

Hardware Read-Write Lock Elision

Pascal Felber





Shady Issa
Paolo Romano




Alexander
Matveev




Multicores are everywhere



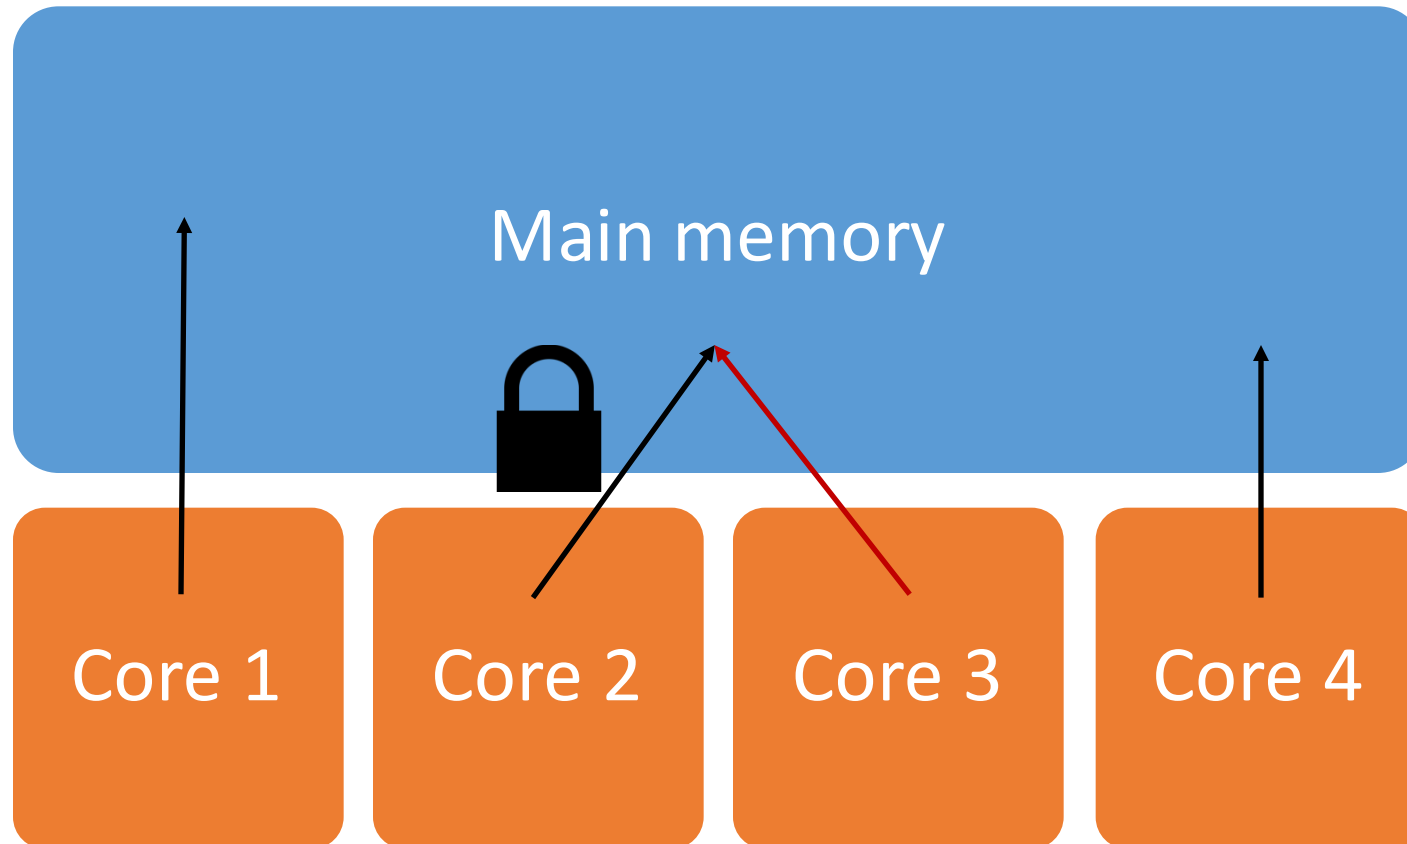
Parallel programming



Parallel programming

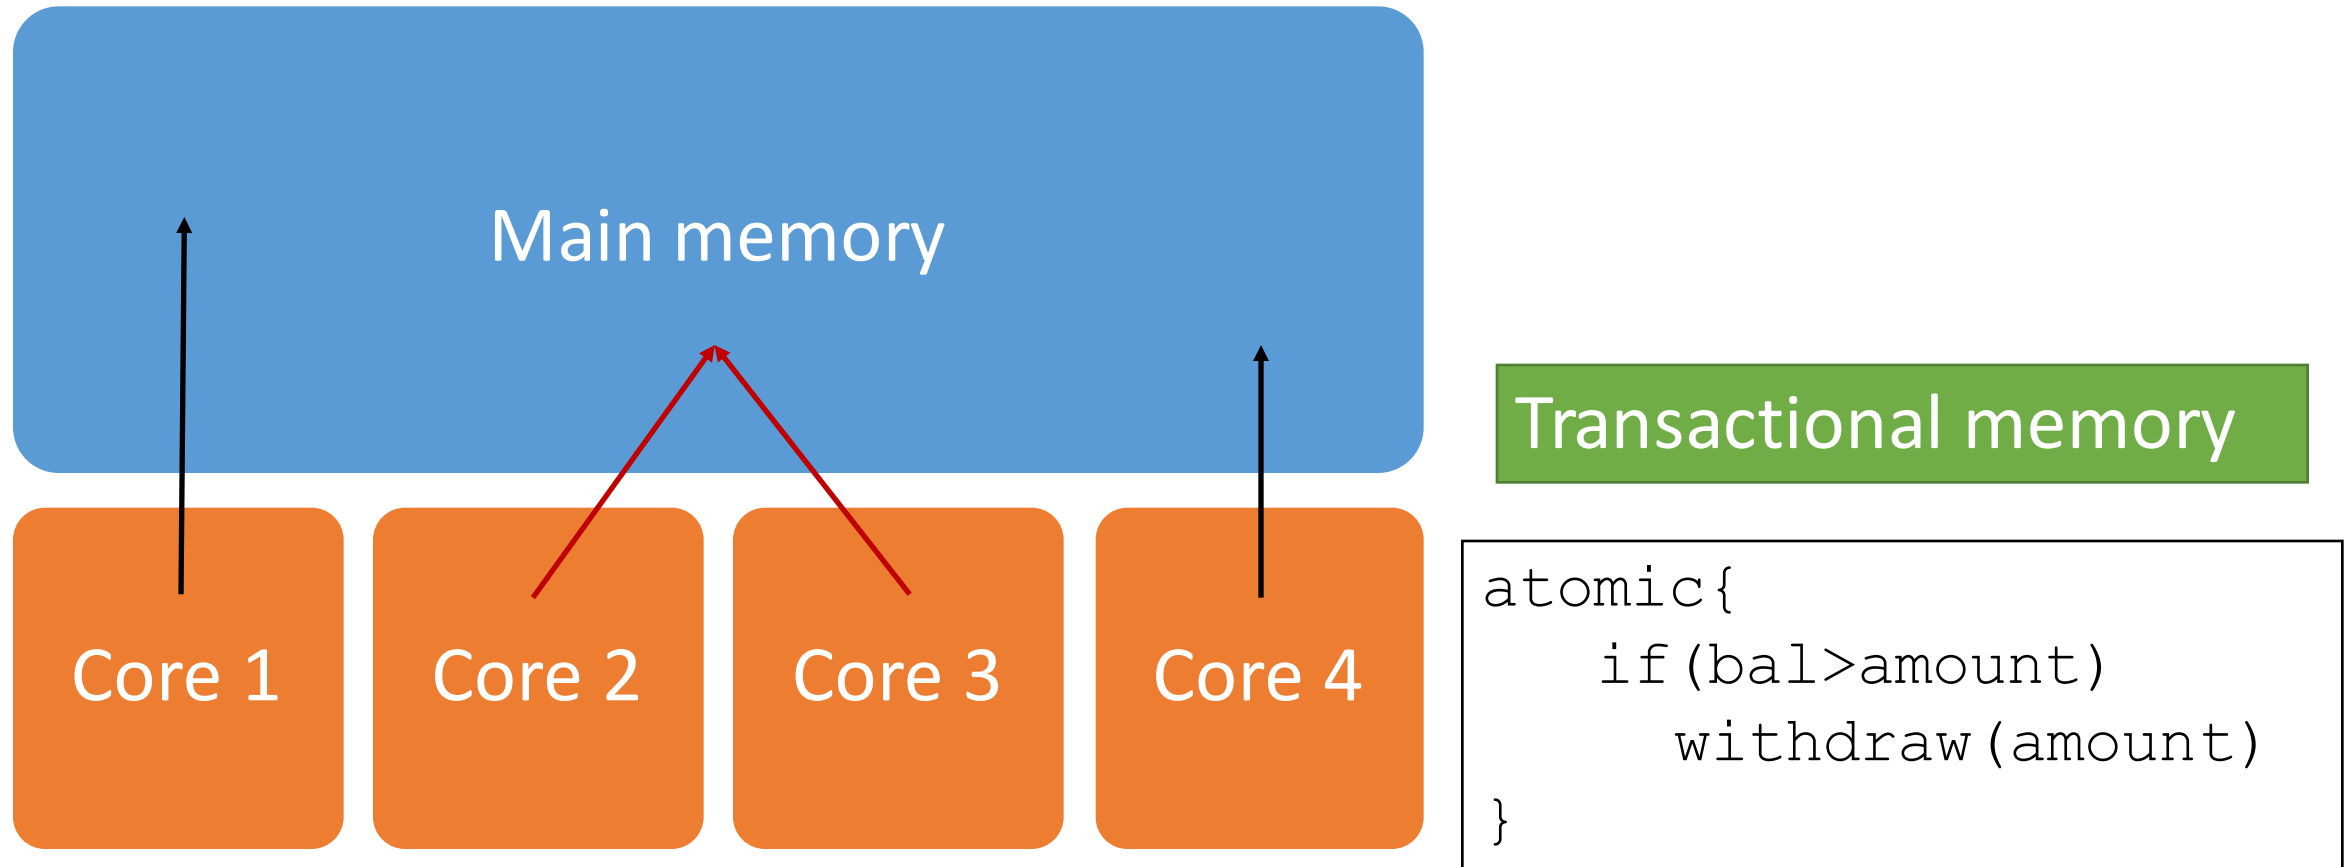


Parallel programming




- complexity
- deadlocks
 - livelocks
 - priority inversions
 - convoy effects


Parallel programming



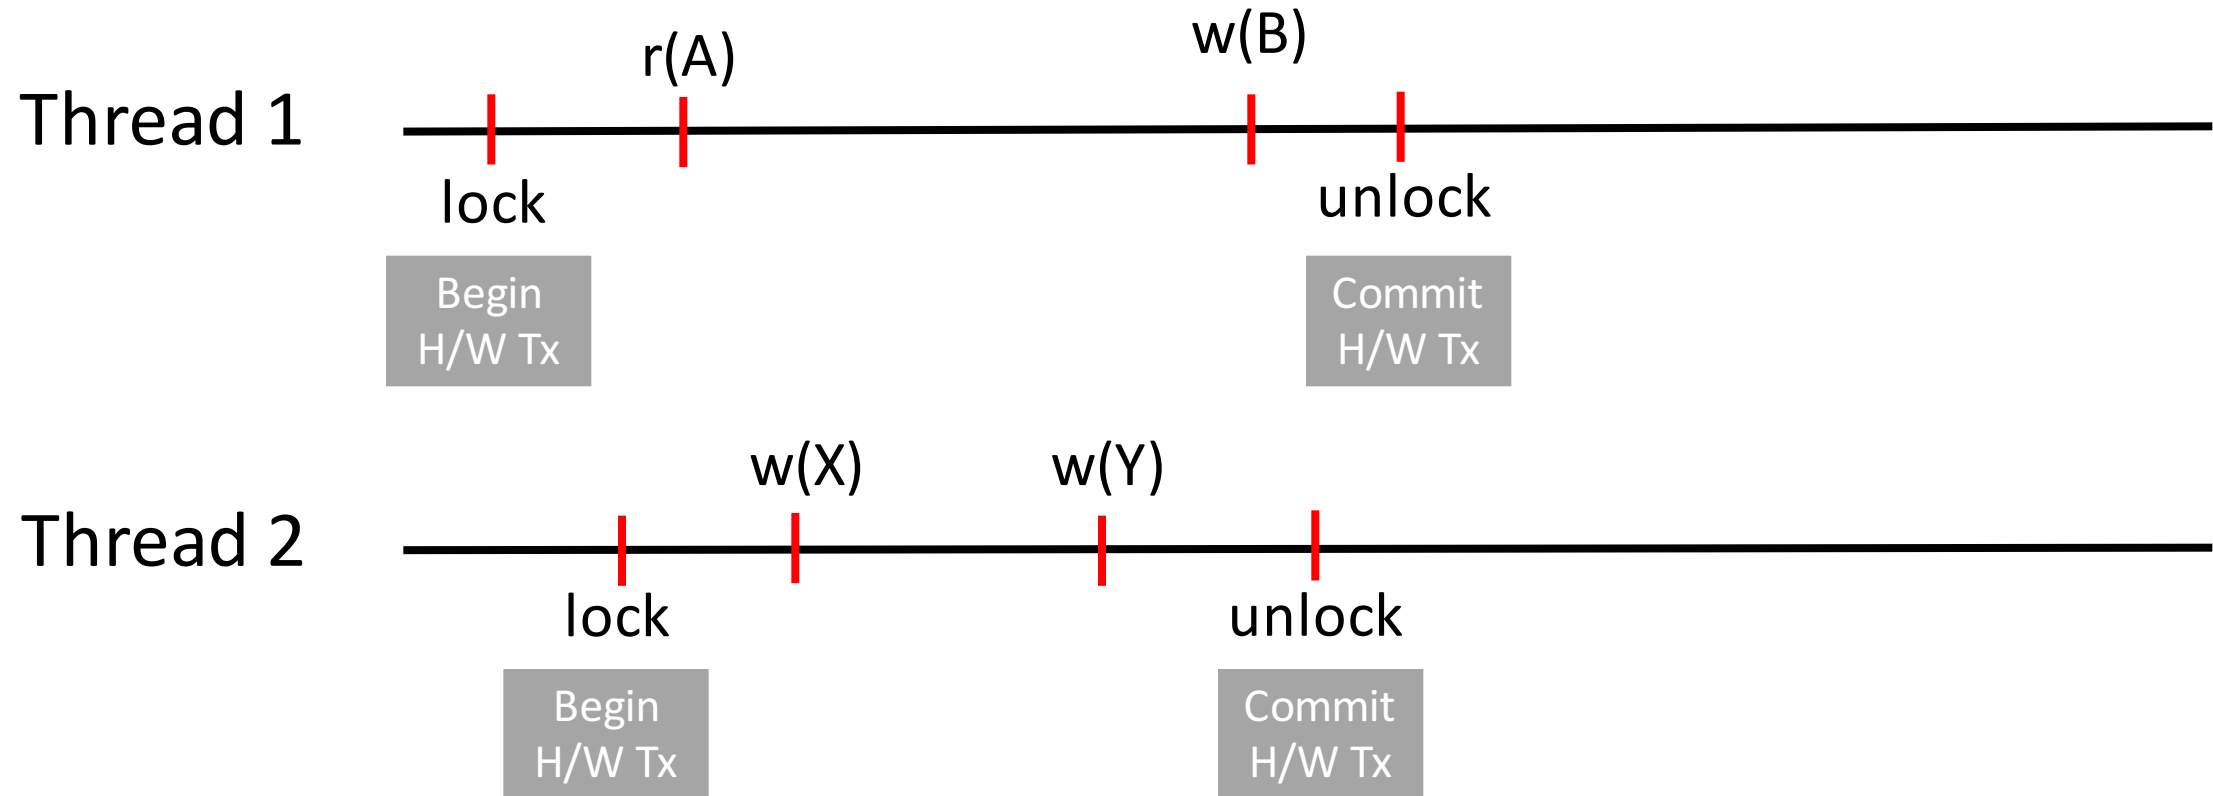
Hardware Transactional Memory



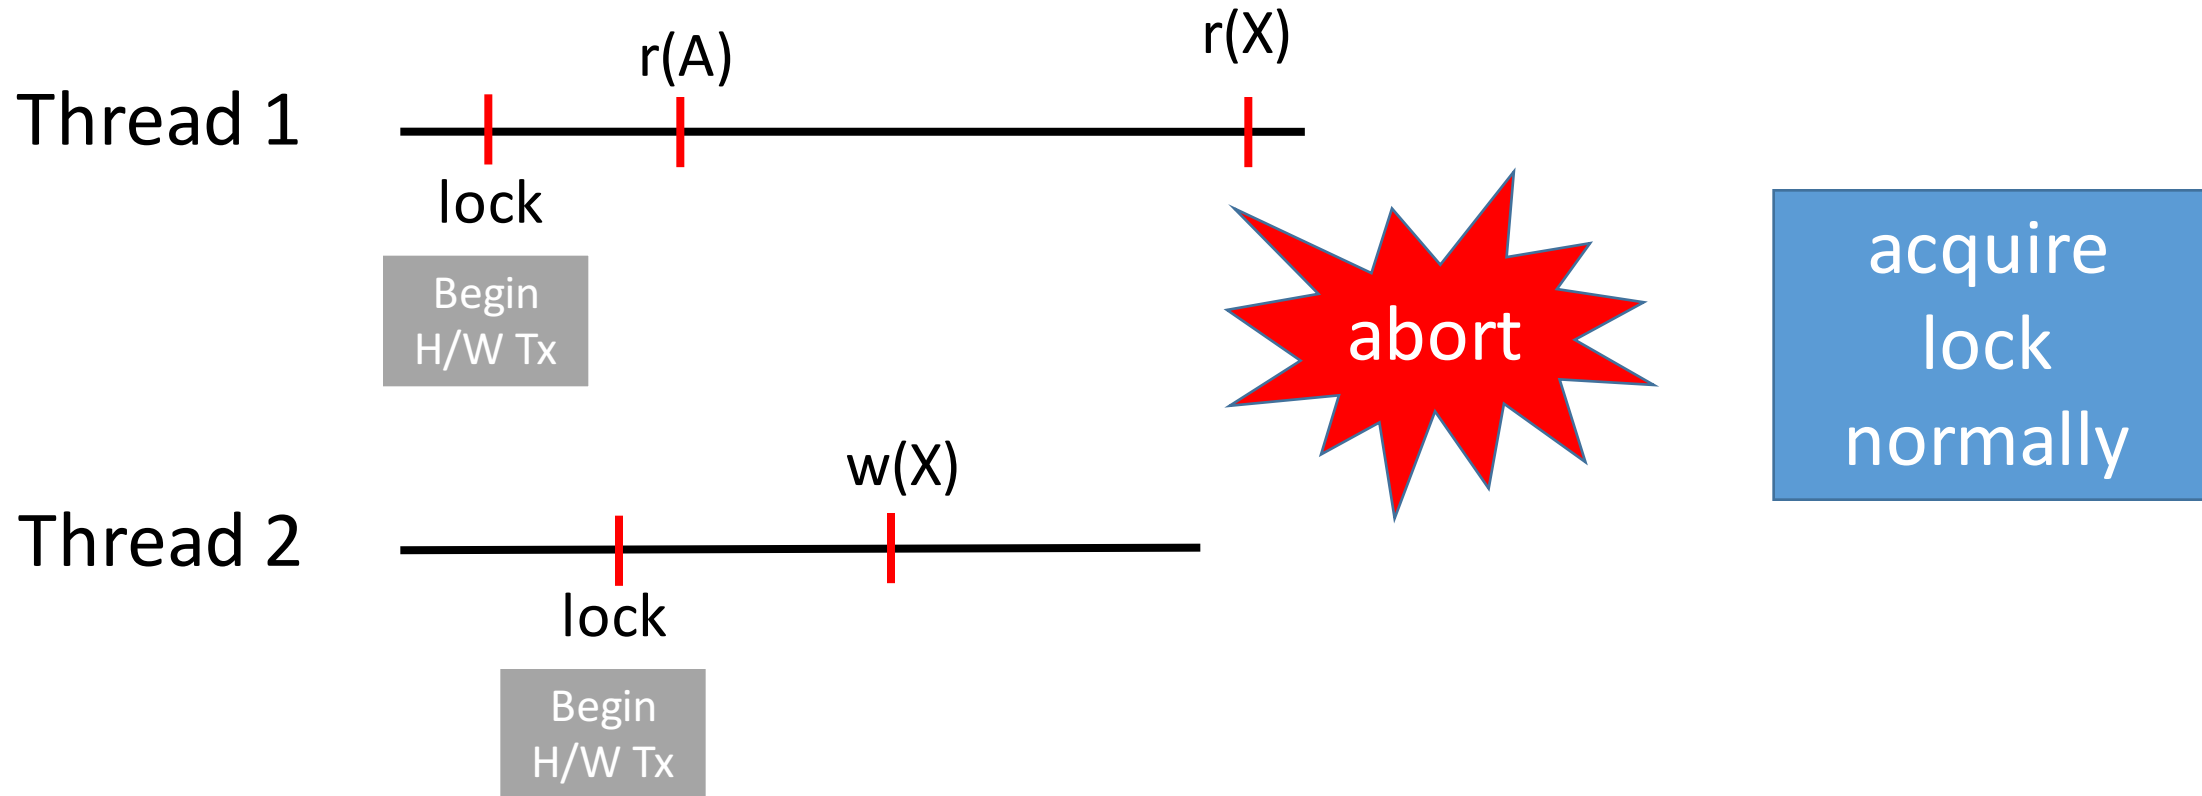
