High-Performance Distributed RMA Locks

PATRICK SCHMID, MACIEJ BESTA, TORSTEN HOEFLER

Need for Efficient large-Scale Synchronization

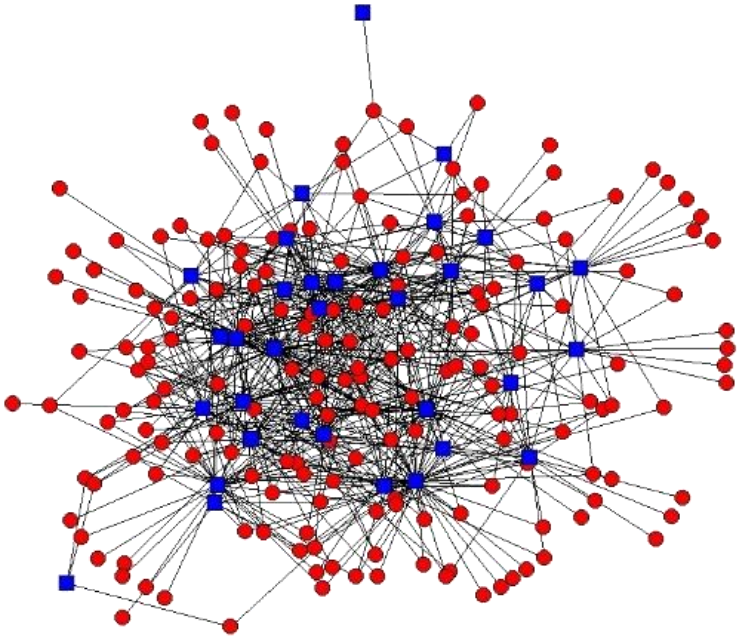

Locks

Inuitive semantics

Various performance penalties

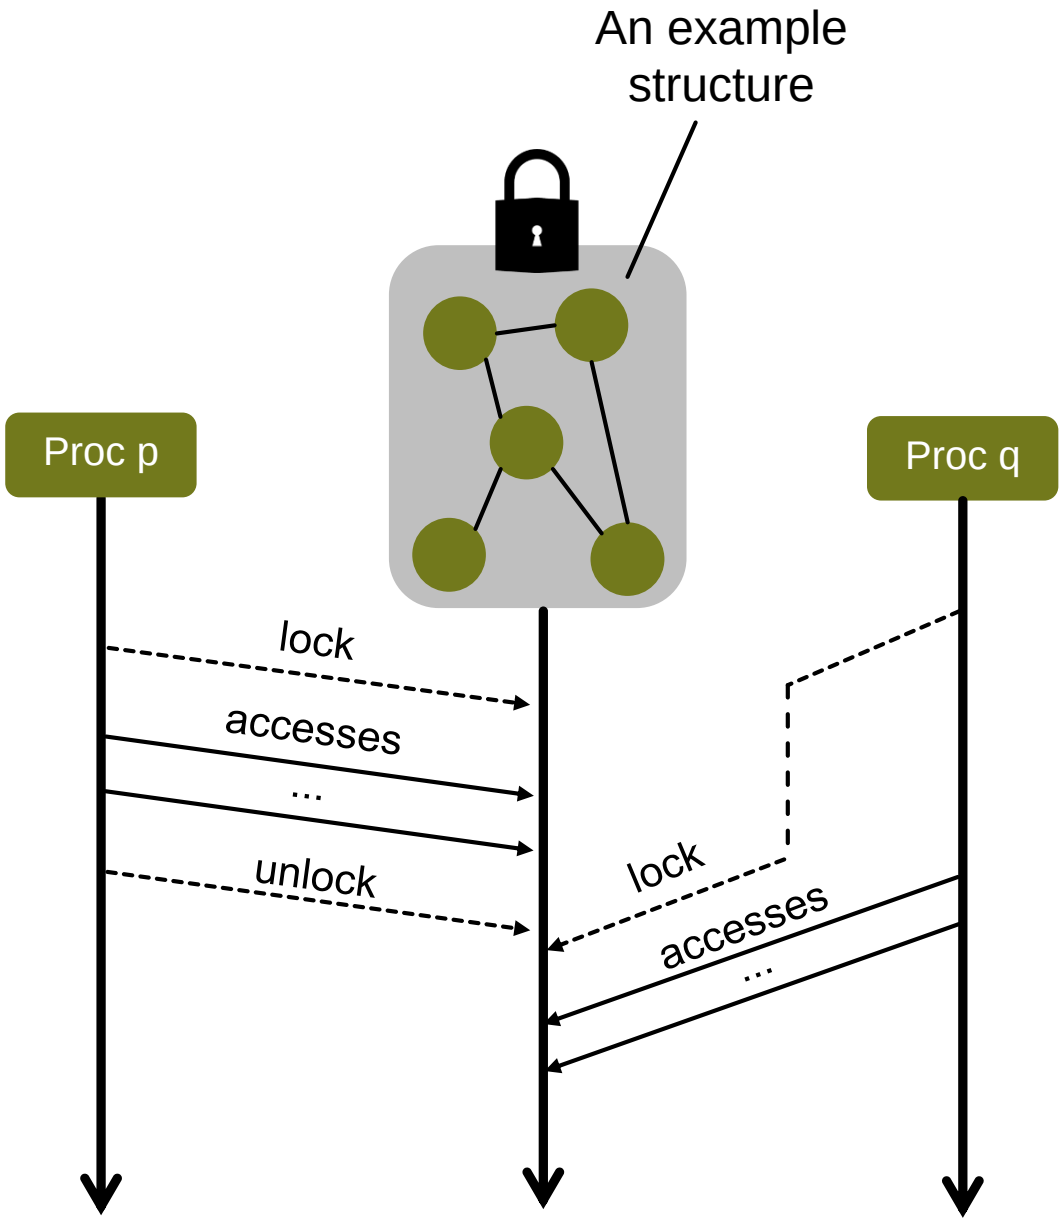

LOCKS: CHALLENGES

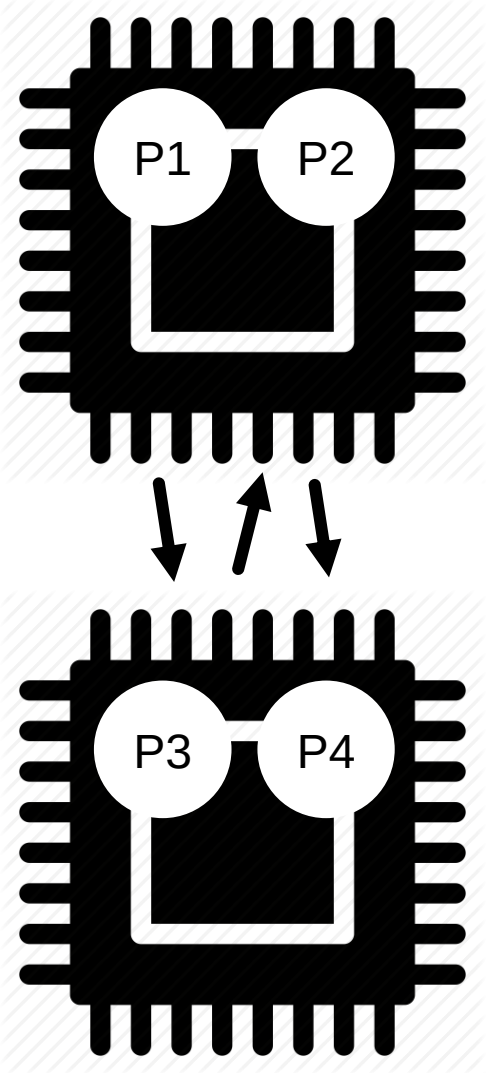

Calciu et al.: NUMA-aware reader-writer locks, PPoPP'13

LOCKS: CHALLENGES

We need intra- and inter-node topology-awareness

We need to cover arbitrary topologies

Locks: Challenges

We need to distinguish between readers and writers

Reader

Reader

Reader

Writer

We need flexible performance for both types of processes

[1] V. Venkataramani et al. Tao: How facebook serves the social graph. SIGMOD'12.

What will we use in the design?

What we will use
MCS Locks

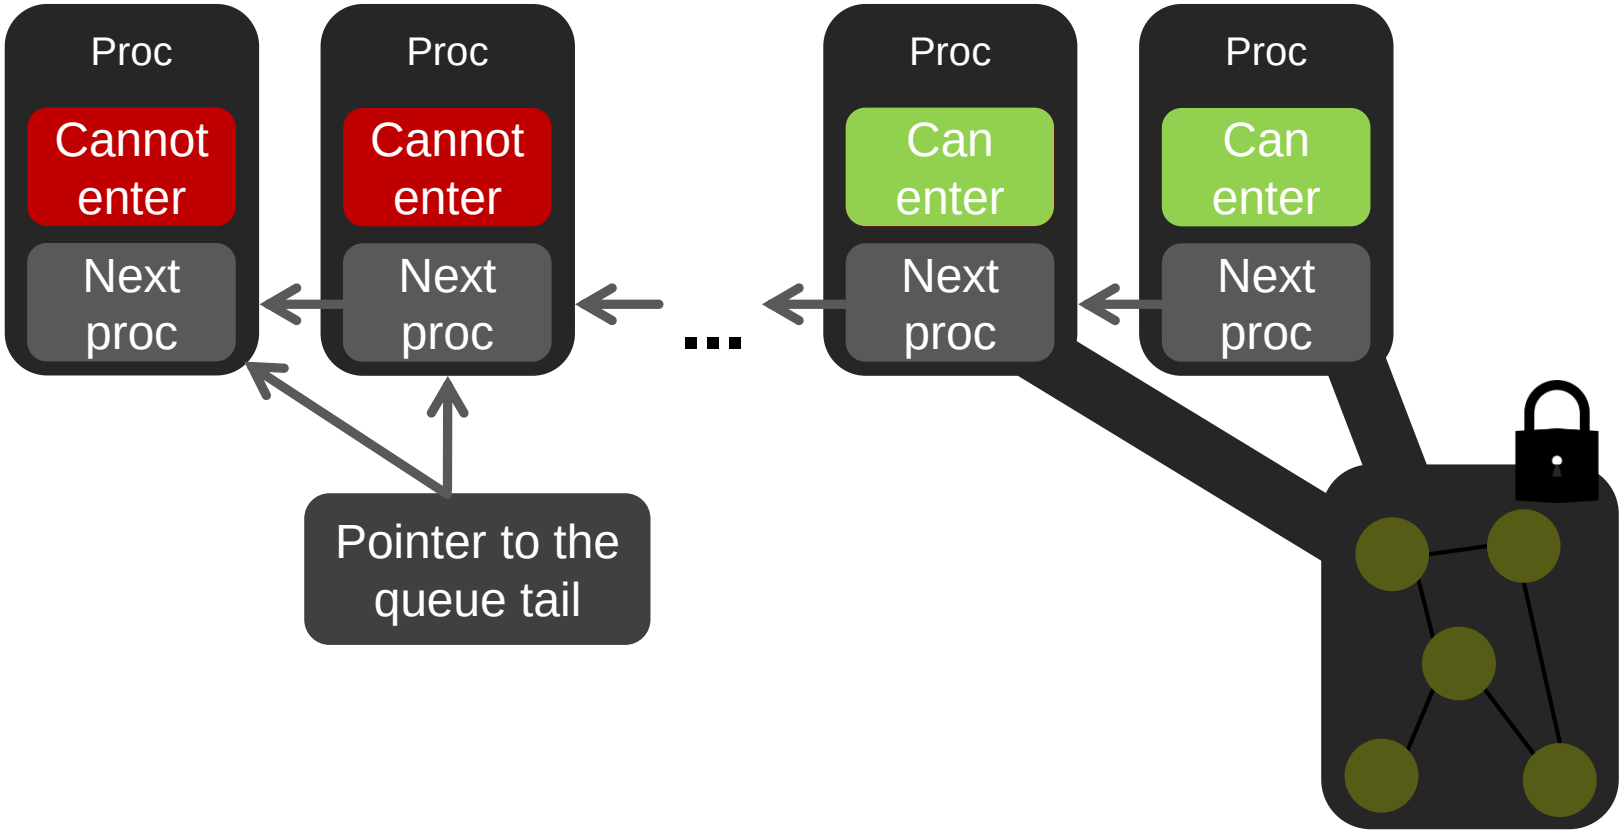

Mellor-Crummey and Scott: Algorithms for Scalable Synchronization on Shared-Memory Multiprocessors, ACM TOCS'91

What we will use
Reader-Writer Locks

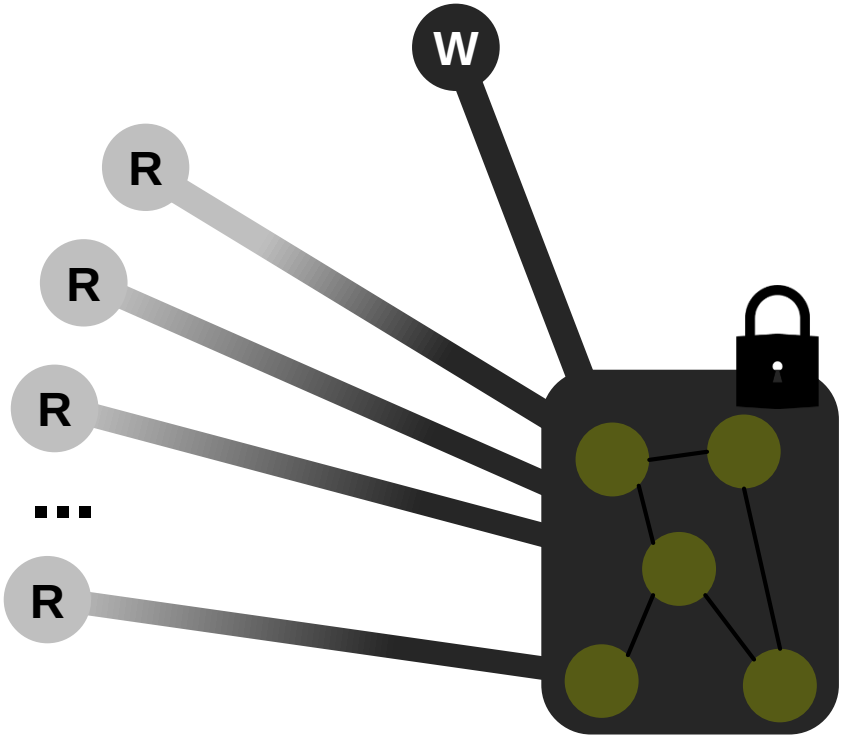

How to manage the design complexity?

How to ensure tunable performance?

What mechanism to use for efficient implementation?

