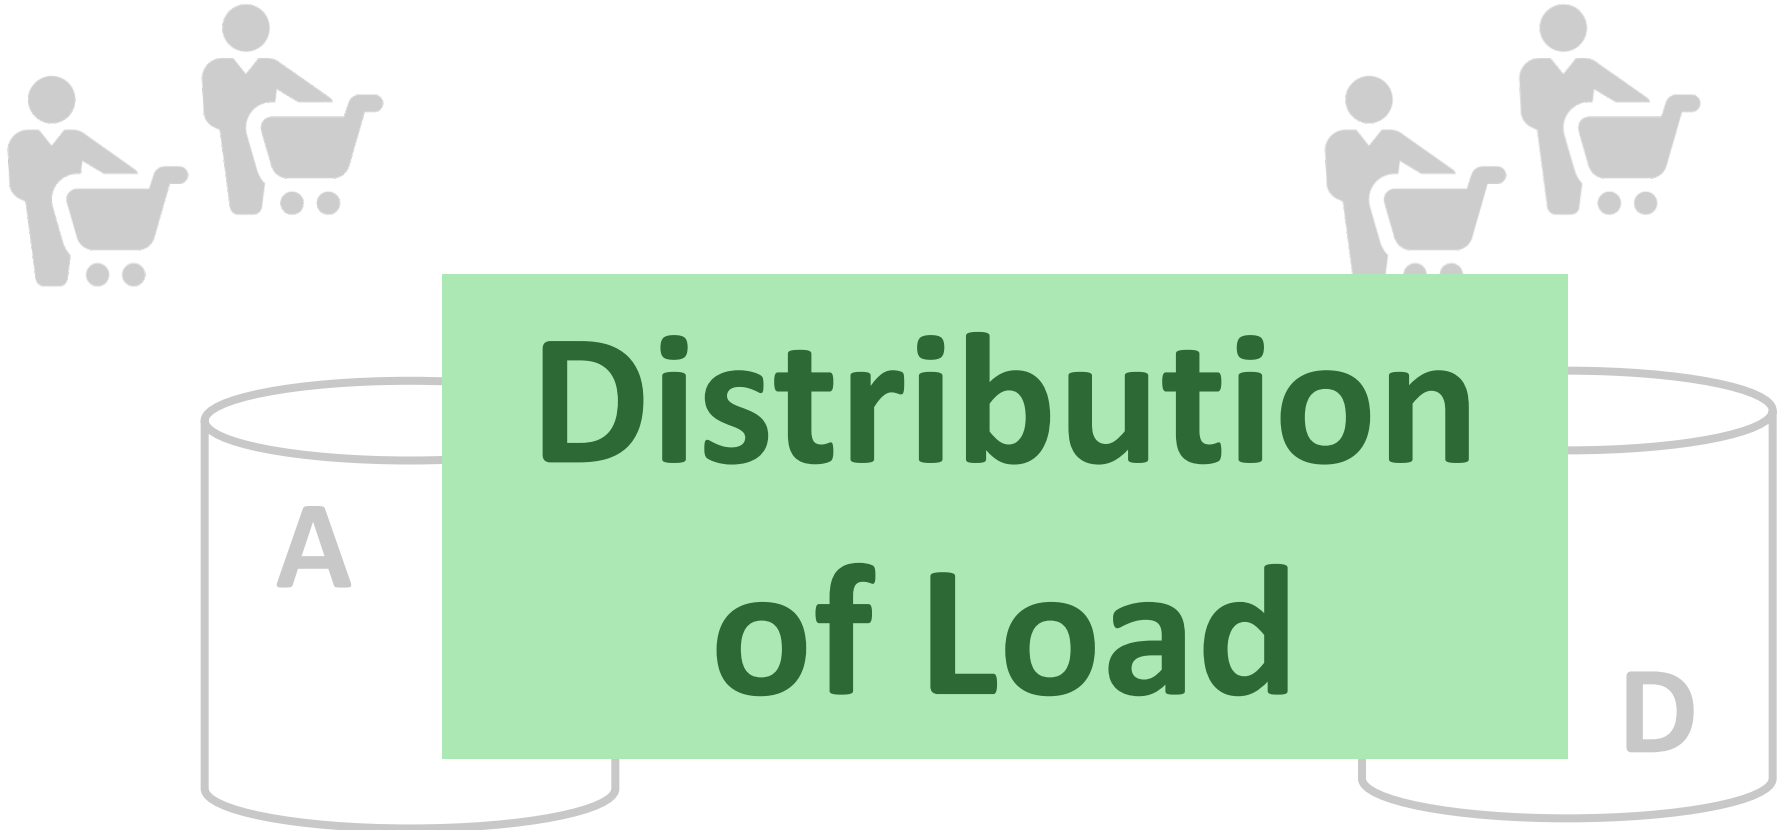

Distribution of Load

Database Partitioning

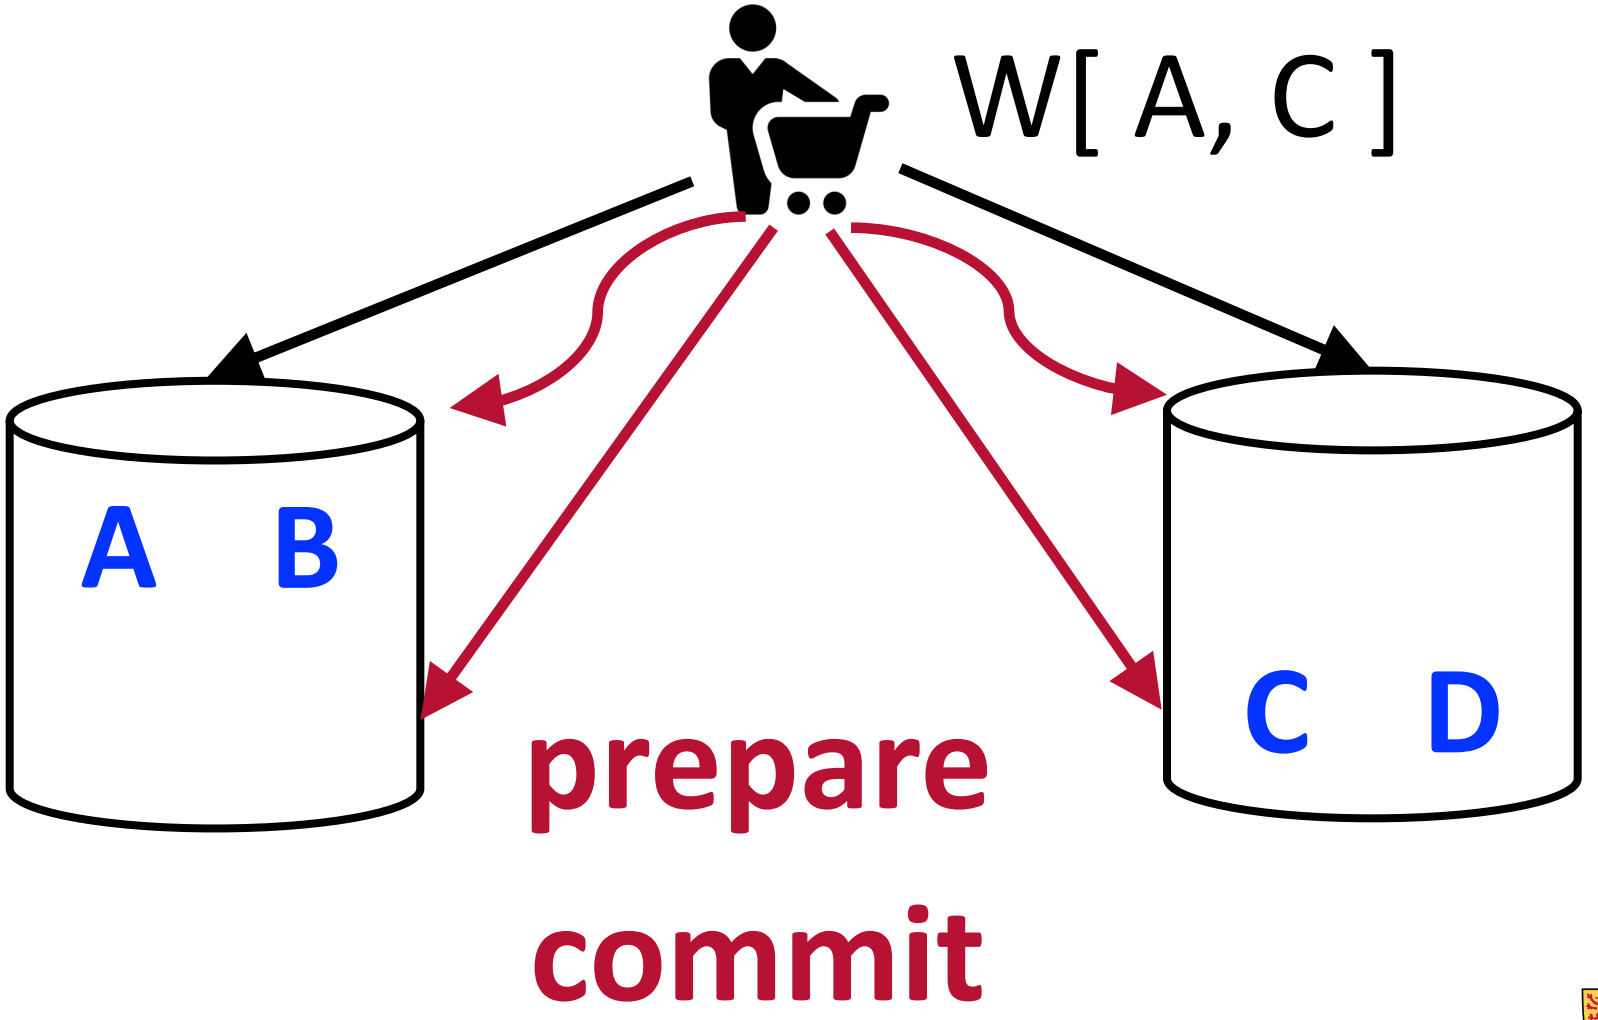

Database Partitioning

Expensive Coordination

commit

Storage Formats

Row Layout

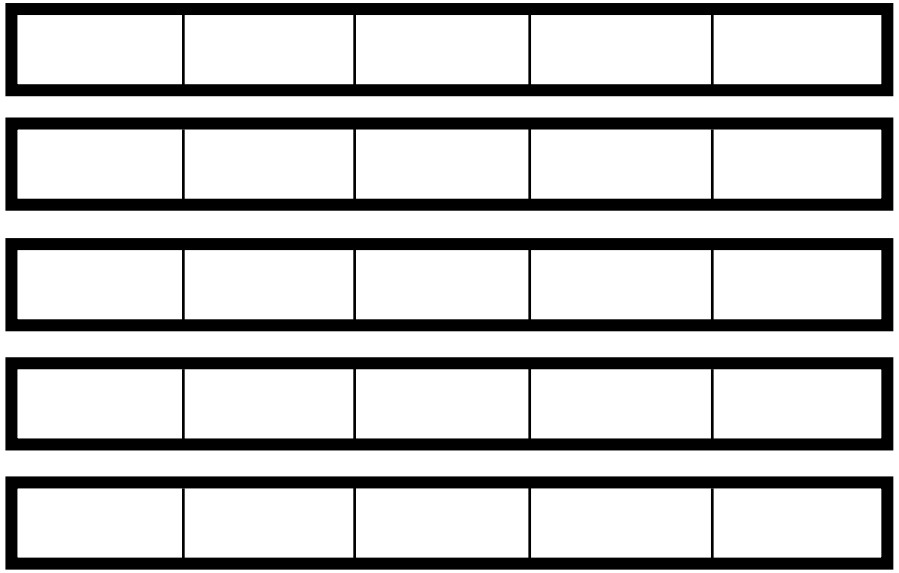

Updates

Columnar Layout

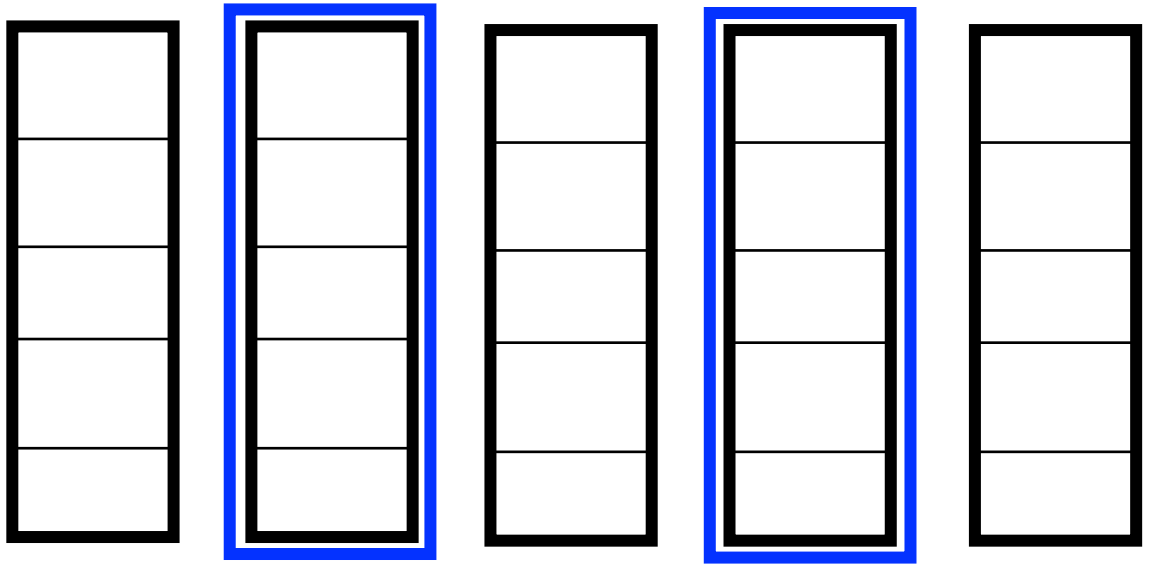

Analytics

UNIVERSITY OF WATERLOO

Storage Formats and HTAP

Row Layout

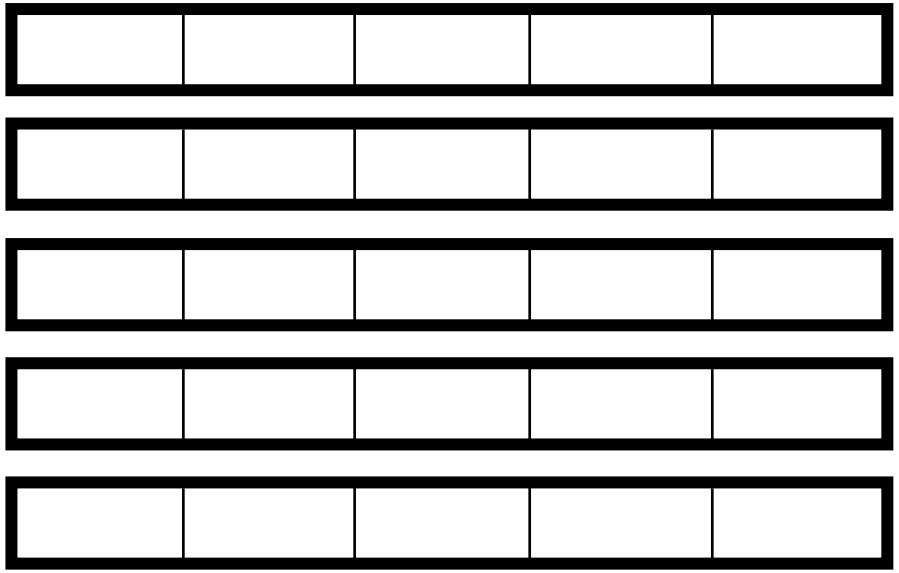

Columnar Layout

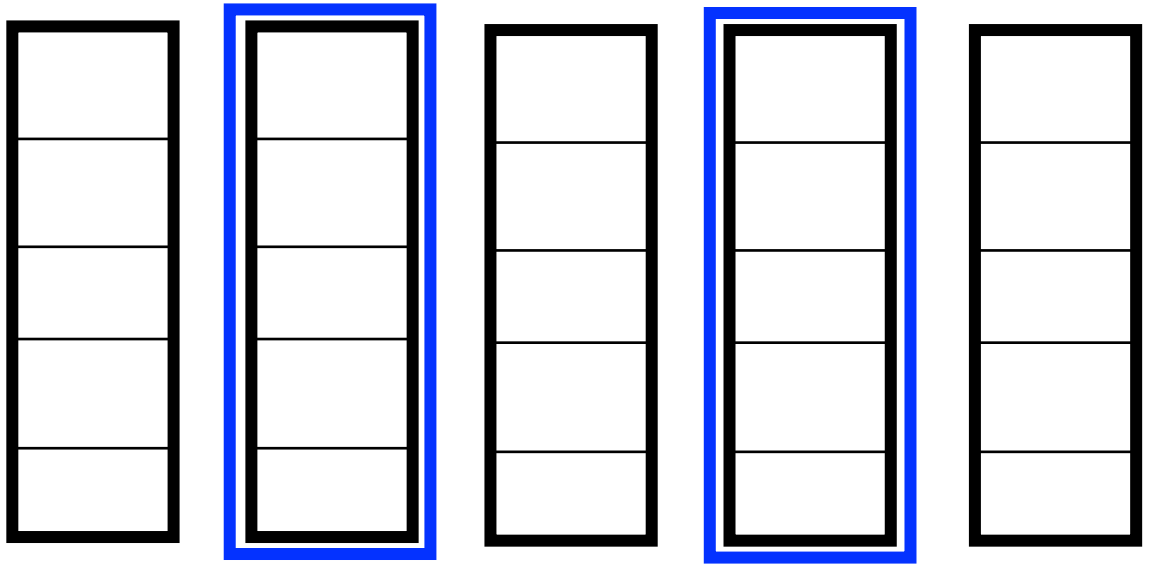

Updates (OLTP)

Analytics (OLAP)

UNIVERSITY OF WATERLOO

Storage Formats and HTAP

Row Layout

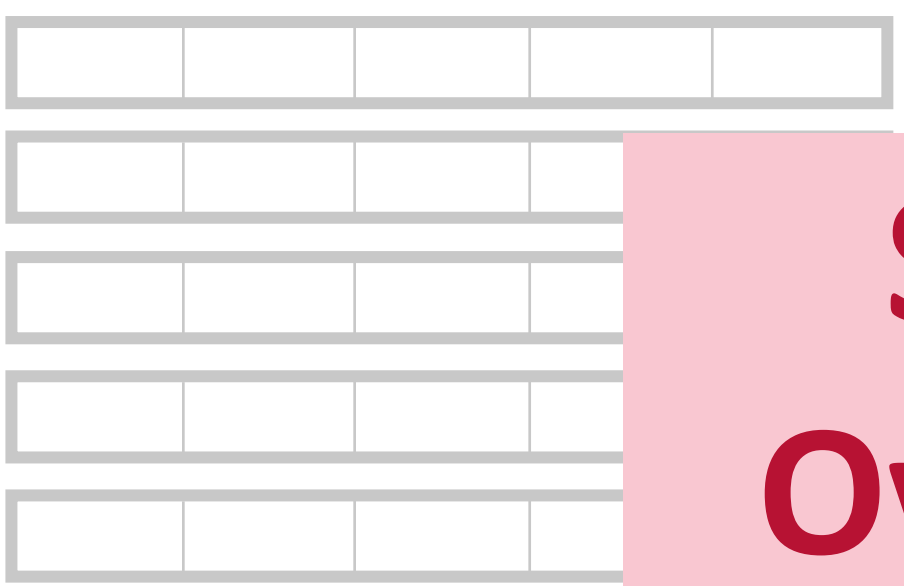

Columnar Layout

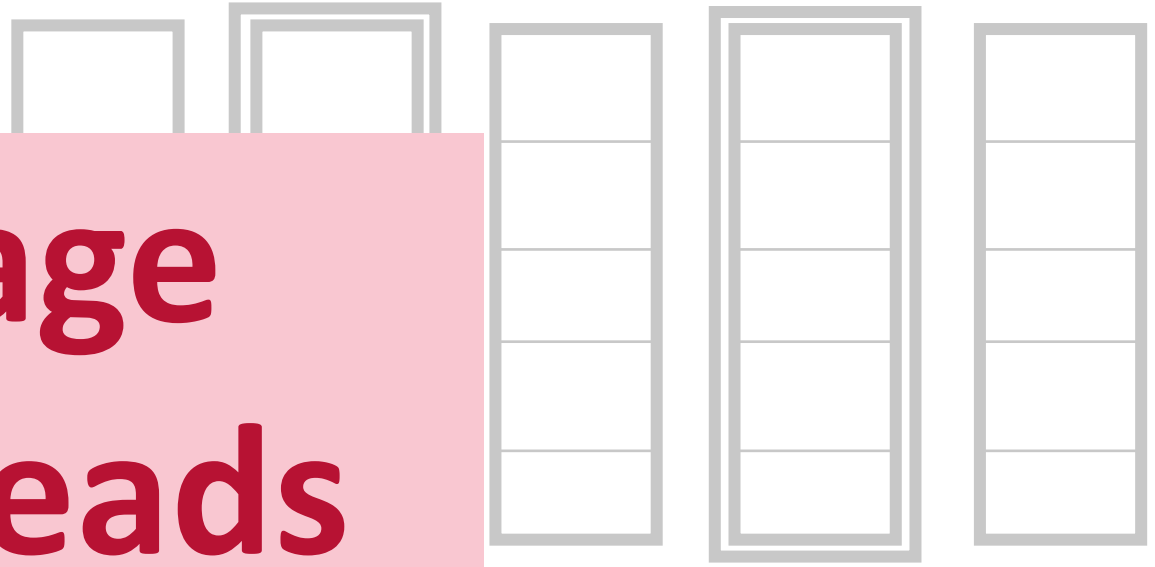

Storage Overheads

Updates (OLTP)

Analytics (OLAP)