Hardware lock elision



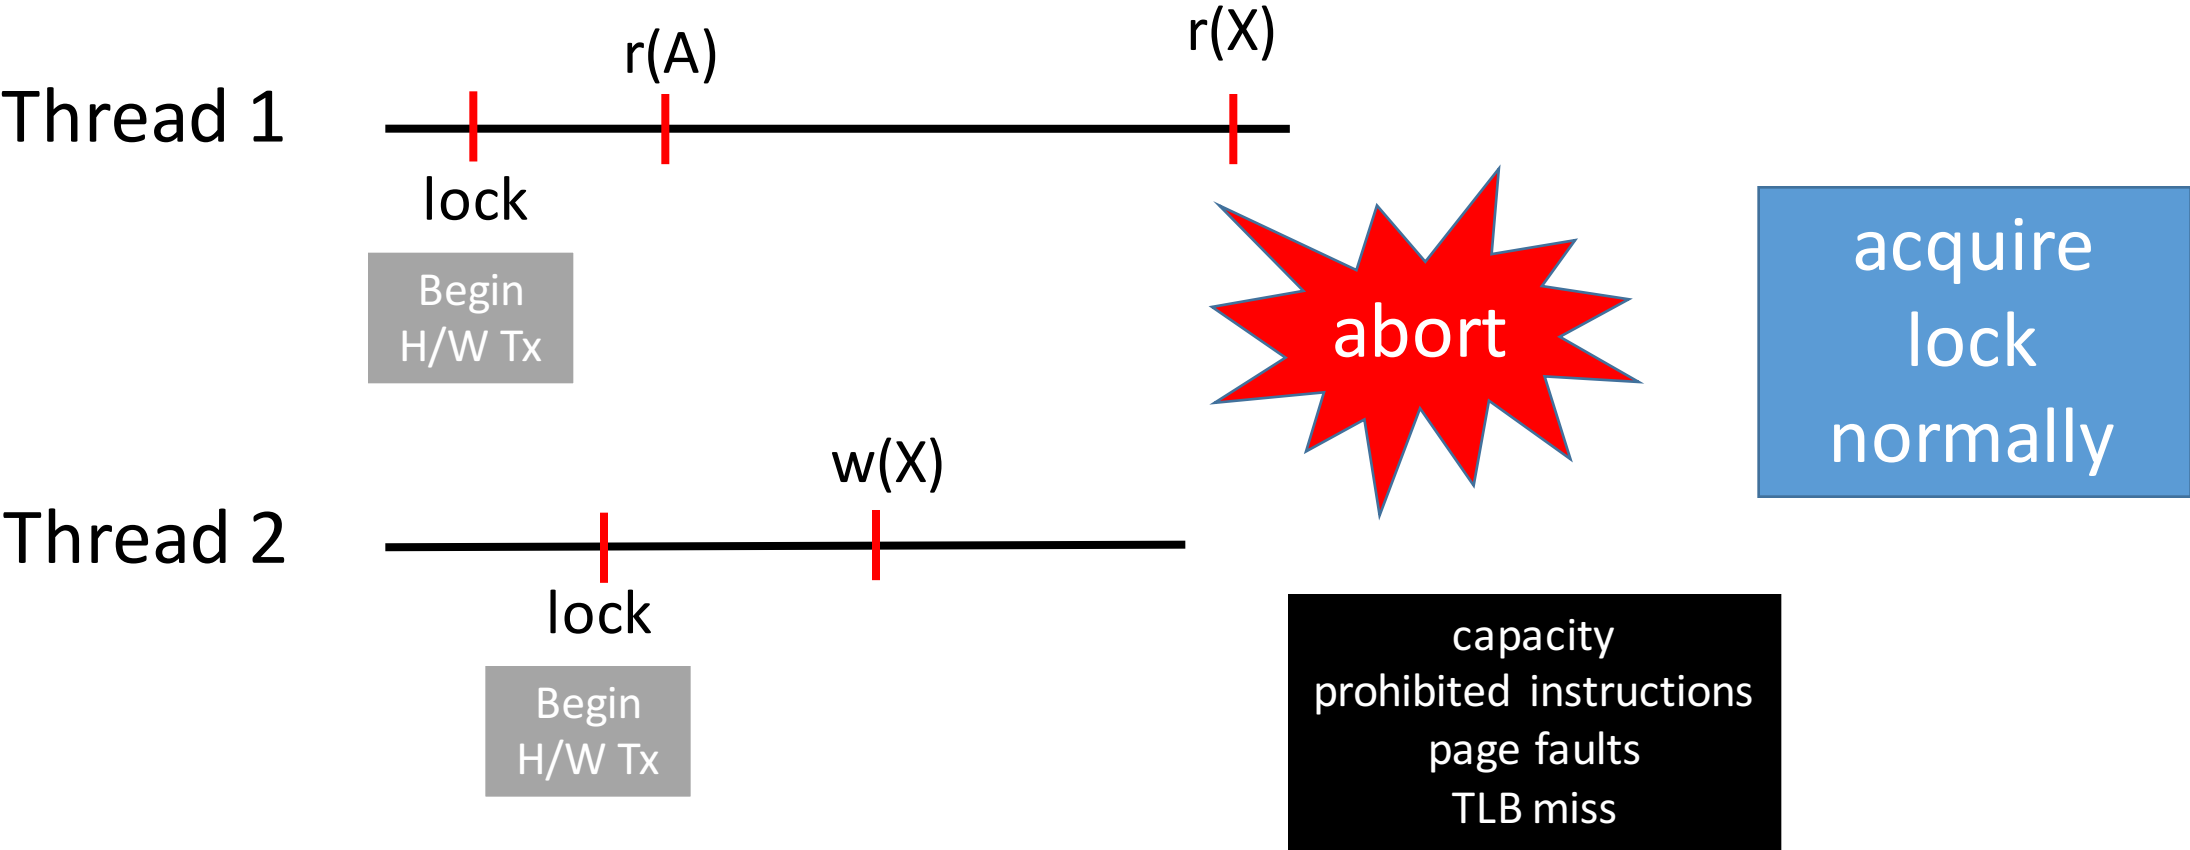
Hardware lock elision



Hardware lock elision



Hardware lock elision



Read-write Locks

read mode

✓ concurrent readers

write mode

✗ sequential writers
✗ blocks readers

read dominated workloads

Hardware read-write lock elision


readers run **without**
instrumentation:

- no H/W TxS
- no S/W tracking of read locations
- no lock acquisition

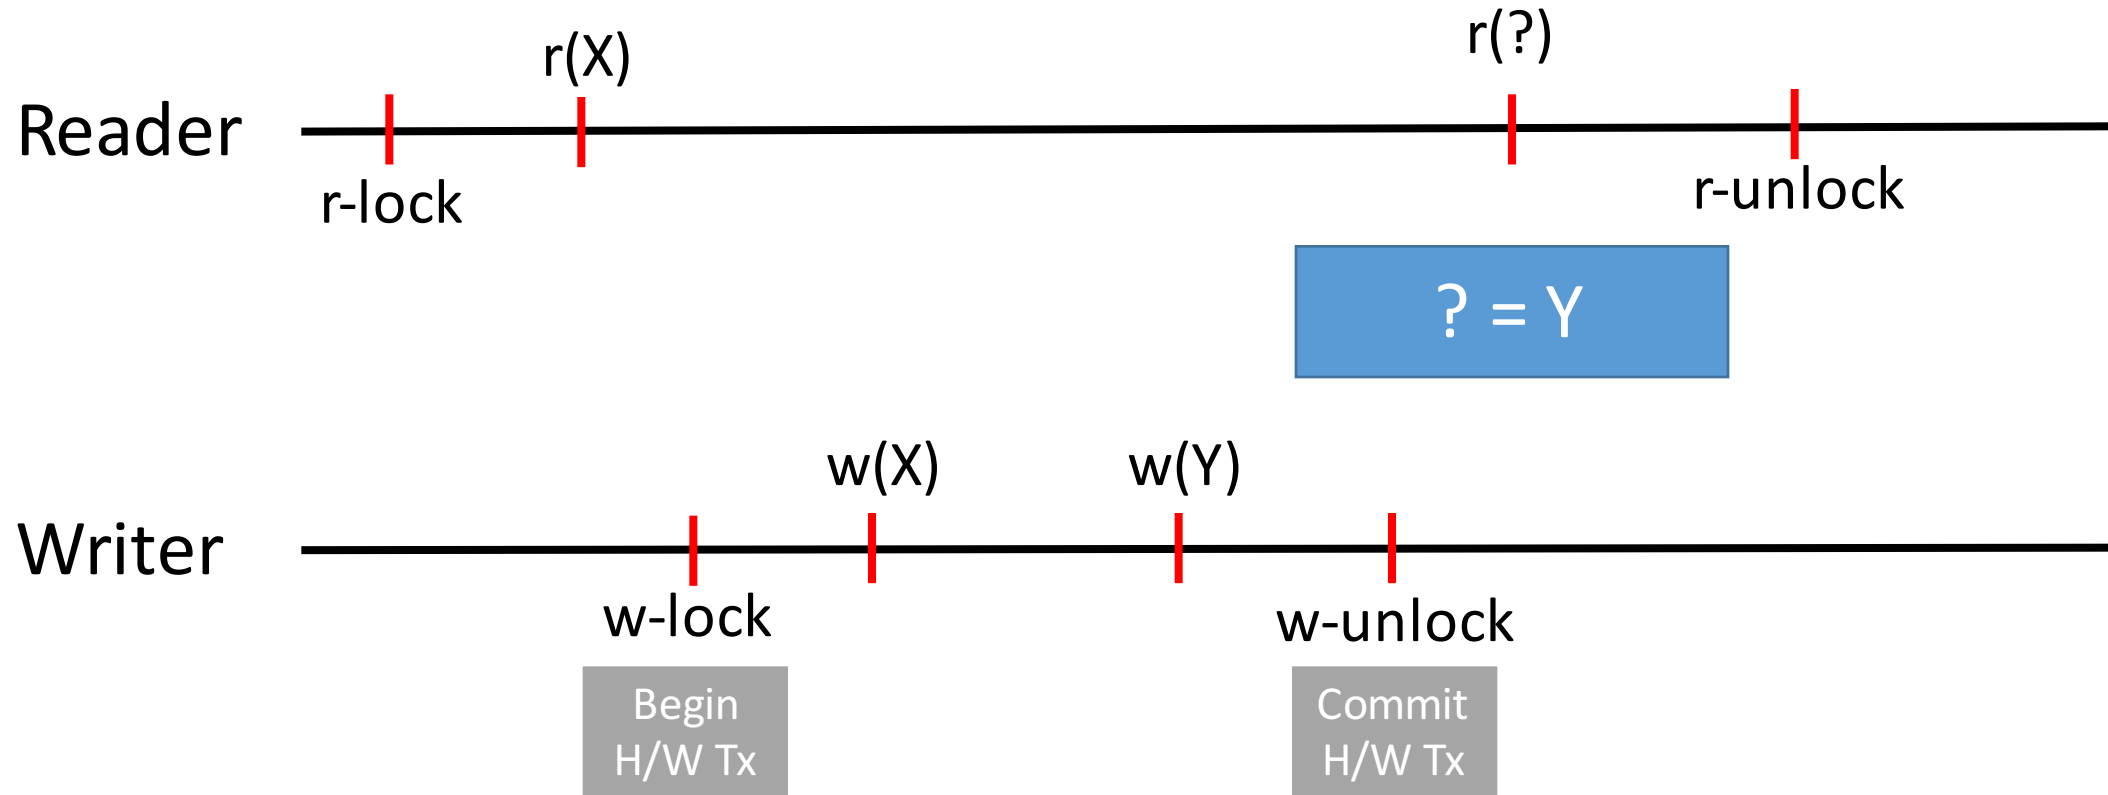
writers run in H/W TxS:

- no lock acquisition
- H/W tracking of read/write locations

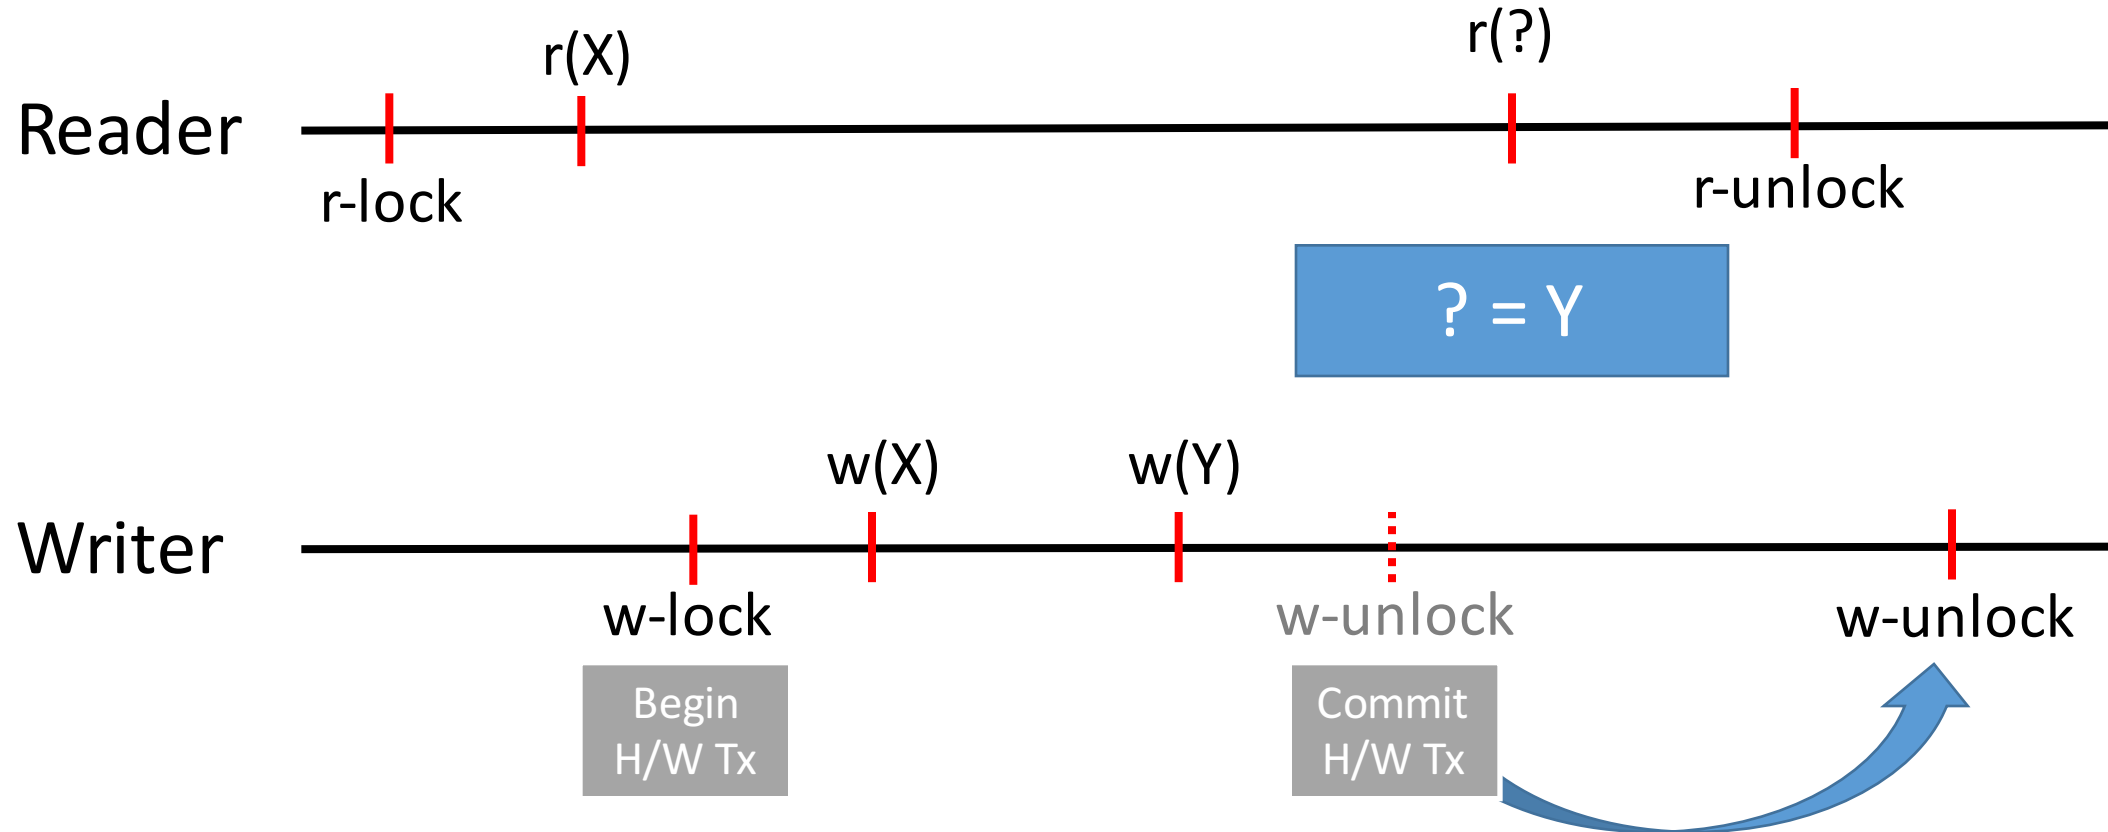
Hardware read-write lock elision



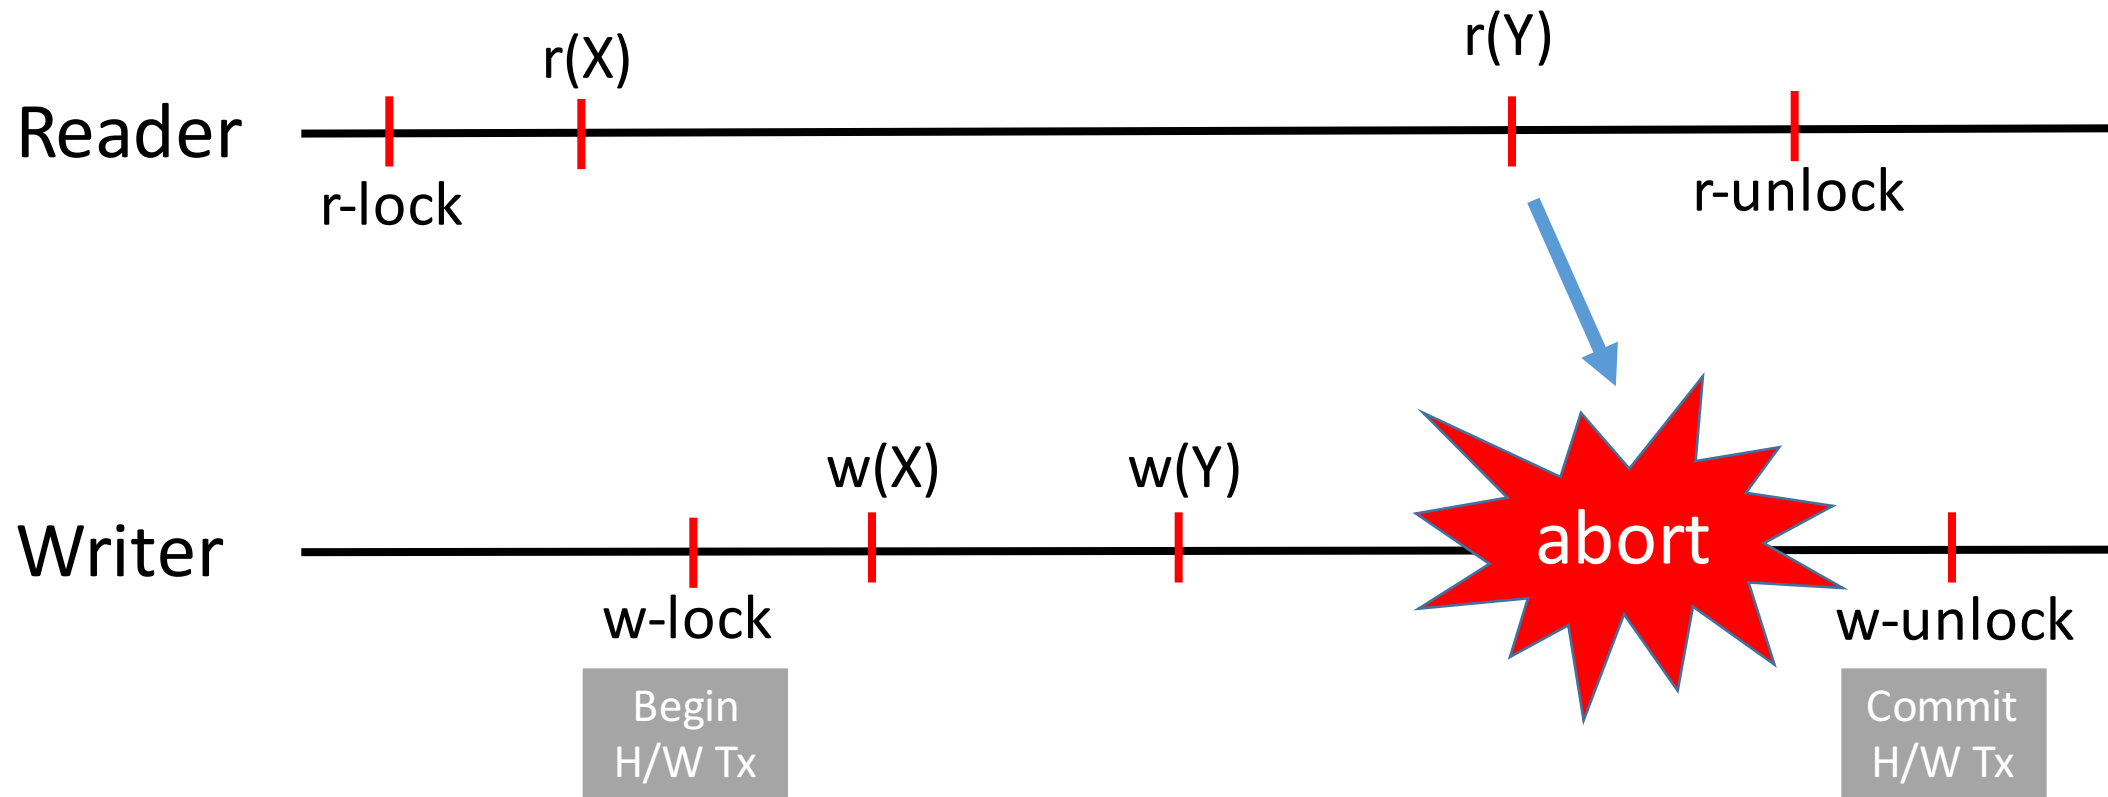
Hardware read-write lock elision



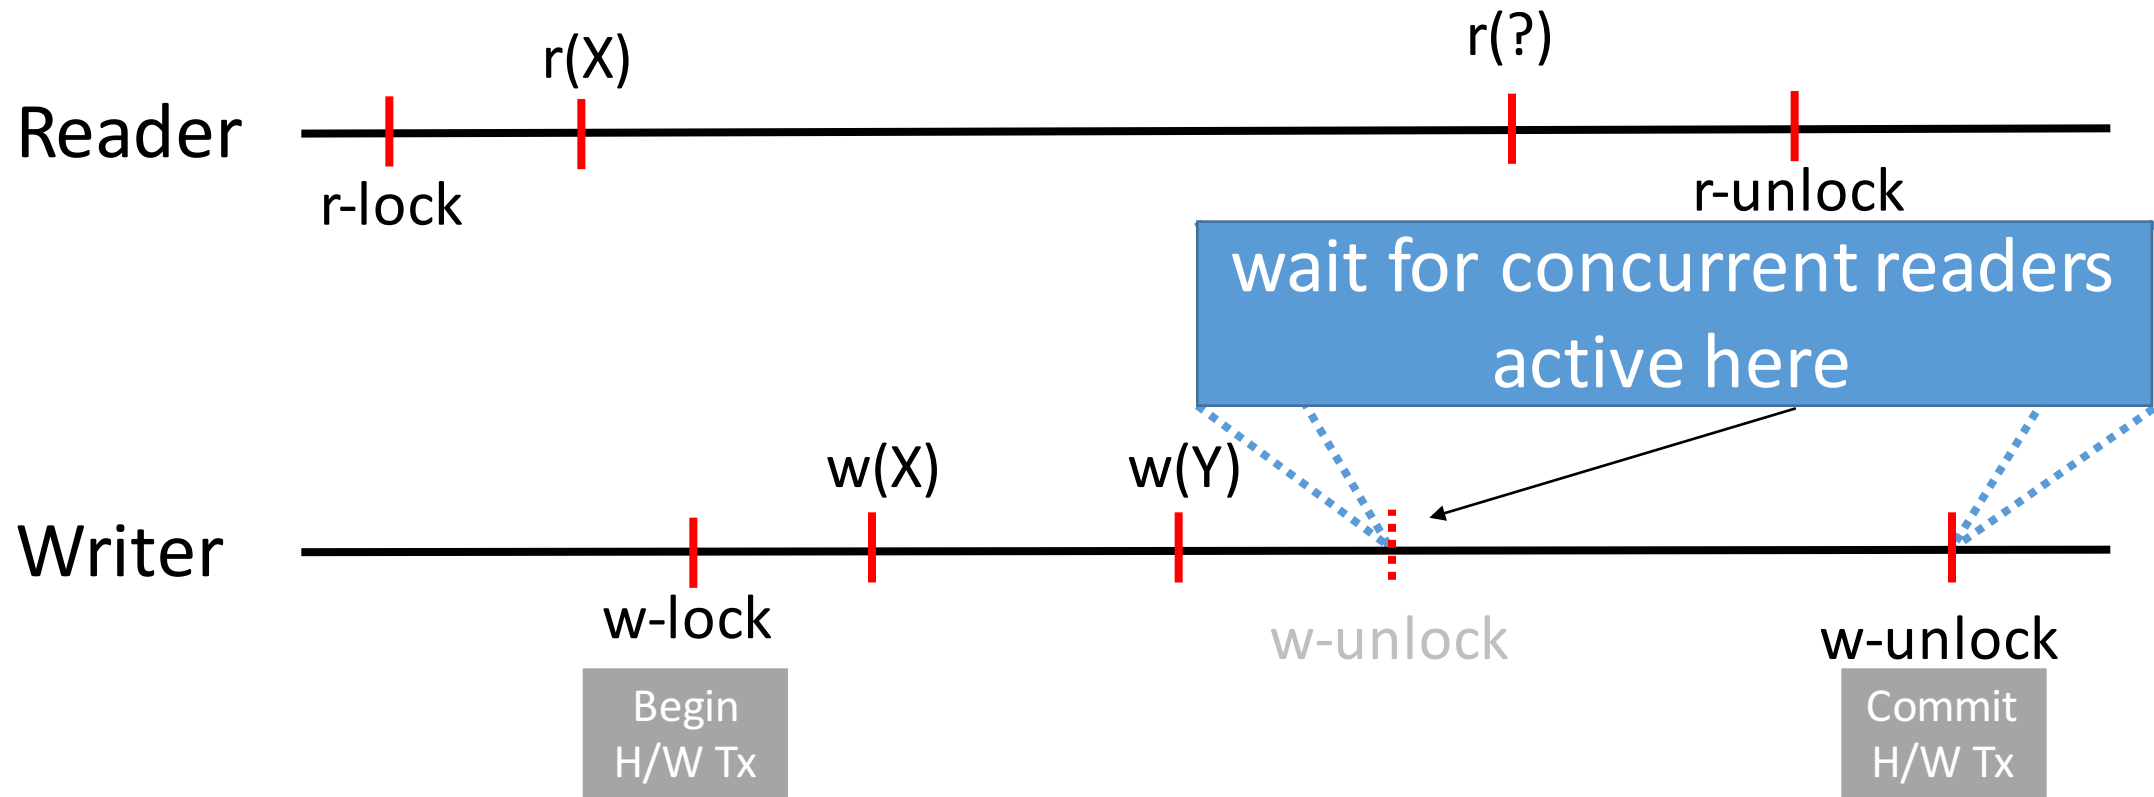
Hardware read-write lock elision



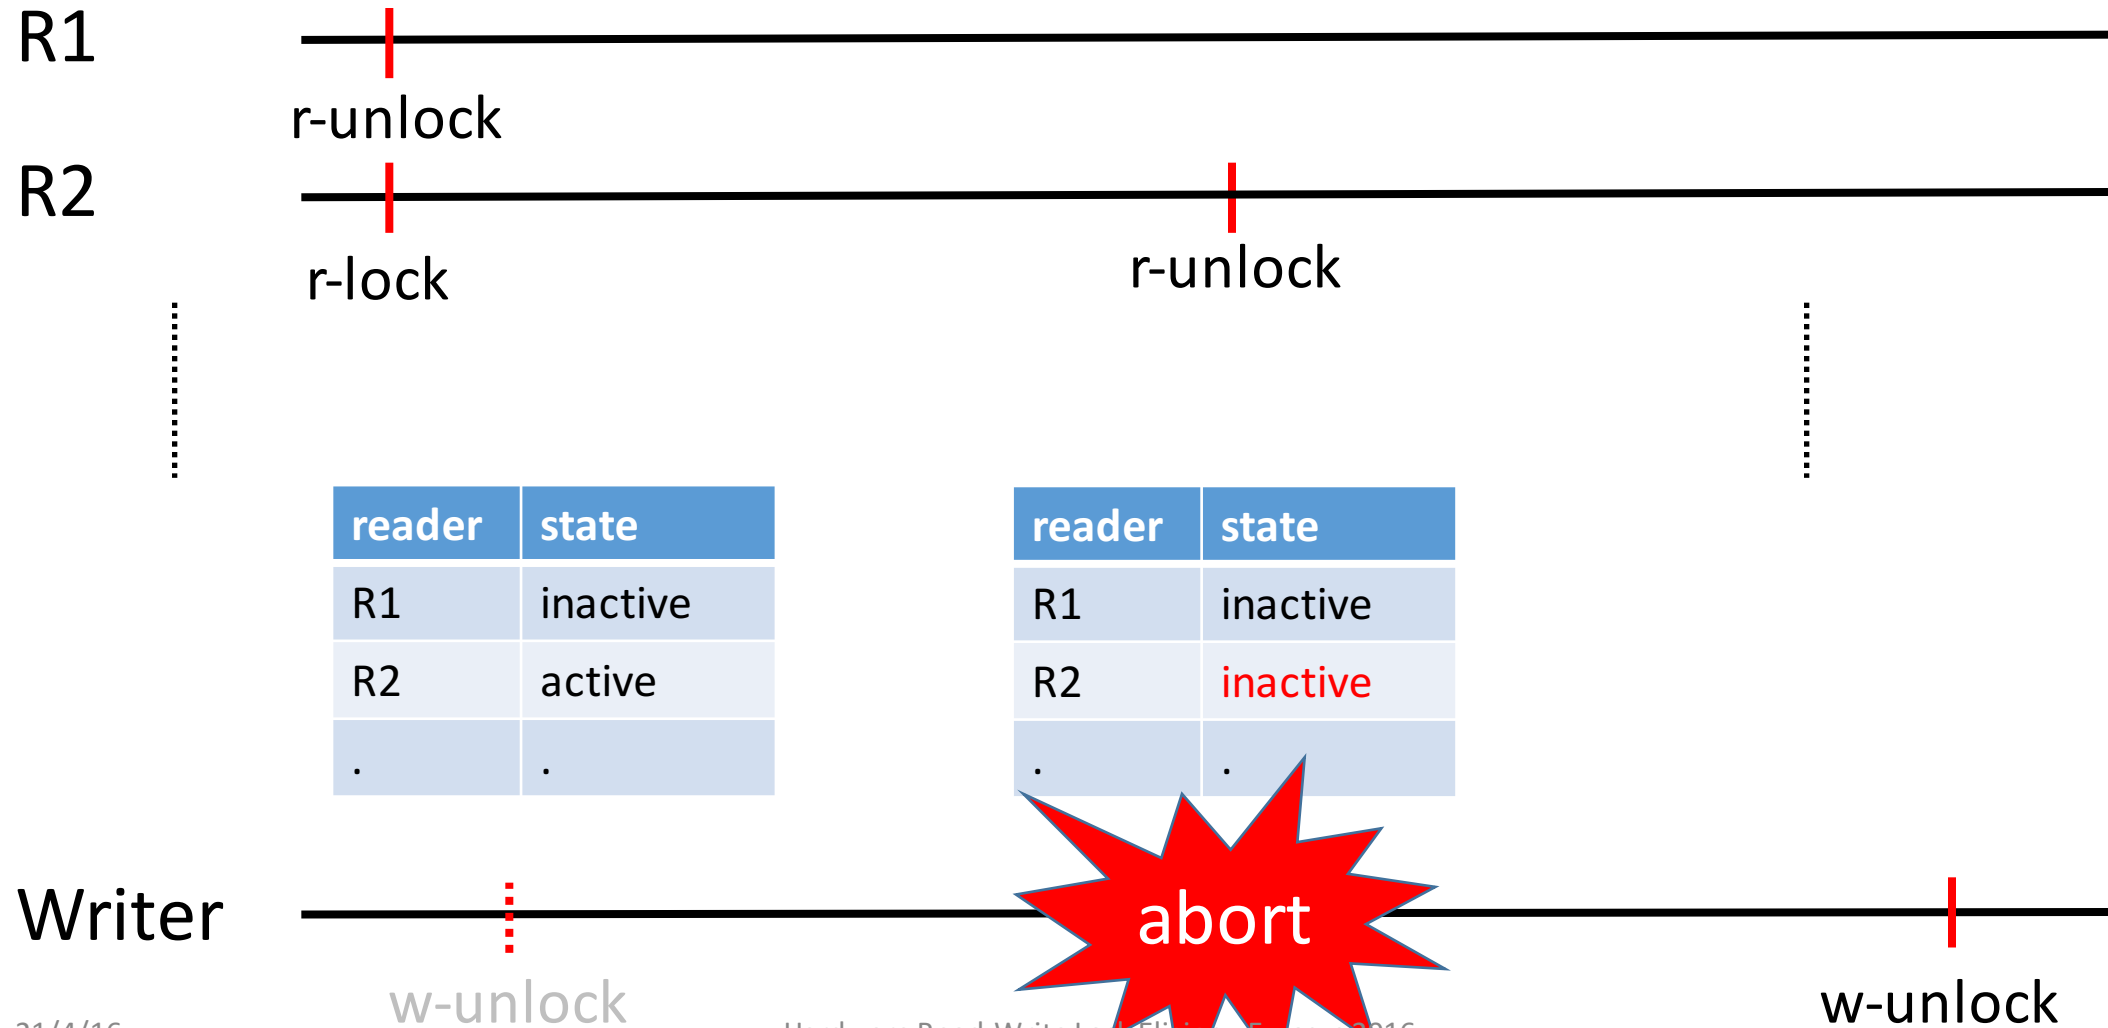
Hardware read-write lock elision



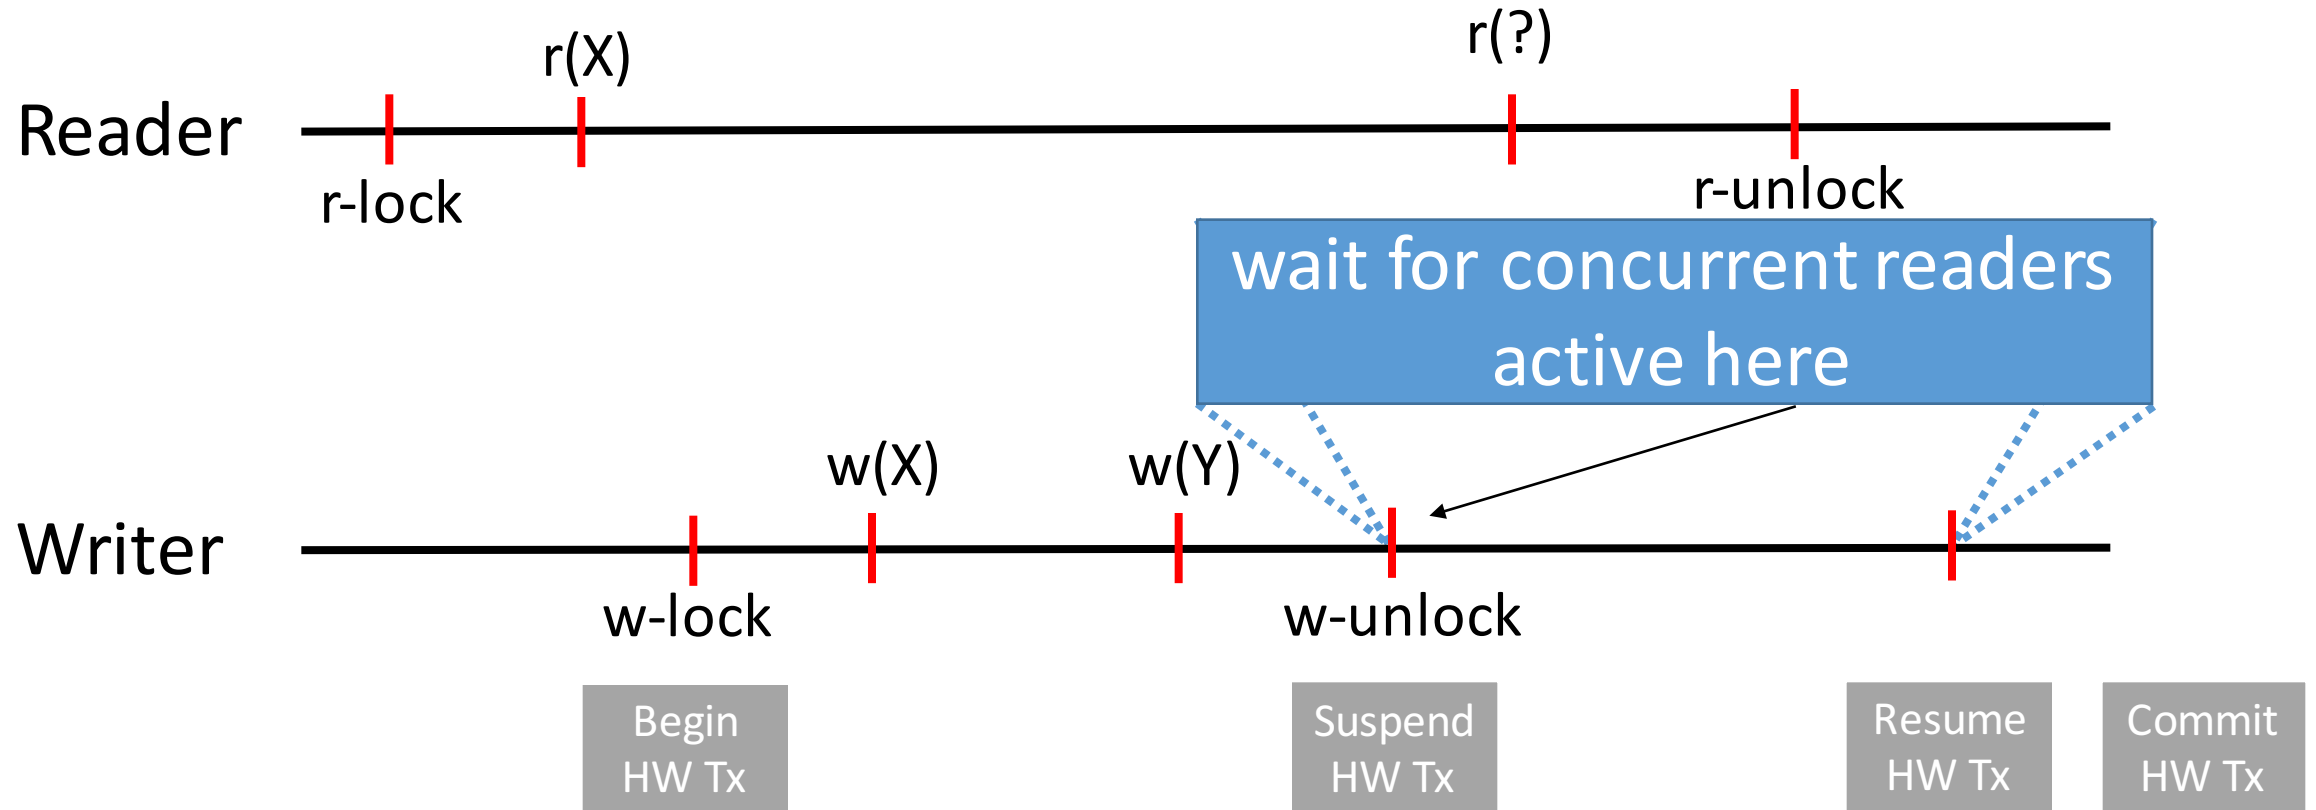