REMOTE MEMORY ACCESS (RMA) PROGRAMMING

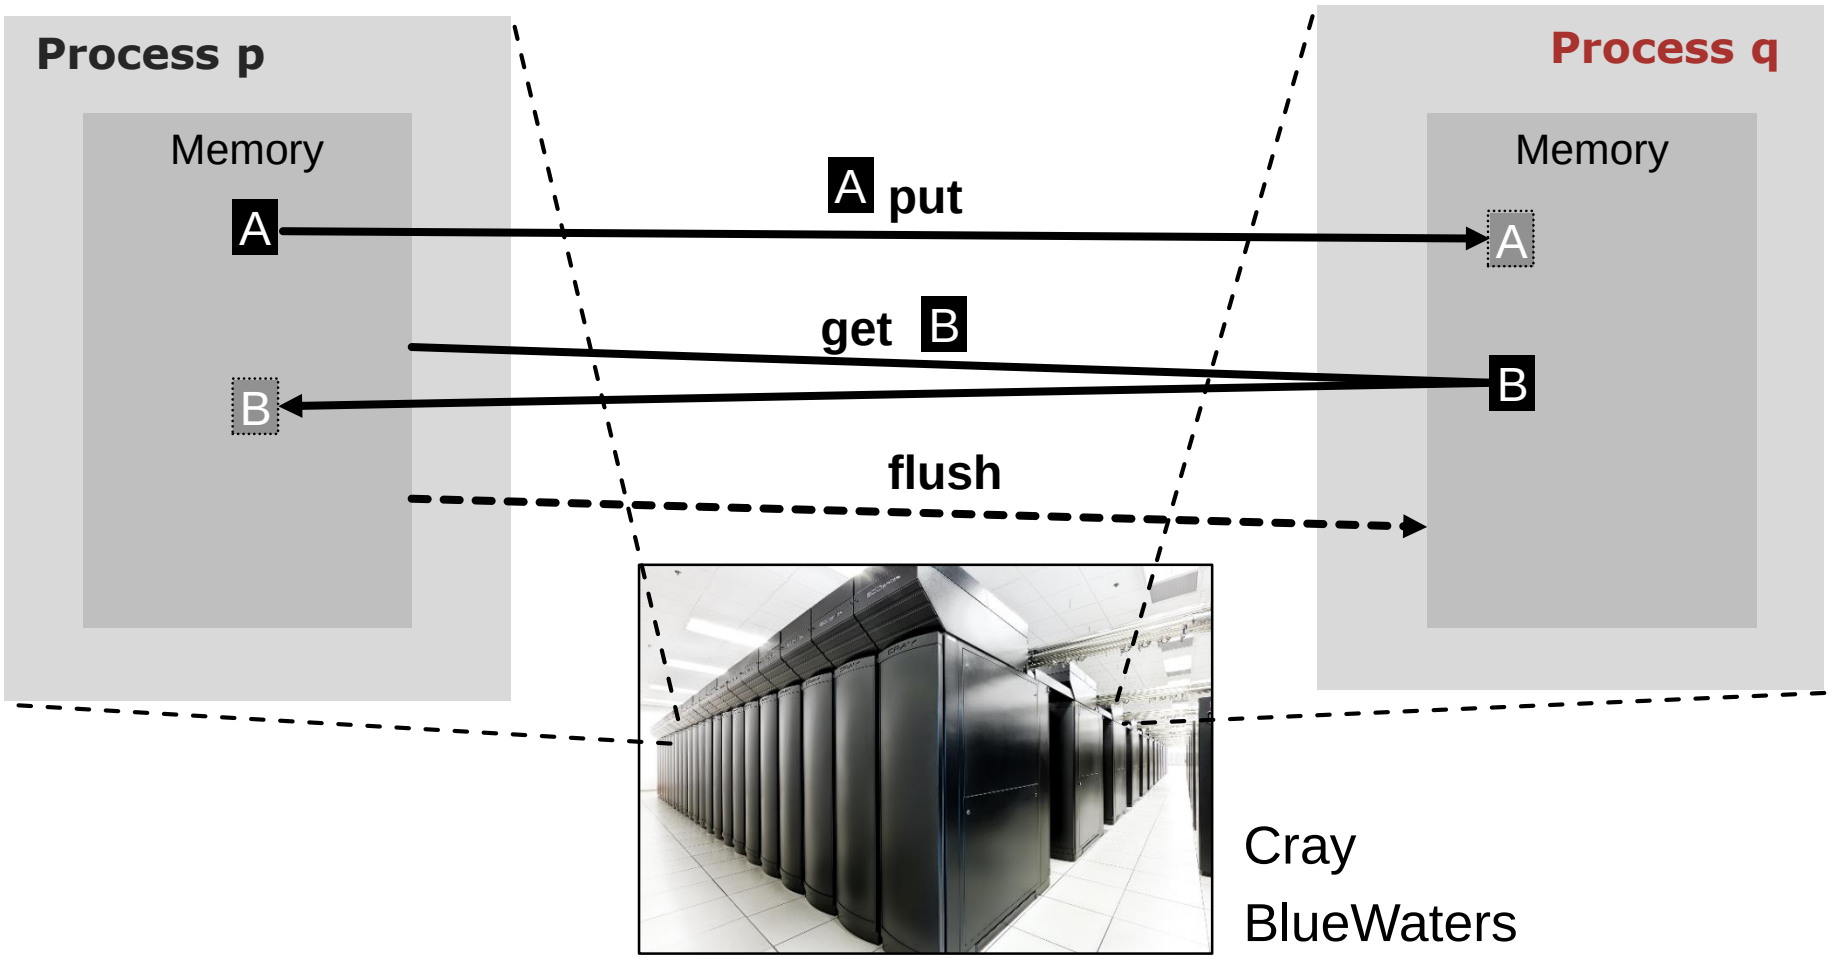

TH, J. Dinan, R. Thakur, B. Barrett, P. Balaji, W. Gropp, K. Underwood: Remote Memory Access Programming in MPI-3, ACM TOPC'15

REMOTE MEMORY ACCESS PROGRAMMING

- Implemented in hardware in NICs in the majority of HPC networks (RDMA support).

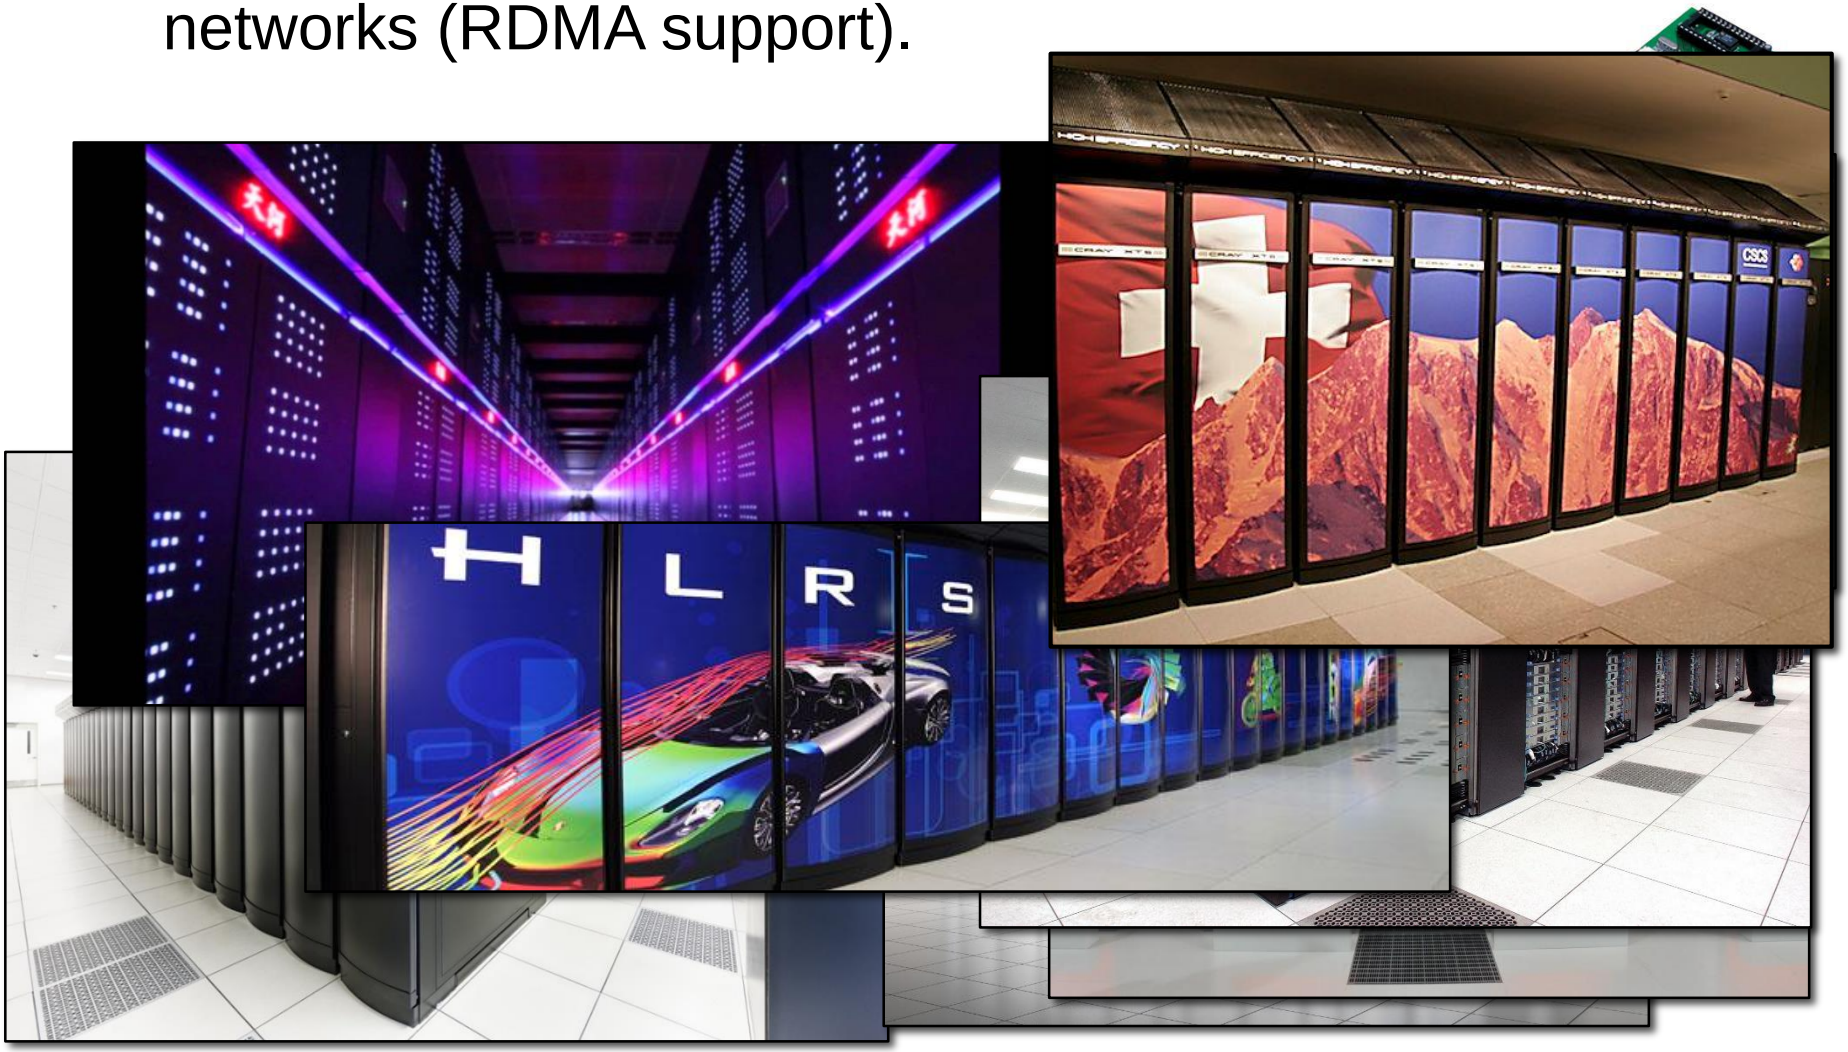

How to manage the design complexity?

How to ensure tunable performance?

What mechanism to use for efficient implementation?

How to manage the design complexity?

Each element has its own distributed MCS queue (DQ) of writers

Readers and writers synchronize with a distributed counter (DC)

MCS queues form a distributed tree (DT)

Modular design

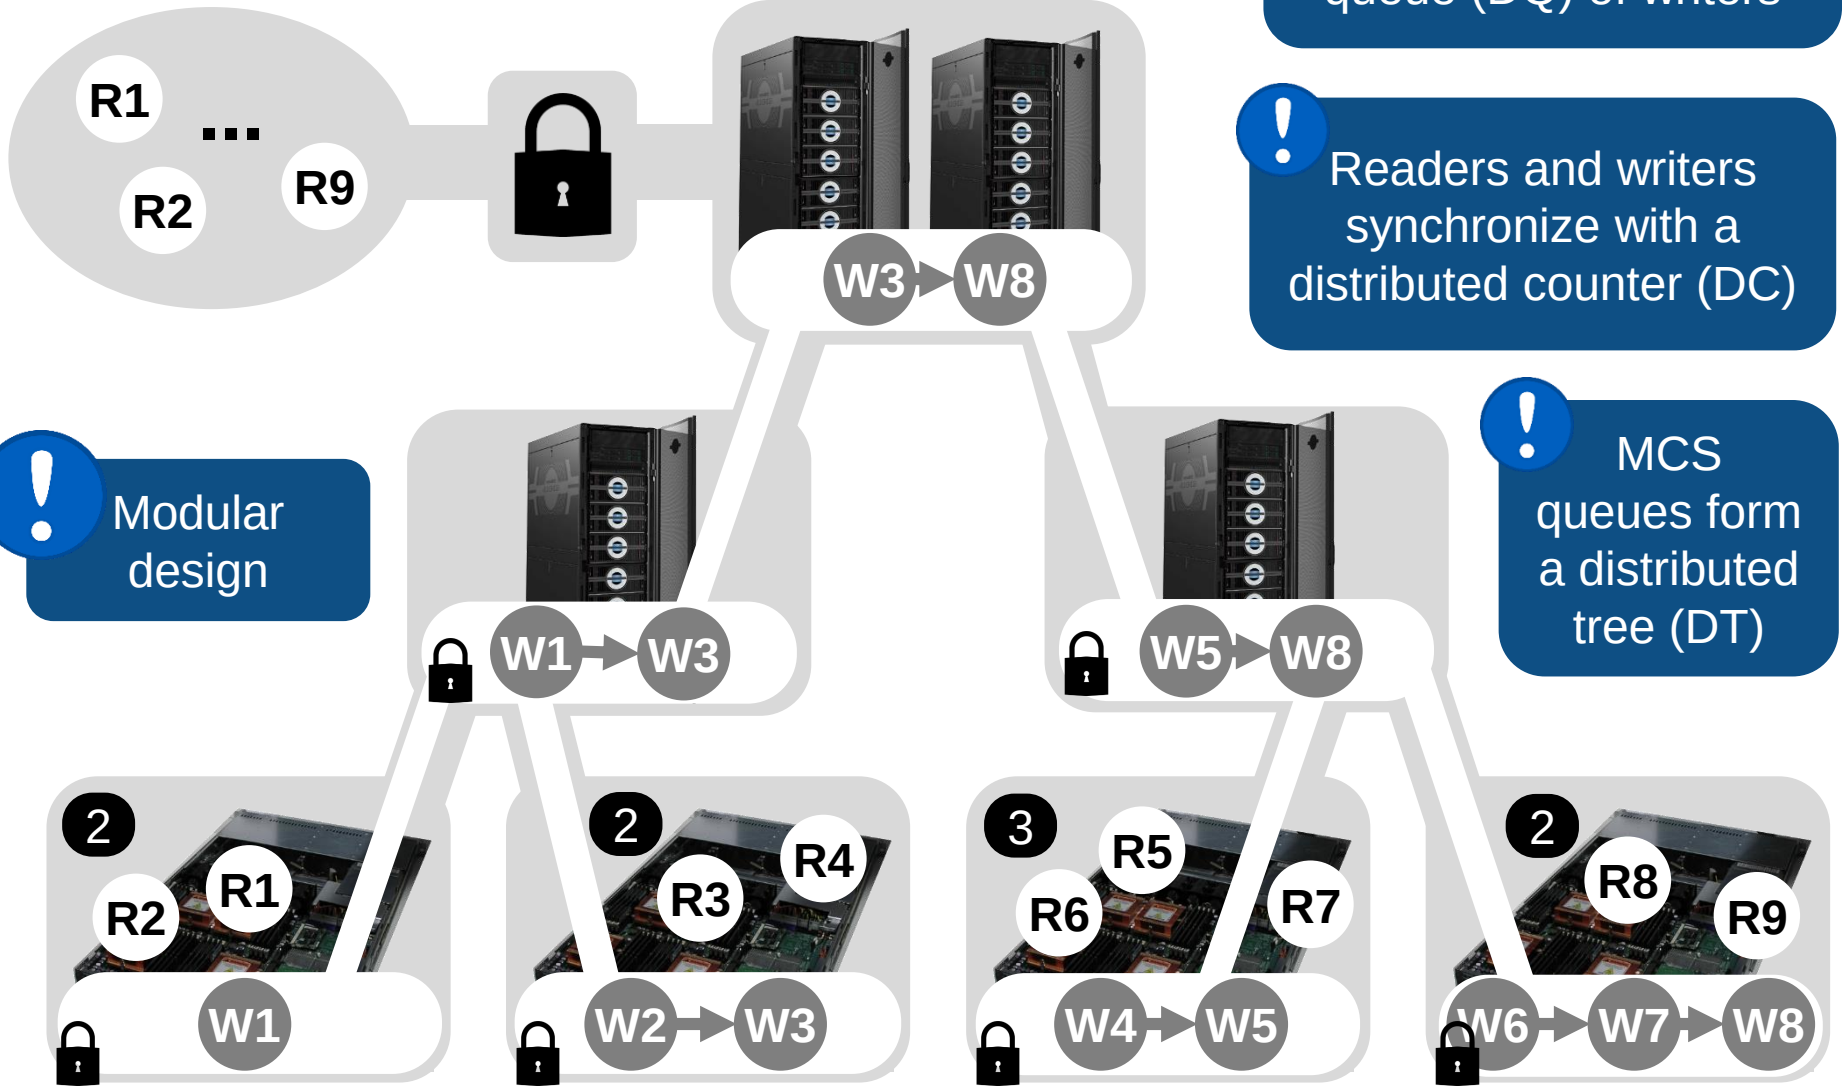

How to ensure tunable performance?