Distributed Databases

How and where to store data?

Trade-offs dependent on workload

Distributed DBMSs must adapt

UNIVERSITY OF WATERLOO

Workloads Can Change

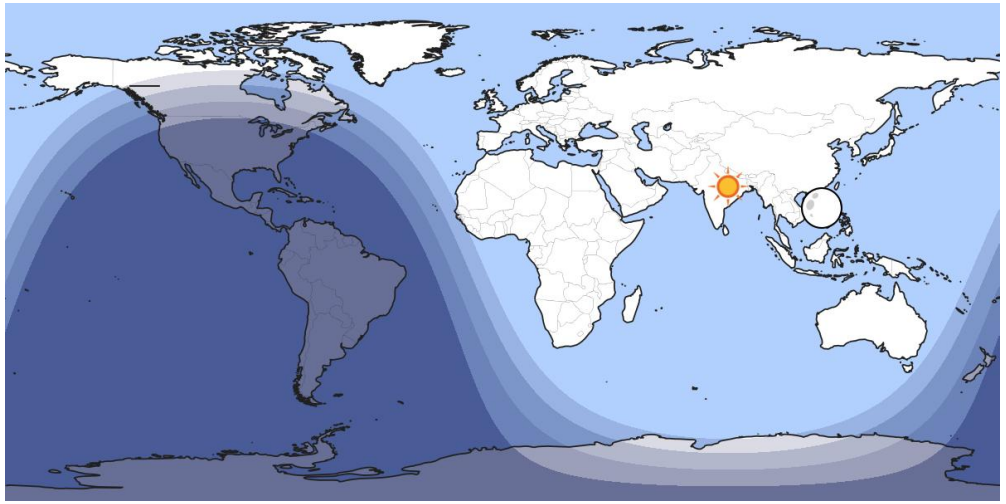

Distributed DBMSs must adapt

Thesis Statement

Automatic **adaptation** of **how & where** data stored

Using online workload information

Improves performance of distributed DBMSs

Thesis Contributions

Automatic **adaptation** of **how & where** data stored

DynaMast (ICDE'20)
Dynamic transfer data mastership to reduces overhead of coordination

MorphoSys (PVLDB'20)
Automatically select physical design: partitioning, & data placement

Proteus (SIGMOD'22) (PVLDB'22)
Adapt data storage (formats & tiers) for HTAP workloads