Hardware read-write lock elision



Hardware read-write lock elision



Hardware read-write lock elision

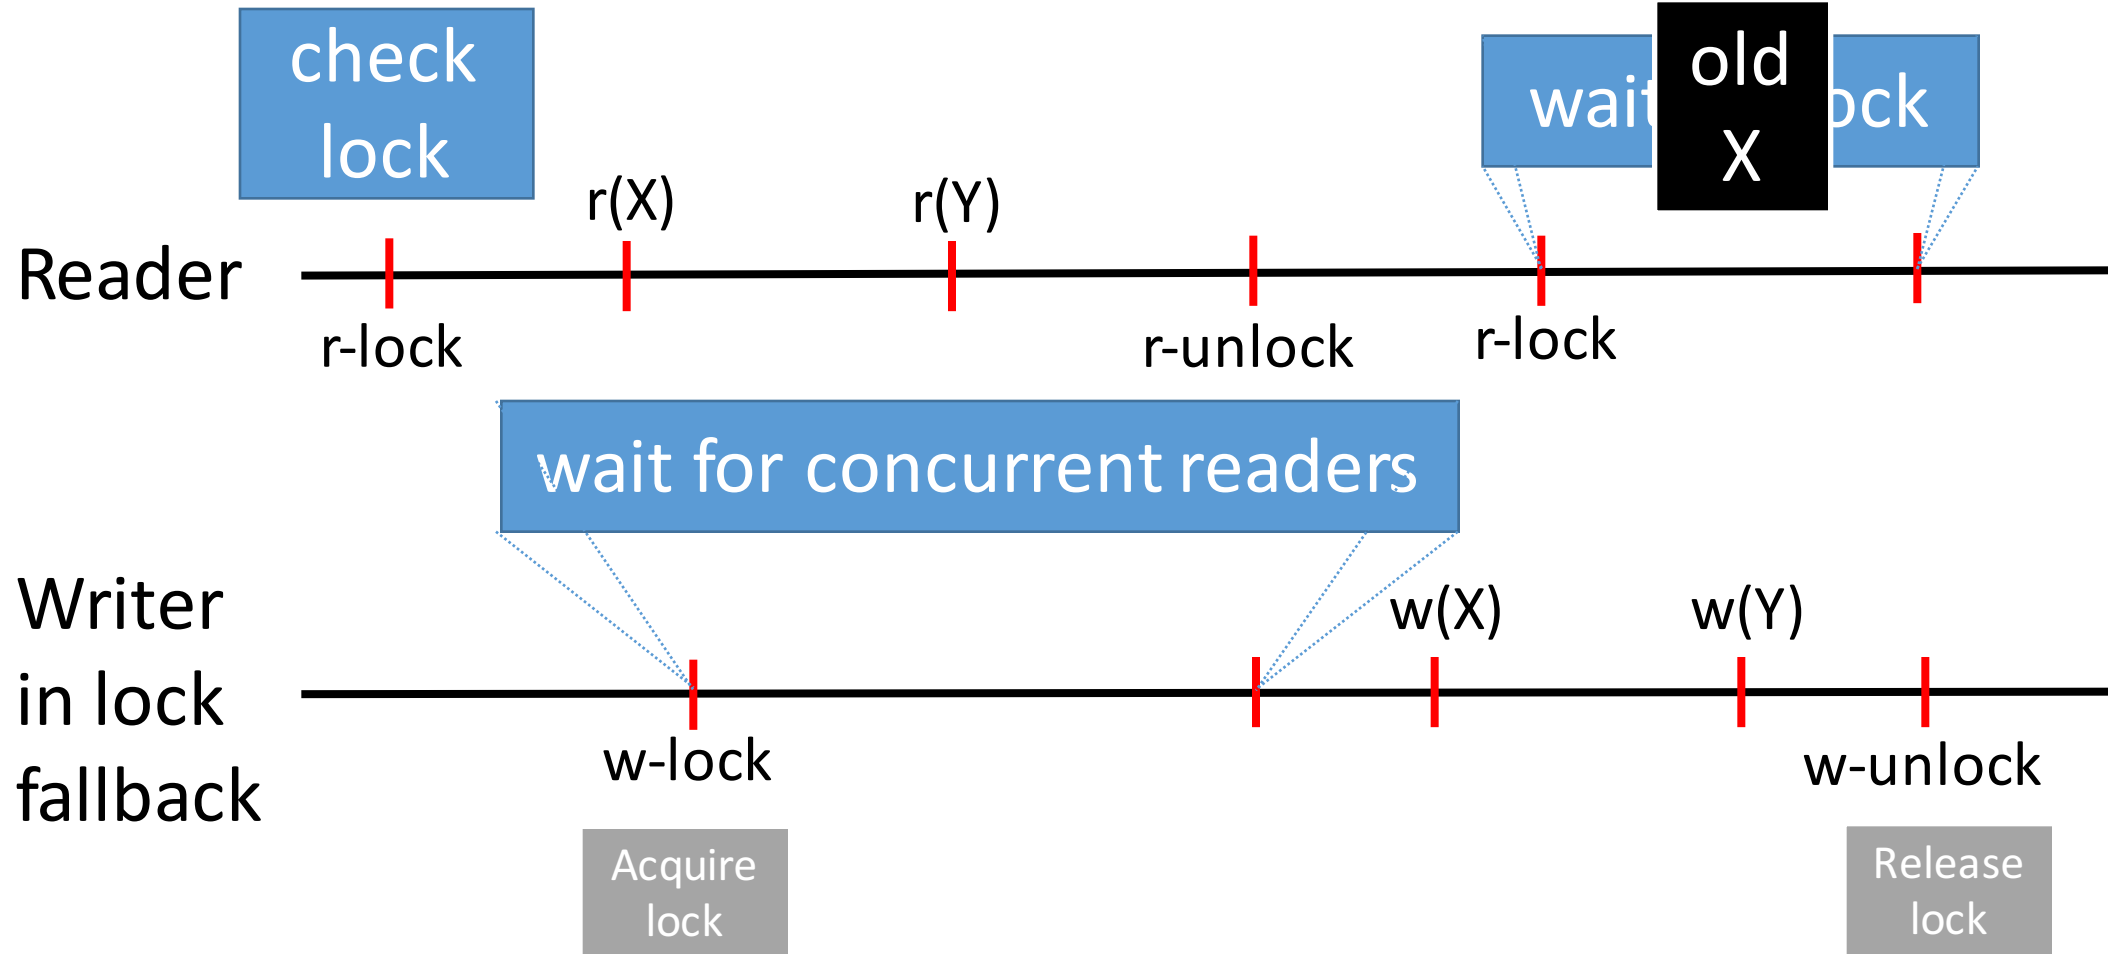


Hardware read-write lock elision

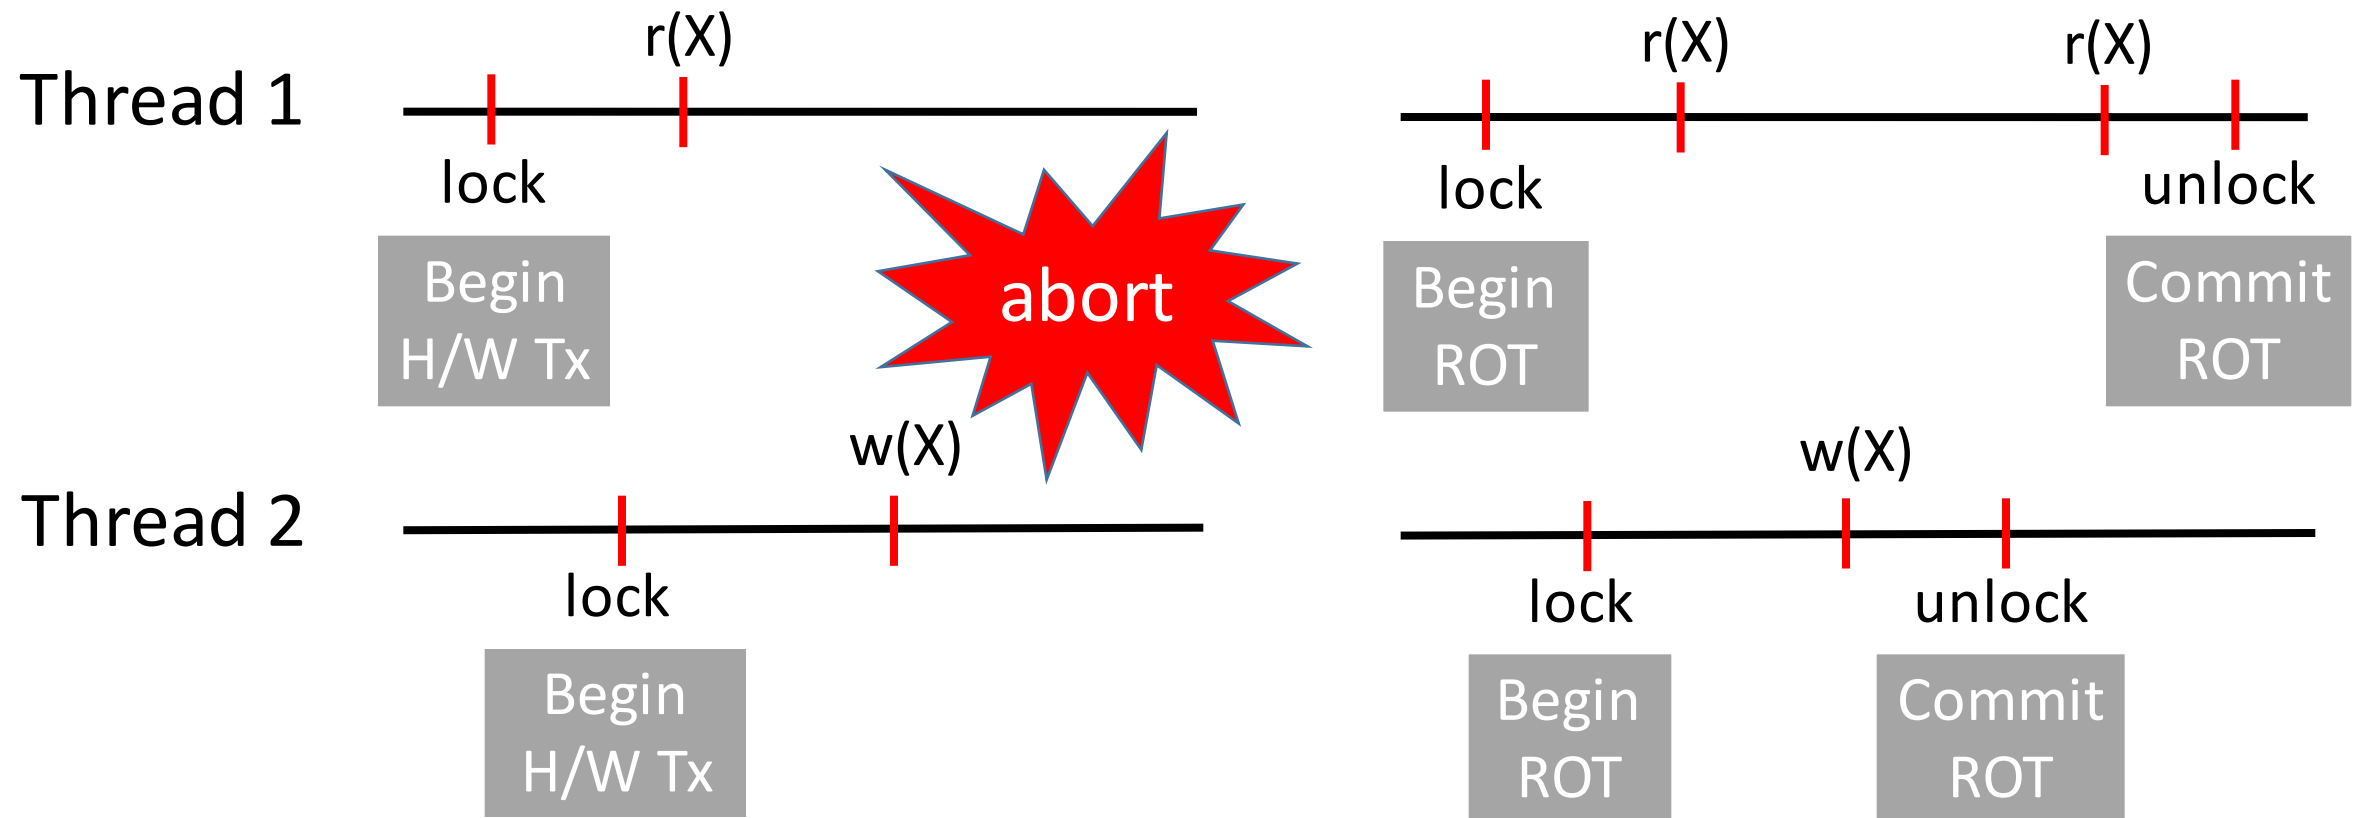
Writes use
H/W TxS

- ✓ concurrency among writers
- ✗ TxS may never commit
 - fallback lock is a must
 - Readers must synchronize with lock holder

