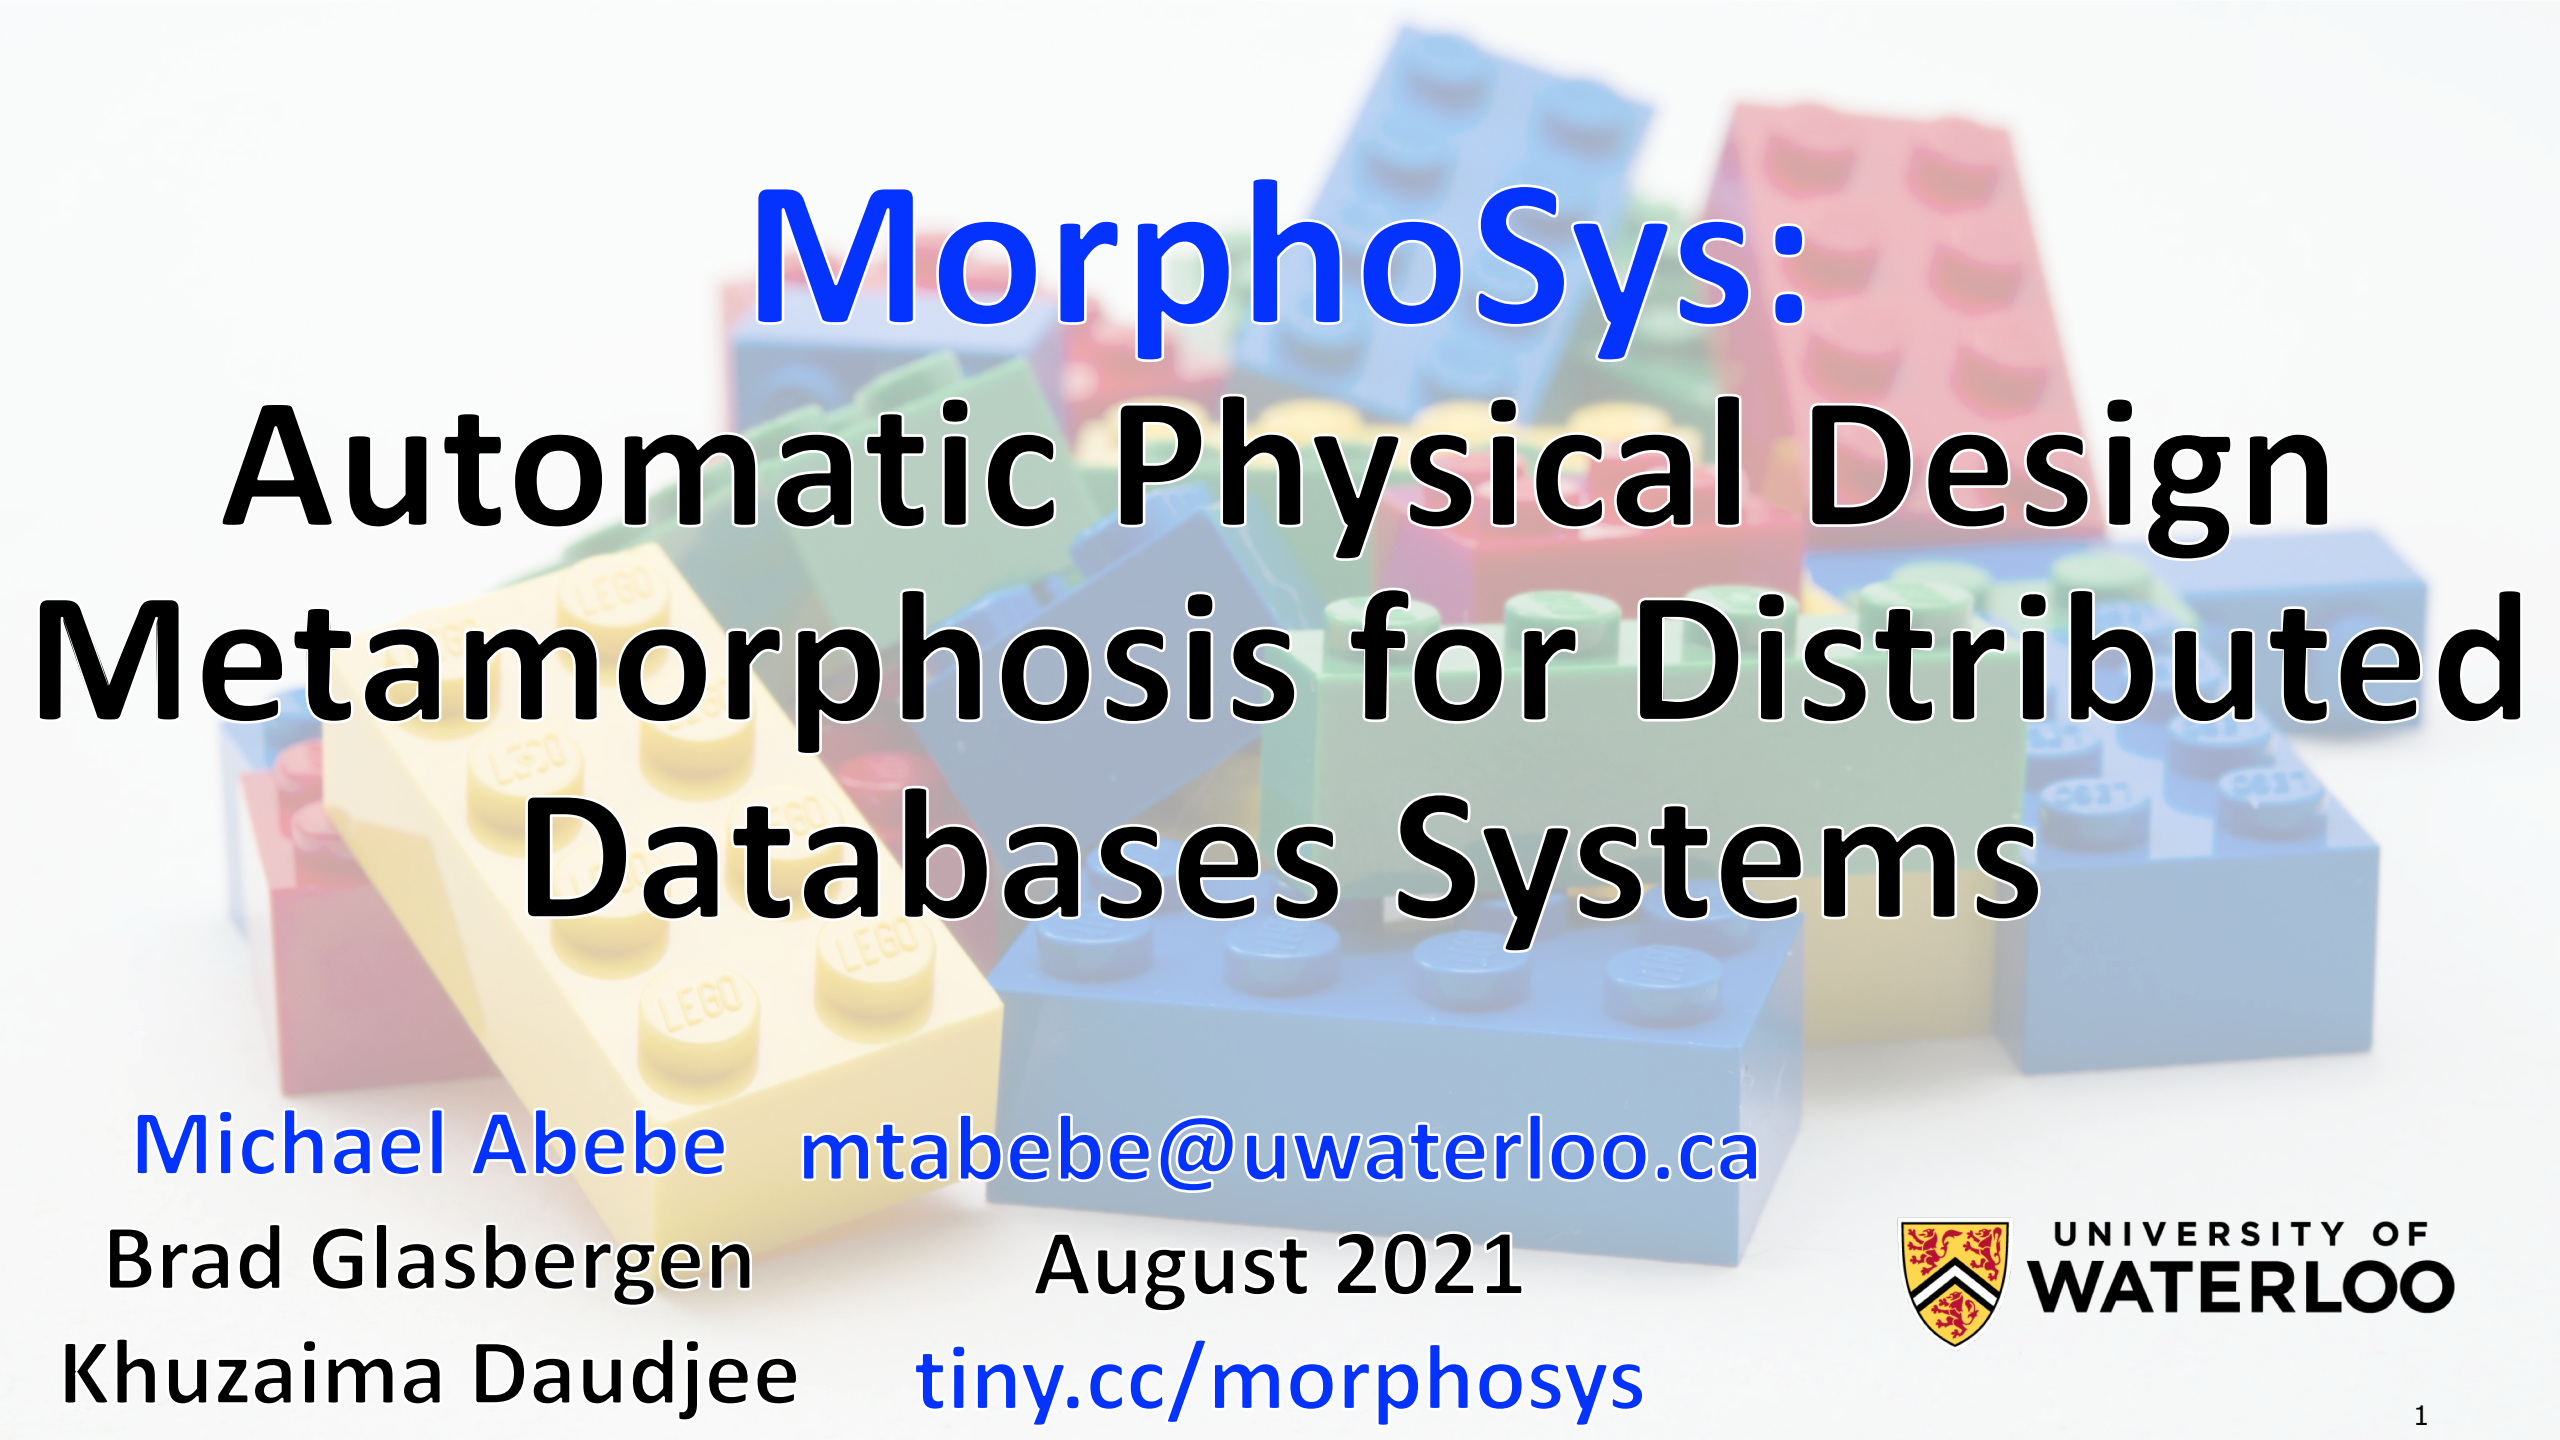

MorphoSys:
Automatic Physical Design Metamorphosis for Distributed Databases Systems

Michael Abebe mtabebe@uwaterloo.ca
Brad Glasbergen
Khuzaima Daudjee

August 2021

tiny.cc/morphosys

UNIVERSITY OF WATERLOO

Distributed Databases

How to distribute data?