UNIVERSITY OF WATERLOO

Architecture

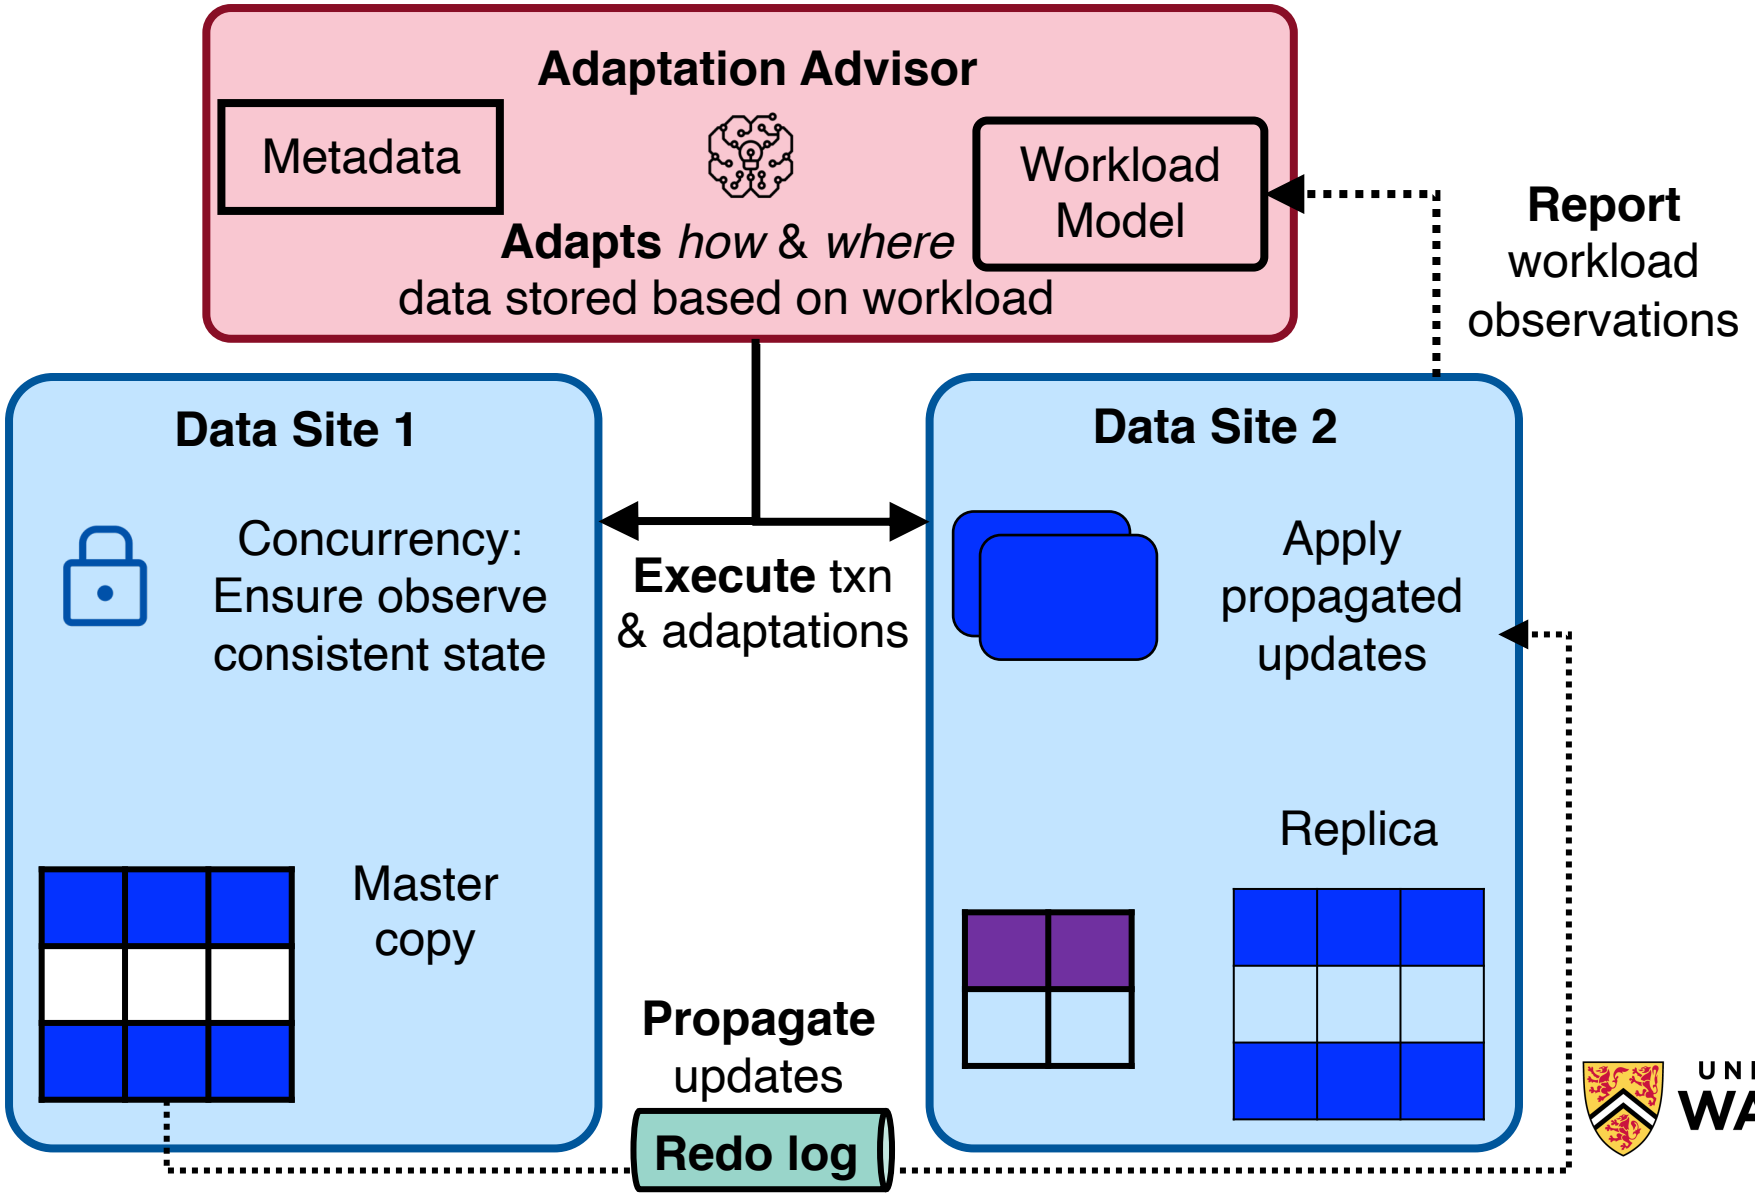

Thesis Contributions

Automatic **adaptation** of **how & where** data stored

DynaMast Dynamic transfer data mastership to reduces overhead of coordination

MorphoSys Automatically select physical design: partitioning, & data placement

Proteus Adapt data storage (formats & tiers) for HTAP workloads

UNIVERSITY OF WATERLOO

Dynamic Mastering

W[A, C]