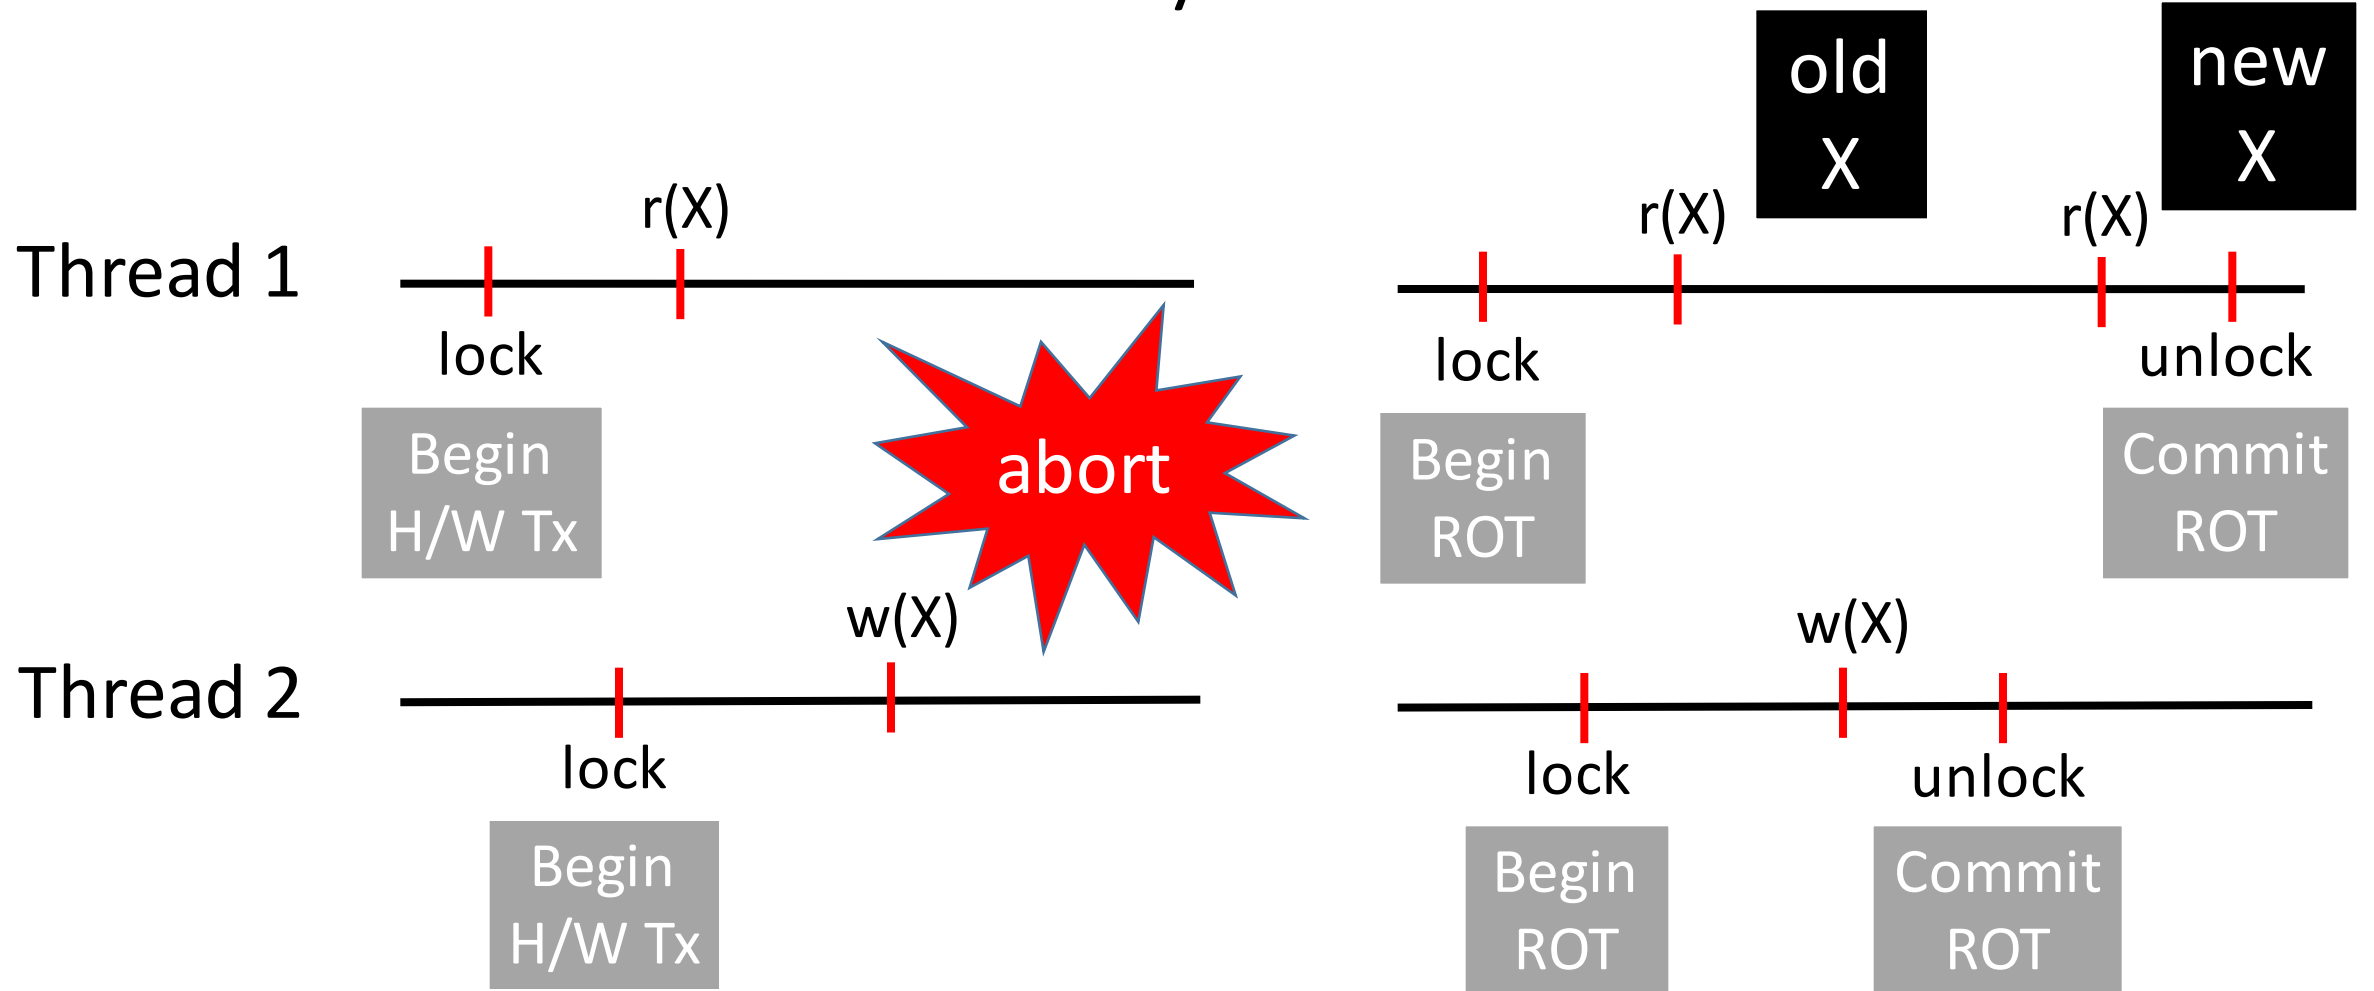
Hardware read-write lock elision



Rollback Only Transactions



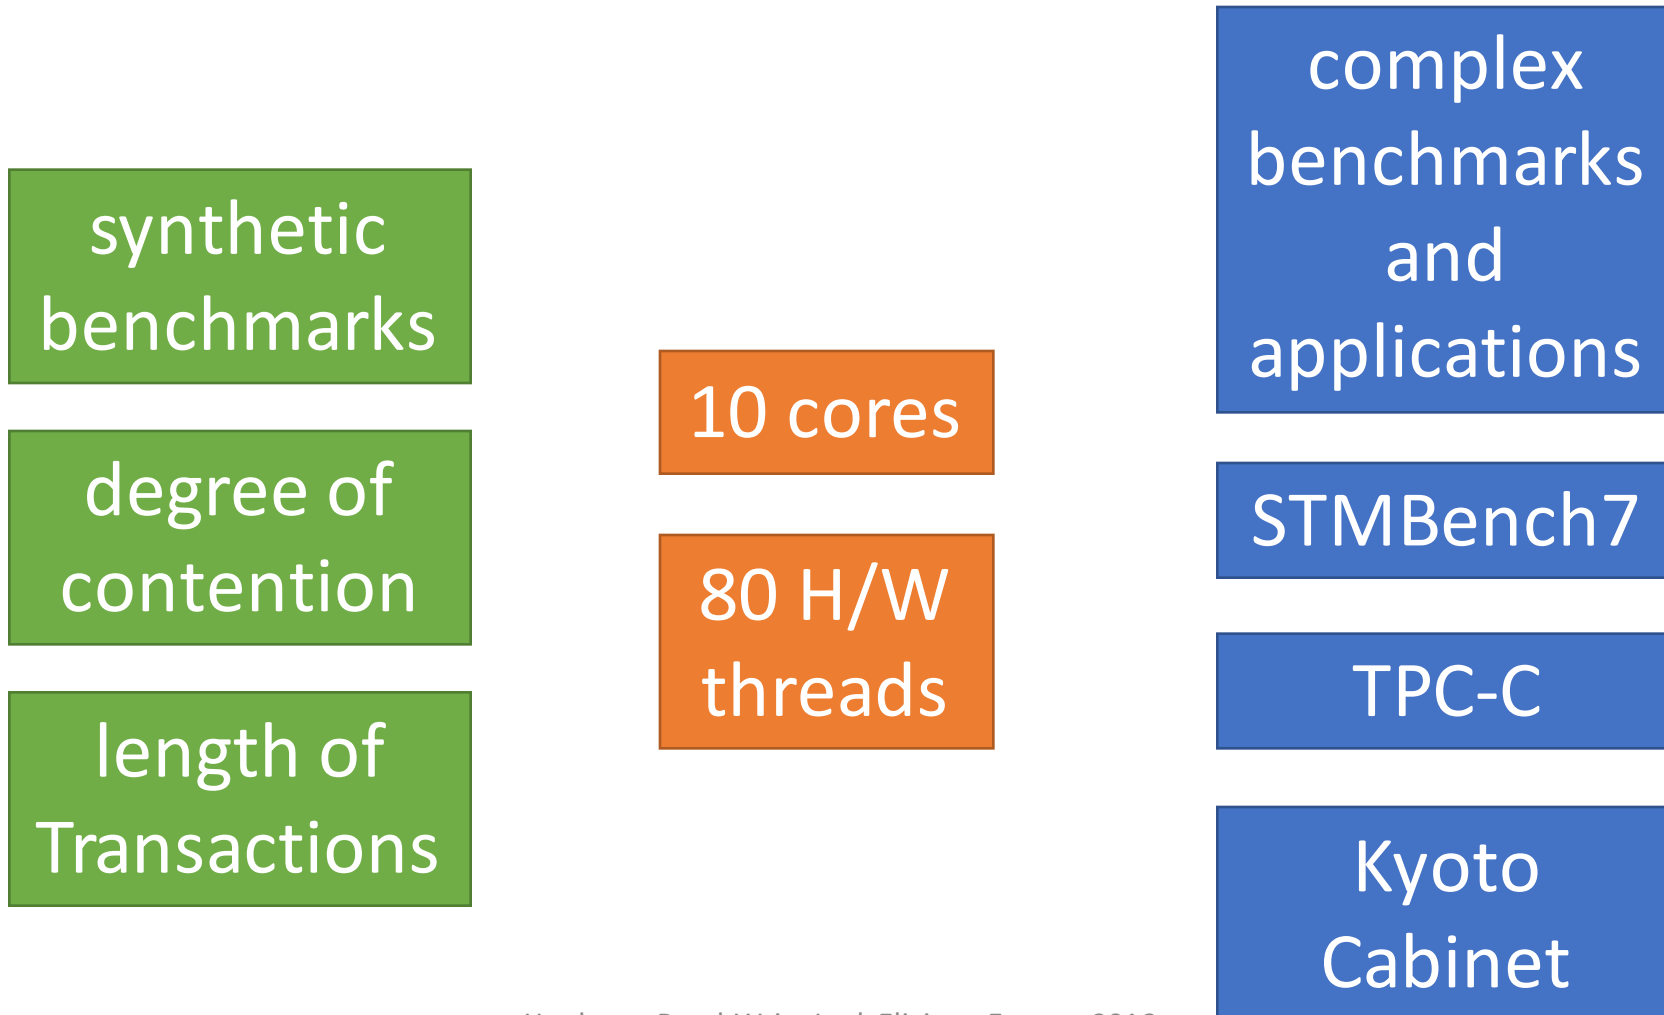
Rollback Only Transactions



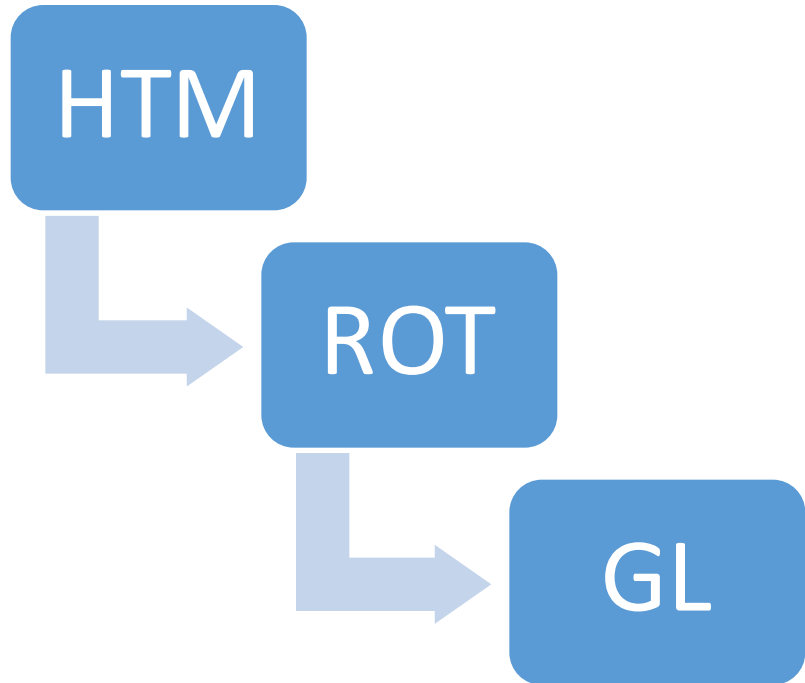
Using Rollback Only Transactions (ROTs)

- ✓ atomic
- ✓ allow larger Tx
- ✓ no need for Suspend/Resume
- ✗ no tracking of reads
- ✗ not serializable
- ✗ single writer

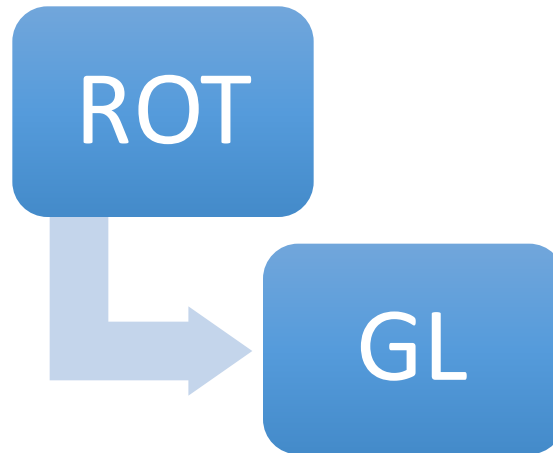
Experiments



Optimistic




Pessimistic





RW-LE^{OPT} 
RW-LE^{PES} 

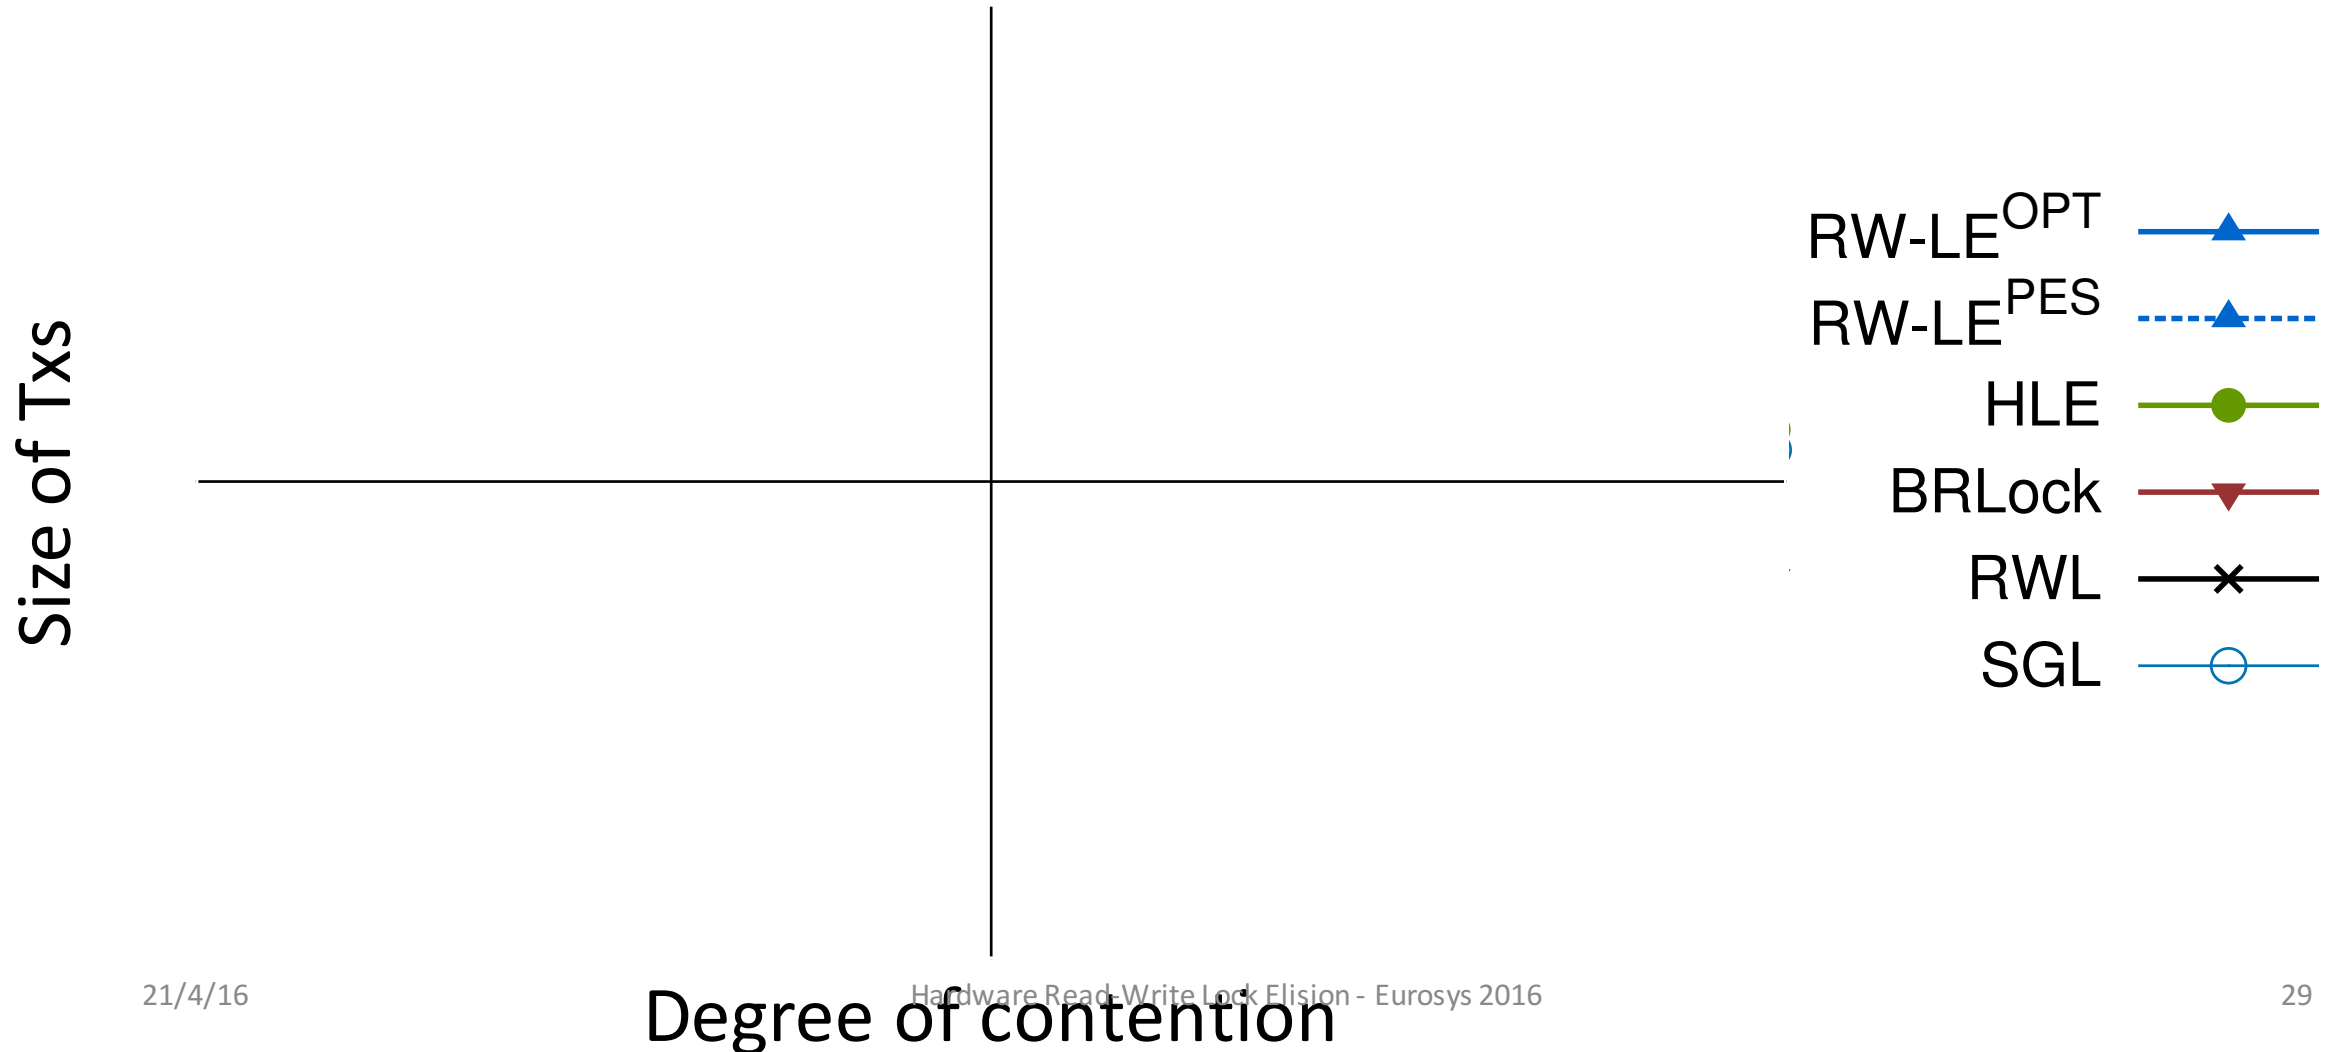
Optimistic



Pessimistic



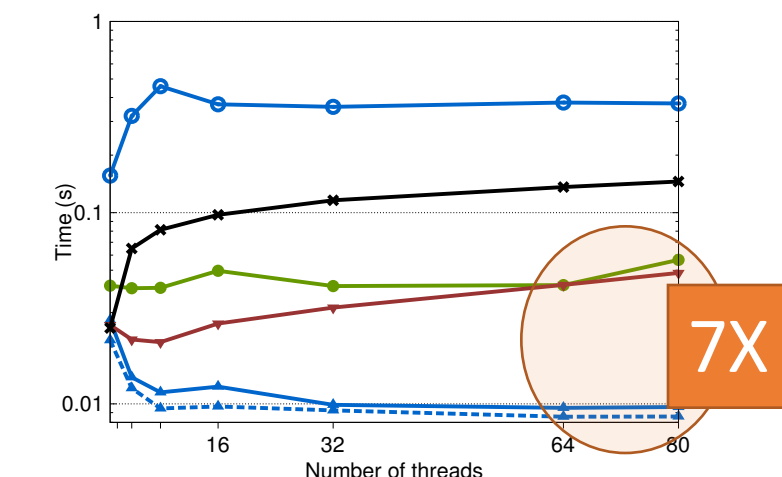
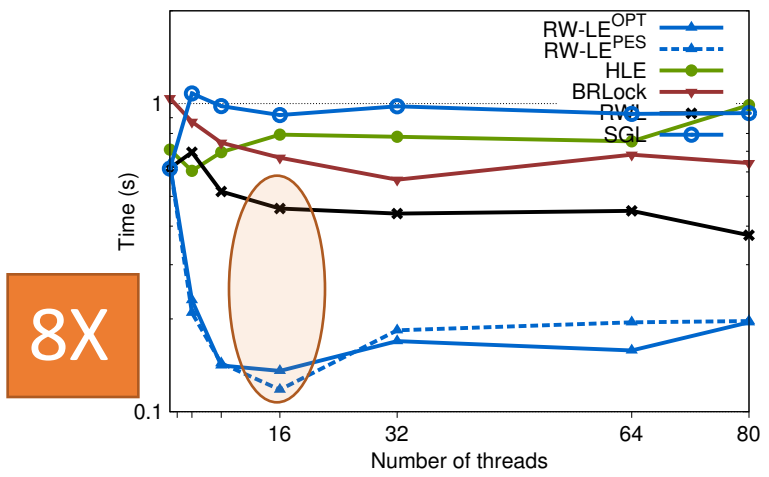
Synthetic benchmarks