Each DQ: fairness vs throughput of writers

DC: a parameter for the latency of readers vs writers

DT: a parameter for the throughput of readers vs writers

A tradeoff parameter for every structure

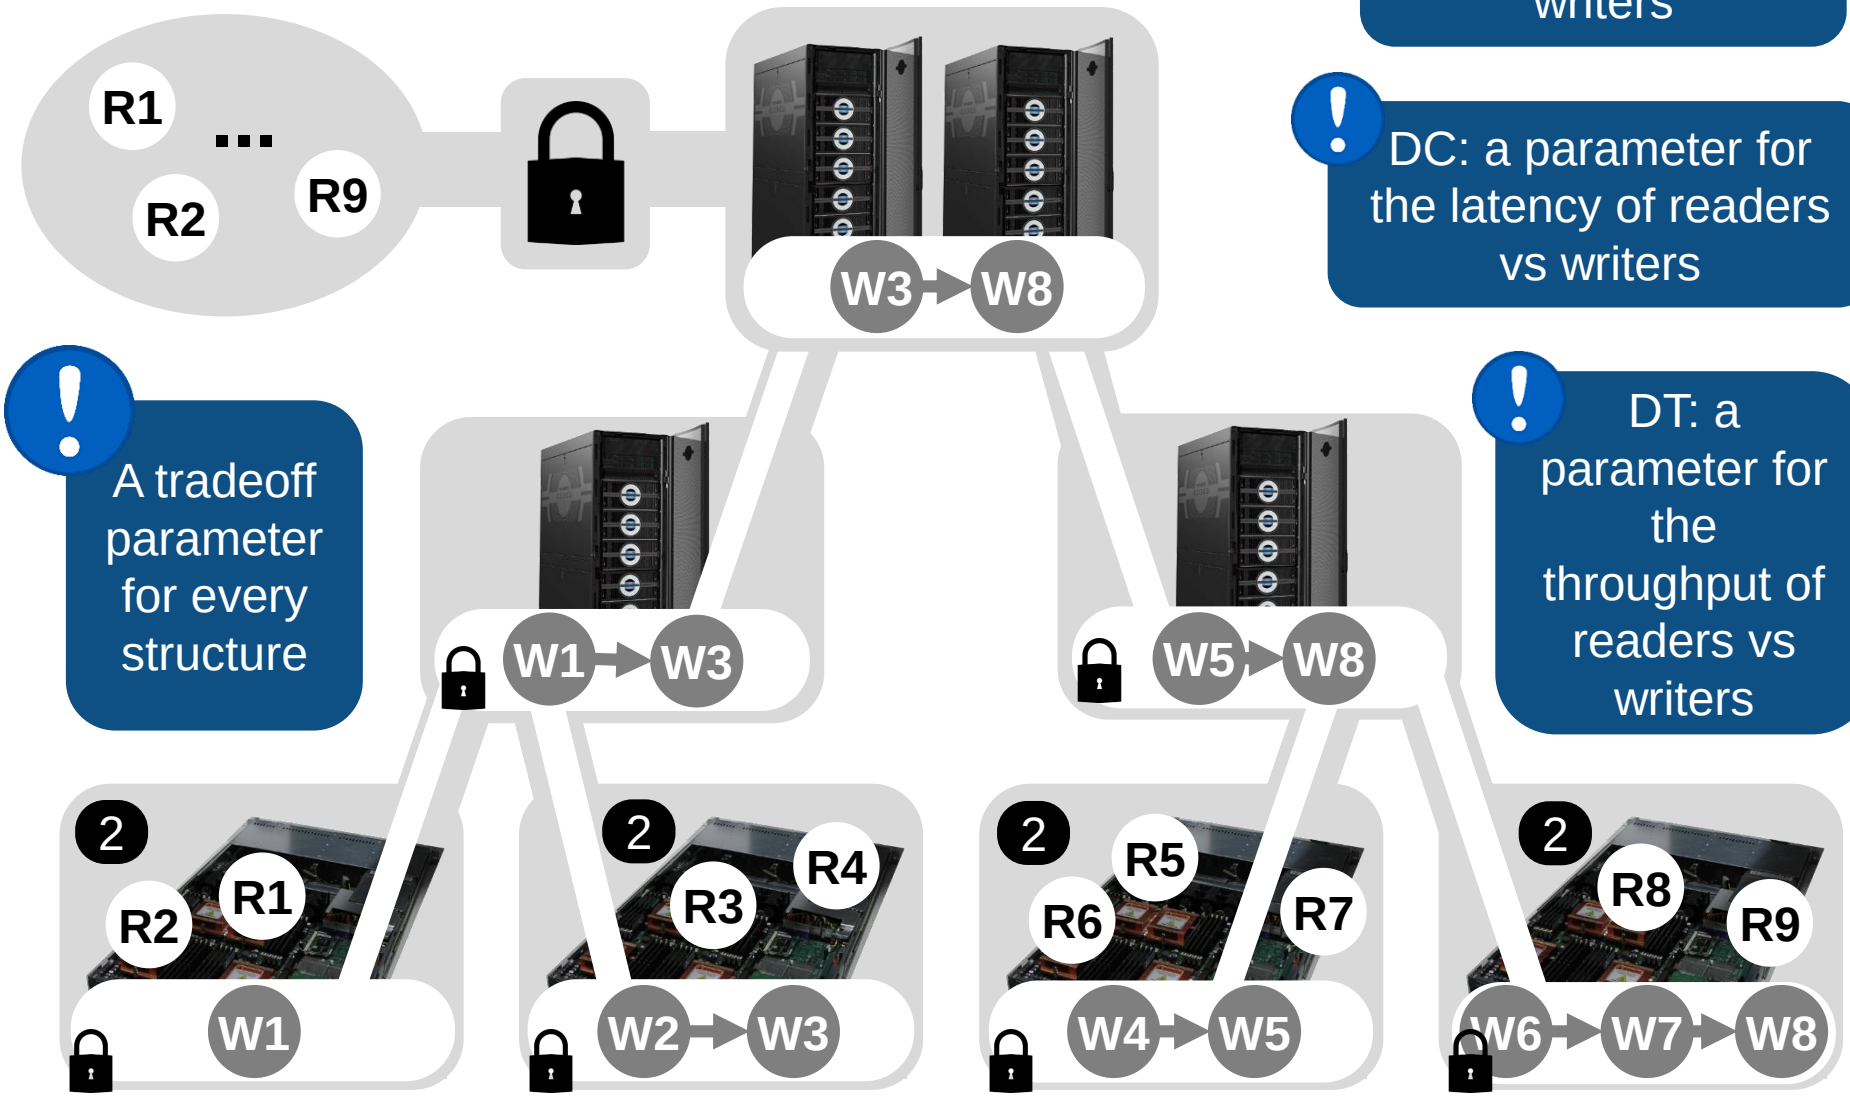

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

DISTRIBUTED MCS QUEUES (DQS)

Throughput vs Fairness

Larger $T_{L,i}$: more throughput at level i. Smaller $T_{L,i}$: more fairness at level i.

Each DQ: The maximum number of lock passings within a DQ at level i, before it is passed to another DQ at i.

$T_{L,i}$

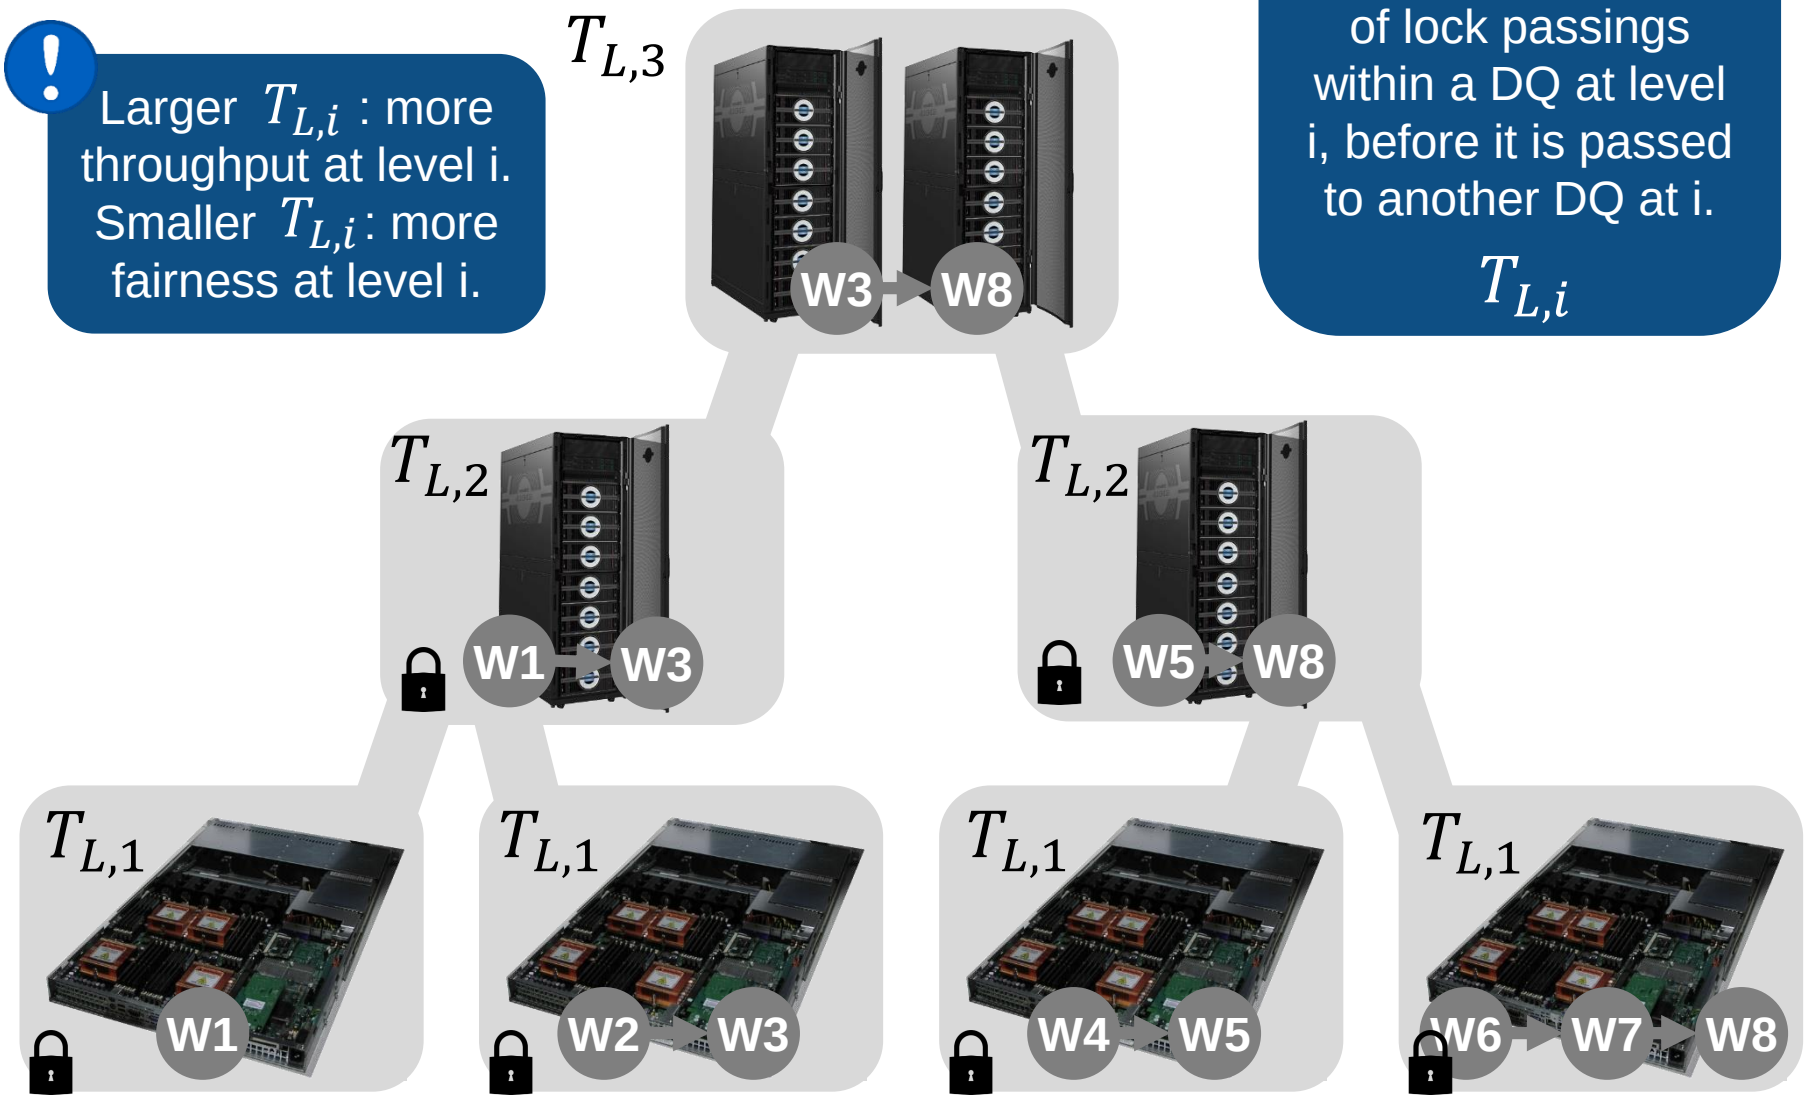

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

Distributed Tree of Queues (DT)

Throughput of readers vs writers

DT: The maximum number of consecutive lock passings within readers (T_R).