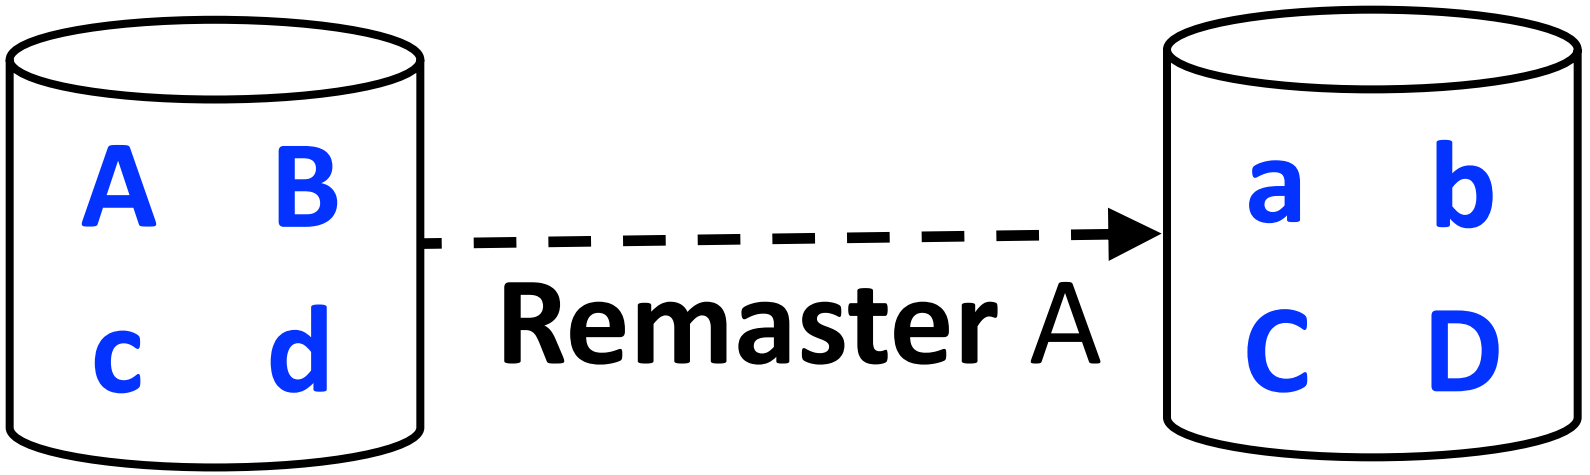

UNIVERSITY OF WATERLOO

Dynamic Mastering

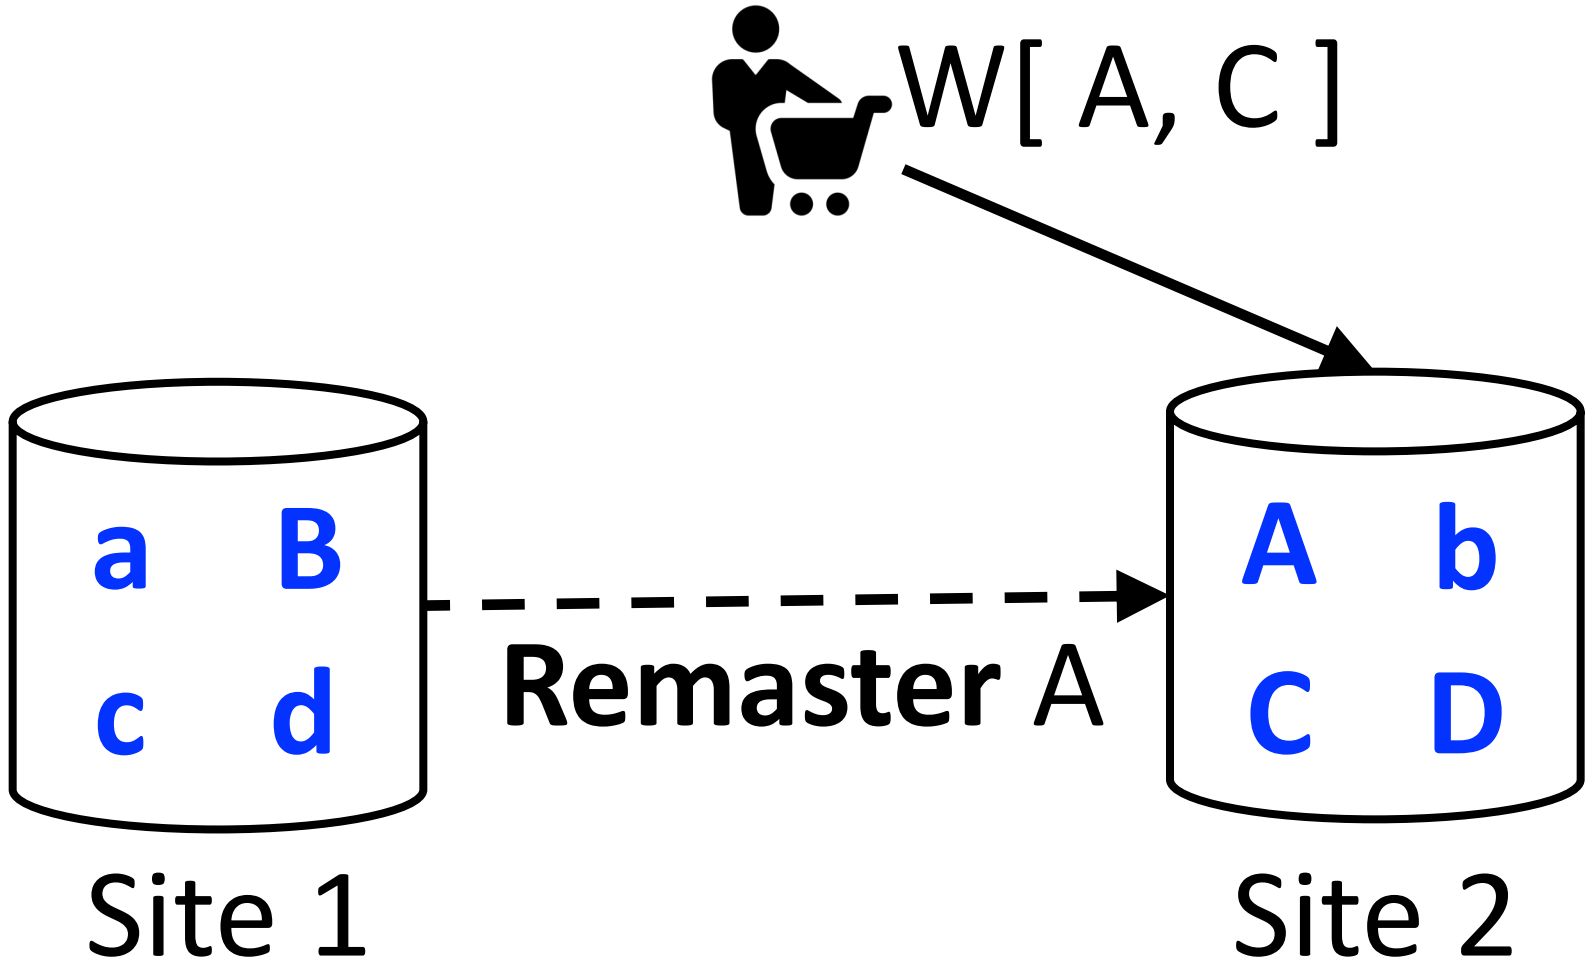

Dynamic Mastering

Outside transaction boundaries

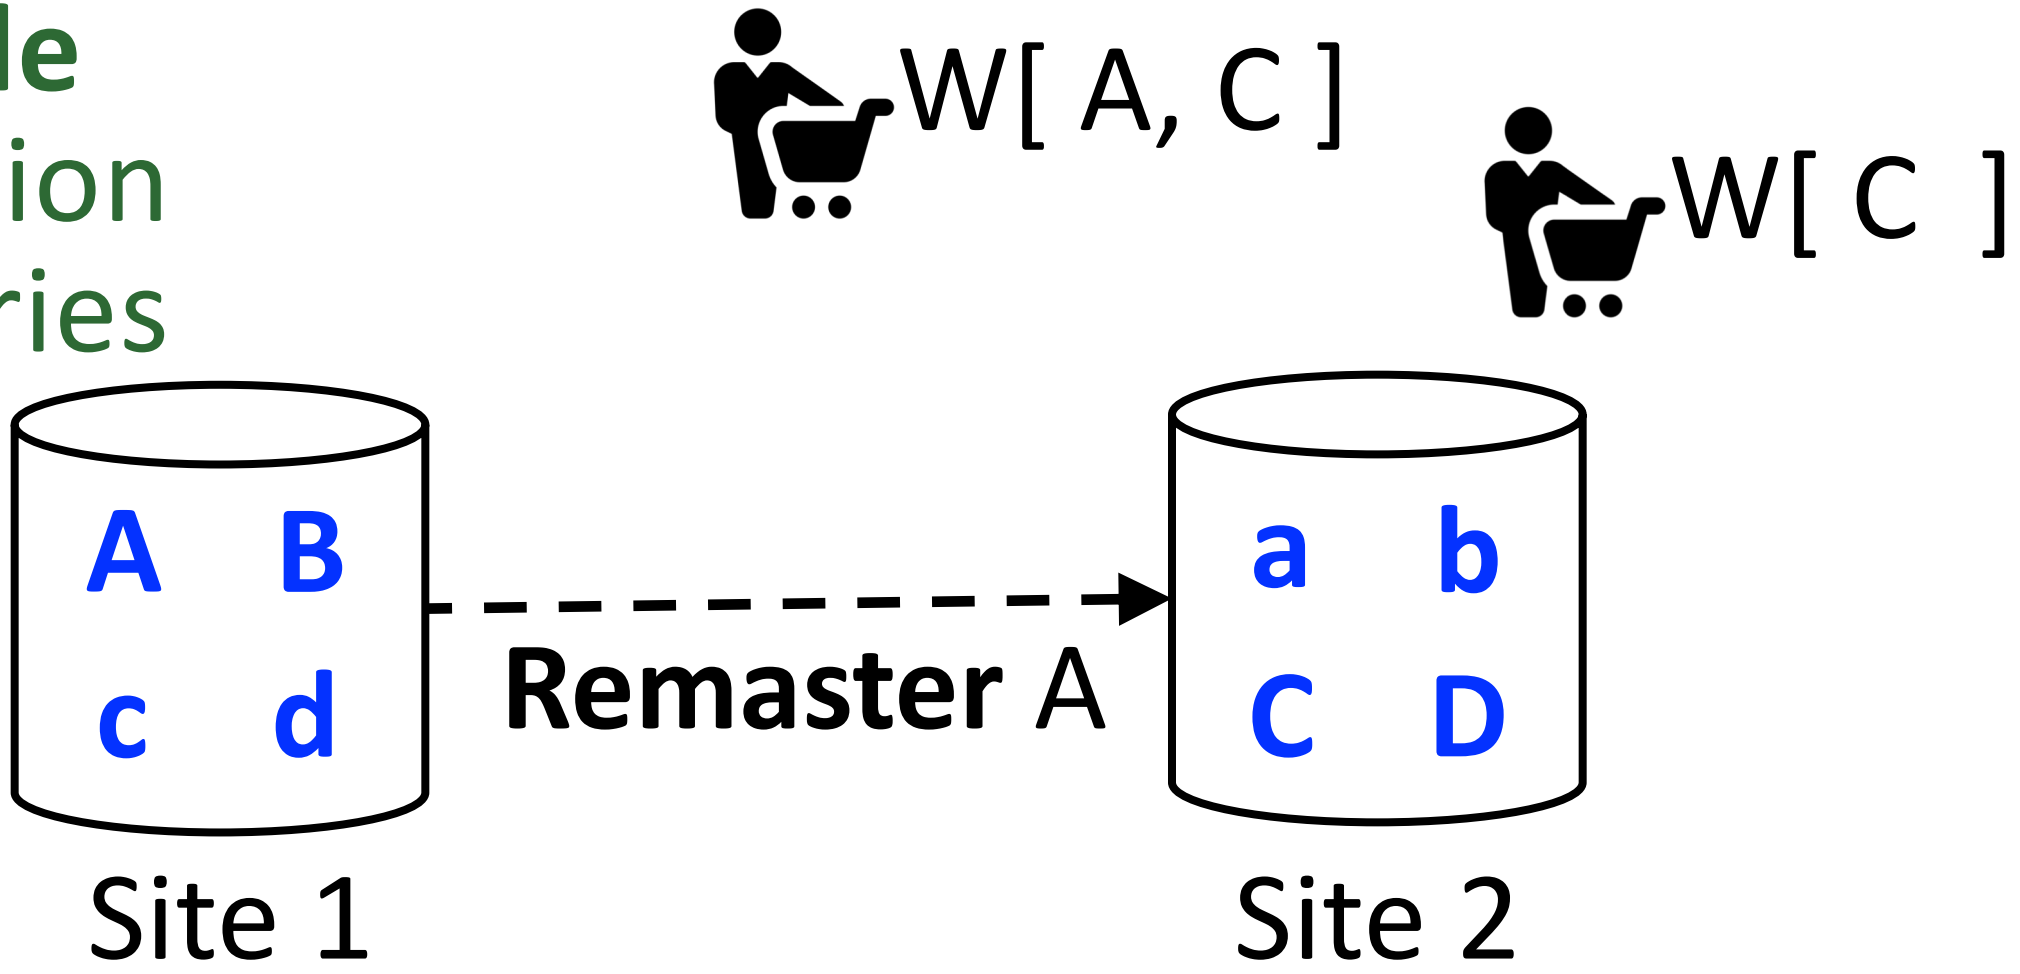

UNIVERSITY OF WATERLOO

Dynamic Mastering

W[B]

W[A, C]