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

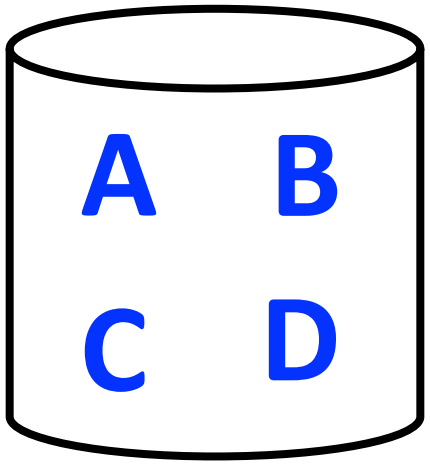

Master

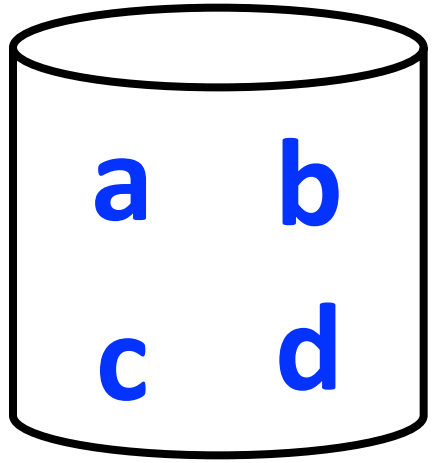

Replica

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

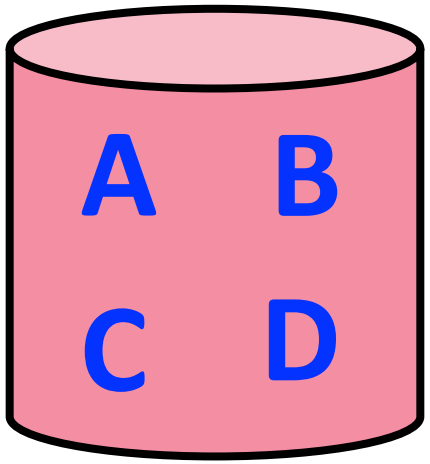

Master

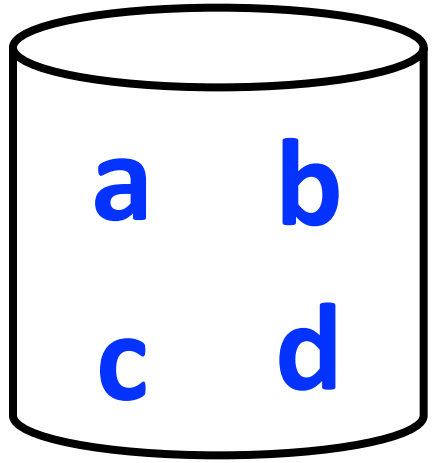

Replica

UNIVERSITY OF WATERLOO

Database Replication

Writers

Readers

Performance bottleneck

Master

Replica

UNIVERSITY OF WATERLOO

Partitioned Databases

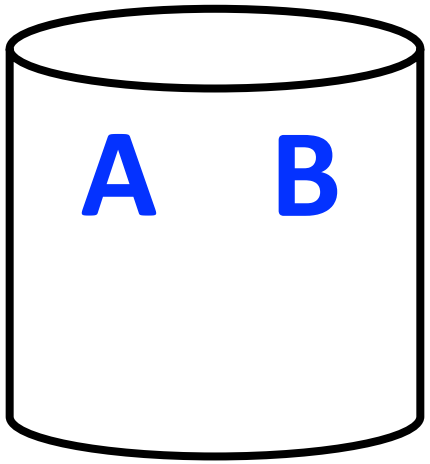

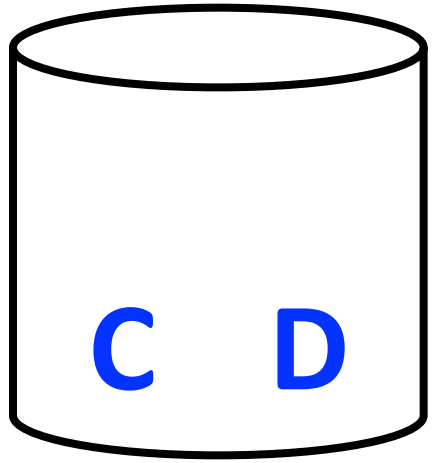

UNIVERSITY OF WATERLOO

Partitioned Databases

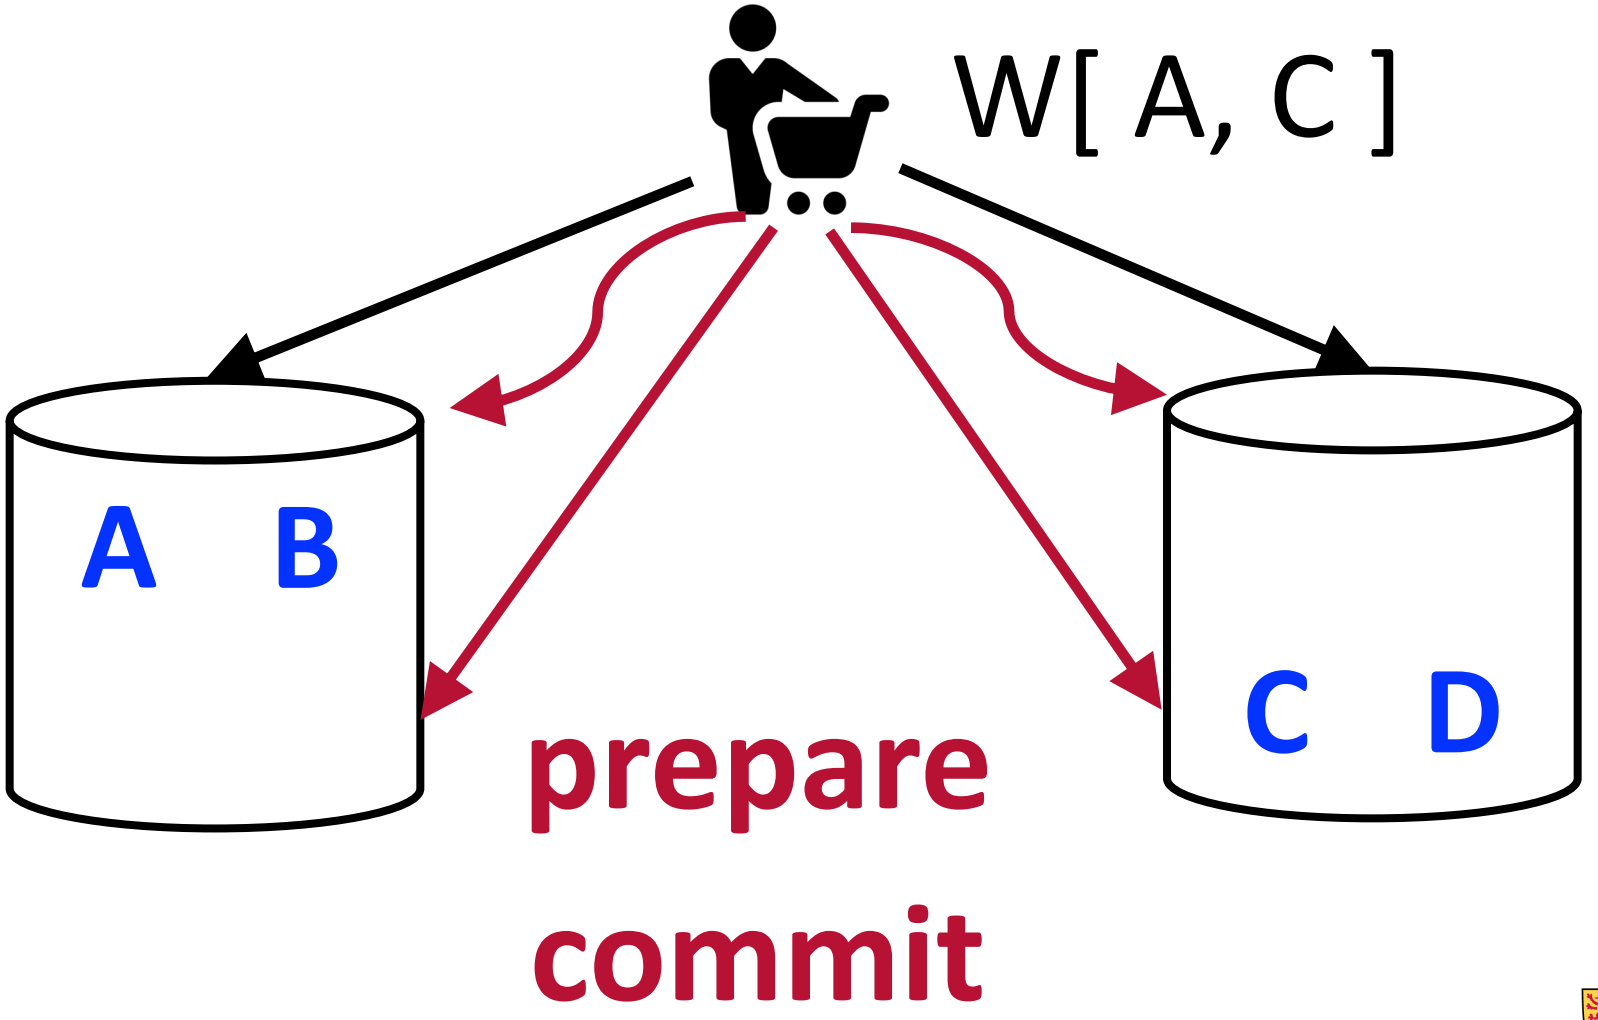

Partitioned Databases

Expensive Coordination

commit

Distributed Databases

How to distribute data?

Distributed Database Physical Design

UNIVERSITY OF WATERLOO

Distributed Database Physical Design

For each **data item**

Where is the **master**?

What nodes **replicate** it?

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

UNIVERSITY OF WATERLOO

Physical designs

Any combination of master data placement, replication, & data partitioning

Physical designs

Any combination of master data placement, replication, & data partitioning

Which physical design?

