

Extending Hardware Transactional Memory Capacity

via

Rollback-Only Transactions and Suspend/Resume

Shady Issa



Pascal Felber



Alexander
Matveev



Paolo Romano



Extending Hardware Transactional Memory Capacity

via

Rollback-Only Transactions and Suspend/Resume

POWER8-TM

Shady Issa



Pascal Felber



UNIVERSITÉ DE
NEUCHÂTEL

Alexander
Matveev



Paolo Romano



TÉCNICO
LISBOA



Transactional Memory

- alternative paradigm for parallel programming
- easy to use
- potential of fine-grained locking performance

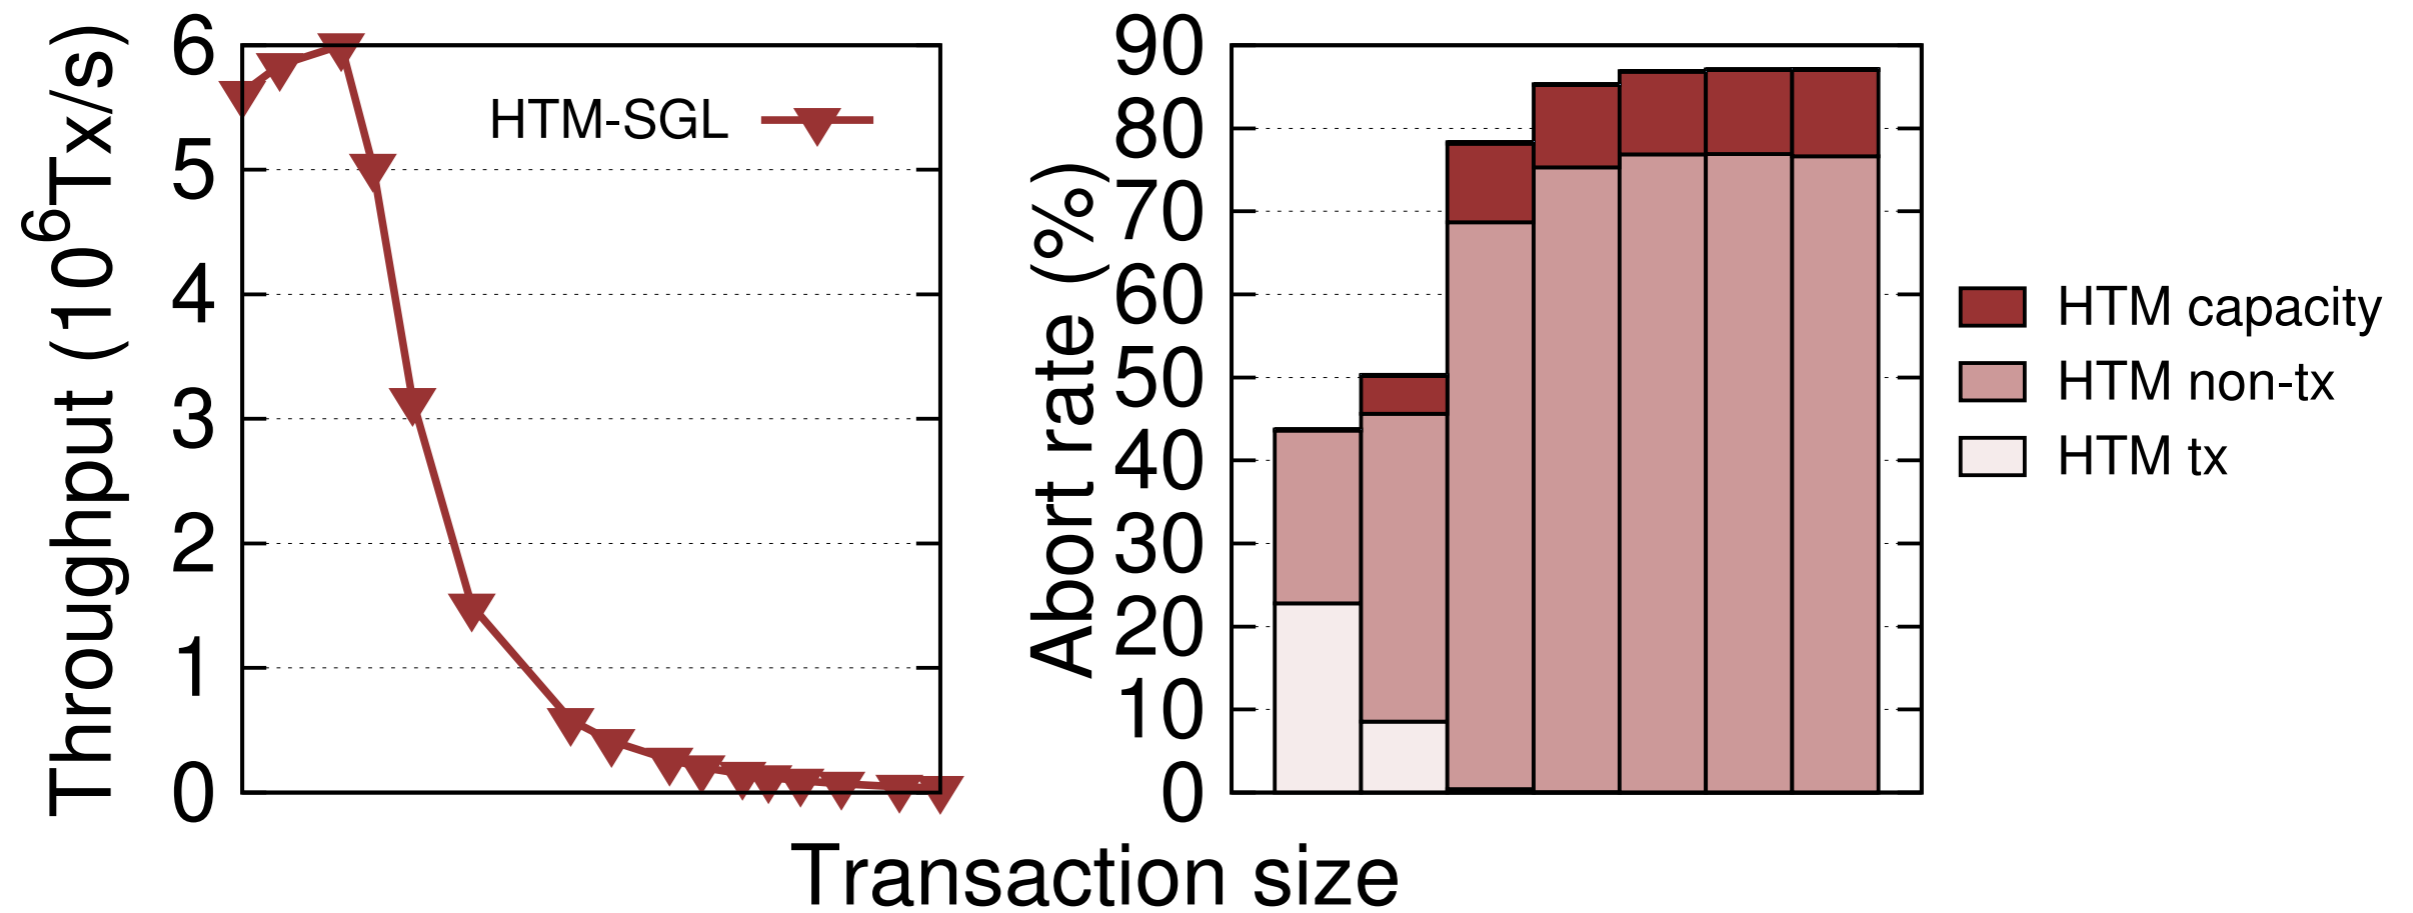
```
withdraw(account, value) {  
  __transaction{  
    if account.balance > value:  
      account.balance -= value;  
      return account.balance;  
    else  
      return -1;  
  }  
}
```