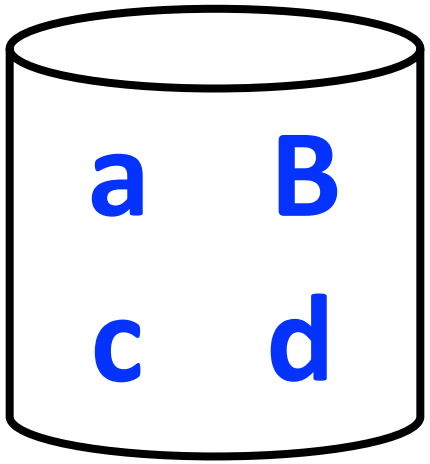

Site 1

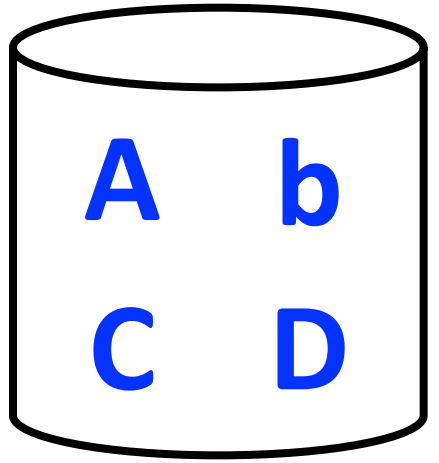

Site 2

UNIVERSITY OF WATERLOO

YCSB with Skew - Throughput

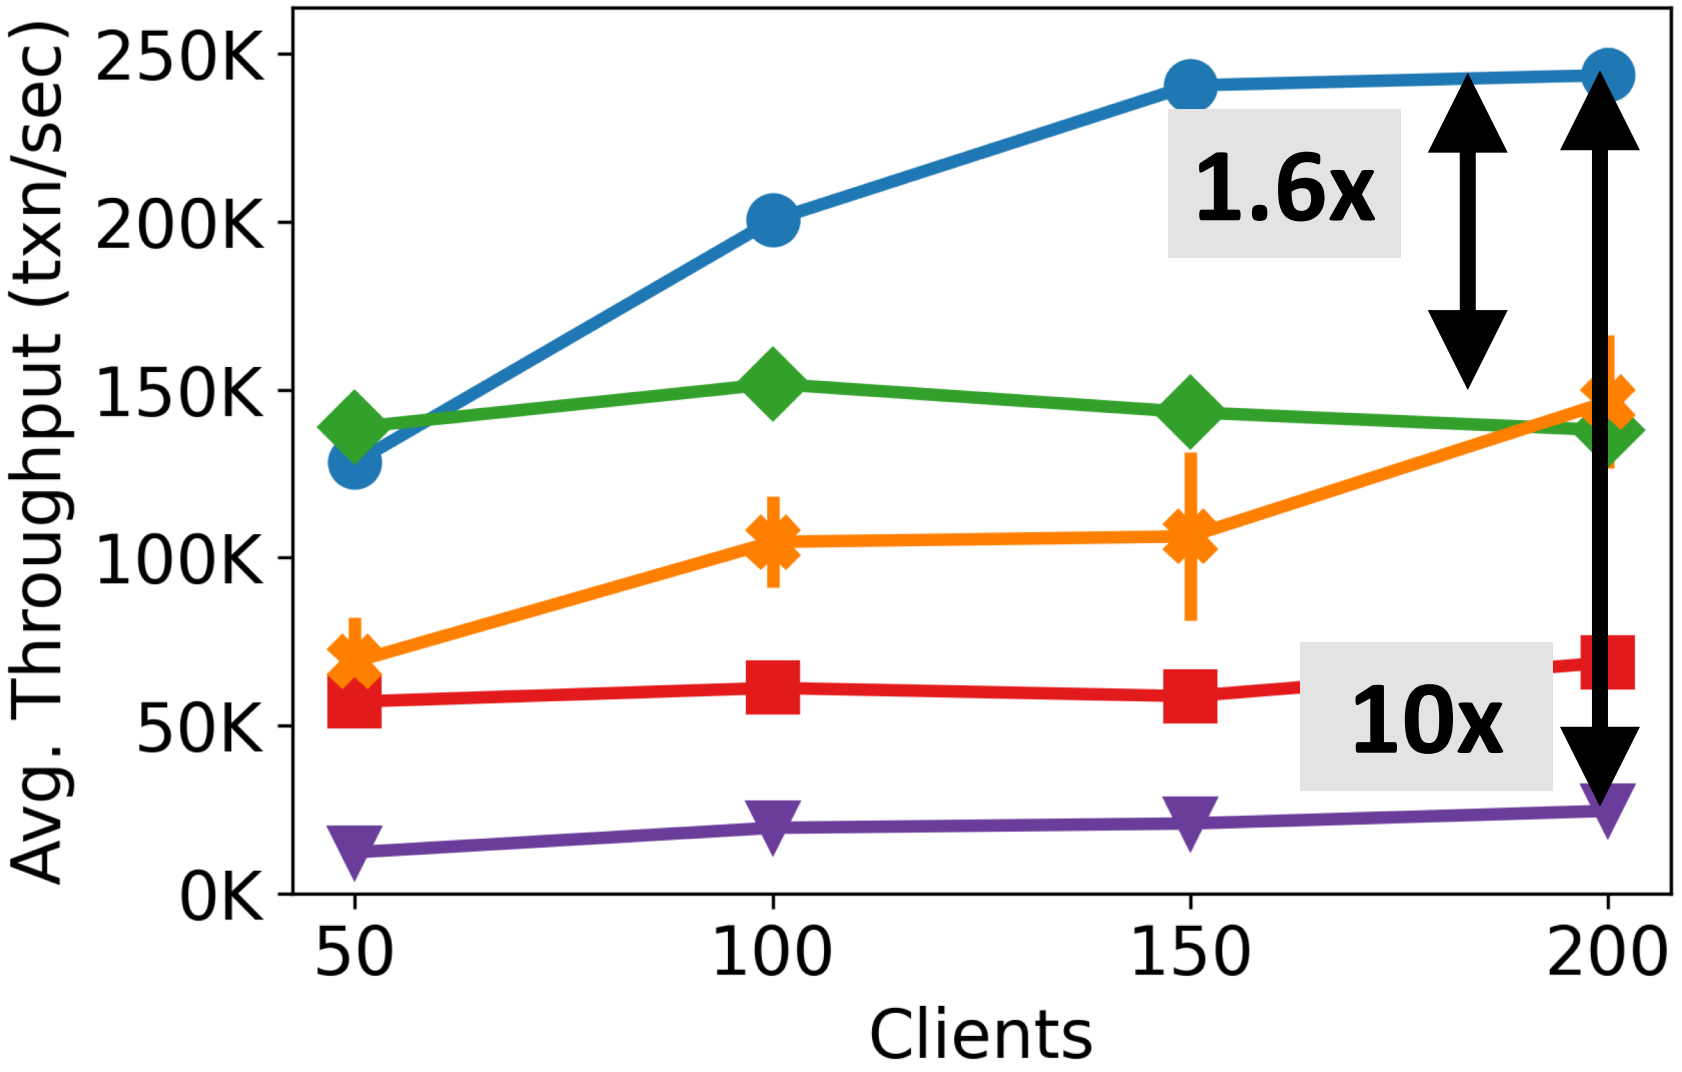

DynaMast

LEAP

Single-Master

Partition-Store

Multi-Master

UNIVERSITY OF WATERLOO

YCSB with Skew - Routing

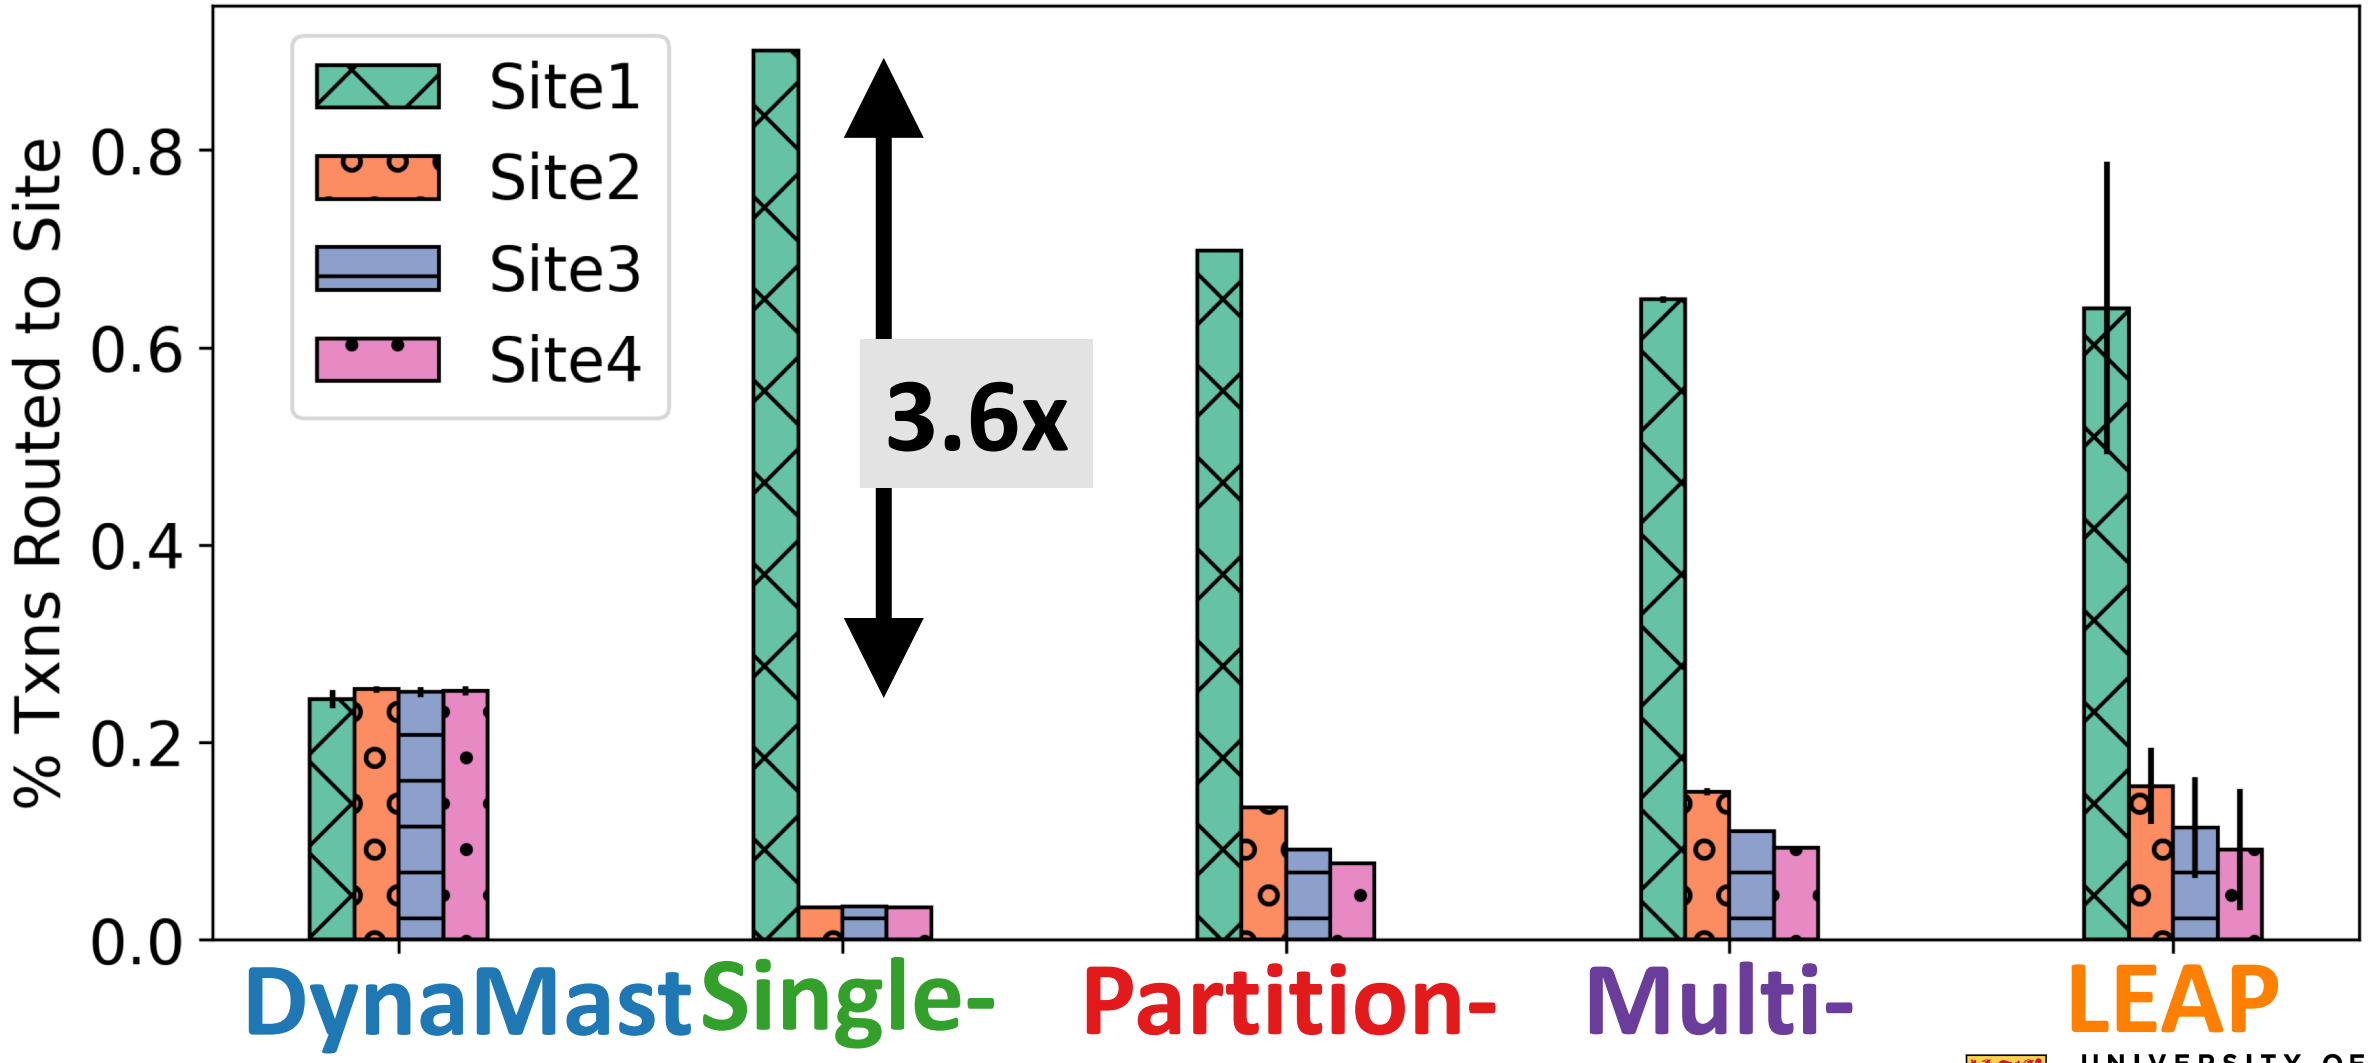

UNIVERSITY OF WATERLOO

Thesis Contributions

Automatic **adaptation** of **how & where** data stored

DynaMast Dynamic transfer data mastership to reduces overhead of coordination

MorphoSys Automatically select physical design: partitioning, & data placement

Proteus Adapt data storage (formats & tiers) for HTAP workloads

UNIVERSITY OF WATERLOO

Distributed DBMS Physical Design

For each **data item**

Where is the **master**?

What nodes **replicate** it?

How is it **grouped (partitioned)** with other data items?