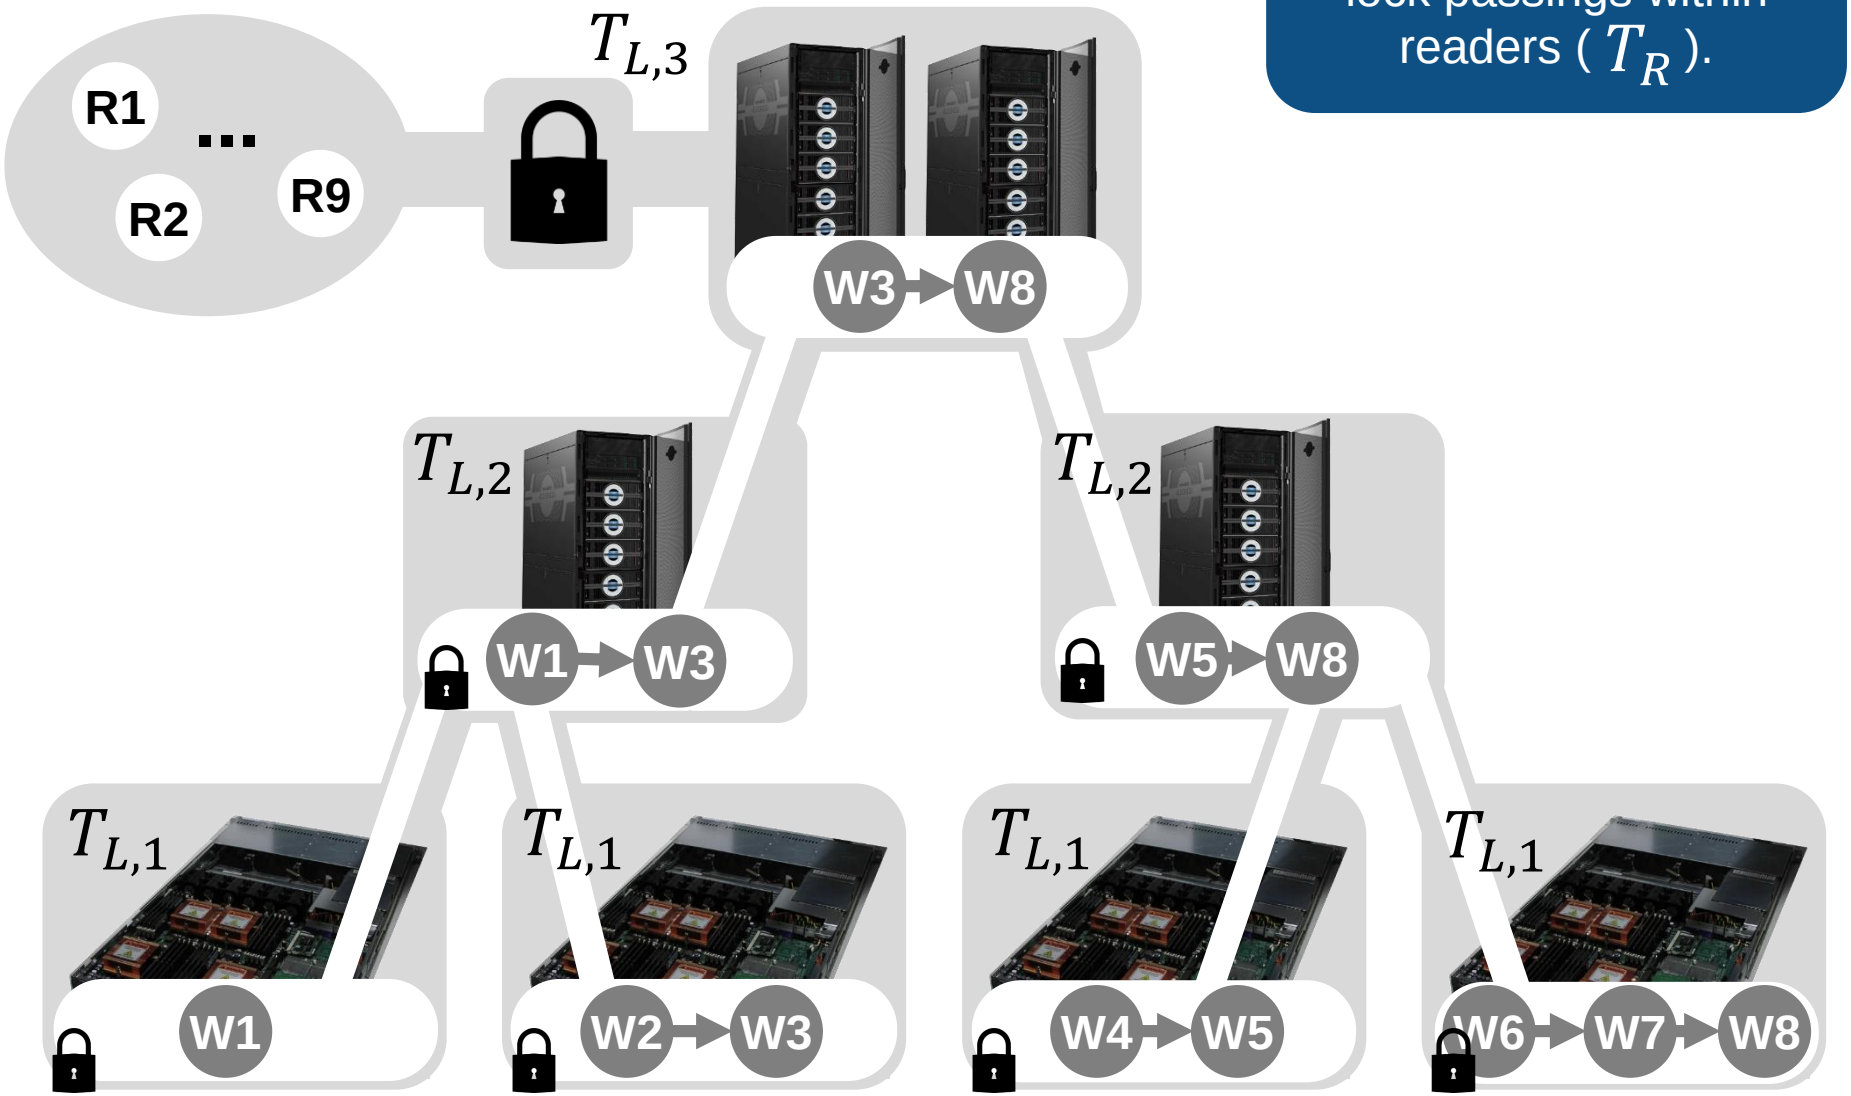

DISTRIBUTED COUNTER (DC)

Latency of readers vs writers

DC: every *k*th compute node hosts a partial counter, all of which constitute the DC.

$$k = T_{DC}$$

A writer holds the lock

b|x|y

Readers that arrived at the CS

Readers that left the CS

$T_{DC} = 1$

$T_{DC} = 2$

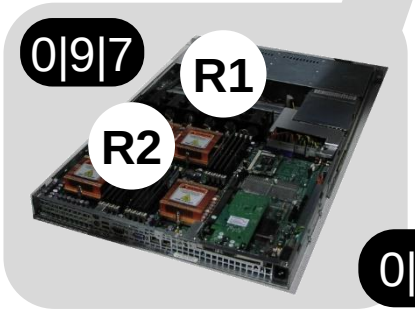

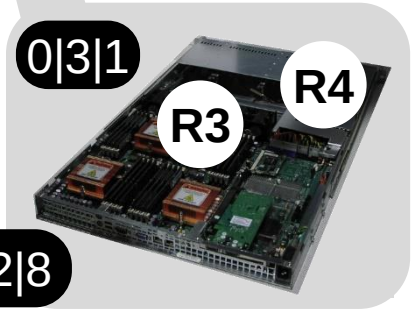

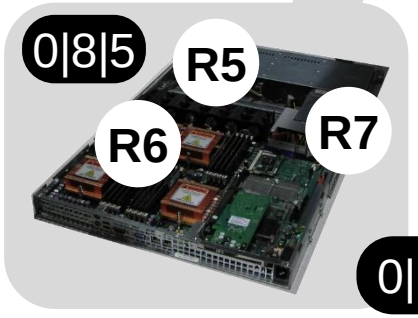

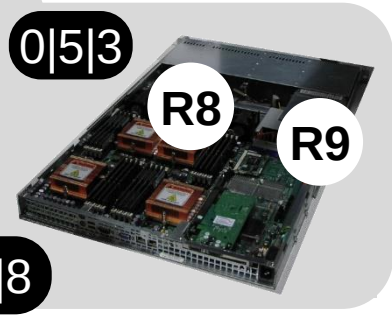

THE SPACE OF DESIGNS

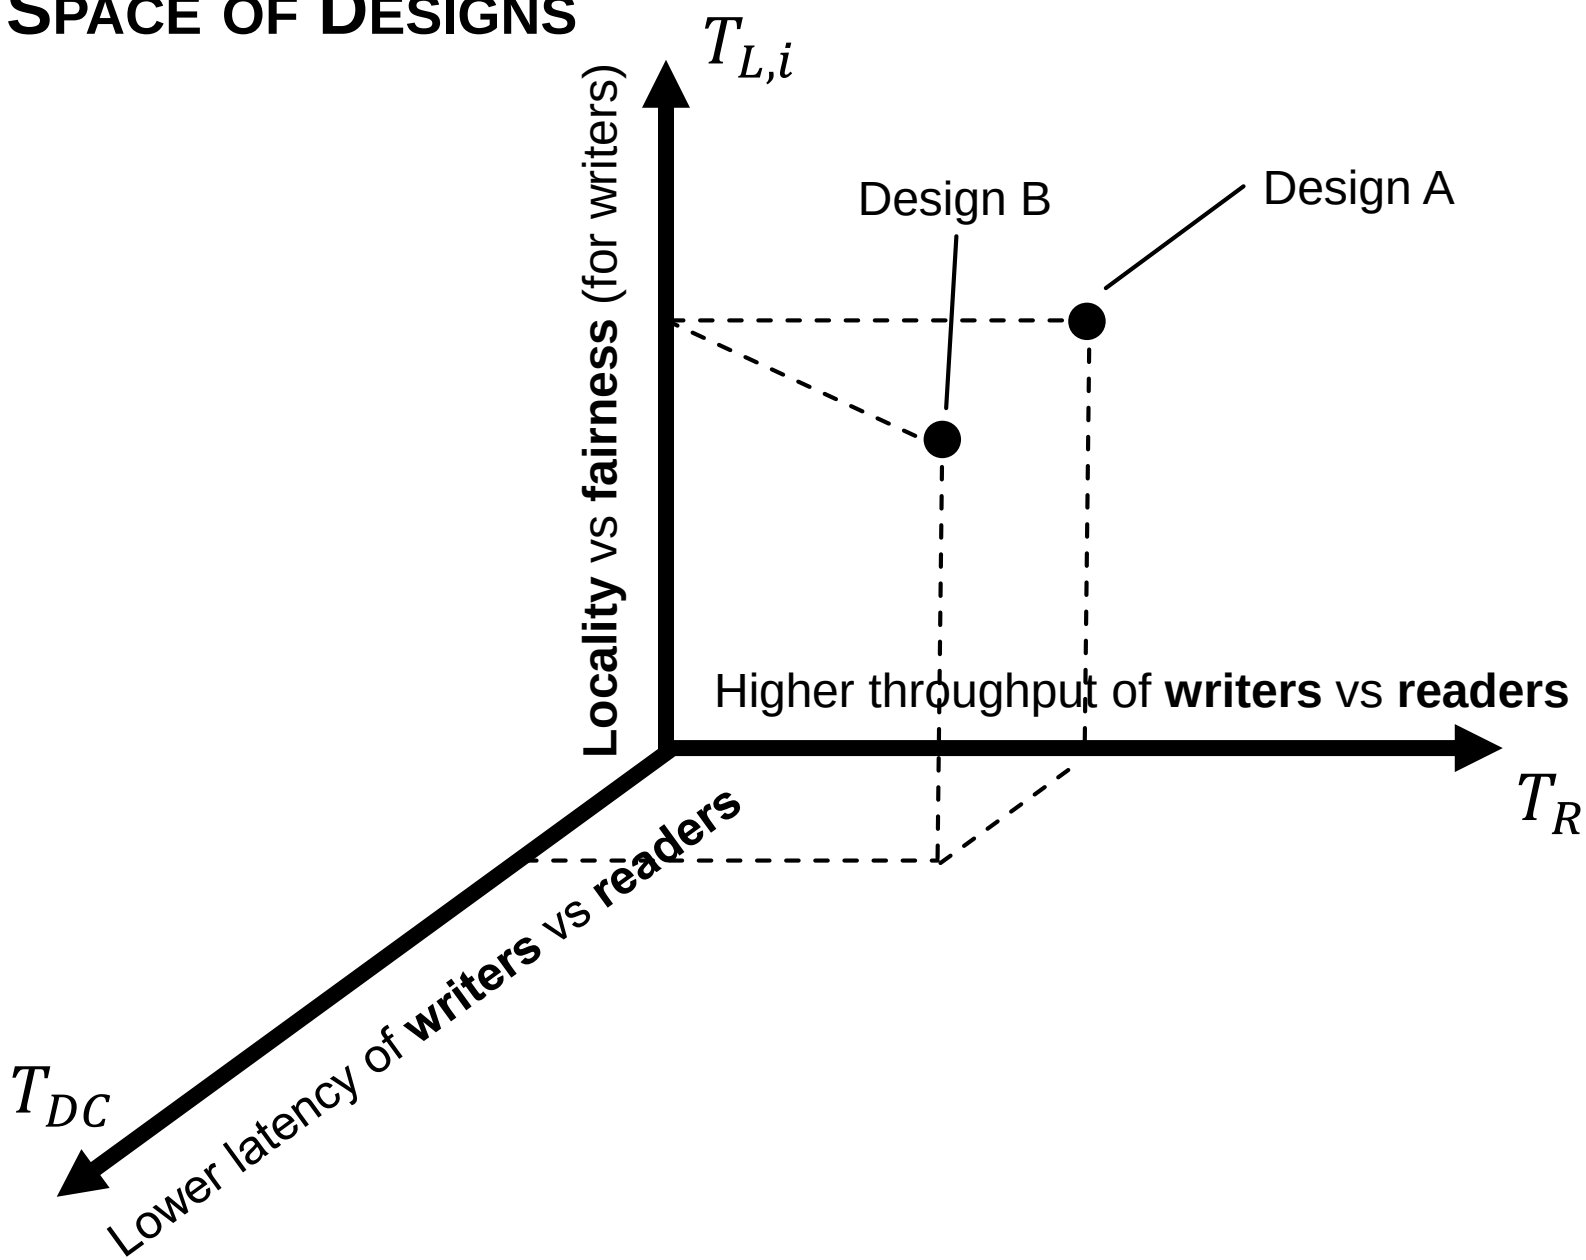

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

Lock Acquire by Readers

A lightweight acquire protocol for readers: only one atomic fetch-and-add (FAA) operation