Transactional memory
implementation

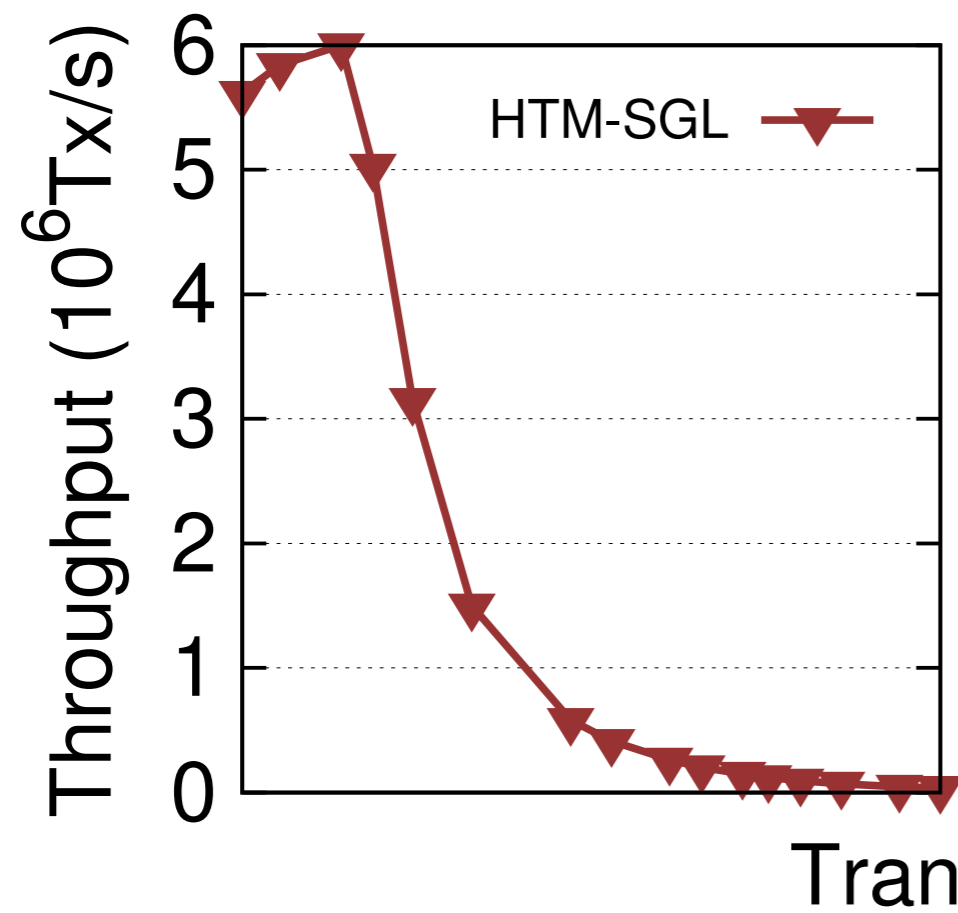
Hardware Transactional Memory

- Intel and IBM processors
- implemented in the cache coherence protocol
- cache line granularity
- best effort
 - S/W fallback is needed