Traditionally: **offline** workload knowledge

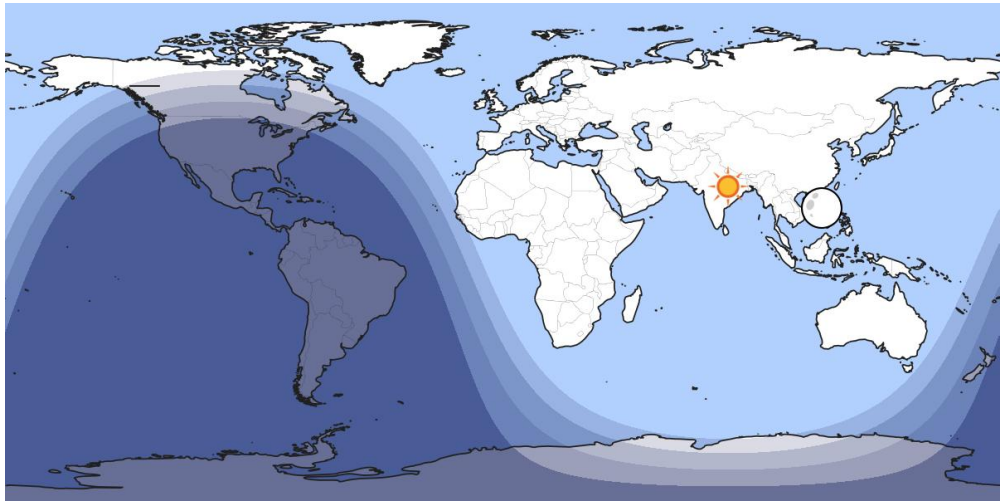

Physical design should change with workload

UNIVERSITY OF WATERLOO

MorphoSys

Automatically chooses a physical design

Automatically adapts the physical design

Aim: improve database system performance

UNIVERSITY OF WATERLOO

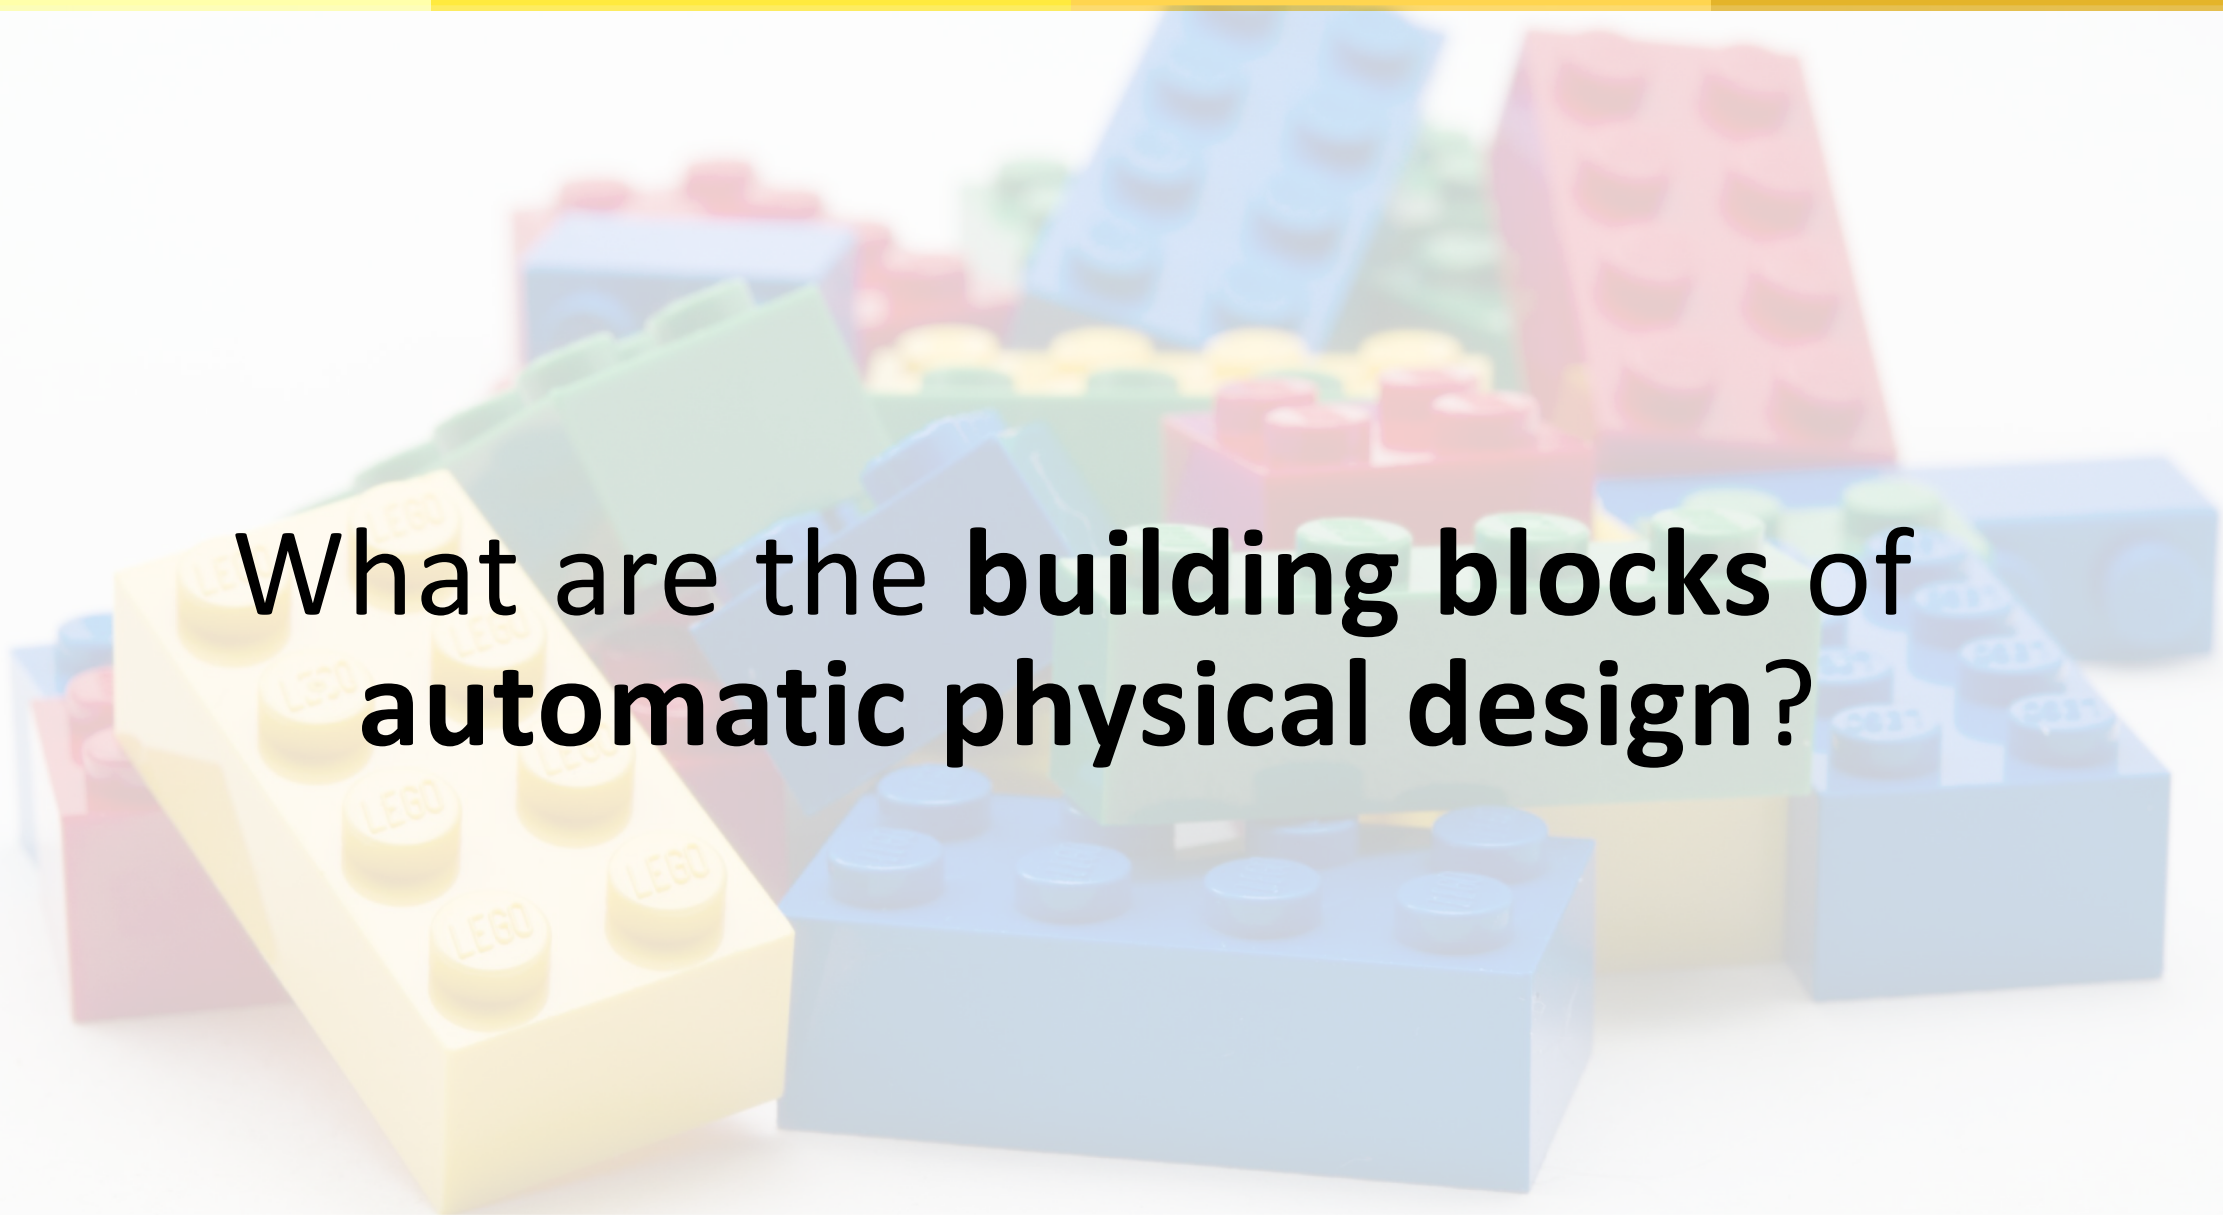

What are the **building blocks** of **automatic physical design**?