Stress test – performance


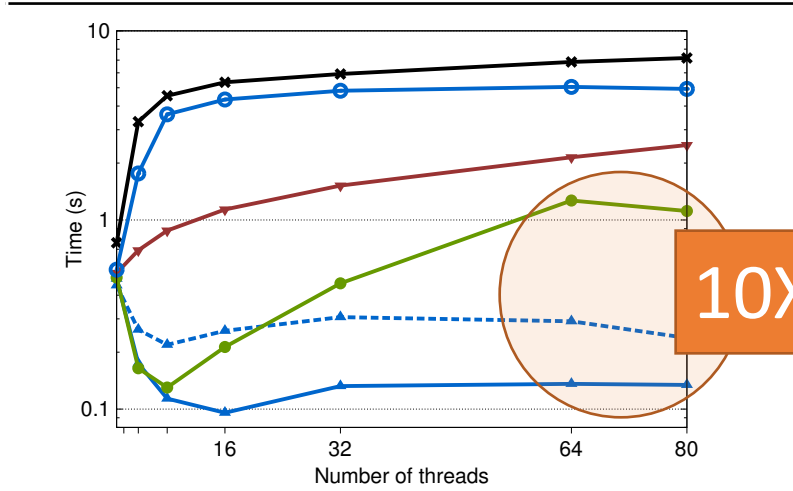
10% writers

Large Tx



- RW-LE^{OPT} —▲—
- RW-LE^{PES} - -▲- -
- HLE —●—
- BRLock —▼—
- RWL —×—
- SGL —○—

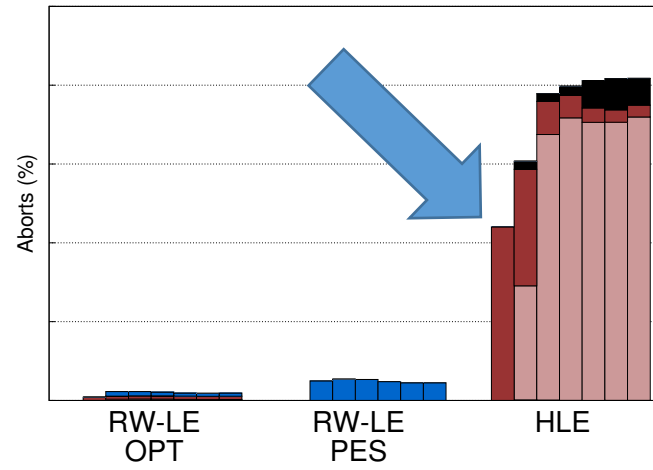
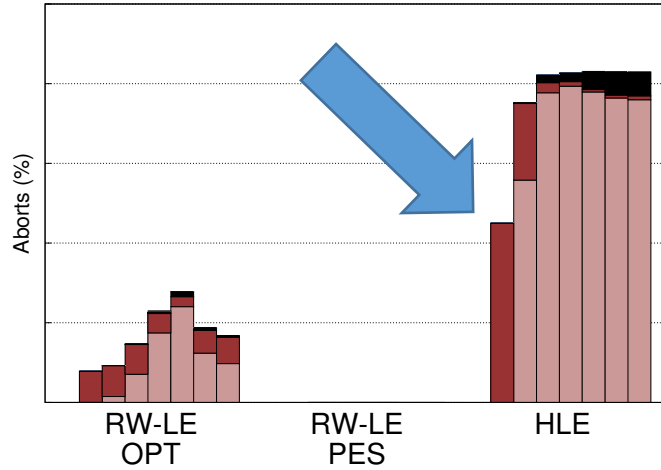
Small Tx




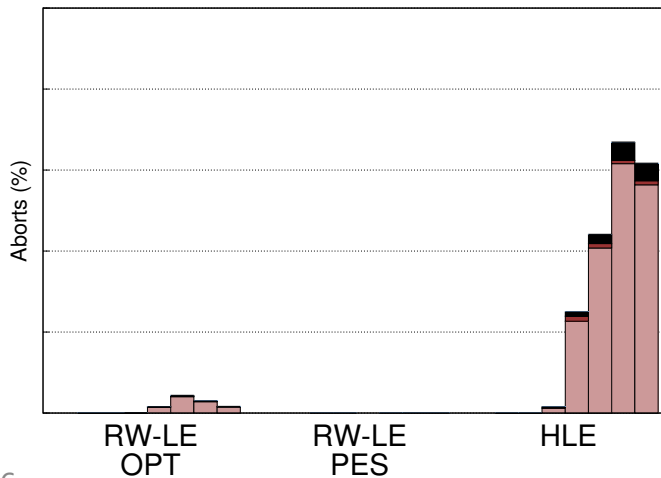
Stress test – abort rate

10% writers

Large Tx



Small Tx


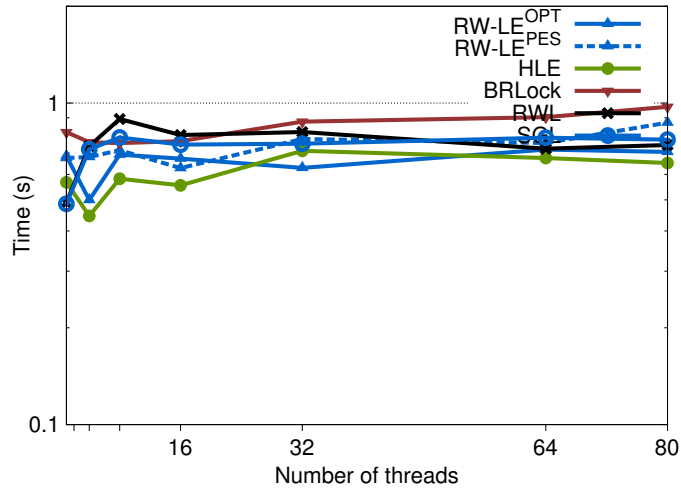


- ROT capacity
- ROT conflicts
- Lock aborts
- HTM capacity
- HTM non-tx
- HTM tx