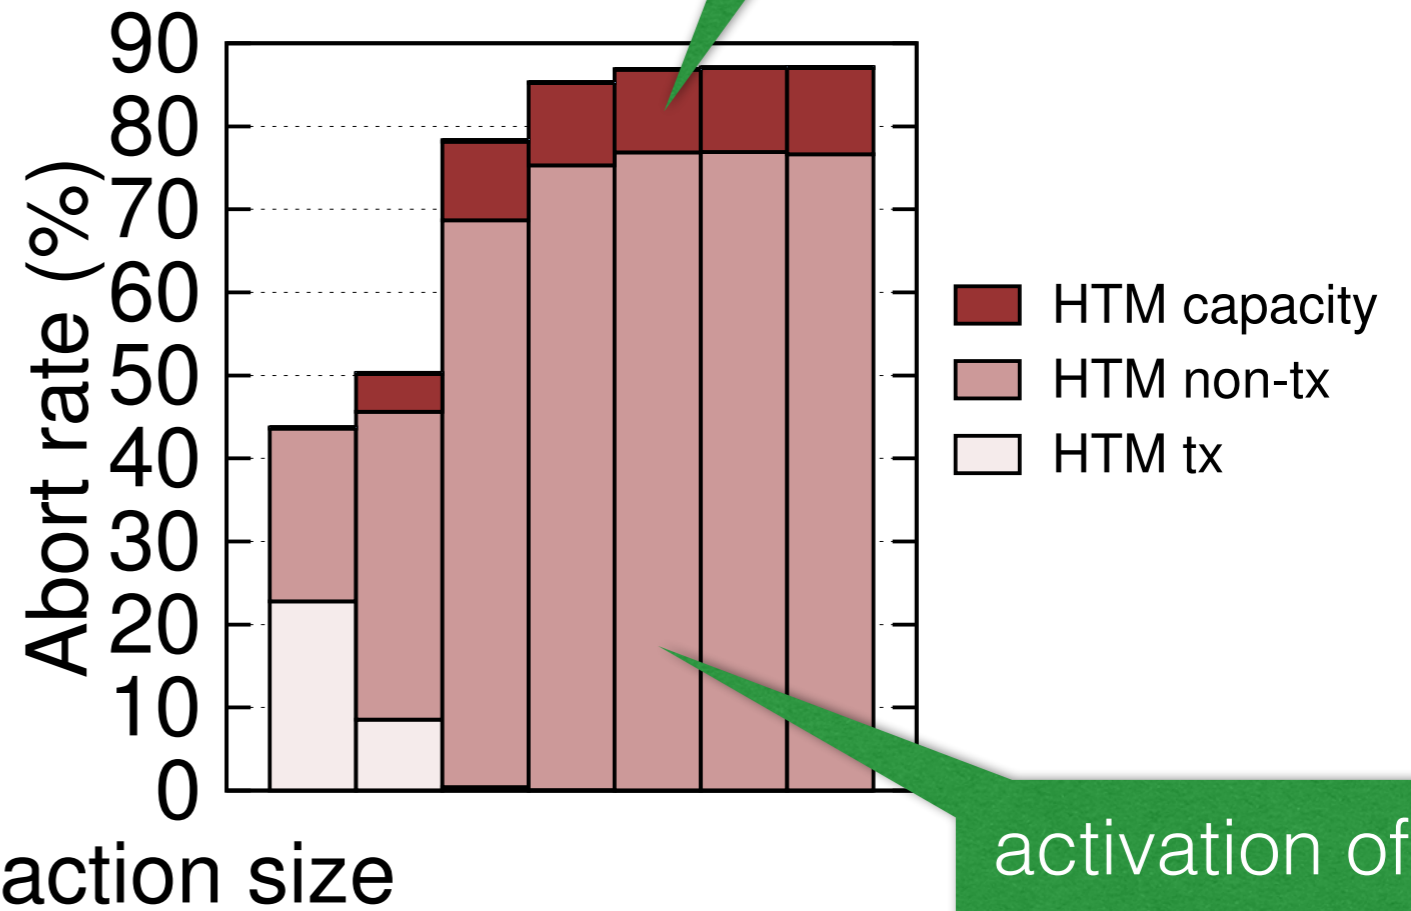

Capacity Limitations



Capacity Limitations



Capacity Limitations



POWER8-TM

- hardware/software co-design
- utilises specific features available in POWER8:
 - suspend/resume
 - ROTs
- to support execution of larger transactions


Rollback-only Transaction

- lightweight transaction type
- updates are applied atomically
- does not track the reads
 - theoretically infinite read-set
- not **serialisable**