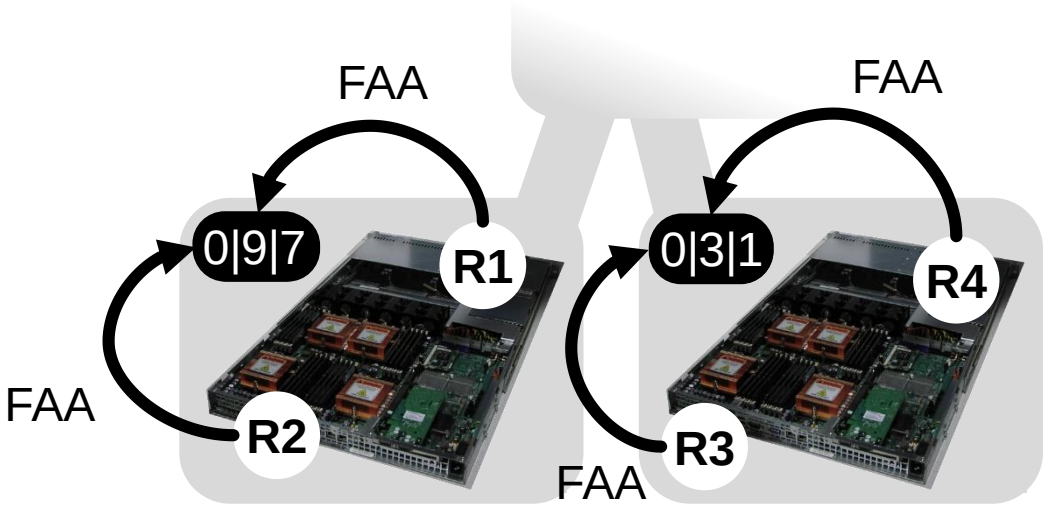

A writer holds the lock

b|x|y

Readers that arrived at the CS

Readers that left the CS

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

LOCK ACQUIRE BY WRITERS

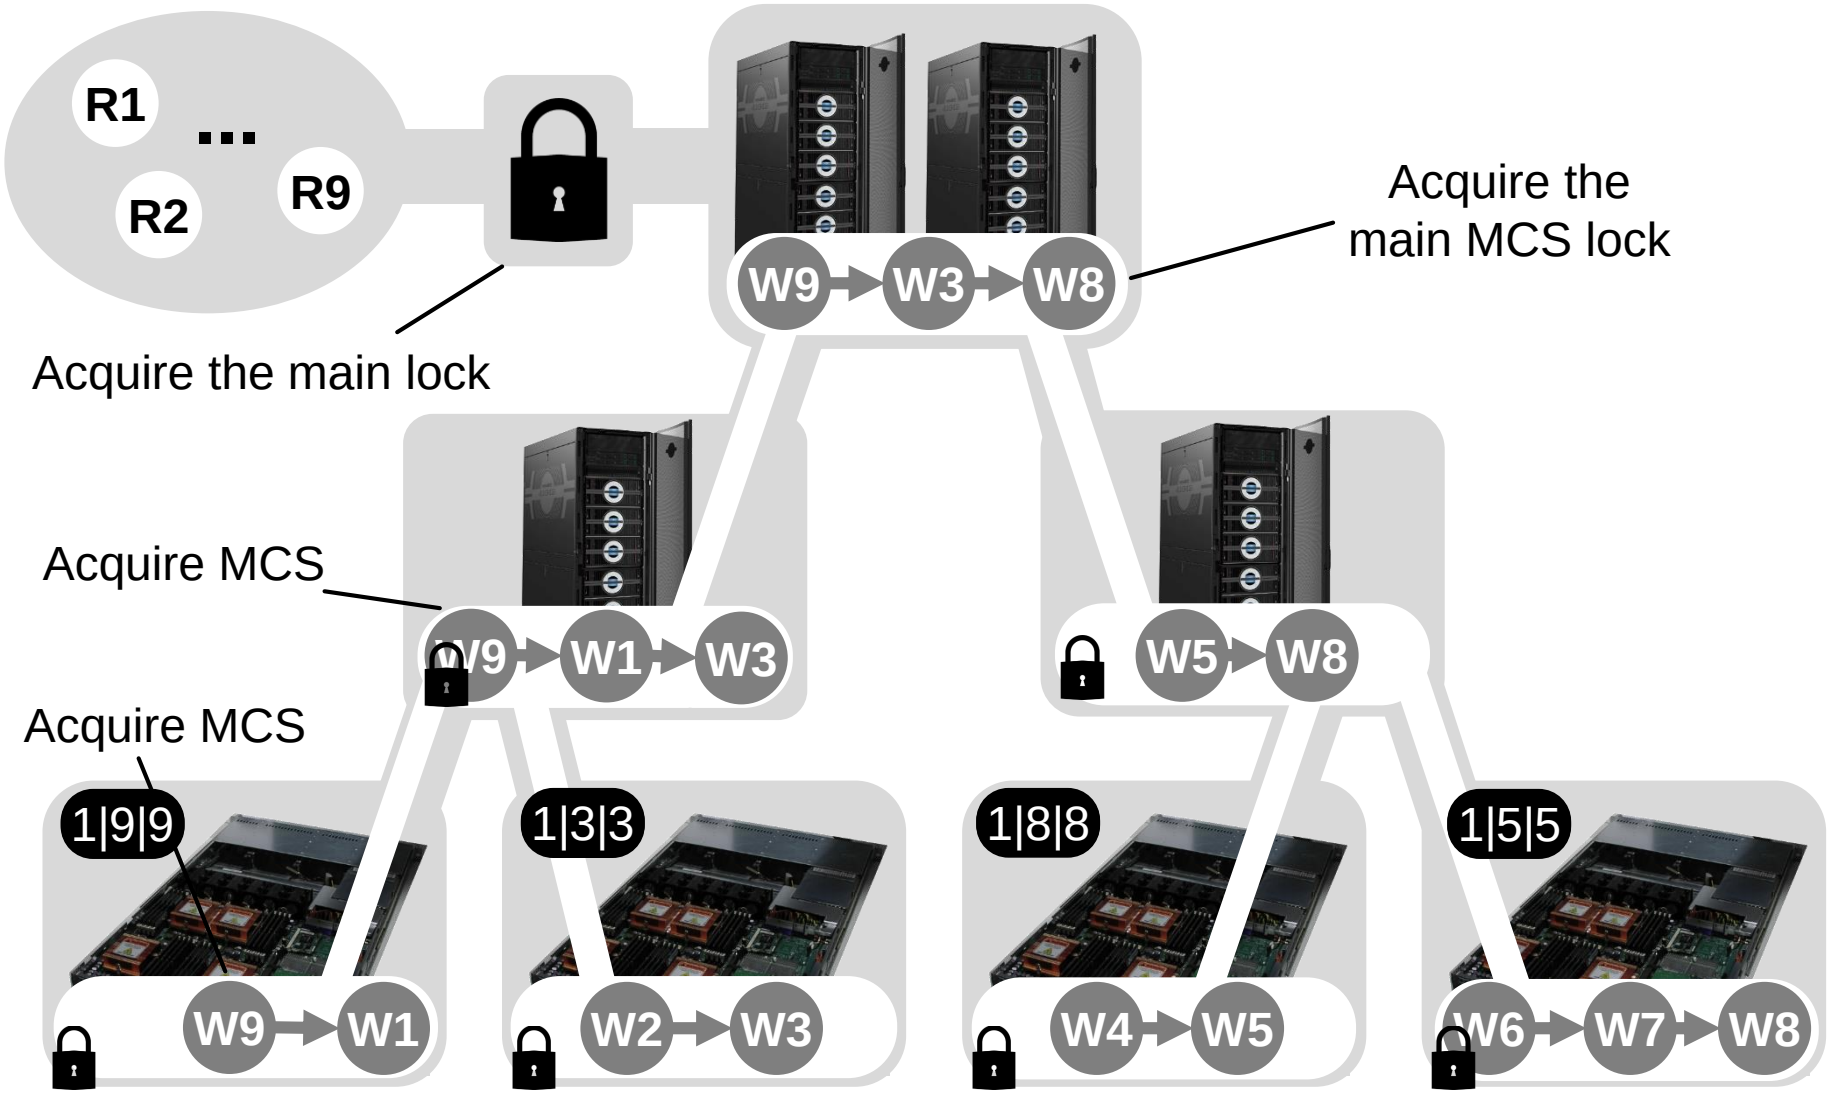

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

EVALUATION

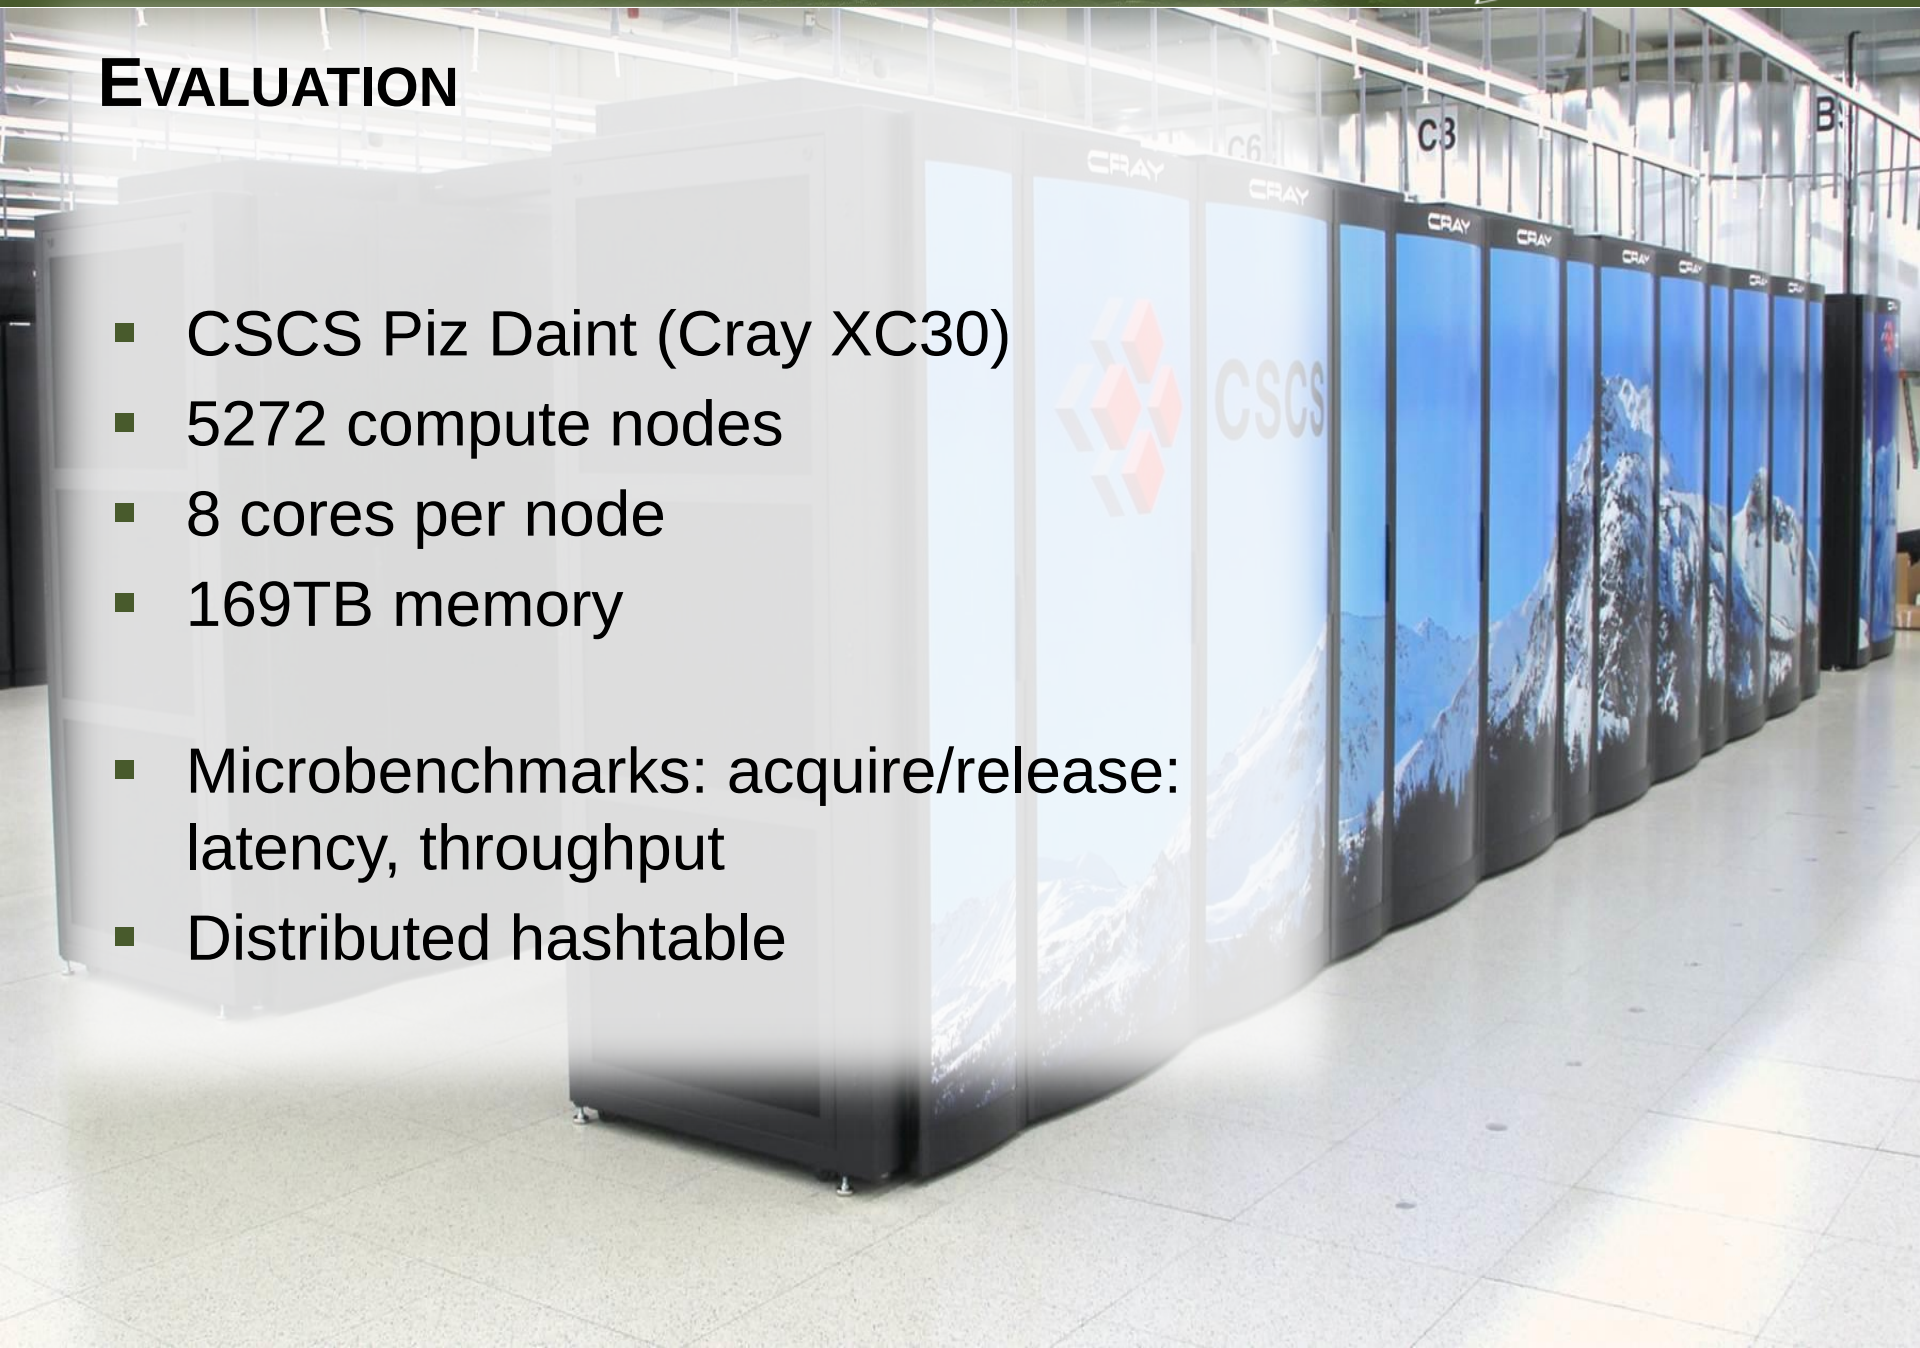

- CSCS Piz Daint (Cray XC30)
- 5272 compute nodes
- 8 cores per node
- 169TB memory