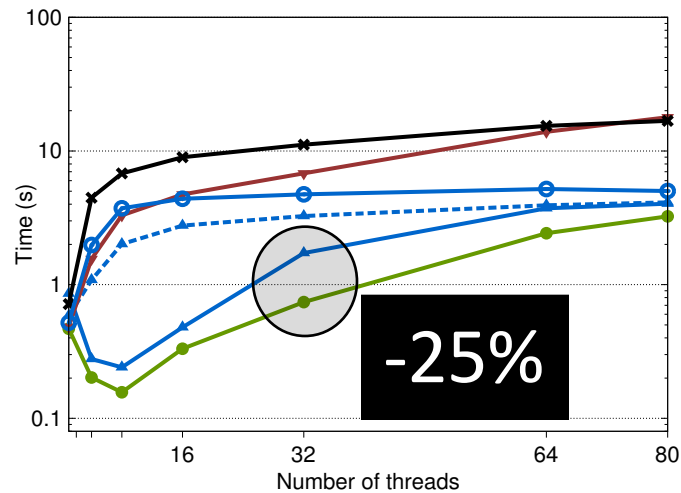
Stress test

90% writers

Large Tx



Small Tx





- RW-LE^{OPT}
- RW-LE^{PES}
- HLE
- BRLock
- RWL
- SGL

-25%

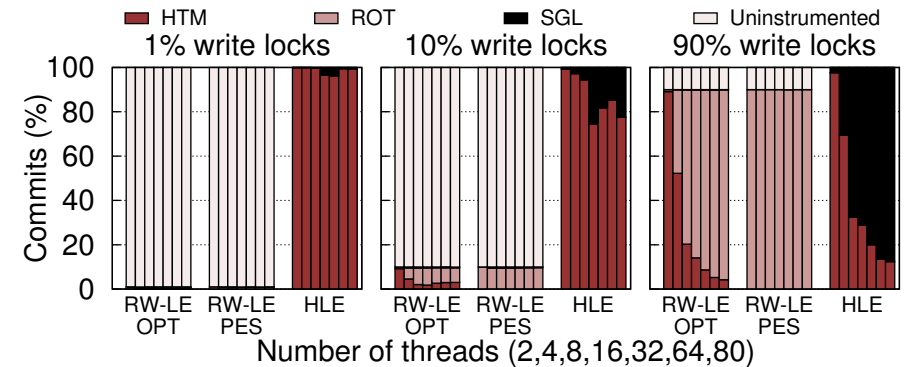

-10%

Synthetic benchmarks

High capacity





Low capacity


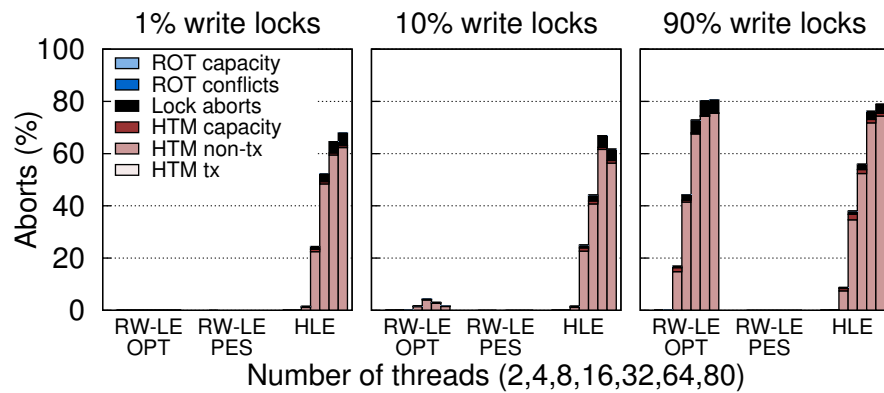


Synthetic benchmarks

High capacity

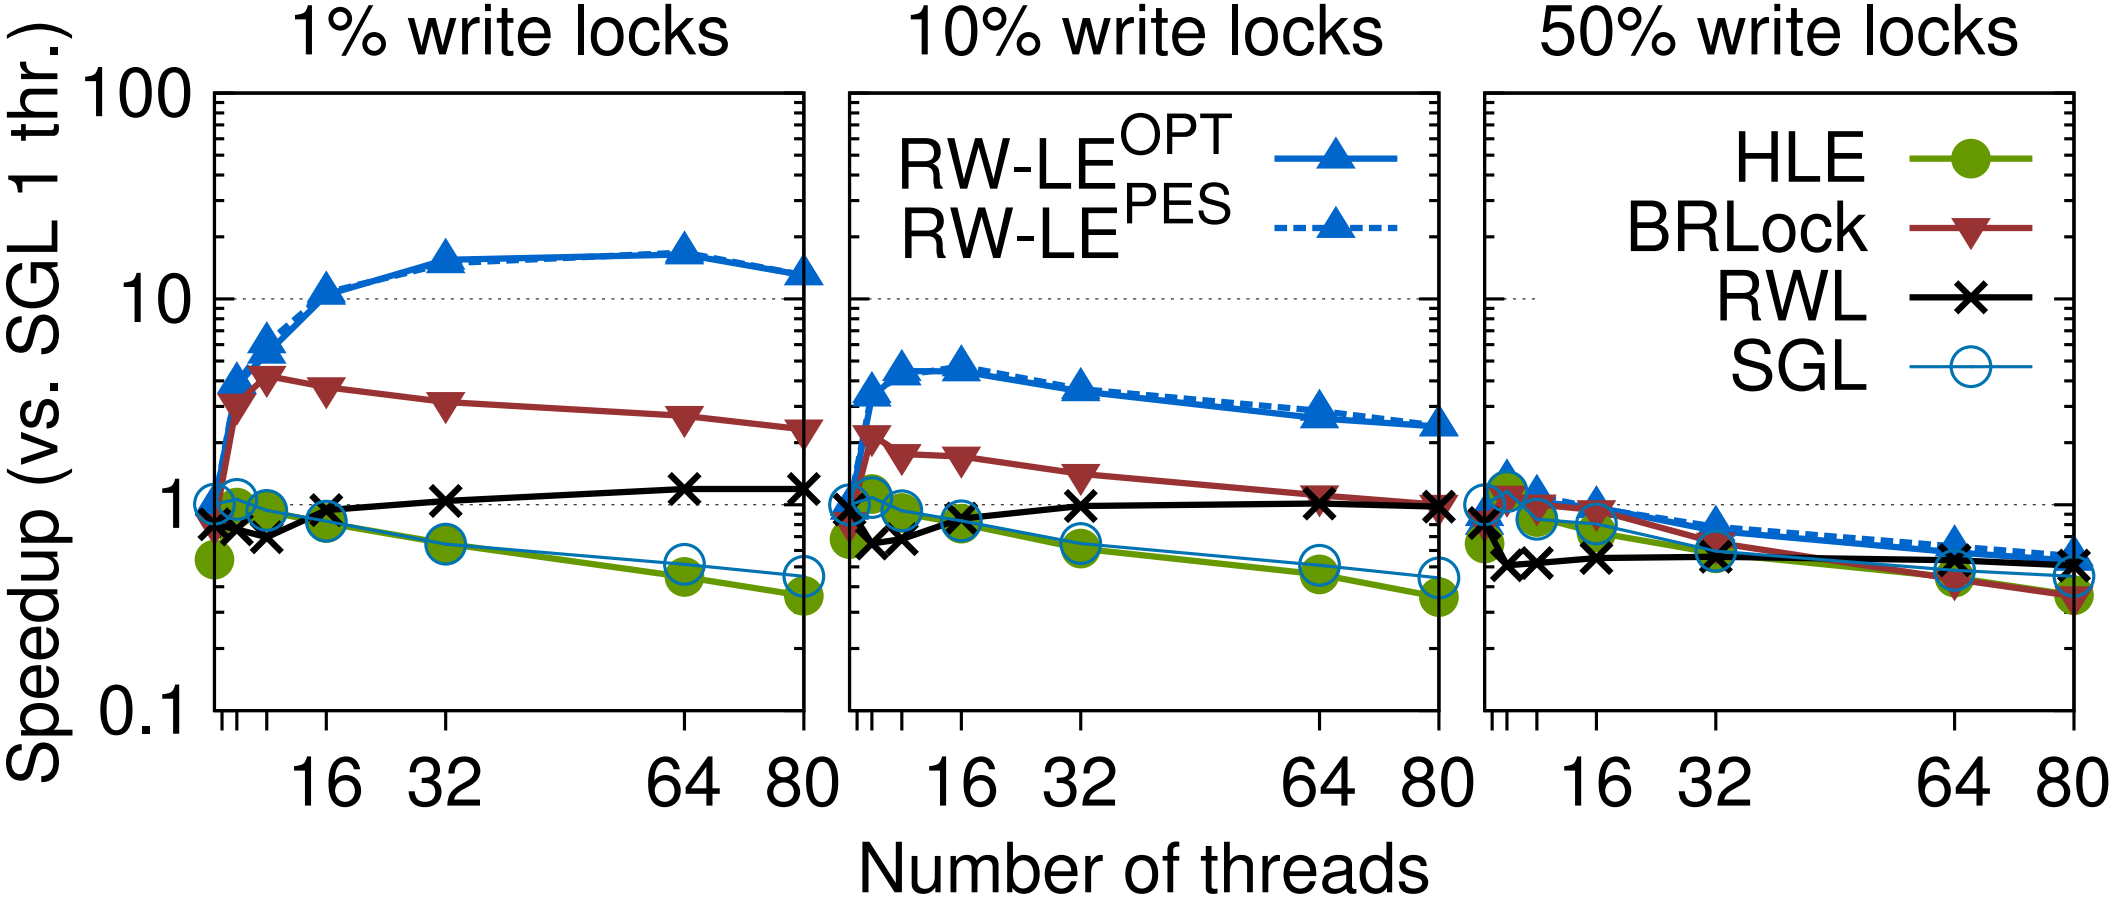


Low capacity

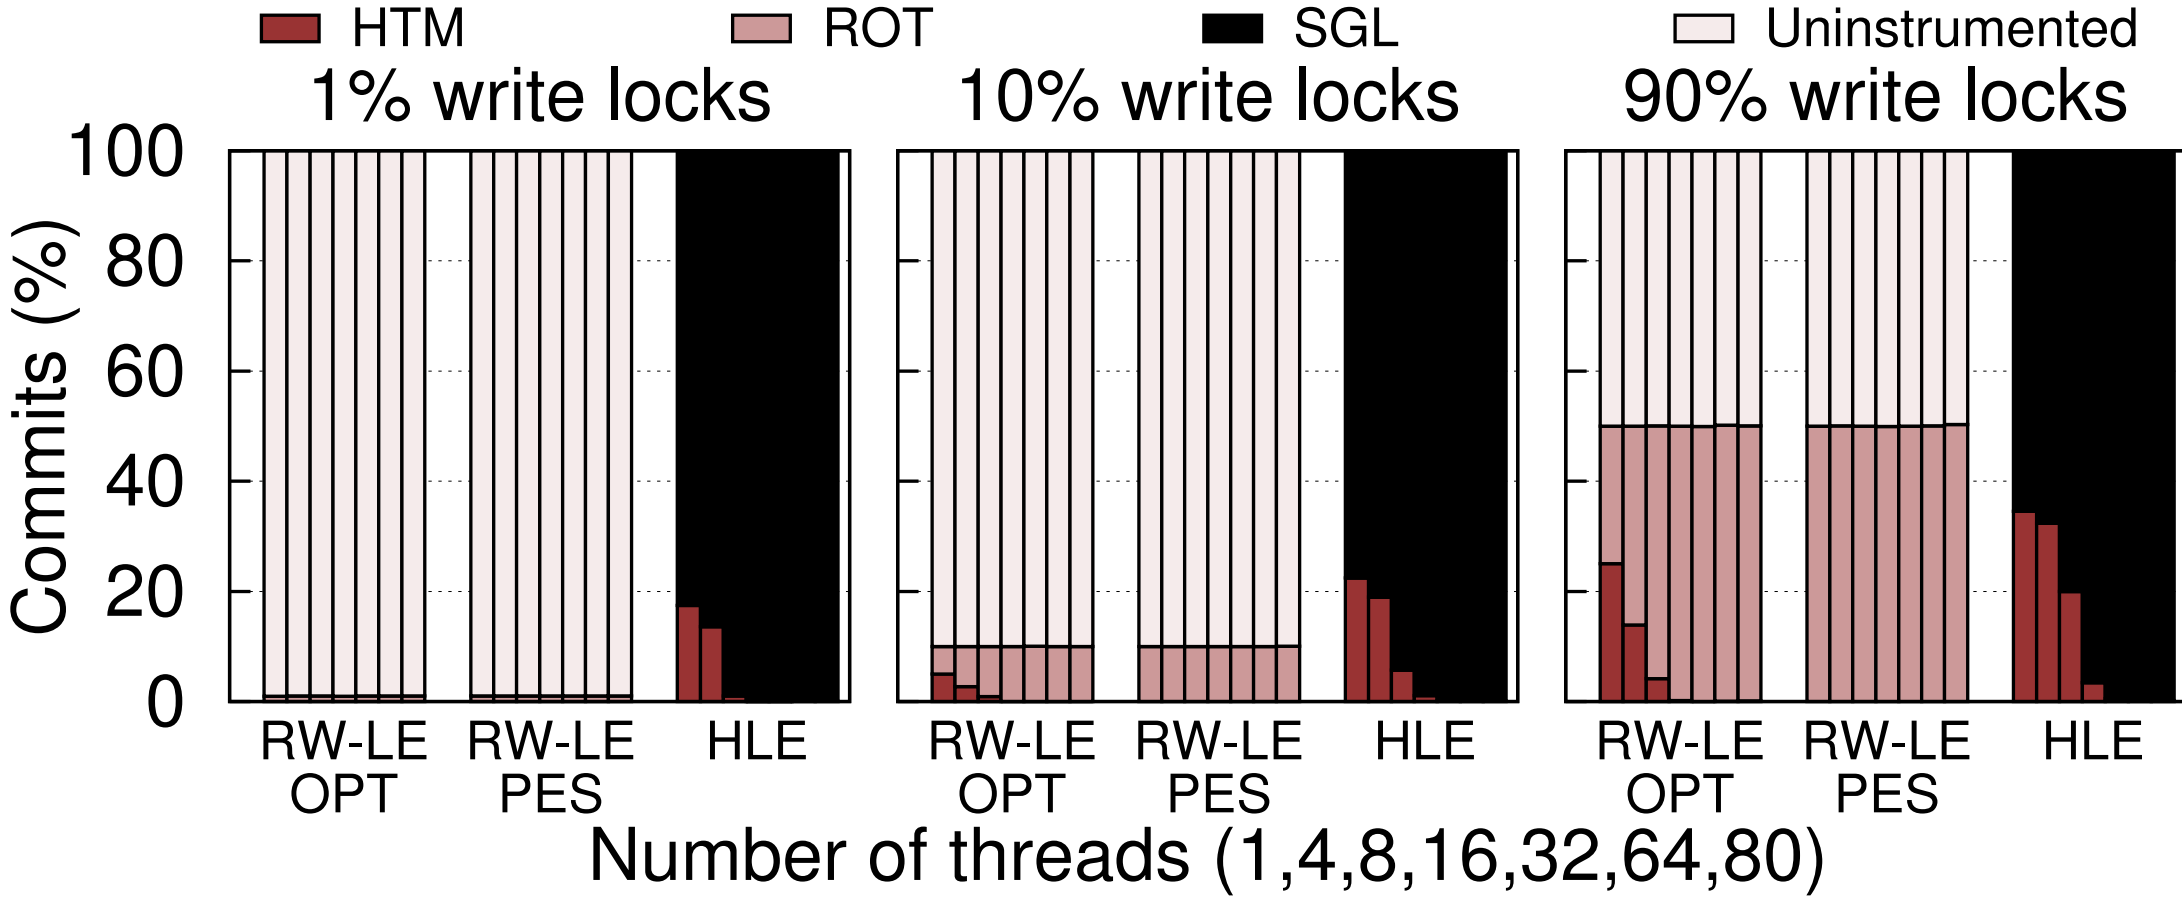


TPC-C