MorphoSys Physical Design Change Operations

Add or **remove replica** of a **partition**

Remaster a **partition**

Split or **merge** **partition(s)**

Making design decisions

Learned cost model quantifies
design change effects

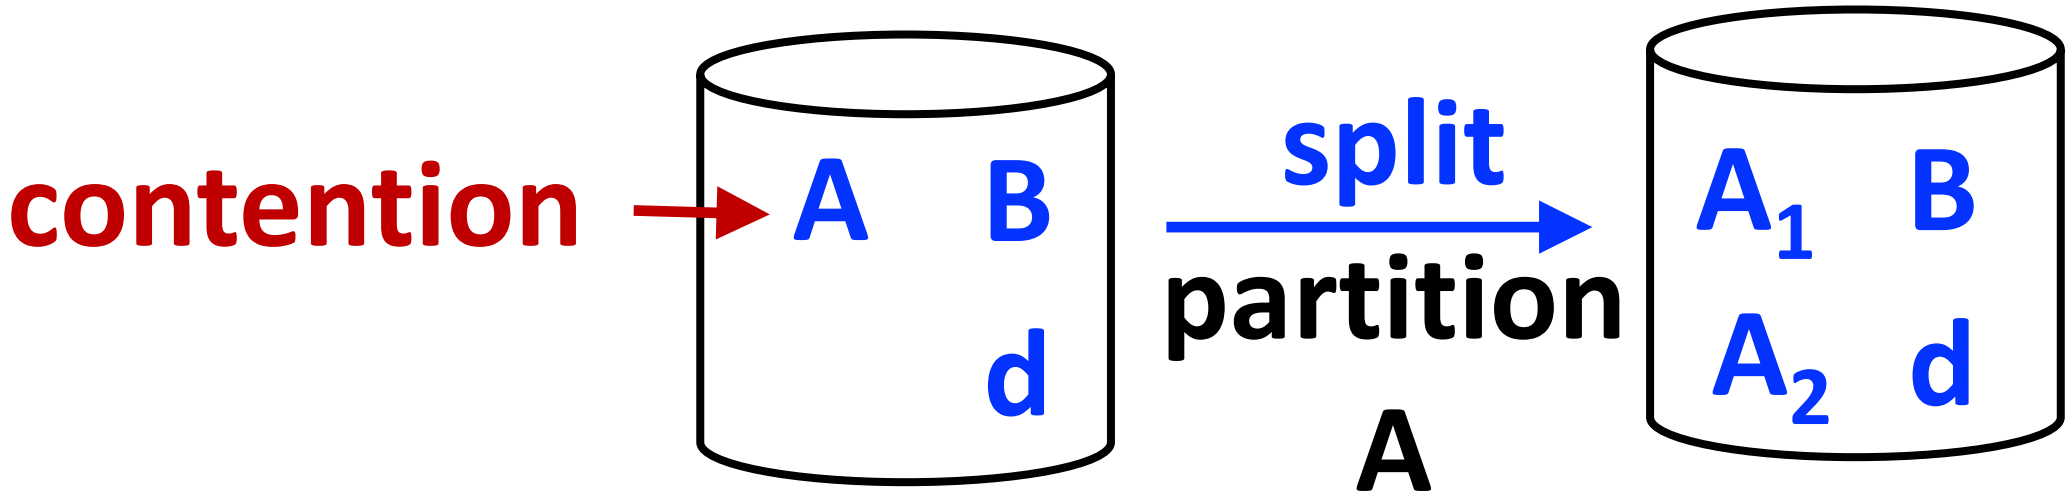

Design change cost < **Expected Benefit**

UNIVERSITY OF WATERLOO

Physical design cost model

Design change cost < **Expected benefit**

Decompose operators into **key costs**

Predict benefit based on workload history

Skewed YCSB - Throughput

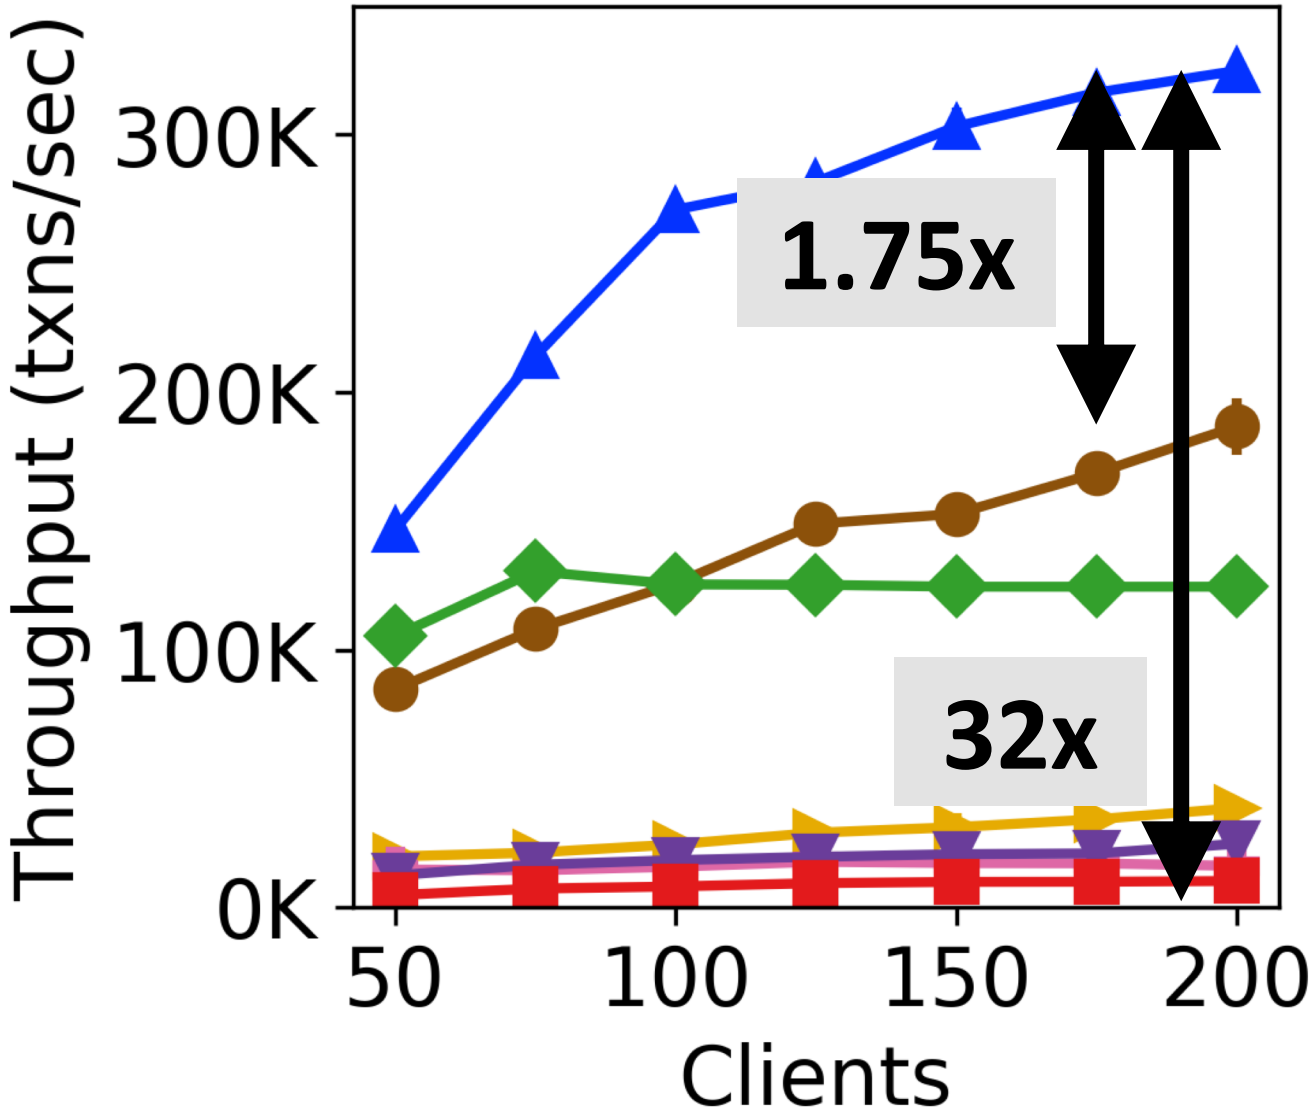

MorphoSys

DynaMast

Single-Master

Clay Multi-Master

VoltDB ADR

UNIVERSITY OF WATERLOO

Number of Replicas

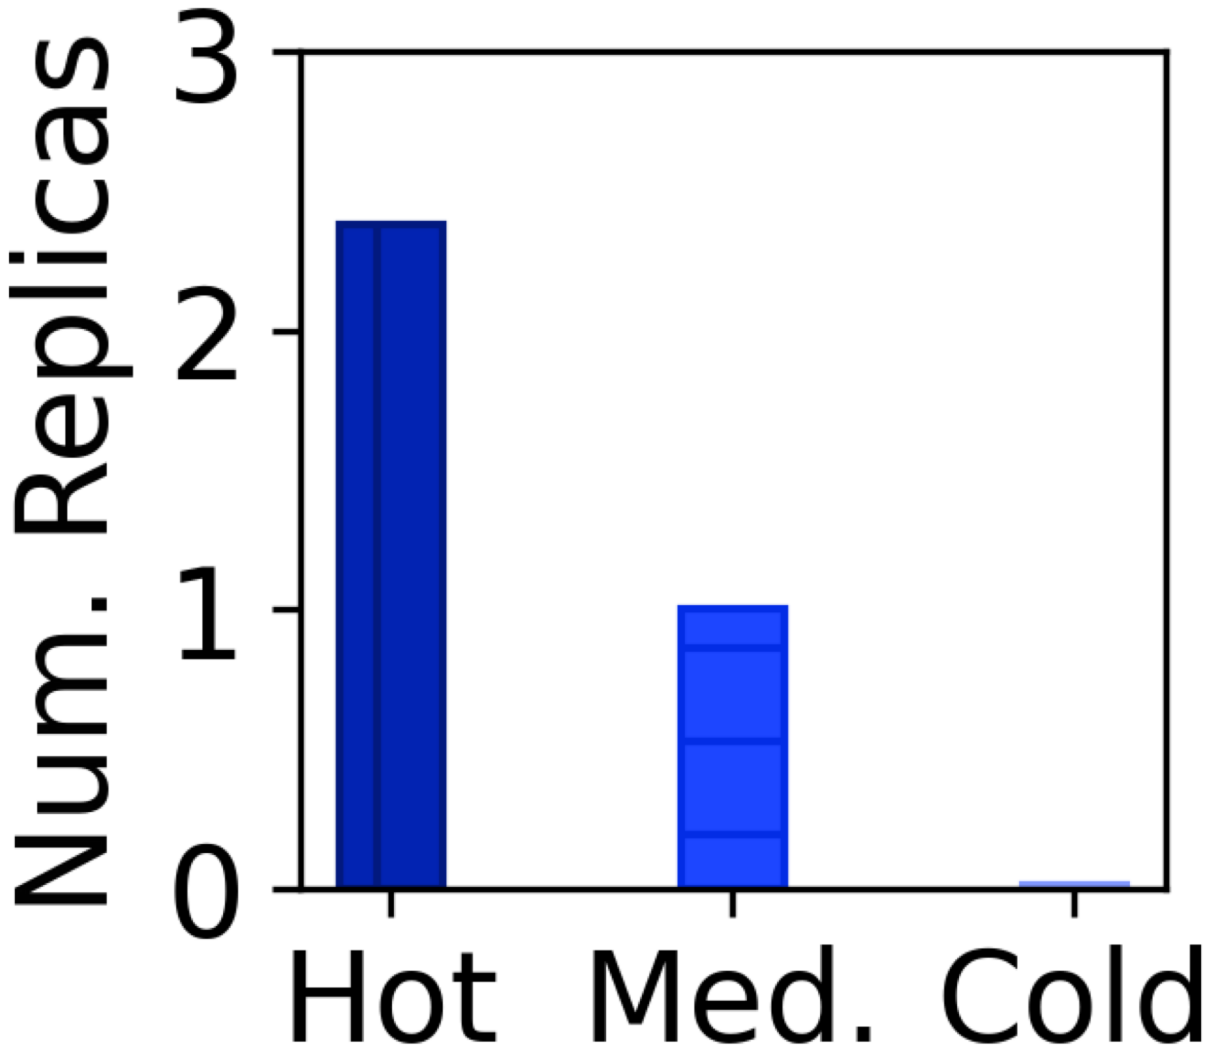