- Microbenchmarks: acquire/release: latency, throughput
- Distributed hashtable

EVALUATION

DISTRIBUTED COUNTER ANALYSIS

0|9|7 0|3|1

0|12|8

Throughput, 2% writers

Single-operation benchmark

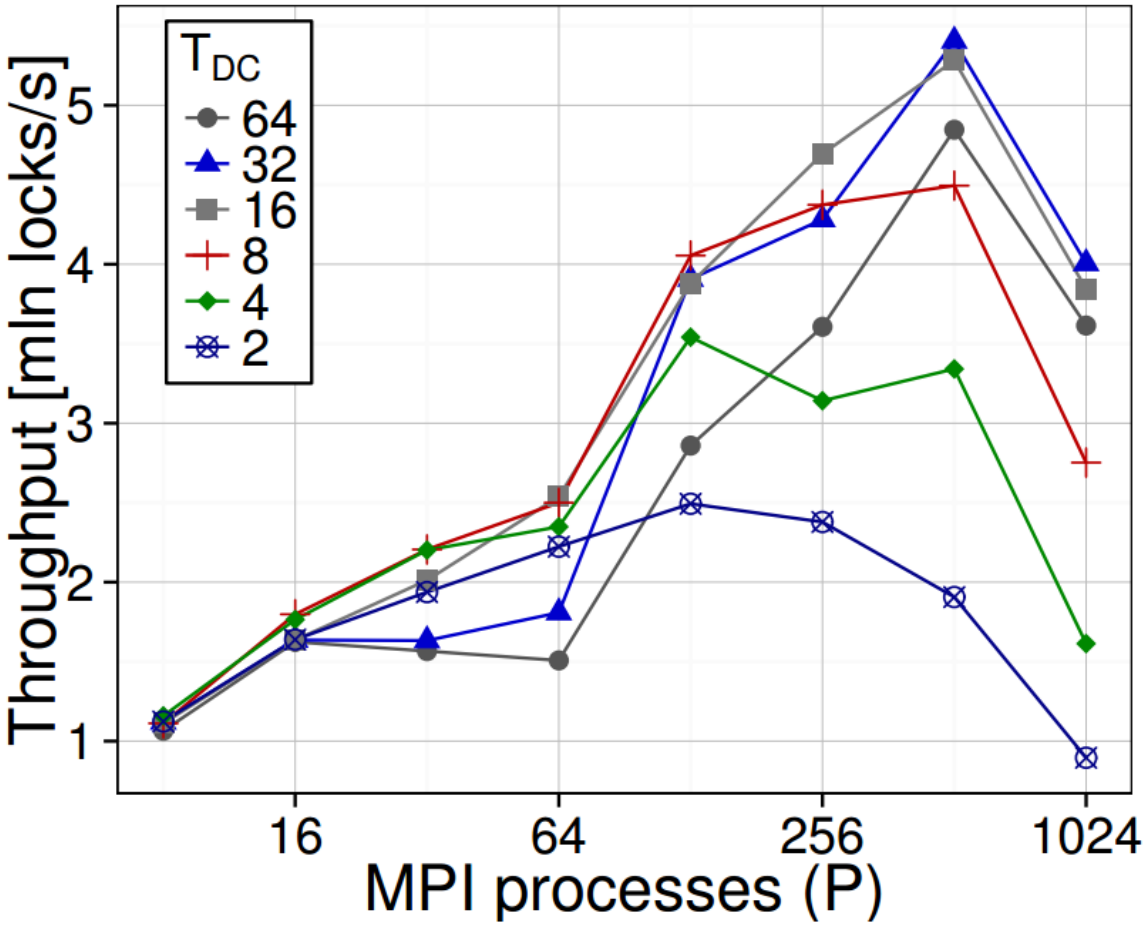

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

EVALUATION
READER THRESHOLD ANALYSIS

Throughput, 0.2% writers,
Empty-critical-section benchmark

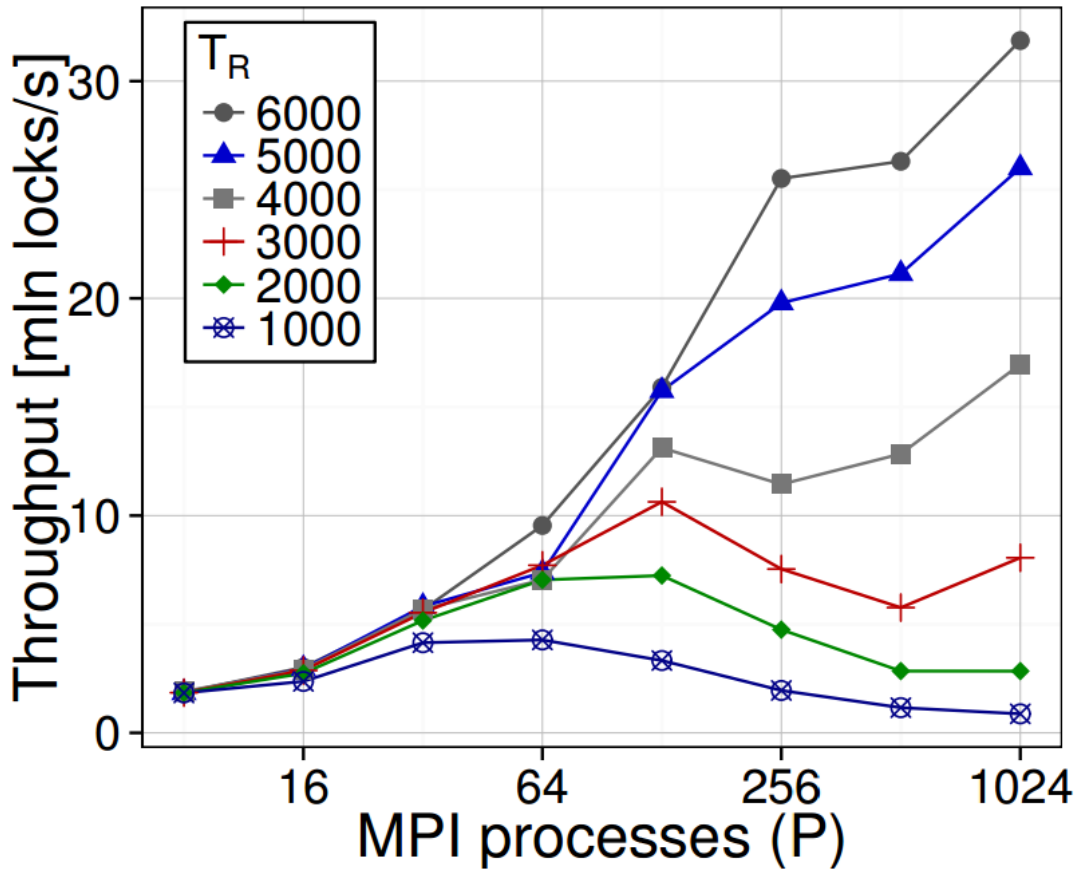

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

EVALUATION
COMPARISON TO THE STATE-OF-THE-ART

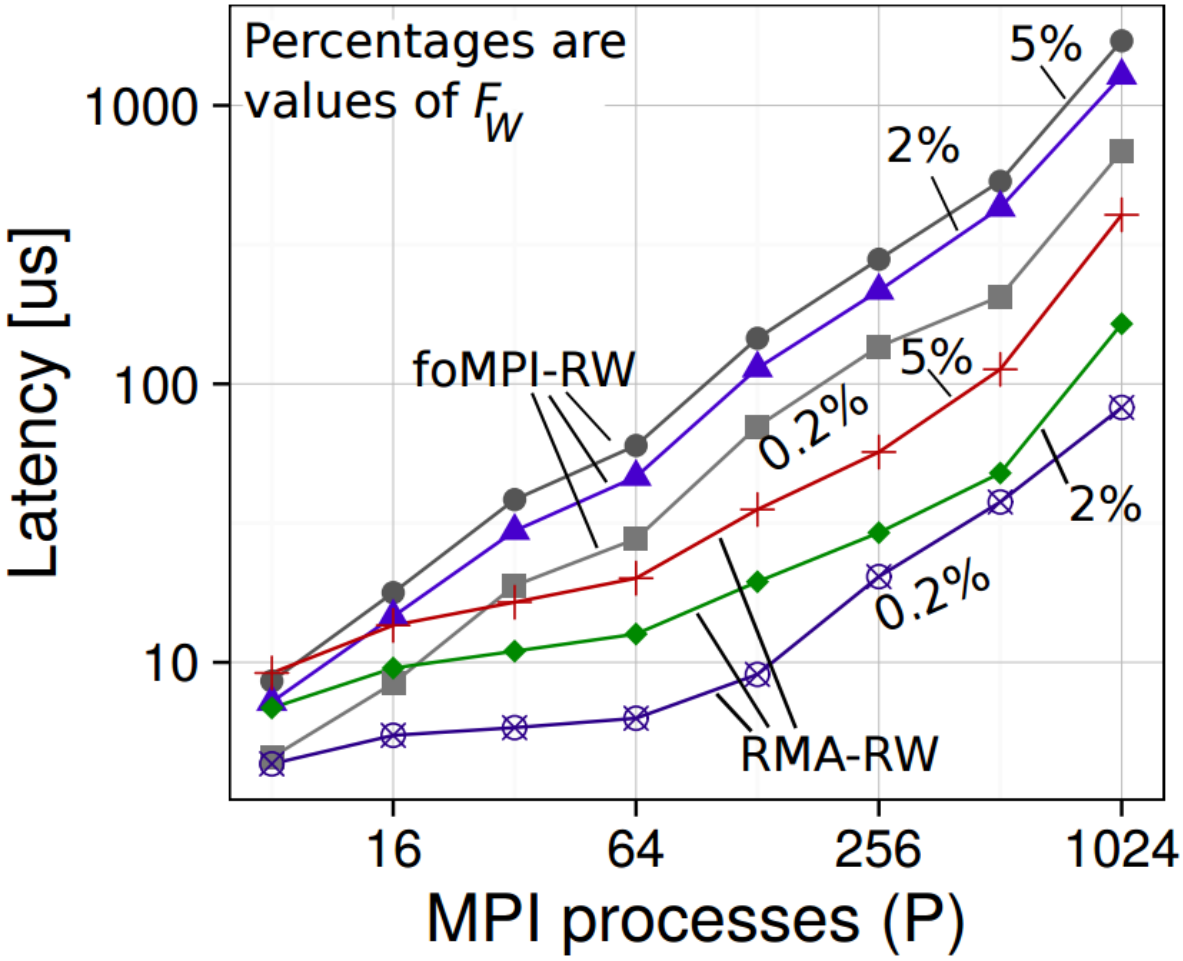

[1] R. Gerstenberger et al. Enabling Highly-scalable Remote Memory Access Programming with MPI-3 One Sided. ACM/IEEE Supercomputing 2013.

EVALUATION

COMPARISON TO THE STATE-OF-THE-ART

Throughput, single-operation benchmark

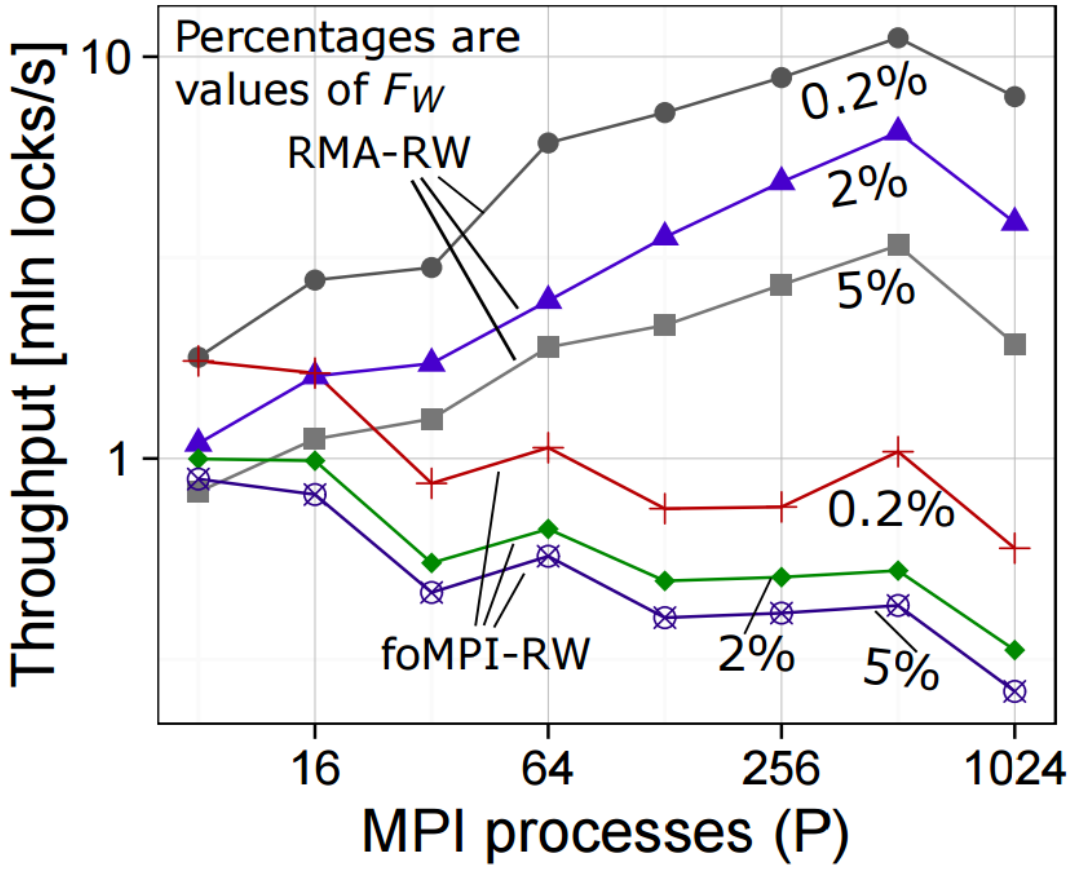

[1] R. Gerstenberger et al. Enabling Highly-scalable Remote Memory Access Programming with MPI-3 One Sided. ACM/IEEE Supercomputing 2013.