UNIVERSITY OF WATERLOO

Dynamic Replication

R[A]

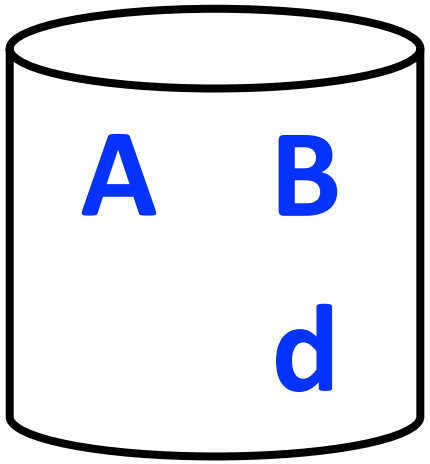

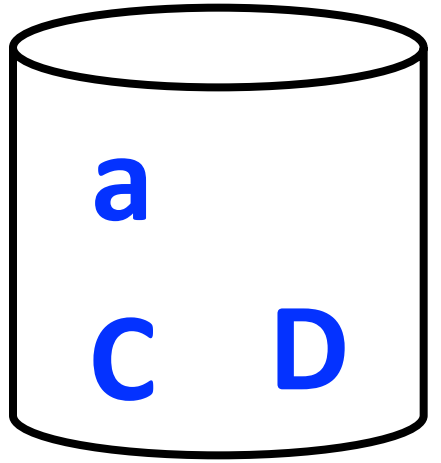

add replica A

Dynamic Replication

R[A]

R[A]

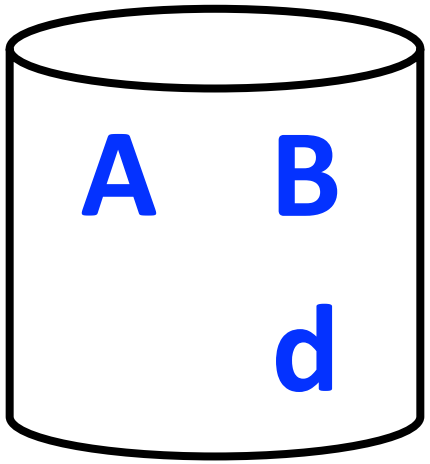

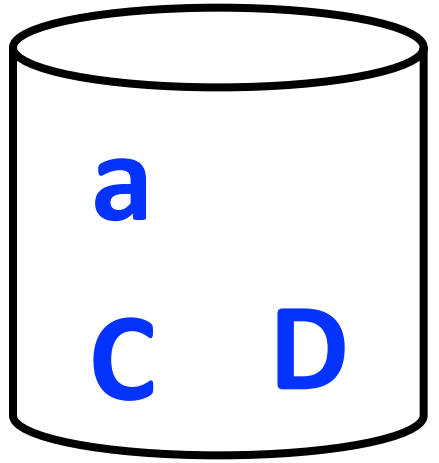

add replica A

Dynamic Replication

Distributes load

add replica A

Dynamic Mastering

W[A, C]