UNIVERSITY OF WATERLOO

Partition Sizes

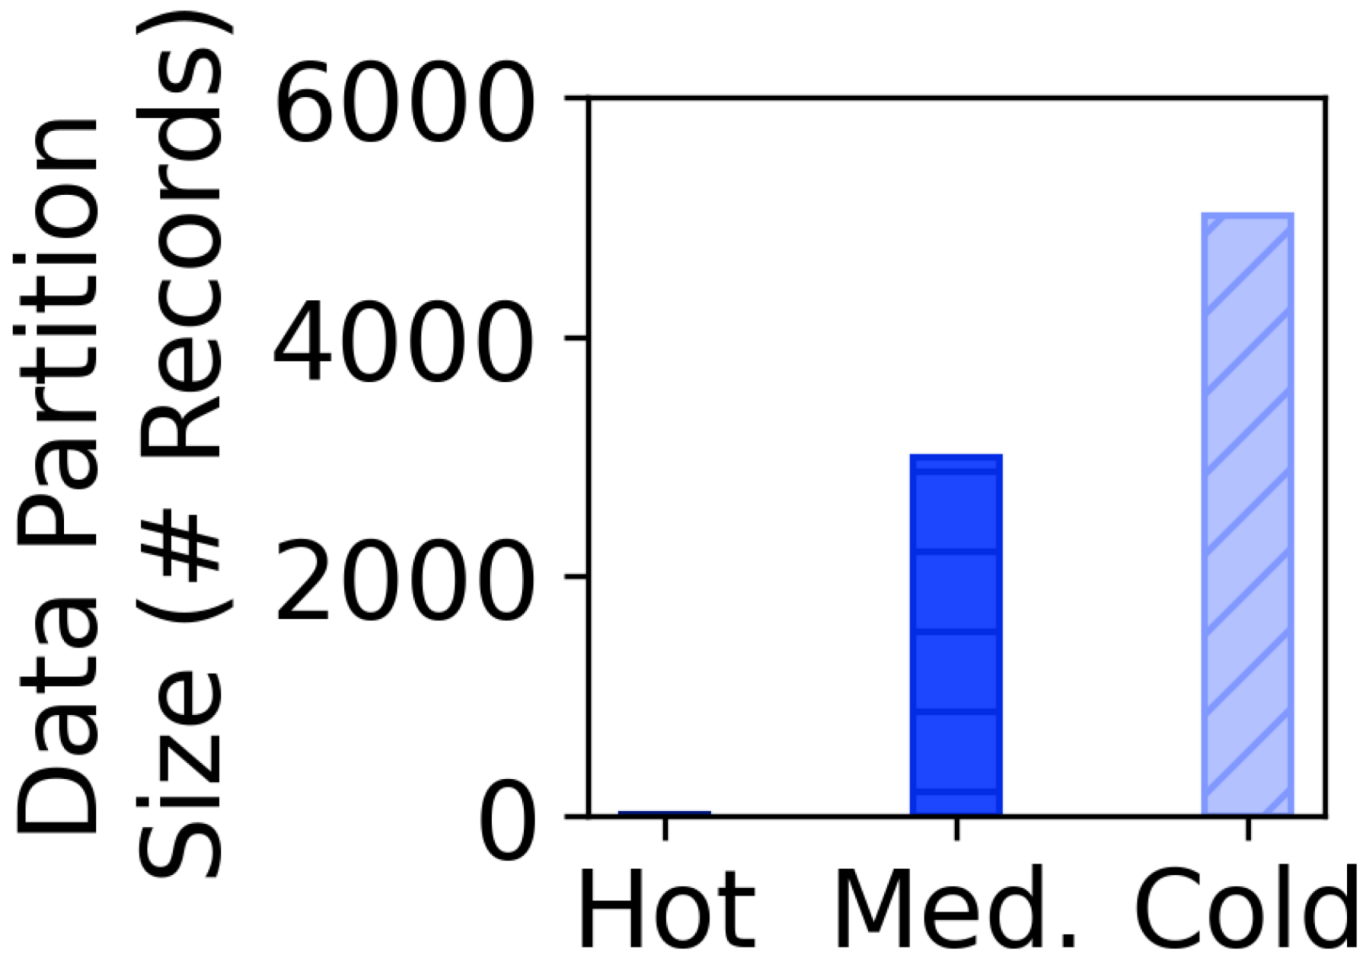

UNIVERSITY OF WATERLOO

Thesis Contributions

Automatic **adaptation** of **how & where** data stored

DynaMast Dynamic transfer data mastership to reduces overhead of coordination

MorphoSys Automatically select physical design: partitioning, & data placement

Proteus Adapt data storage (formats & tiers) for HTAP workloads

UNIVERSITY OF WATERLOO

Performance Trade-Off

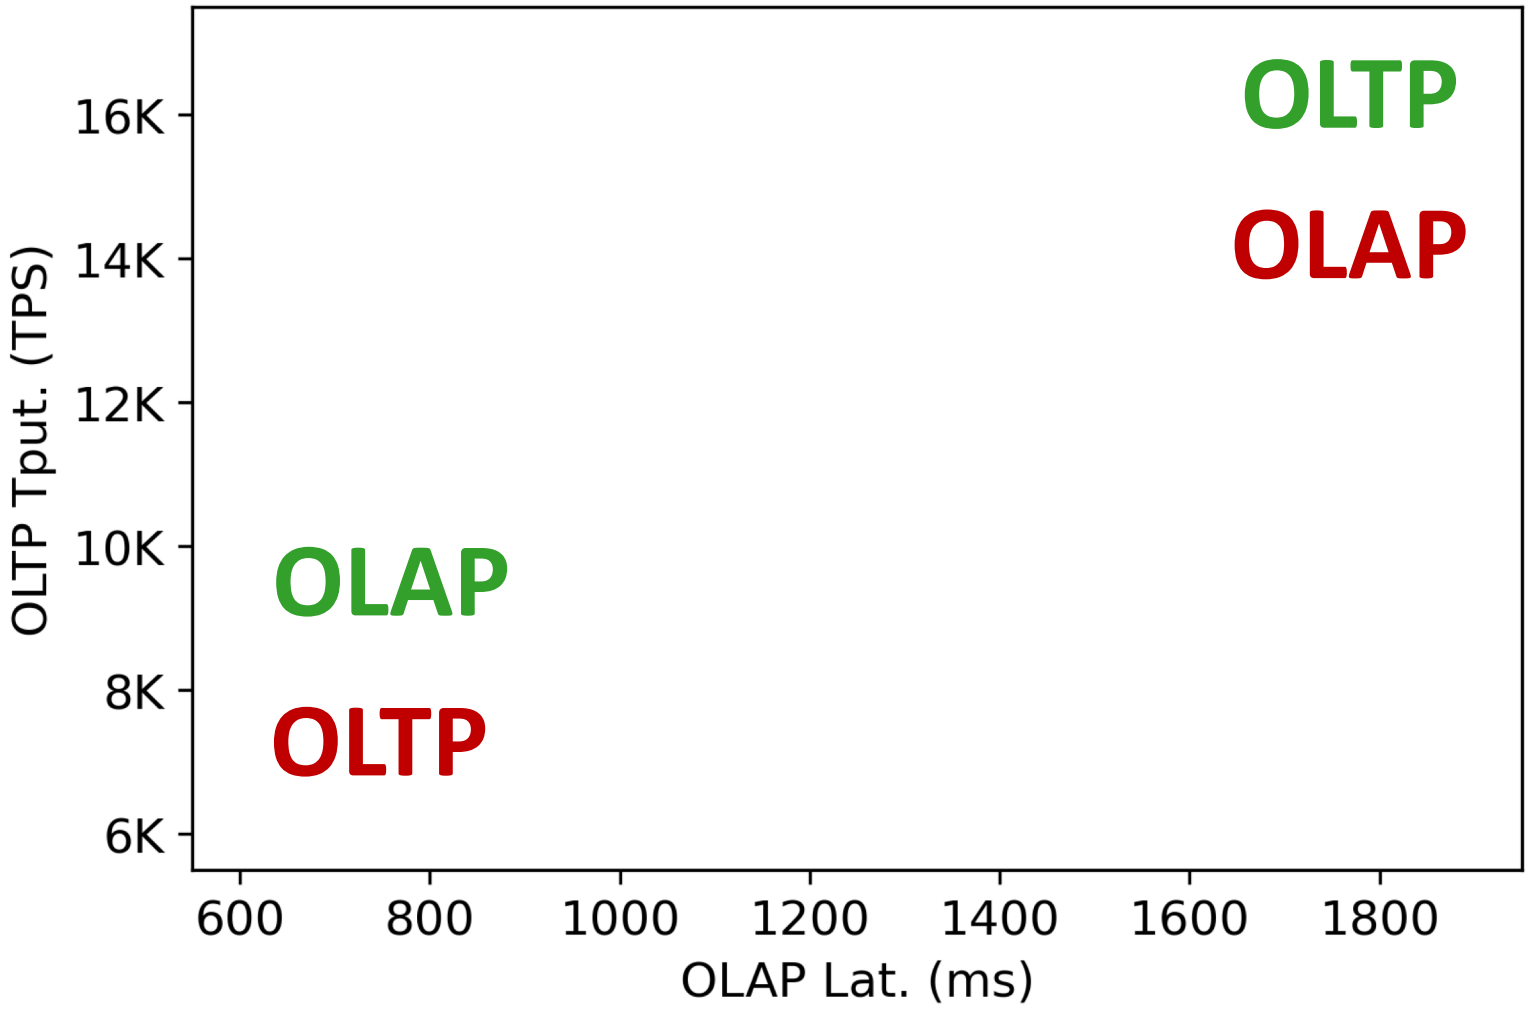

UNIVERSITY OF WATERLOO

Performance Trade-Off

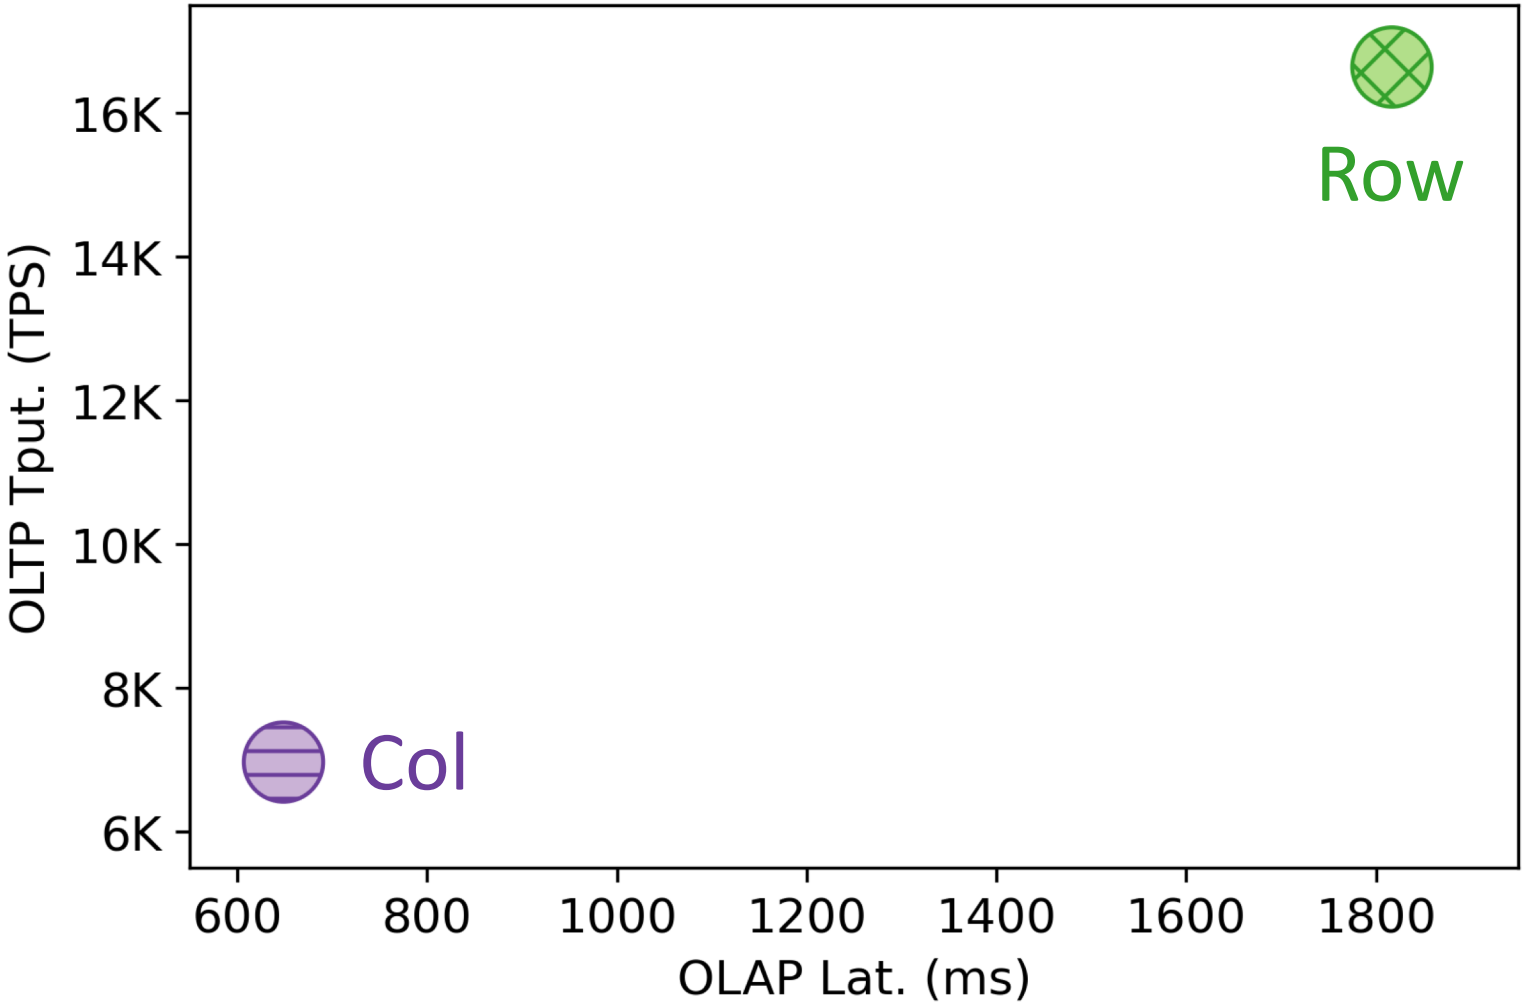

UNIVERSITY OF WATERLOO

Proteus Decisions

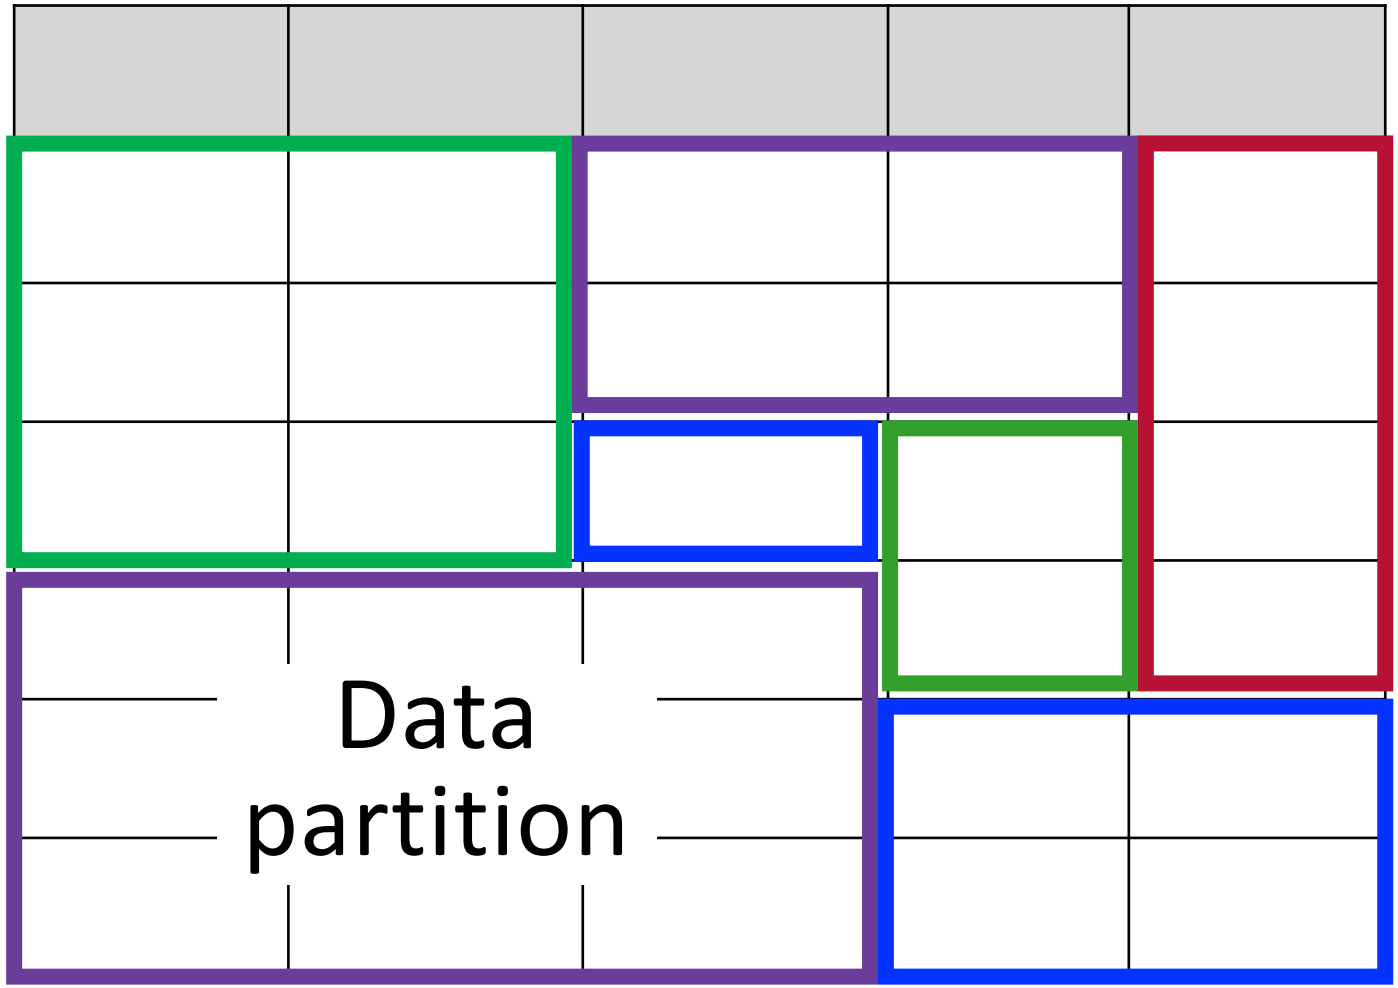

Storage layout

Master/replica(s)