EVALUATION
DISTRIBUTED HASHTABLE

20% writers

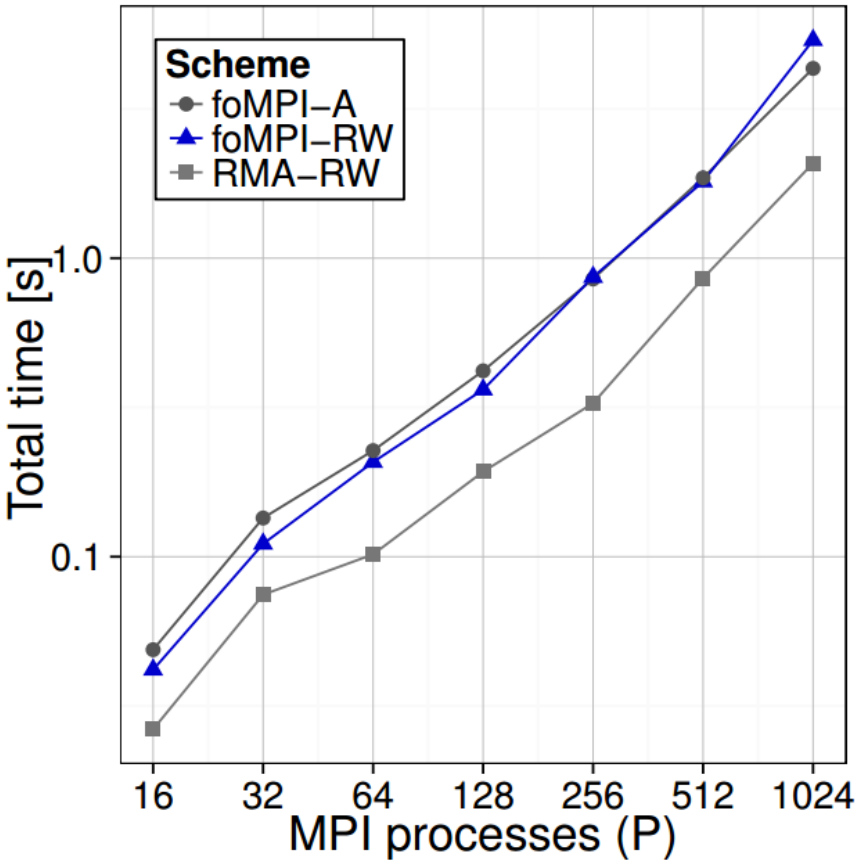

10% writers

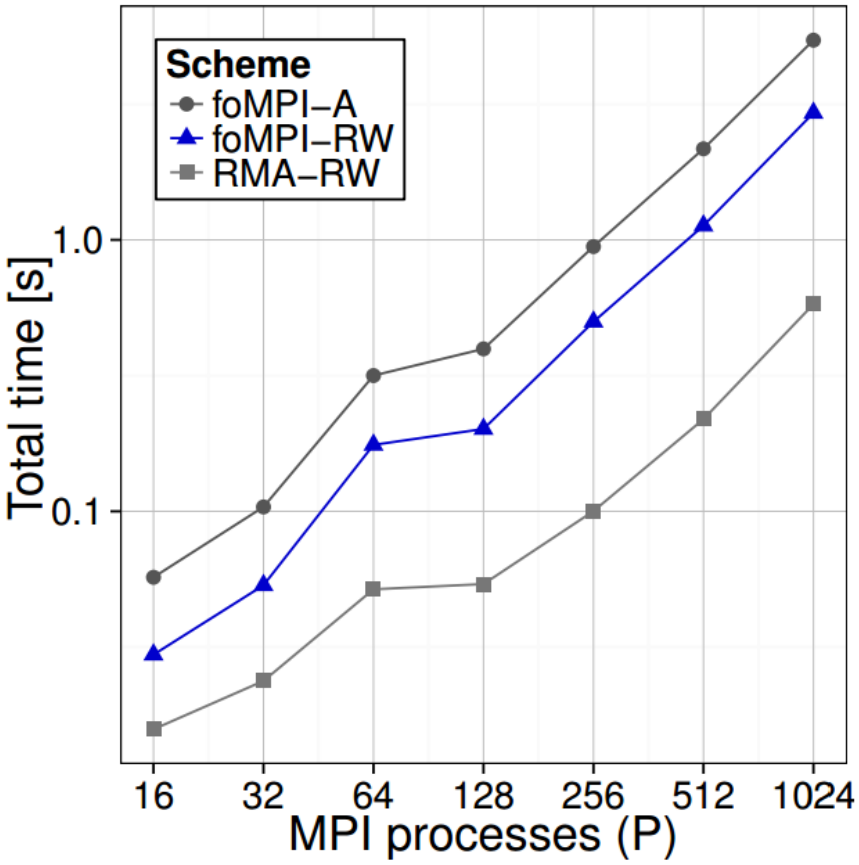

[1] R. Gerstenberger et al. Enabling Highly-scalable Remote Memory Access Programming with MPI-3 One Sided. ACM/IEEE Supercomputing 2013.

EVALUATION
DISTRIBUTED HASHTABLE

2% of writers

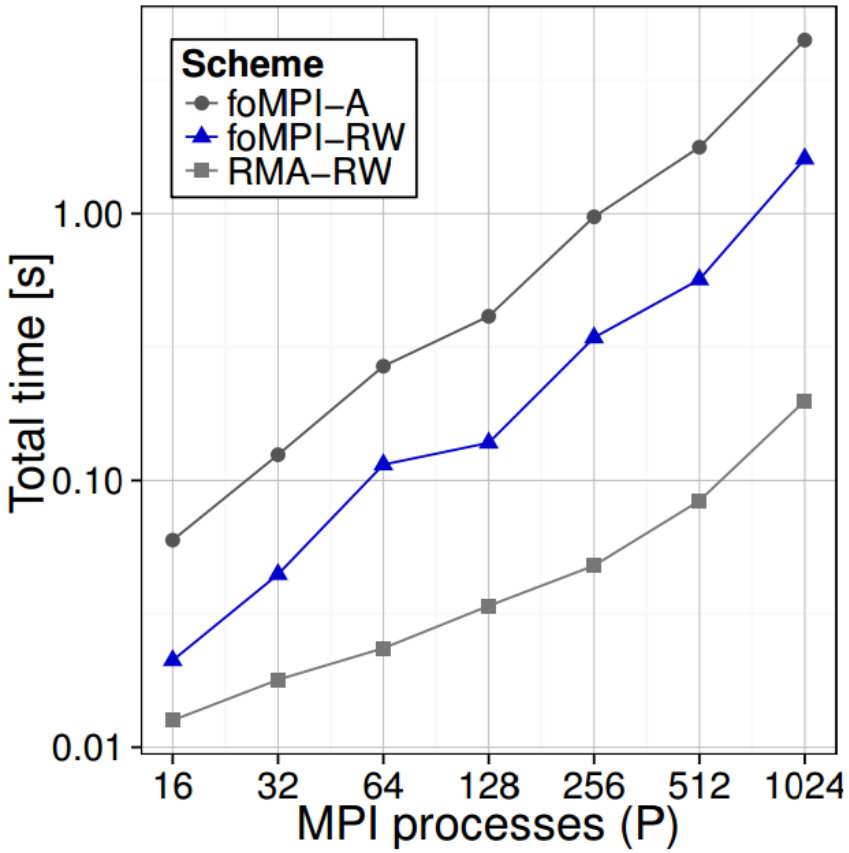

0% of writers

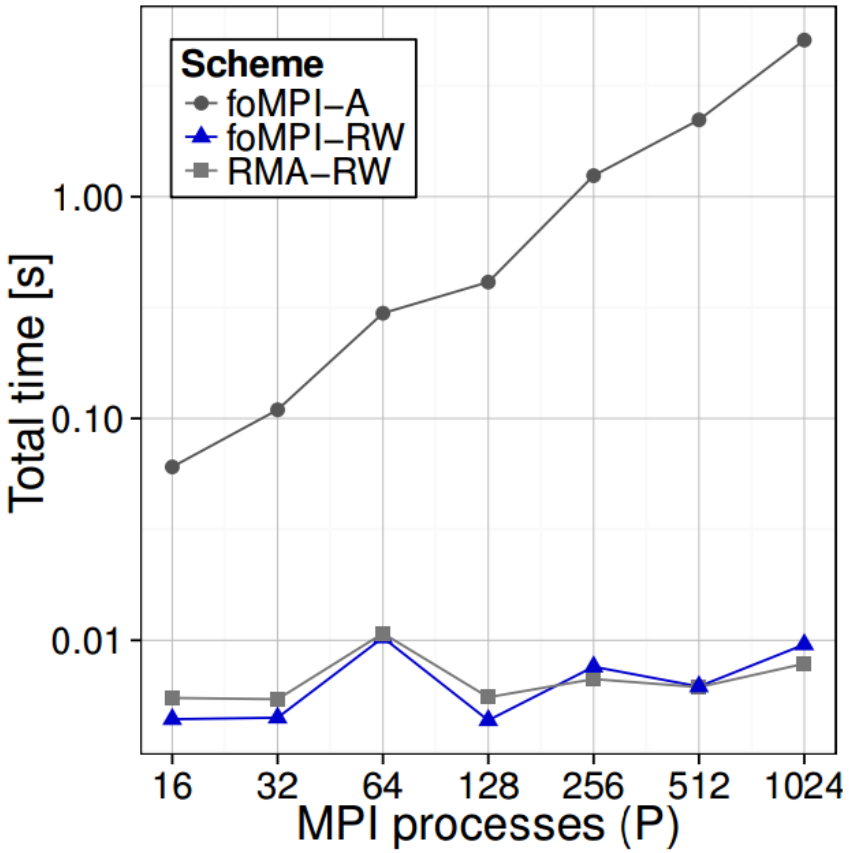

[1] R. Gerstenberger et al. Enabling Highly-scalable Remote Memory Access Programming with MPI-3 One Sided. ACM/IEEE Supercomputing 2013.

Another application area - Databases

- **MPI-RMA for distributed databases?**

Hash-Join

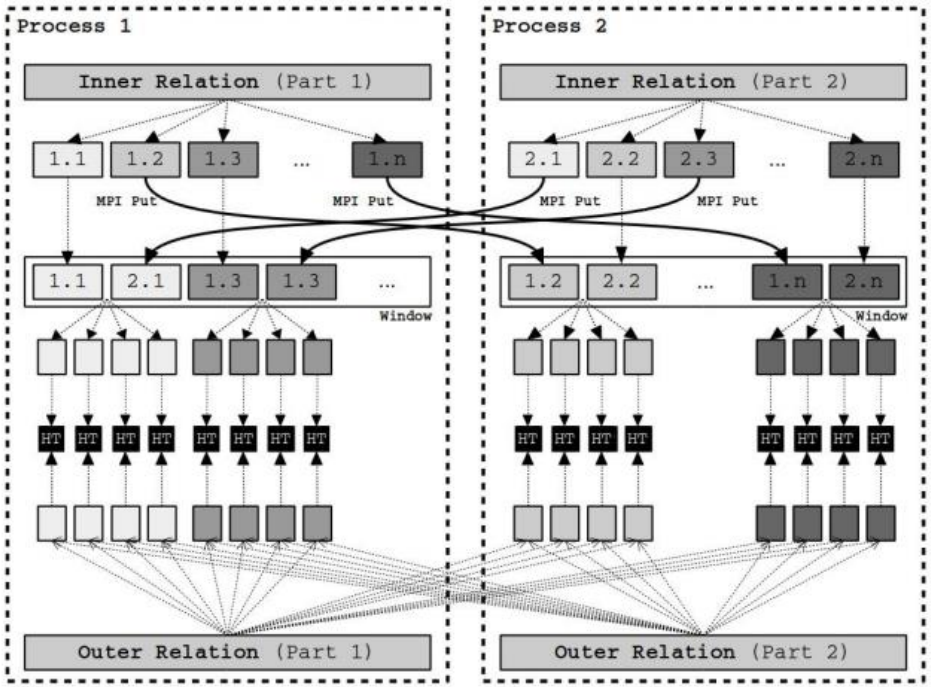

Sort-Join

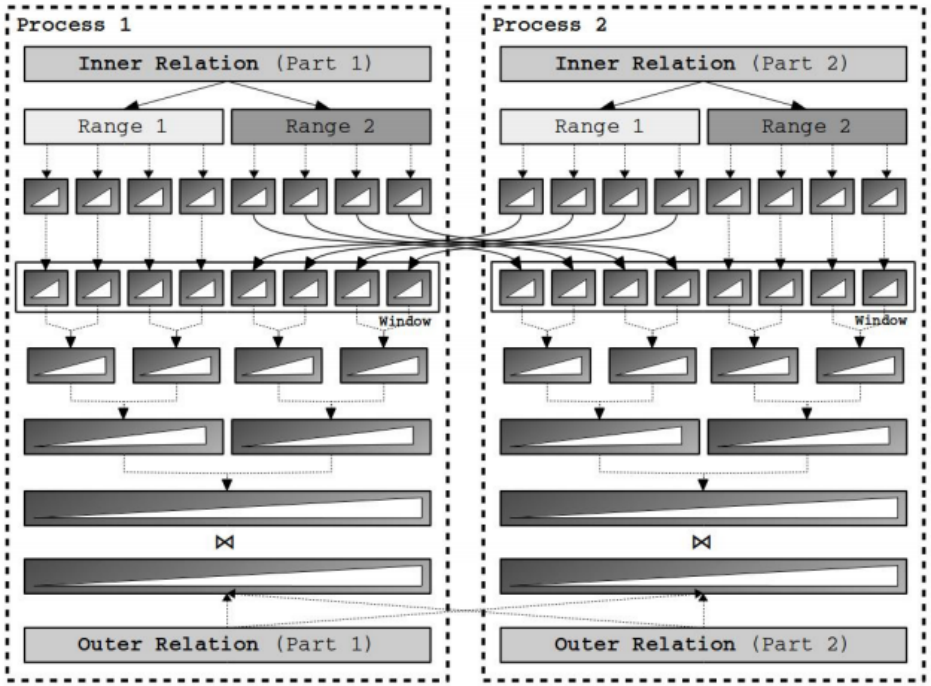

C. Barthels, et al., TH: Distributed Join Algorithms on Thousands of Cores presented in Munich, Germany, VLDB Endowment, Aug. 2017

Another application area - Databases

- **MPI-RMA for distributed databases on Piz Daint**

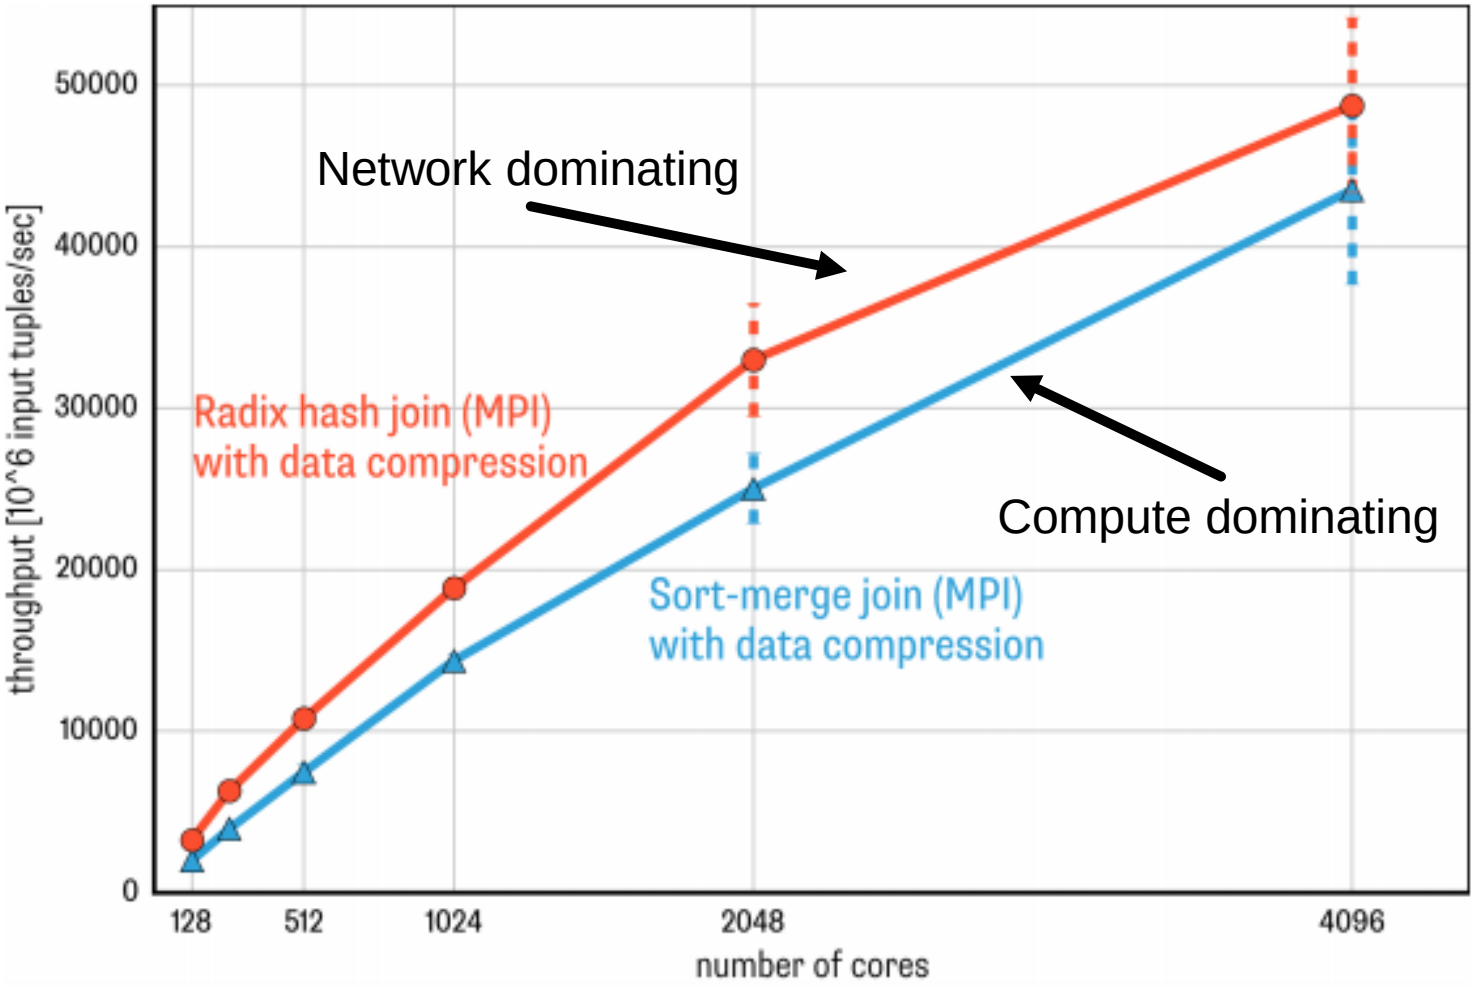

C. Barthels, et al., TH: Distributed Join Algorithms on Thousands of Cores presented in Munich, Germany, VLDB Endowment, Aug. 2017