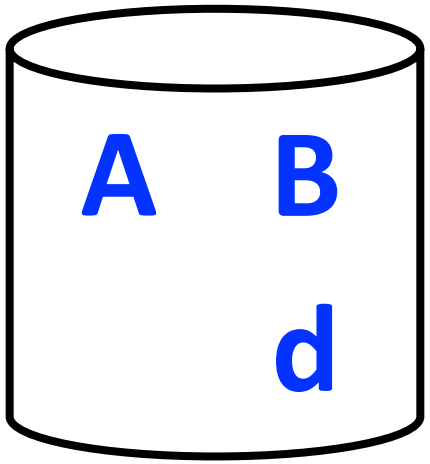

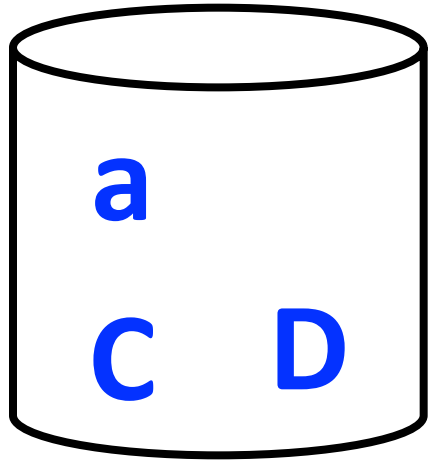

remaster A

Dynamic Mastering

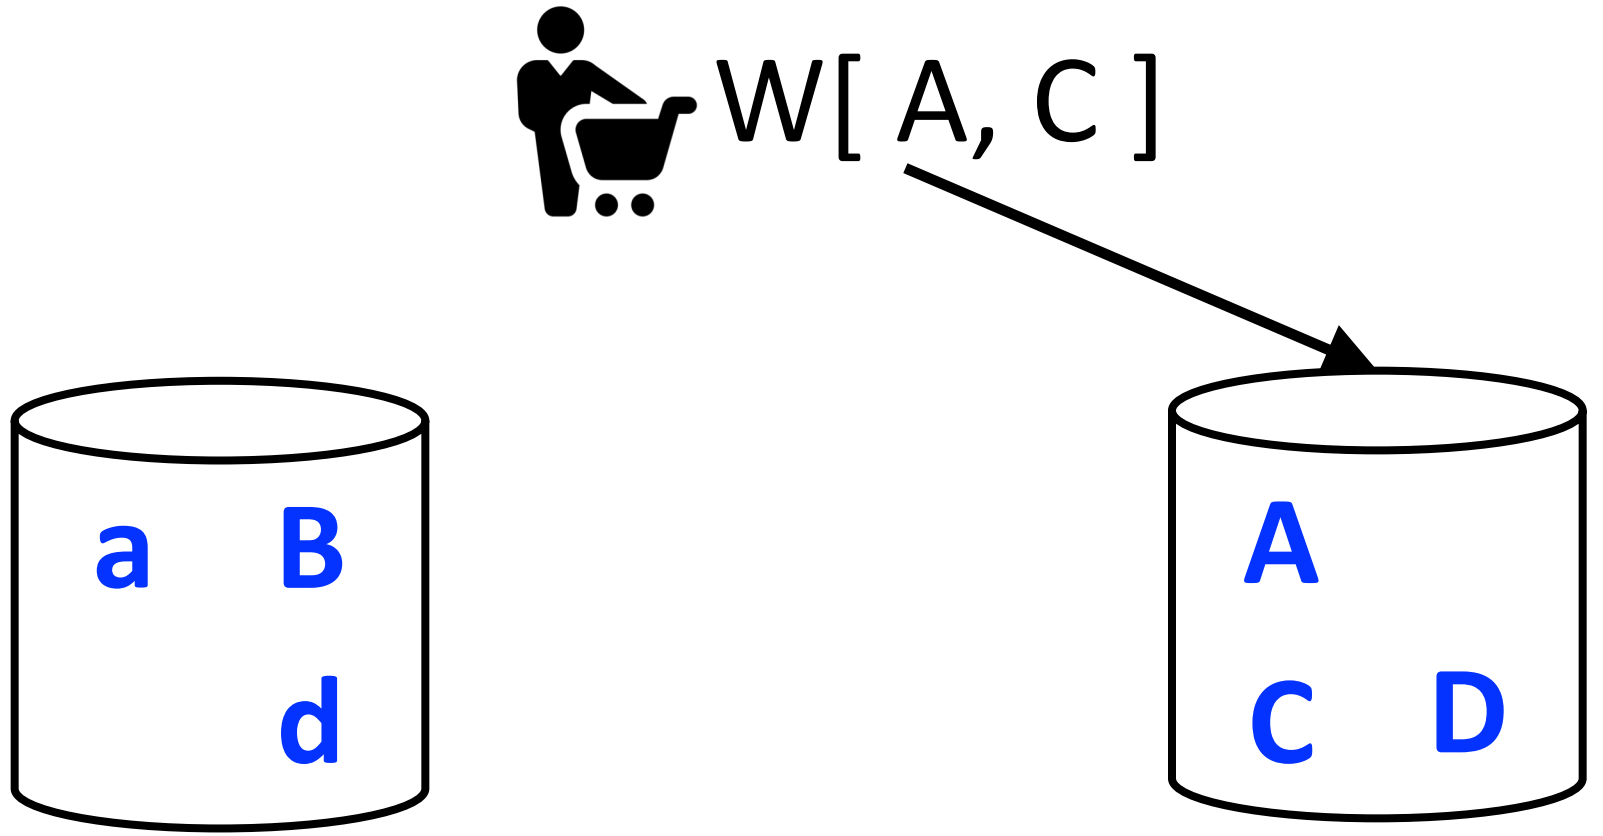

Dynamic Mastering

Single site transactions

remaster A

Dynamic Partitioning

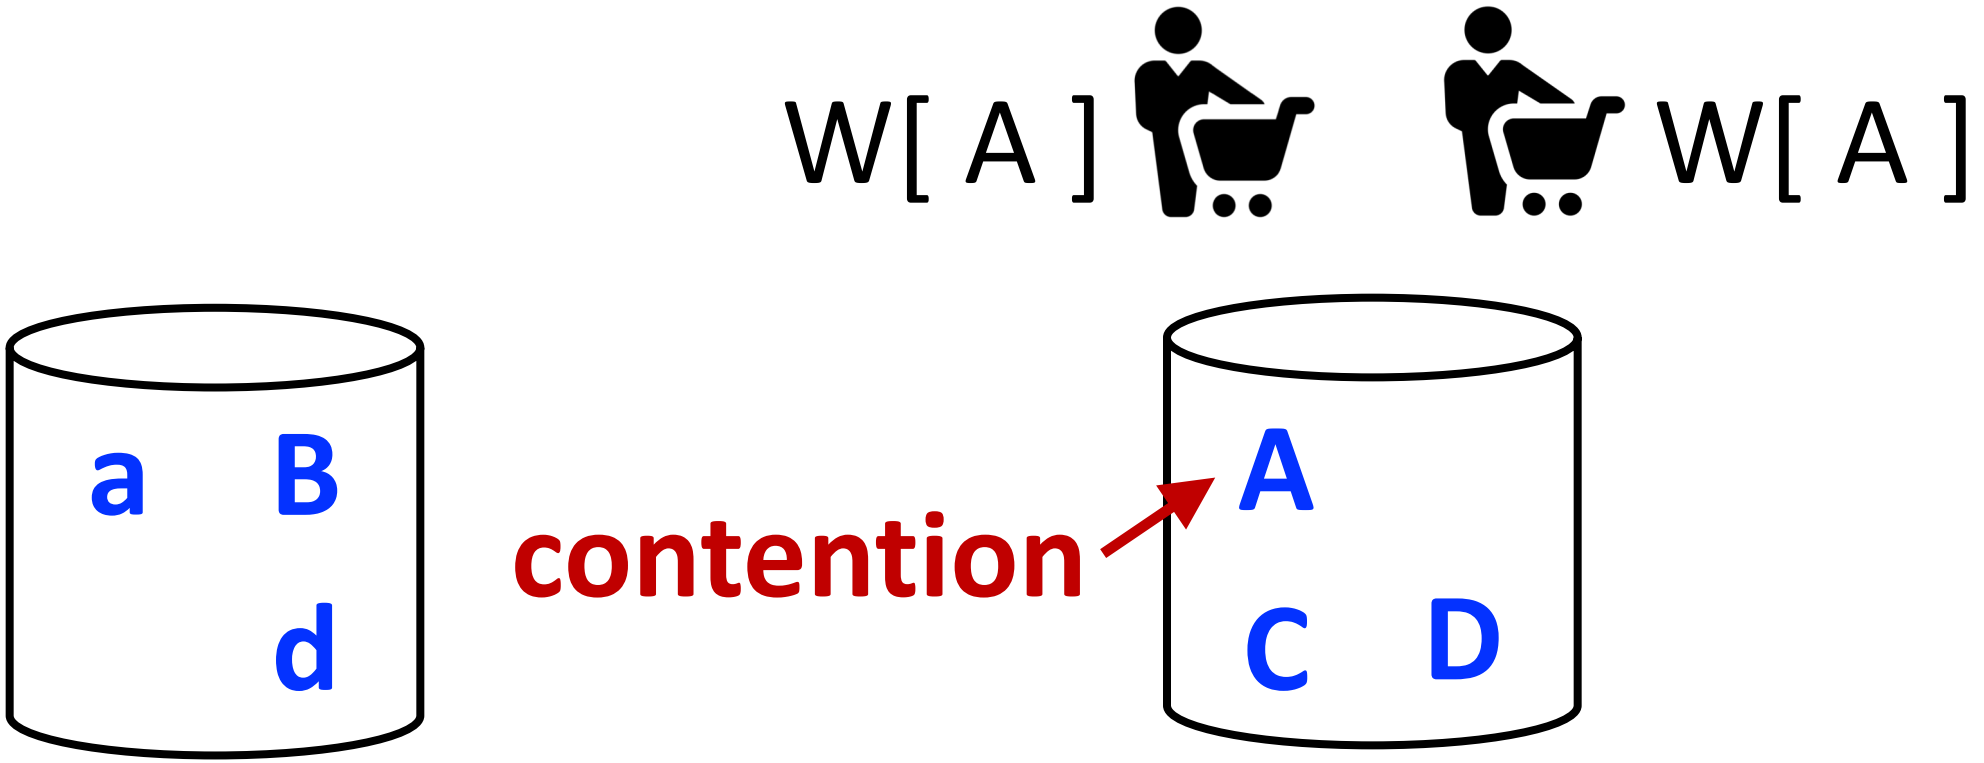

Dynamic Partitioning

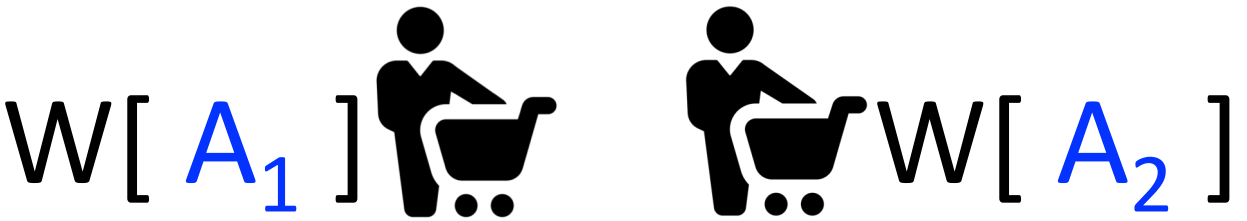

split partition A

Dynamic Partitioning

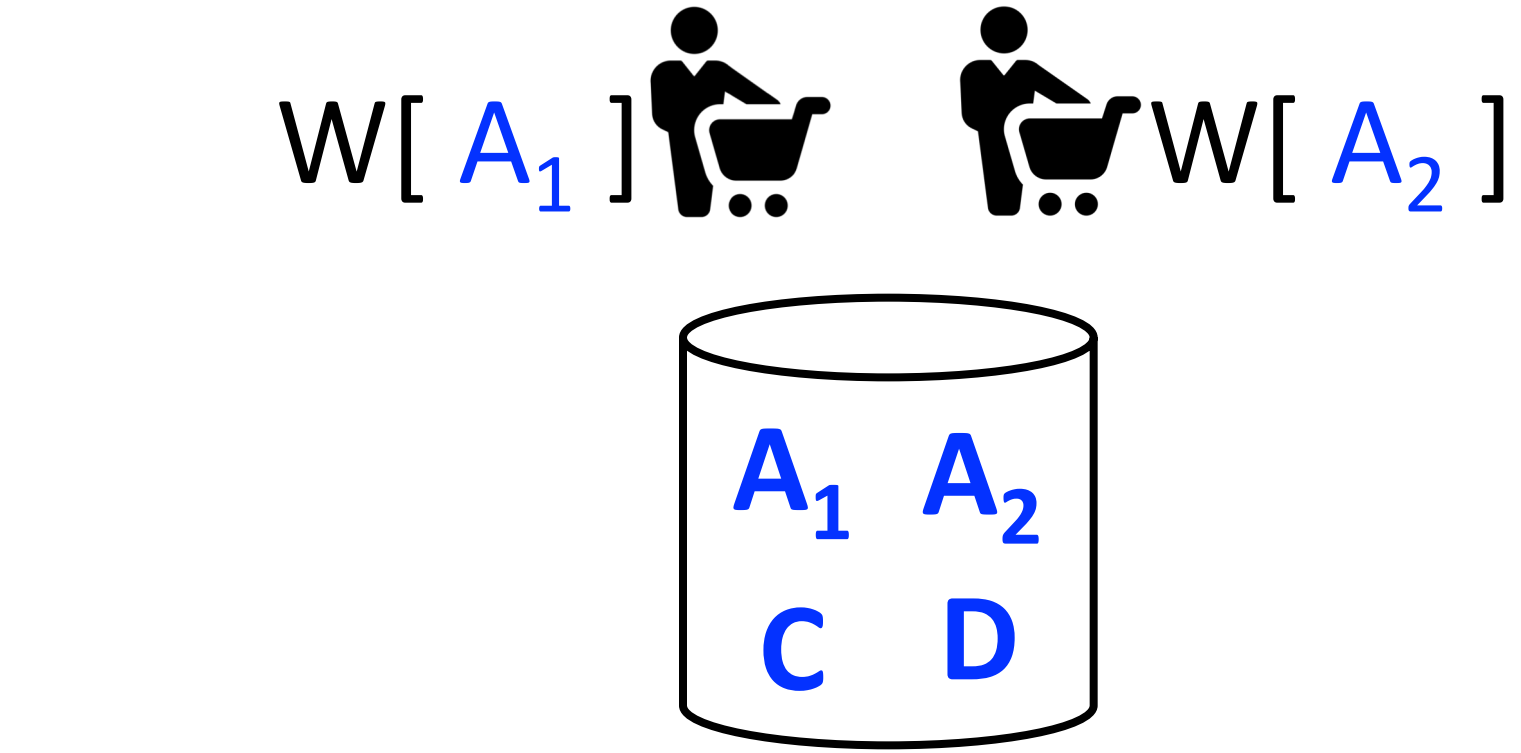

split partition A

Dynamic Partitioning

split partition A

MorphoSys Physical Design Change Operations

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

Remaster a **partition**

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

Challenges of Automatic Physical Design