6X



TPC-C




TPC-C




STMbench7

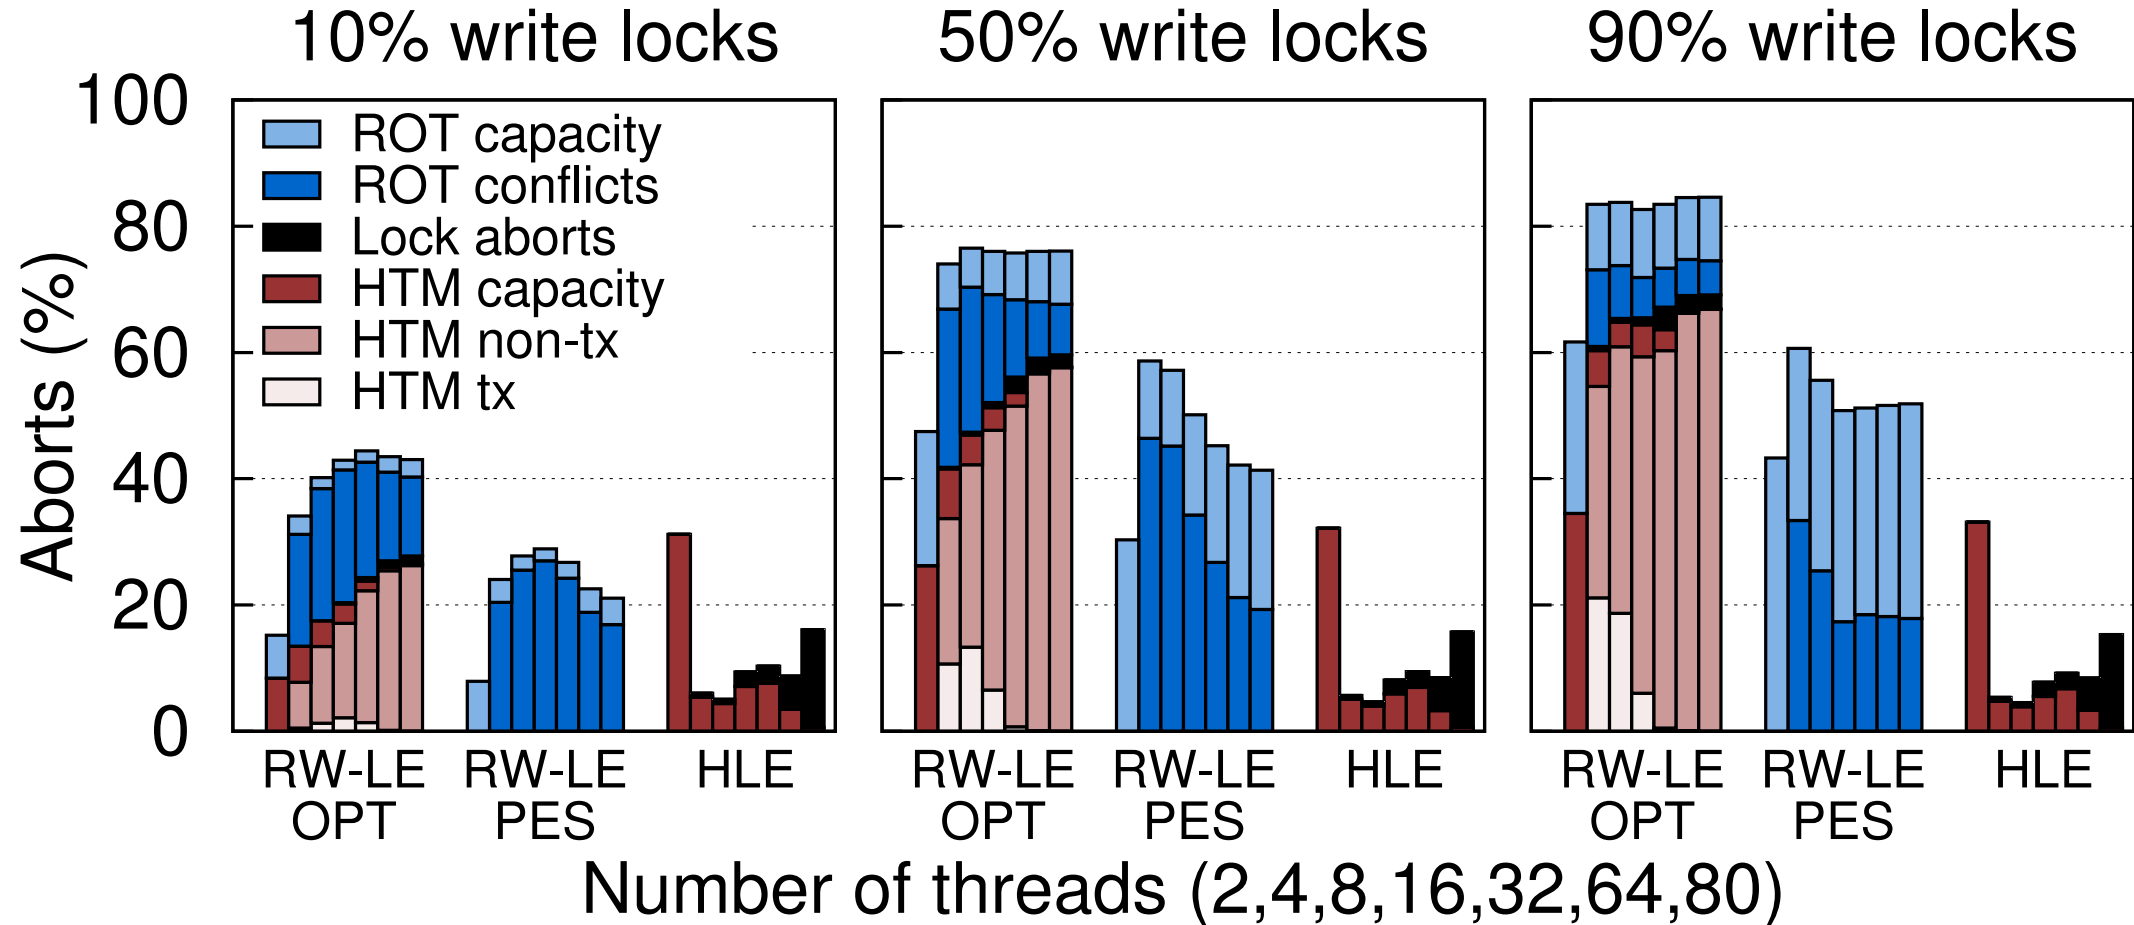
4X



STMbench7

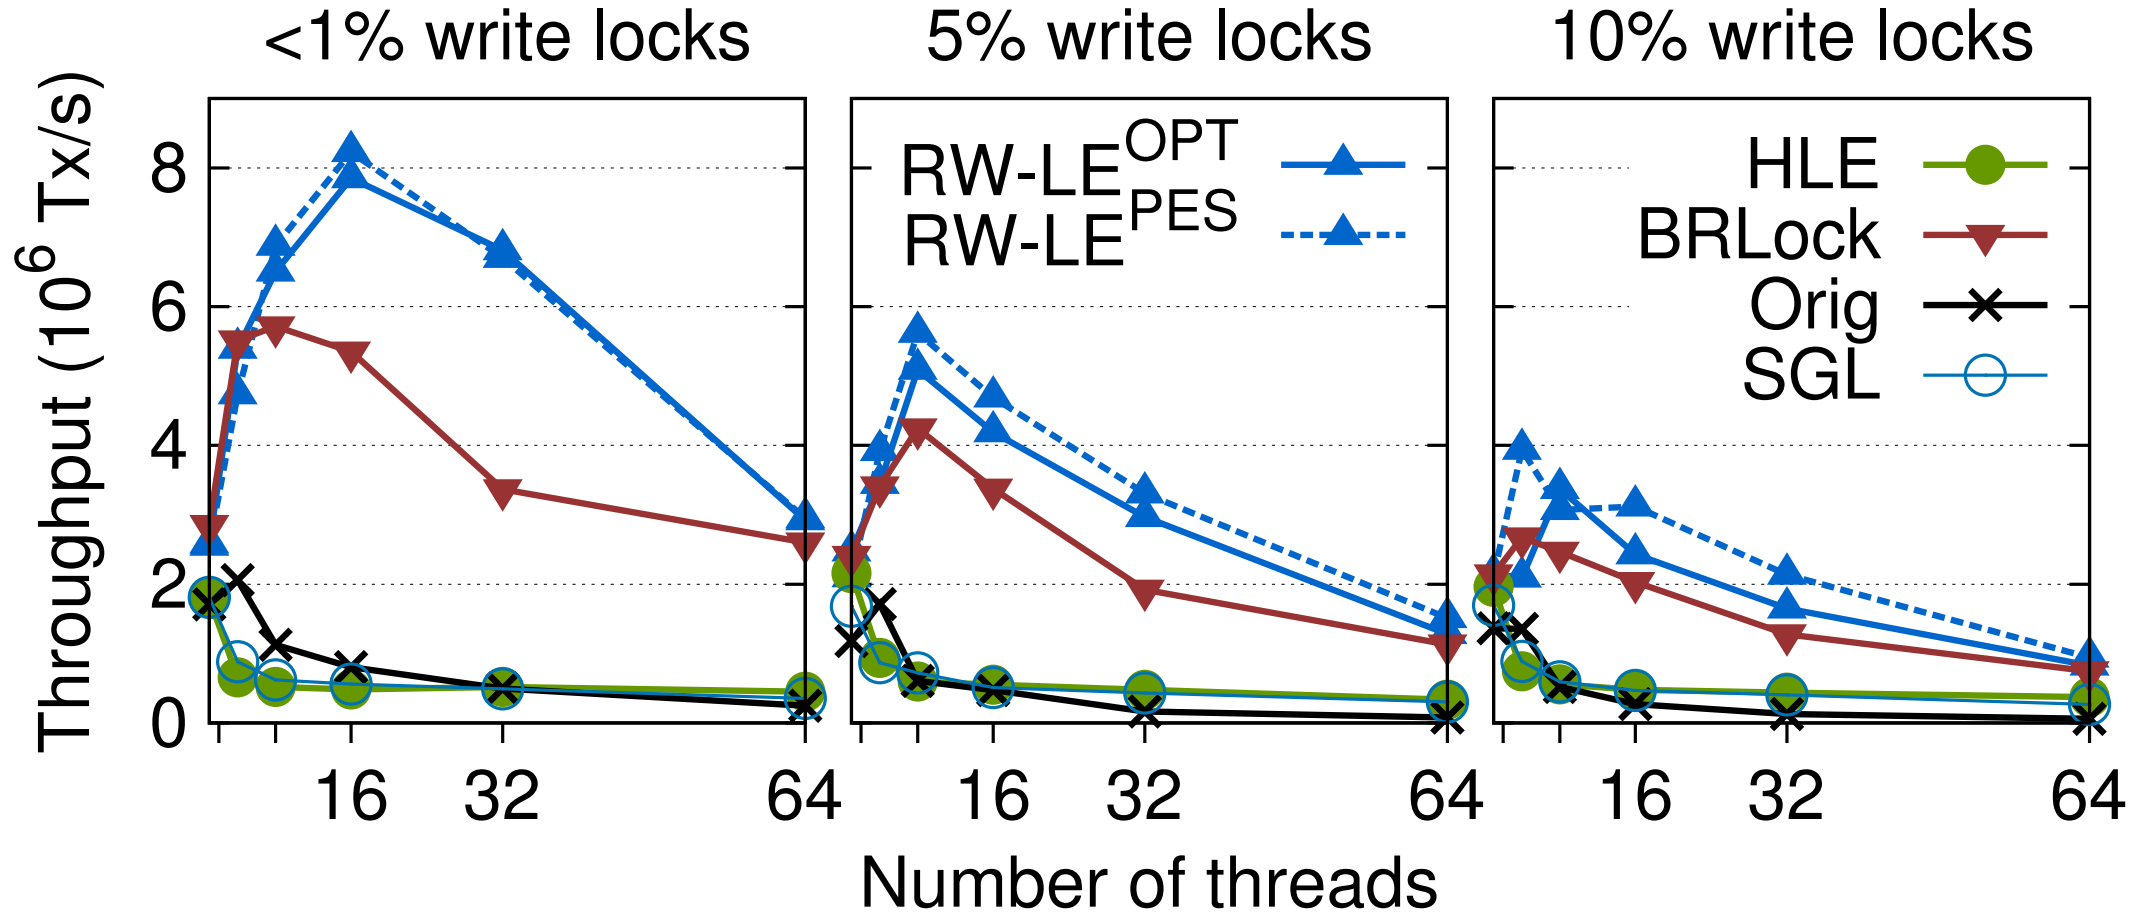


STMbench7




Kyoto Cabinet


2X



Kyoto Cabinet



Kyoto Cabinet



Conclusions

readers without
instrumentation

writers using
H/W Tx

Conclusions

readers without
instrumentation

writers using
H/W Tx

10X

Conclusions

readers without
instrumentation

writers using
H/W Tx



suspend/resume

ROTs

10x