Txn execution

How to **partition?**

When & **what** to change

UNIVERSITY OF WATERLOO

Transactions in Proteus

Breakdown transaction into **physical** operators

SELECT book, **SUM**(qnt) **GROUP BY** book

Row layout

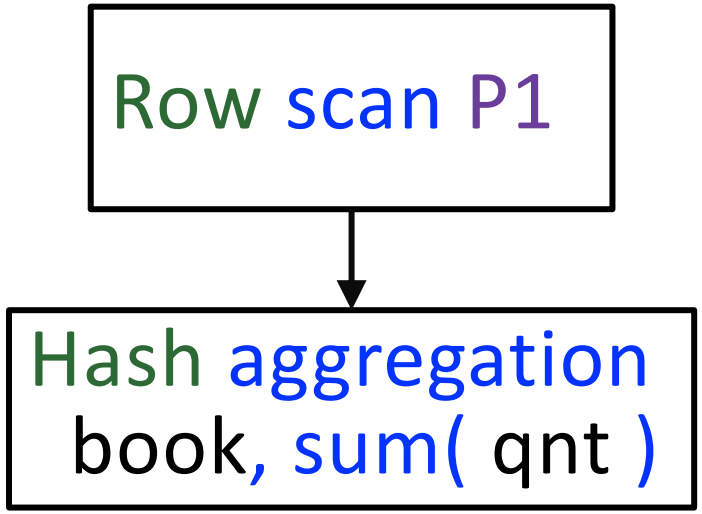

Logical Plan

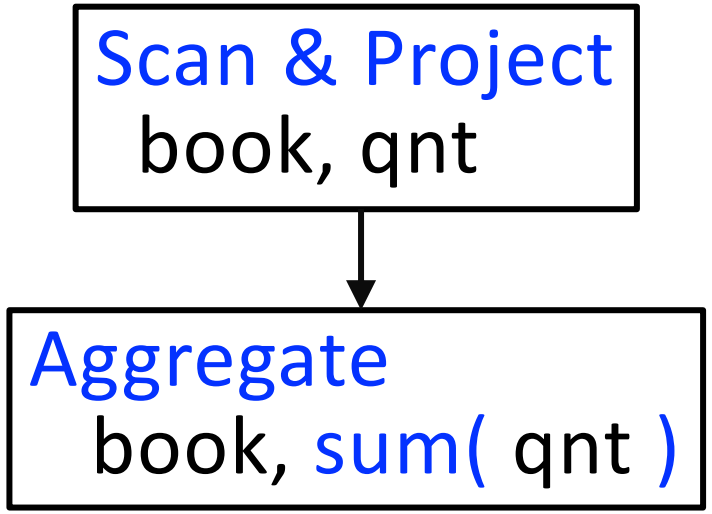

Sorted column layout

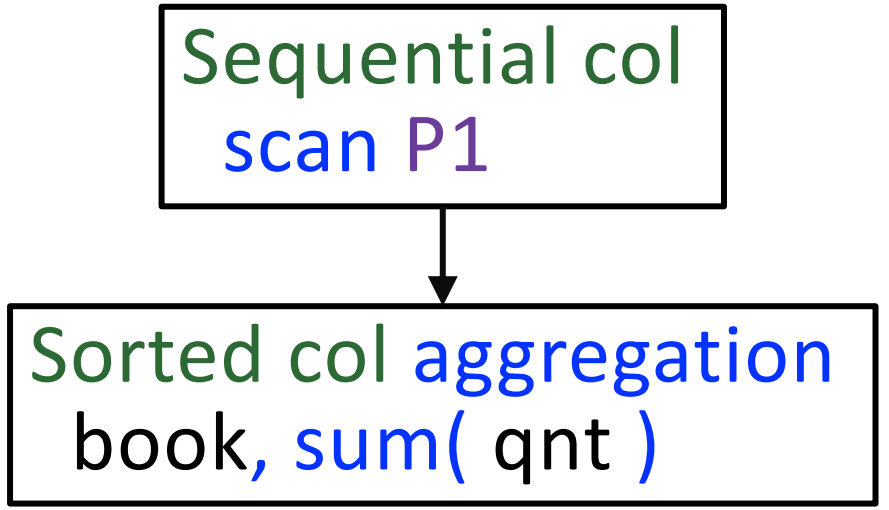

UNIVERSITY OF WATERLOO

Storage-Aware Operators

Per layout implementation of operators

Operate **directly** over columnar, sorted, compressed data

Predict physical operator latency

Cardinality

Data Width

Est Selectivity

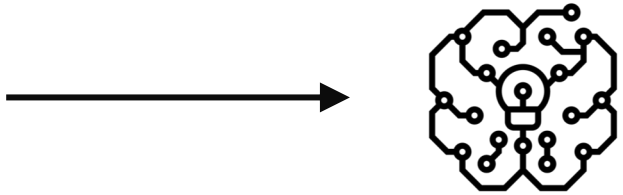

Predicted Latency

Seq col scan

UNIVERSITY OF WATERLOO

Storage-Aware Operators

Per layout implementation of operators

Operate **directly** over columnar, sorted, compressed data

Predict physical operator latency

Cardinality

Data Width

Est Selectivity

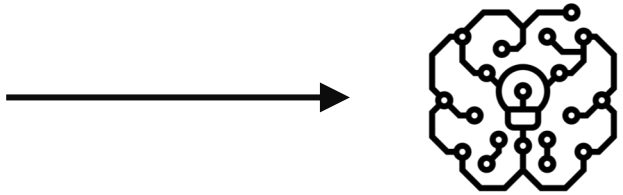

Predicted Latency

Row scan

UNIVERSITY OF WATERLOO

Likelihood of a Transaction

Data **accesses** to **storage**
often follow **predictable** pattern

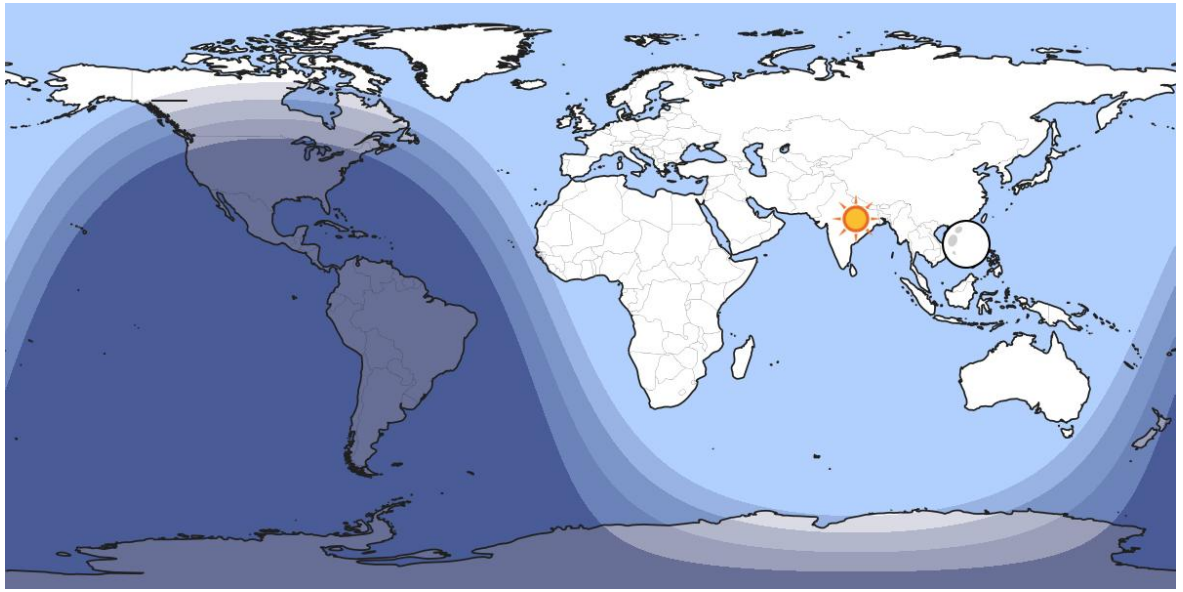

Likelihood of a Transaction

Data **accesses** to **storage**
often follow **predictable** pattern

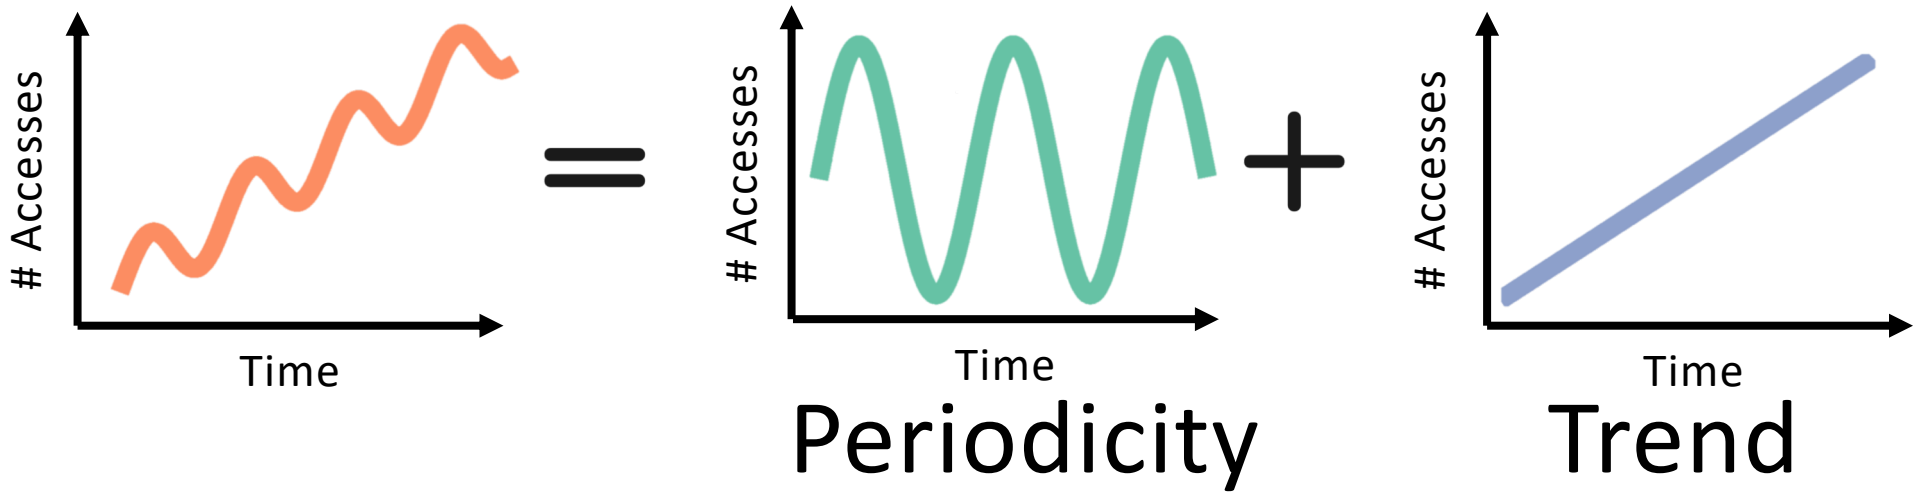

CH BenCHmark

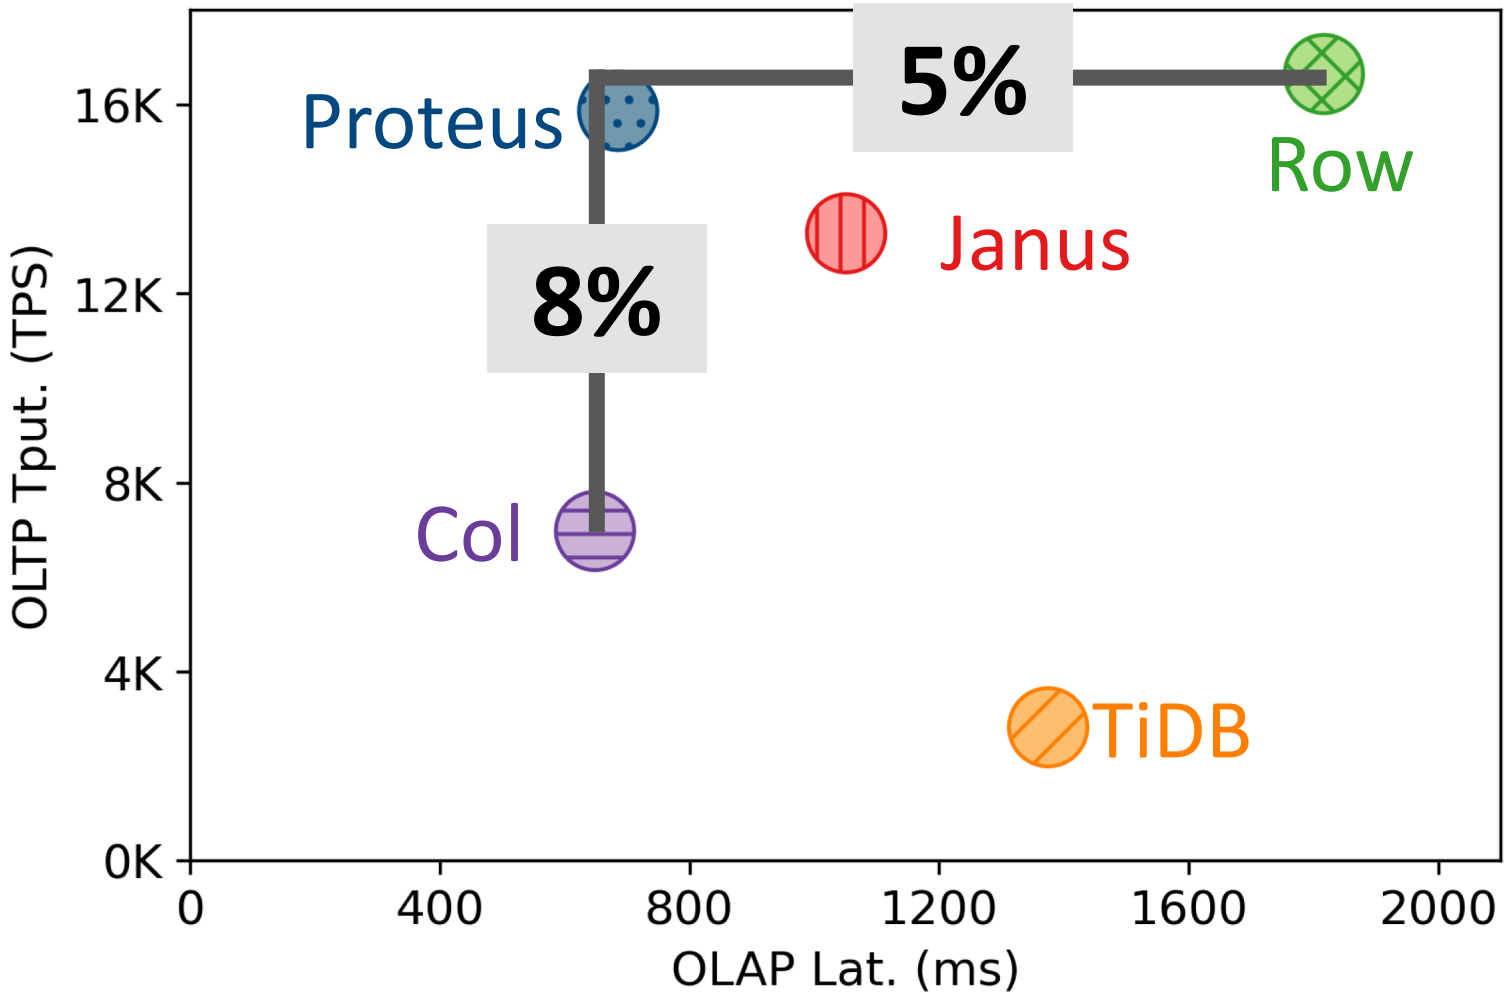

UNIVERSITY OF WATERLOO

Distributed DBMSs are widely used

Distributed DBMSs must adapt

Adaptation of **how** & **where** data stored **improves performance**

DynaMast **MorphoSys** **Proteus**

UNIVERSITY OF WATERLOO

Extra Slides