How to execute both transactions and **design changes efficiently**

How to decide **which physical design operations** to use, and **when**

Efficient Execution

Perform all **operations** at the **partition level**

Decouple partition **reads** and **writes**

Partition based **multi-version** concurrency control

tiny.cc/morphosys

UNIVERSITY OF WATERLOO

Making design decisions

Learned cost model quantifies
design change effects

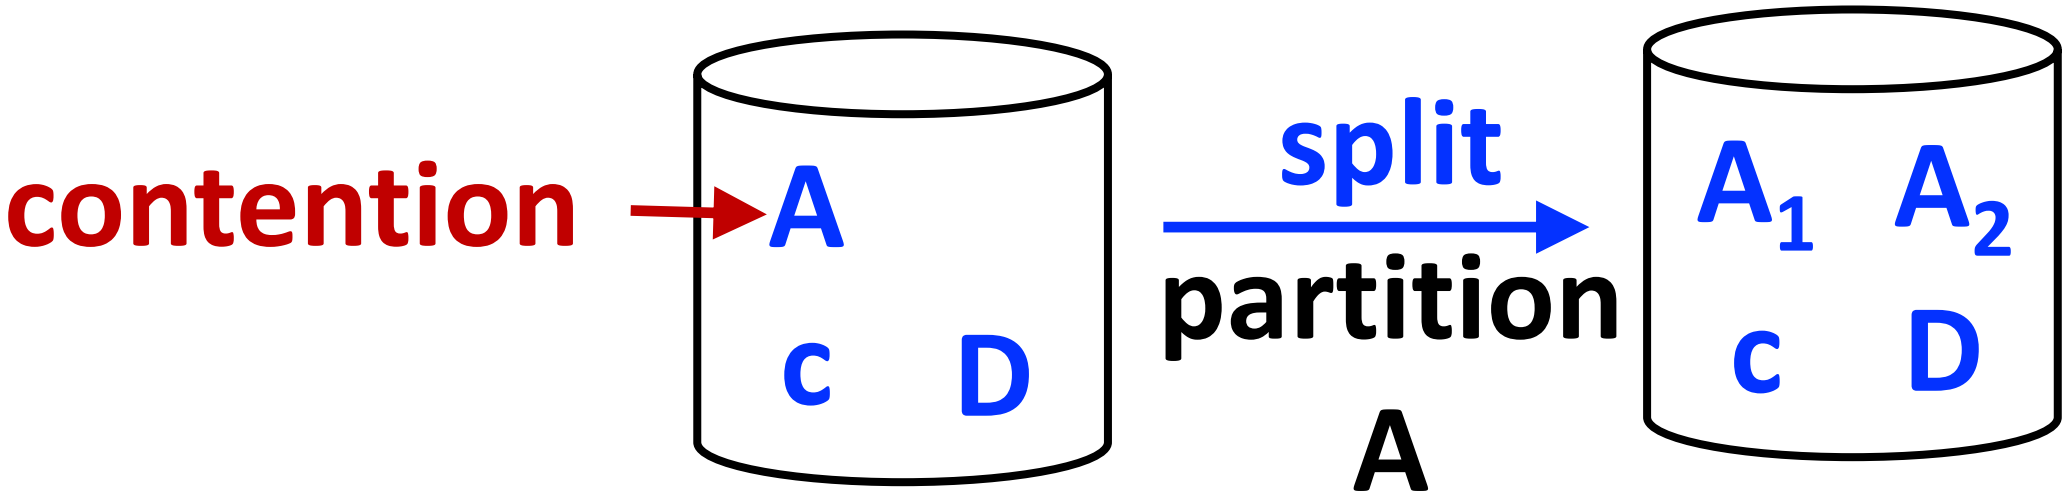

Design change cost < Expected Benefit

UNIVERSITY OF WATERLOO

Physical design cost model

Design change cost < **Expected benefit**

Decompose operators into **key costs**

Predict benefit based on workload history

How well does MorphoSys work?