Another application area - Databases

- **MPI-RMA for distributed databases on Piz Daint**

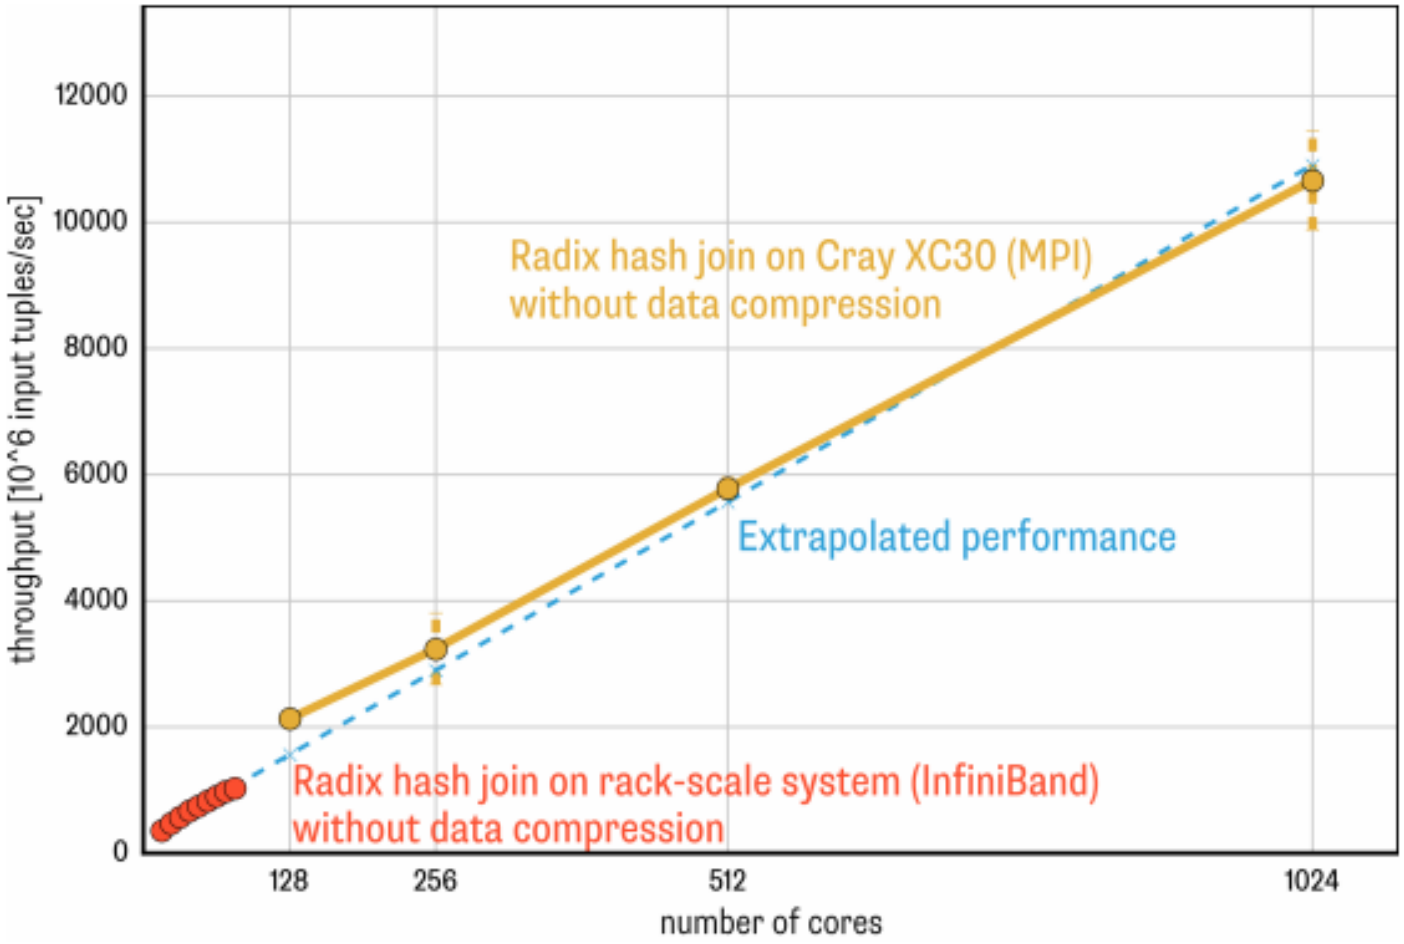

C. Barthels, et al., TH: Distributed Join Algorithms on Thousands of Cores presented in Munich, Germany, VLDB Endowment, Aug. 2017

OTHER ANALYSES

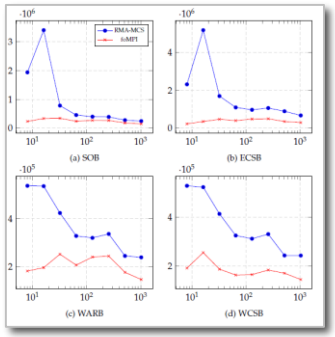

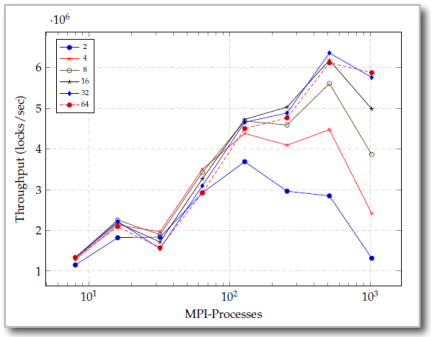

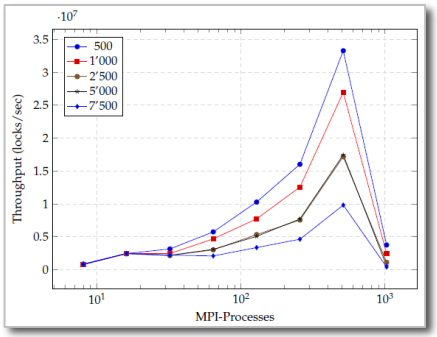

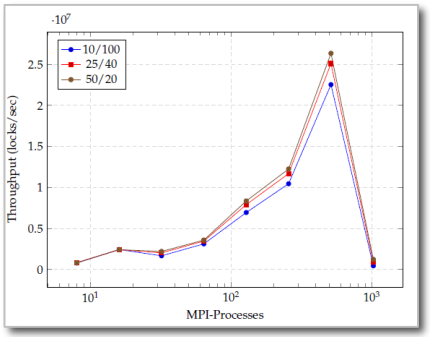

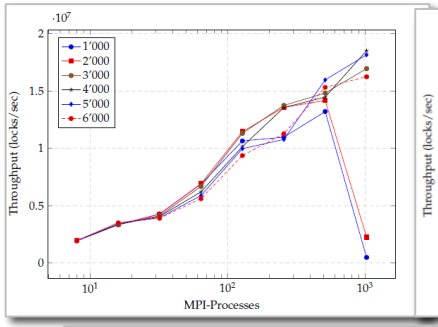

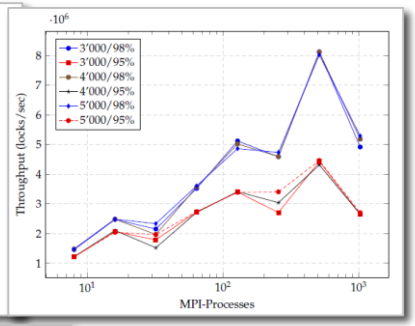

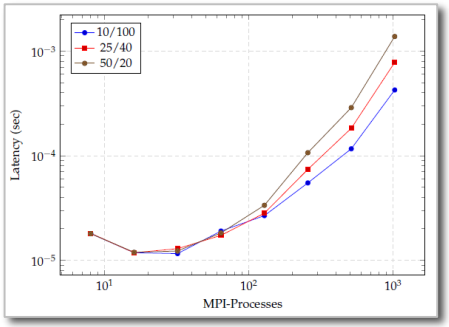

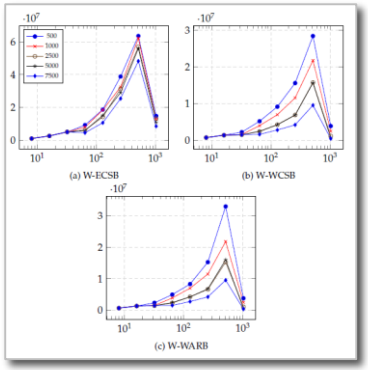

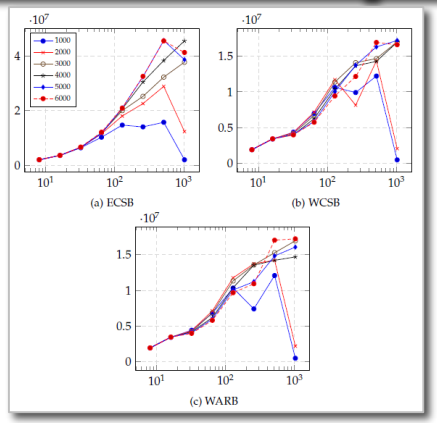

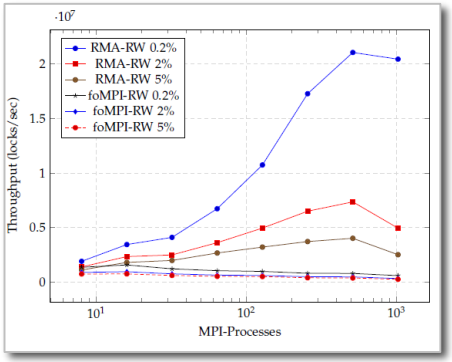

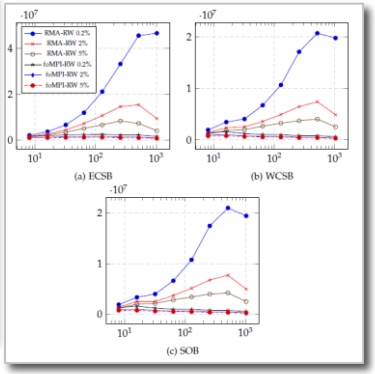

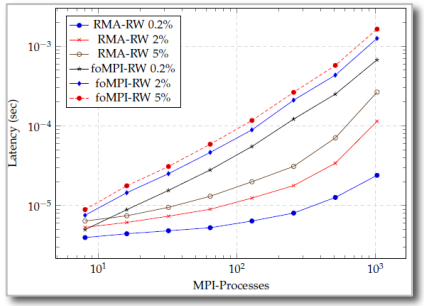

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper

CONCLUSIONS

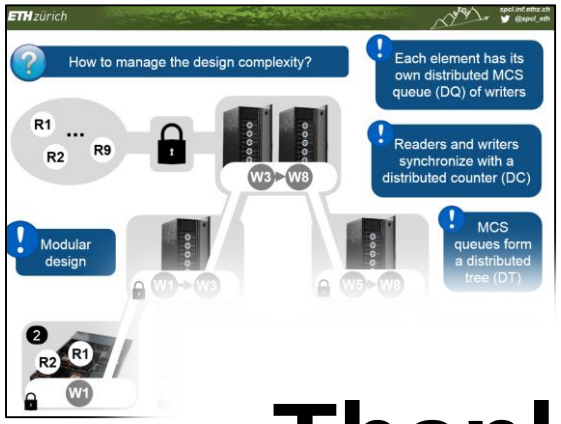

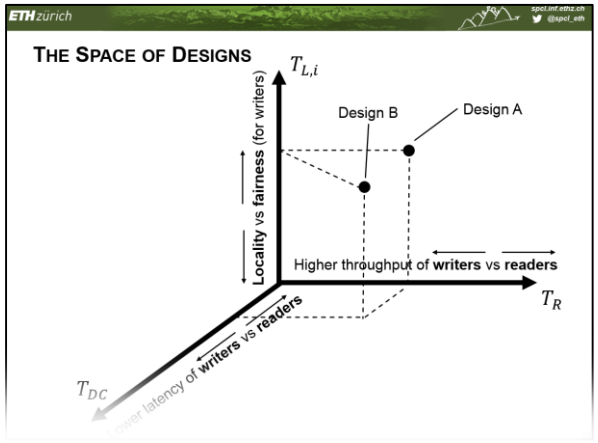

Modular d… …ible
corre… …es

Thank you for your attention

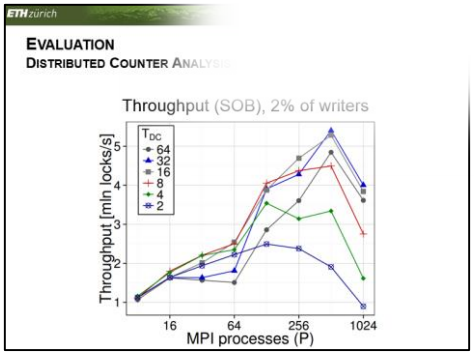

Improves latency and throughput over state-of-the-art

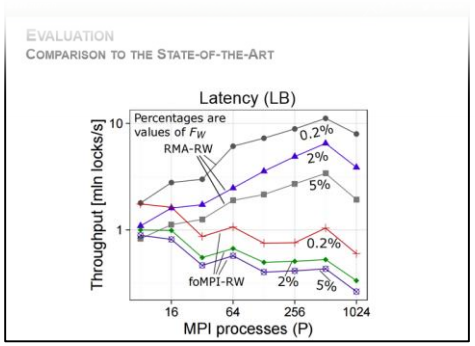

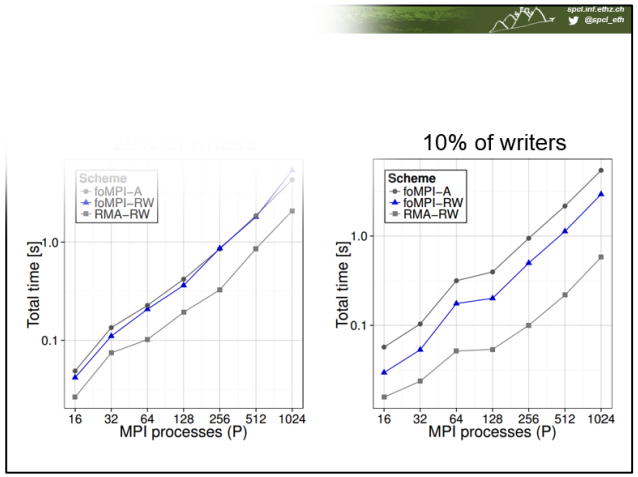

Accelerates distributed hashtabled

P. Schmid, M. Besta, TH: High-Performance Distributed RMA Locks, ACM HPDC'16, best paper