Comparisons

Single-Master
Multi-Master
VoltDB

} *Static Designs*

DynaMast
Adaptive Replication (ADR)
Clay

} *Dynamic Designs*

UNIVERSITY OF WATERLOO

Skewed YCSB - Throughput

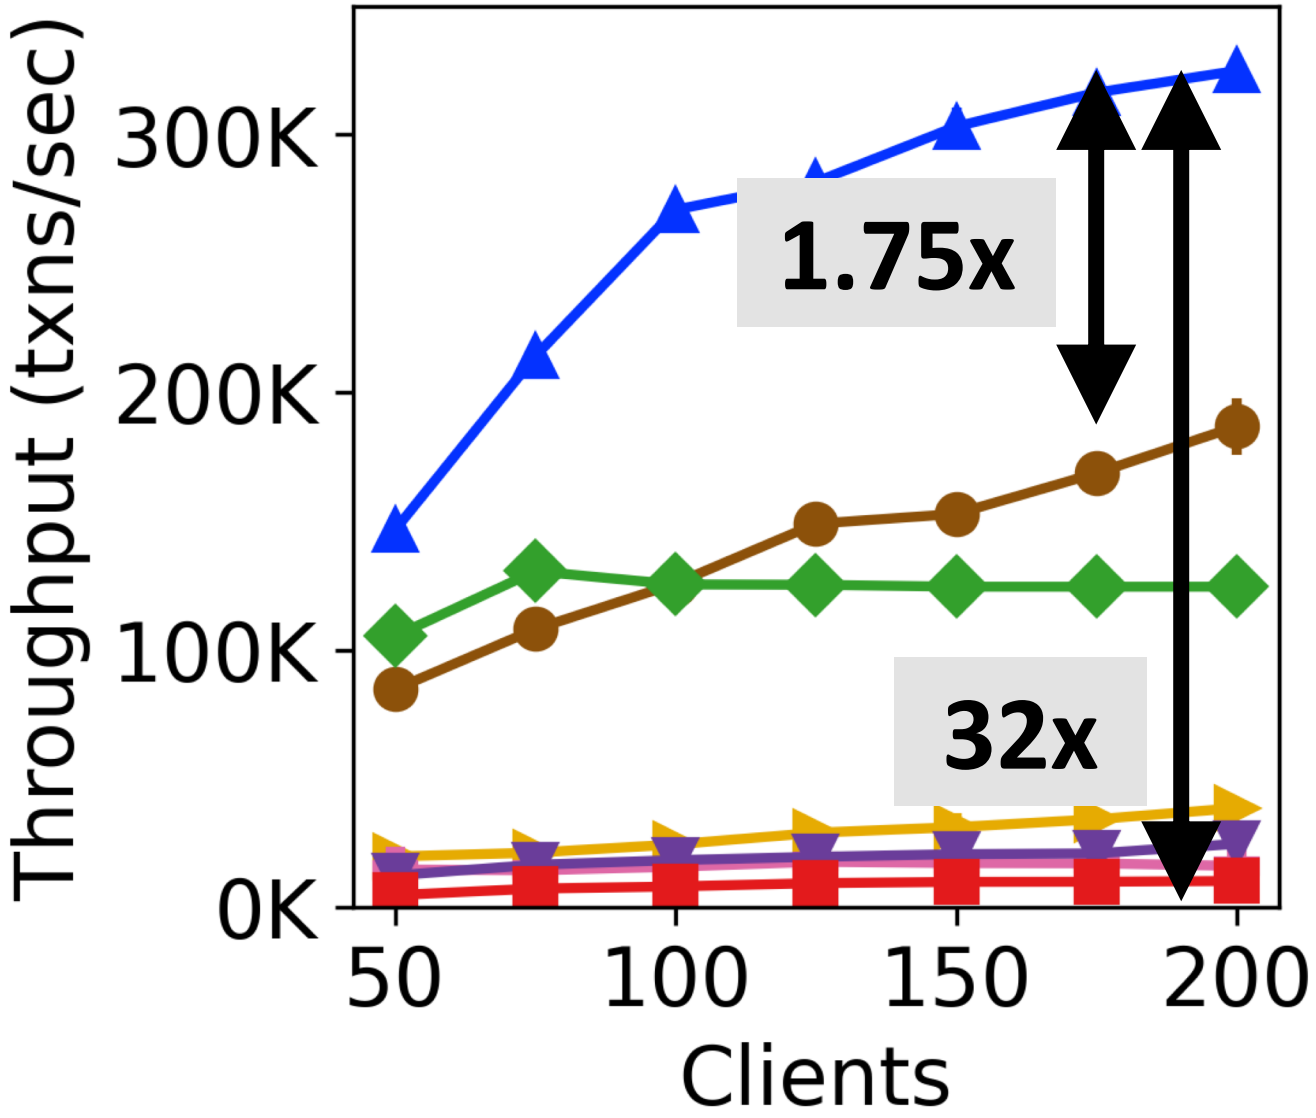

MorphoSys

DynaMast

Single-Master

Clay Multi-Master

VoltDB ADR

UNIVERSITY OF WATERLOO

Number of Replicas

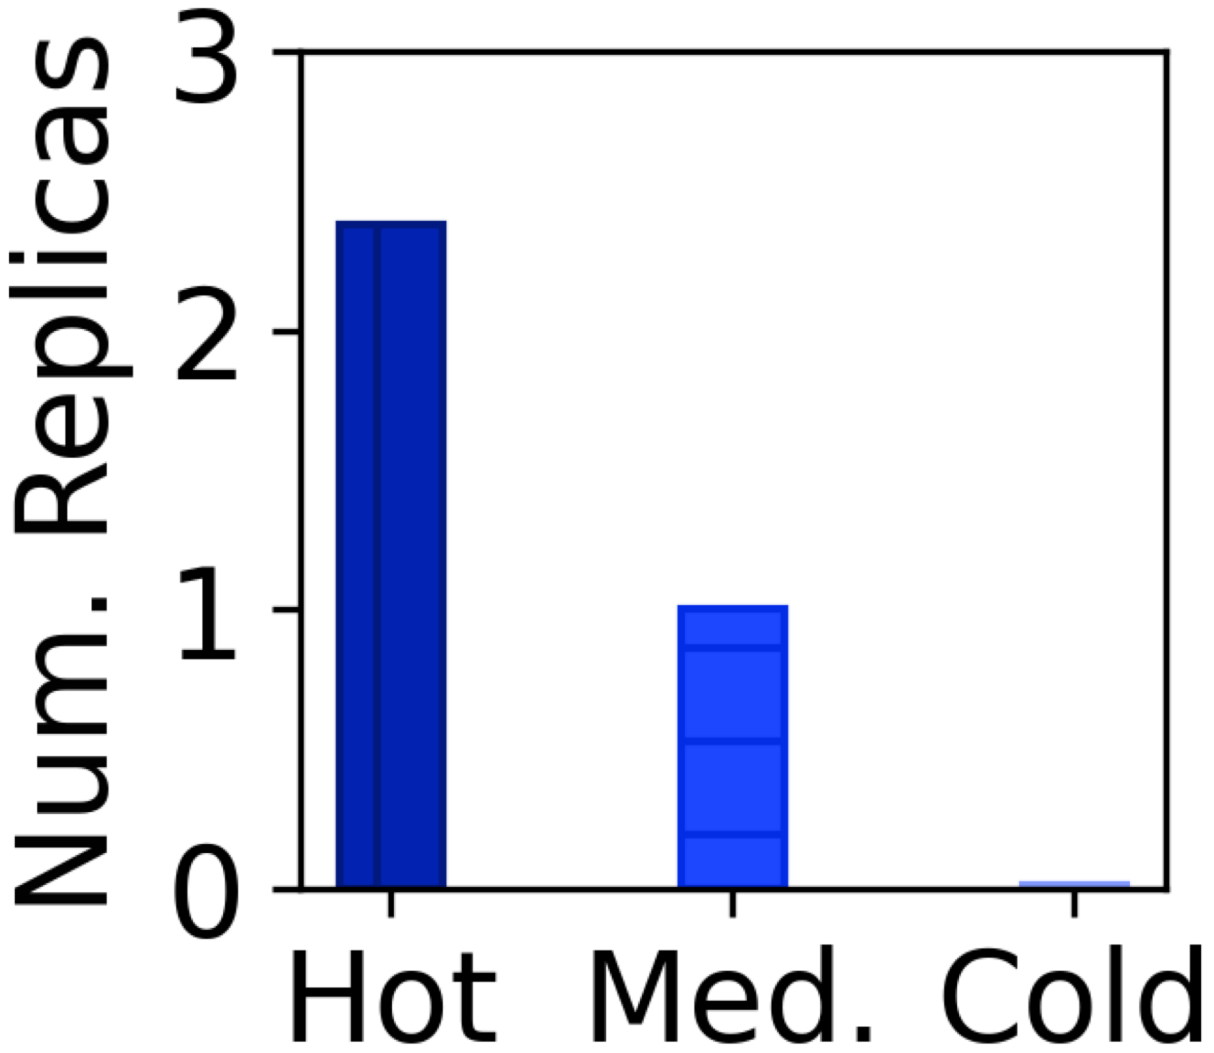

UNIVERSITY OF WATERLOO

Partition Sizes

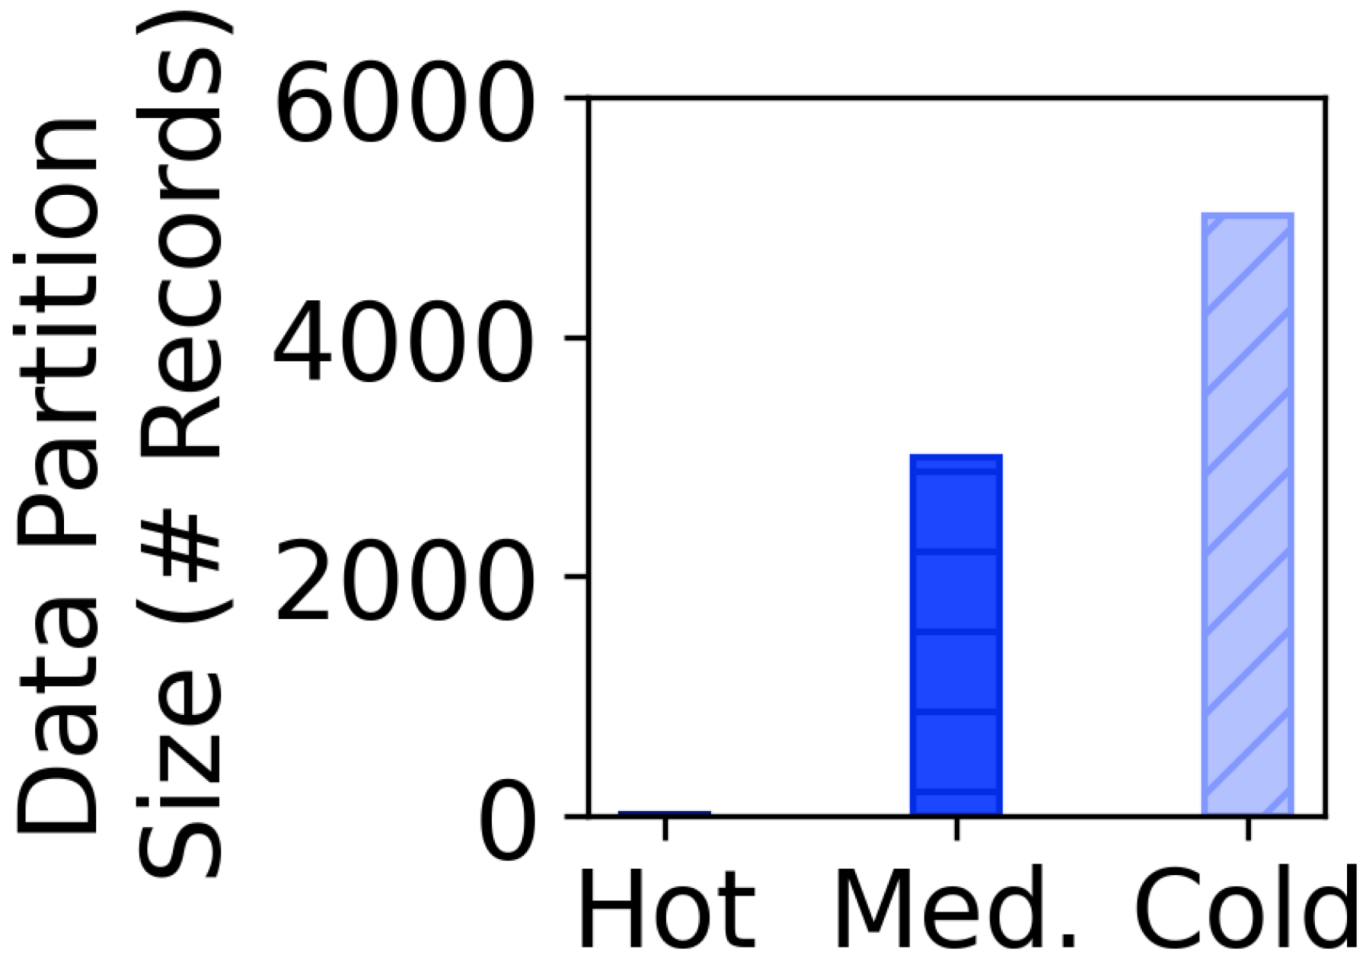

UNIVERSITY OF WATERLOO

Adapts to Workload Changes

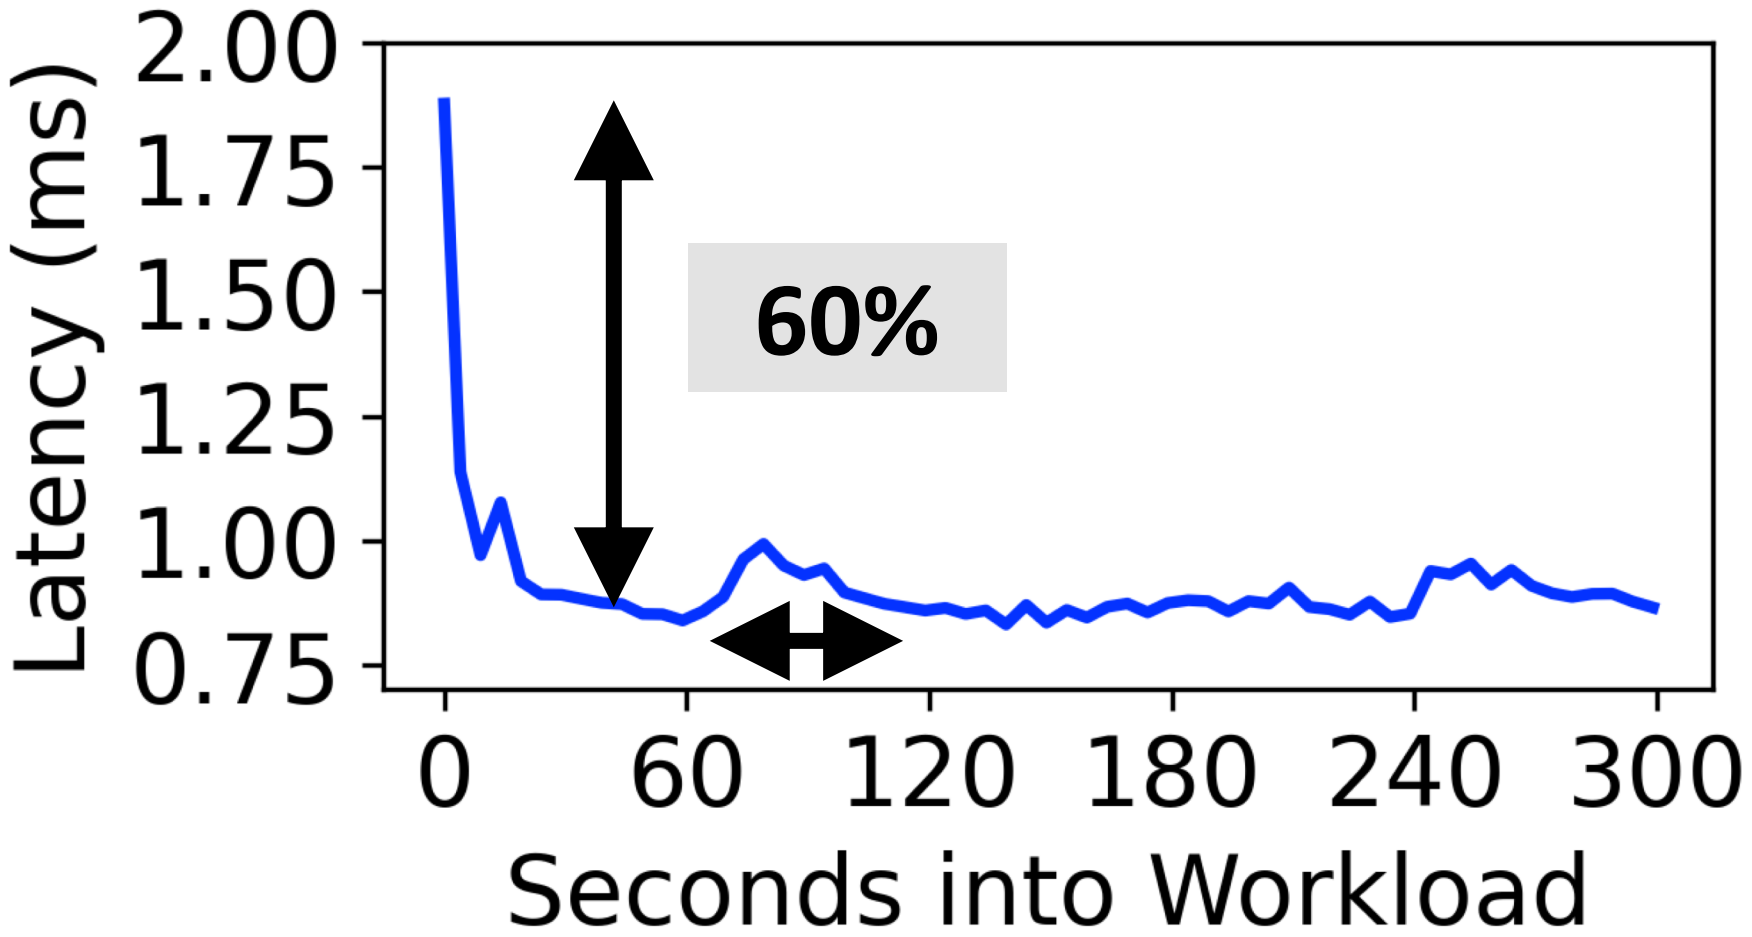

UNIVERSITY OF WATERLOO

More Details in the Paper

Formalism of concurrency control

Role in replica maintenance &
design change operator execution

Generating design change plans

Learned cost functions &
building a workload model

Additional Experiments

UNIVERSITY OF WATERLOO

MorphoSys Takeaways

Automatic physical design changes for distributed databases

tiny.cc/morphosys

Efficiently execute using **partition level operations**

Learned **cost model quantifies** physical design

UNIVERSITY OF WATERLOO