ROTs

$X = 0$


Thread 1



inconsistent
value

$X = 0$


Thread 2



ROTs

$X = 0$

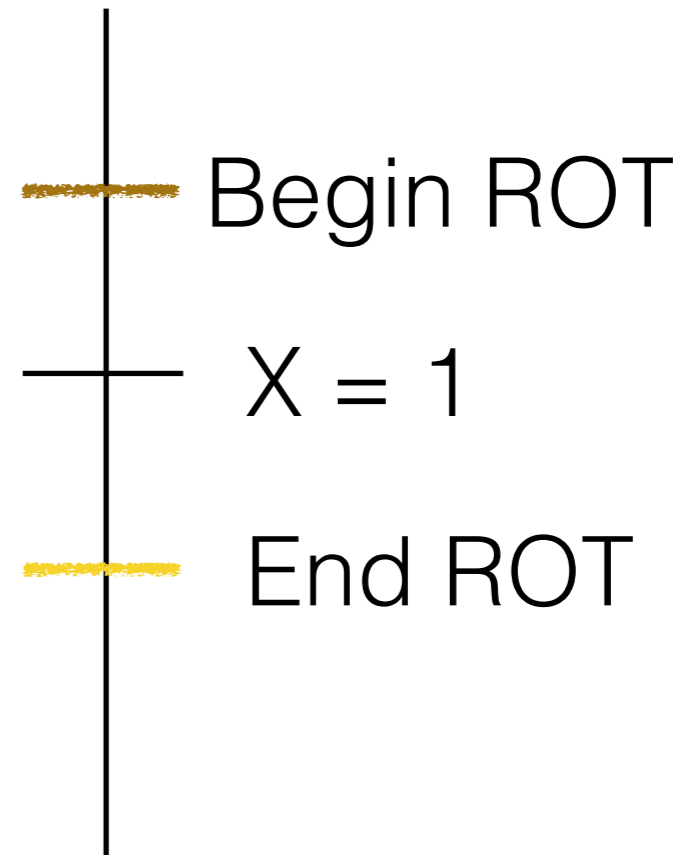
Thread 1



inconsistent value

$X = 0$

Thread 2




WAR

ROTs


$X = 0$

Thread 1



$X = 0$


Thread 2



ROTs

$X = 0$

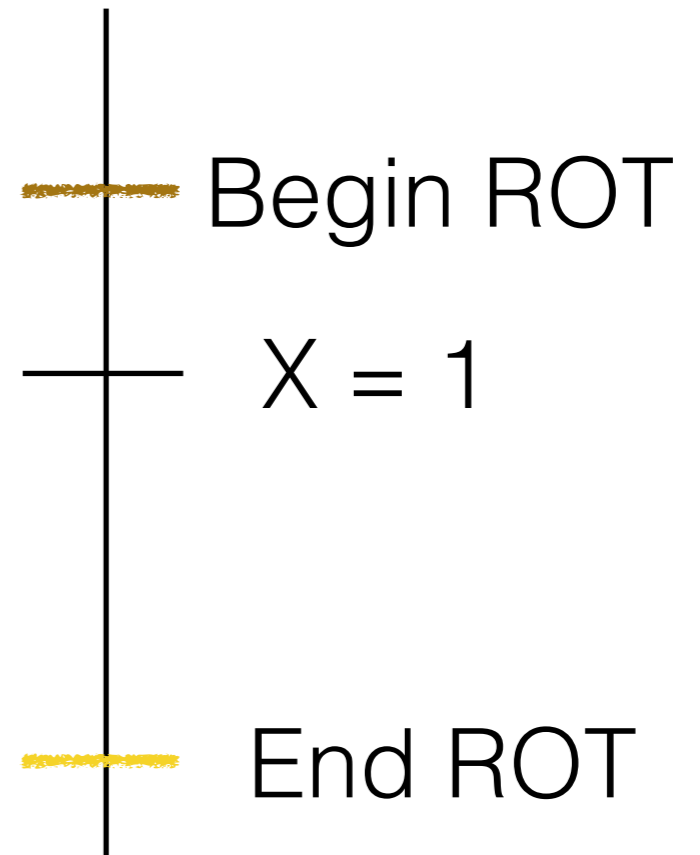
Thread 1



consistent

$X = 0$

Thread 2




new value can only appear now

ROTs

$X = 0$


Thread 1



consistent

$X = 0$

Thread 2

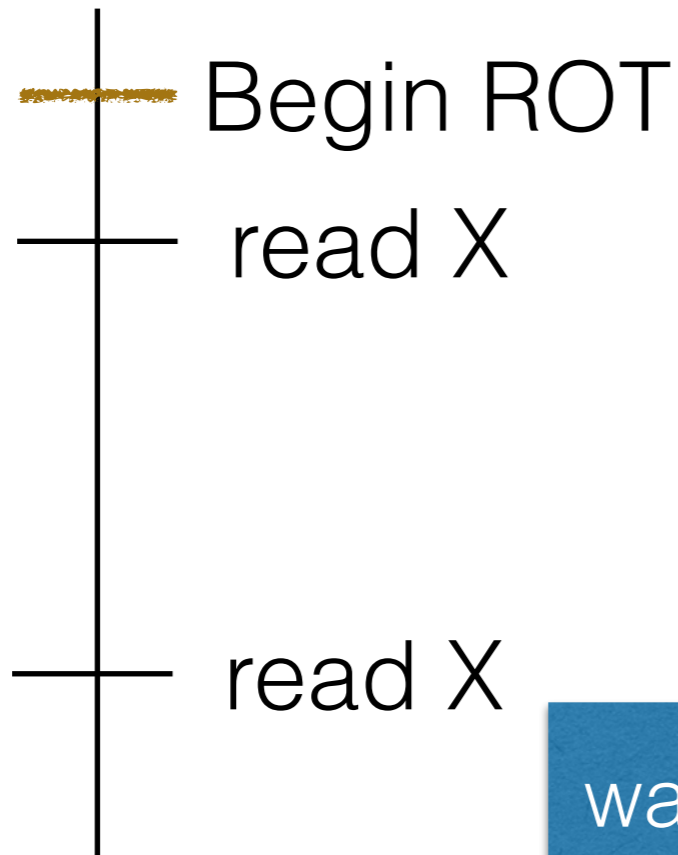


new value can only appear now

ROTs


$X = 0$

Thread 1



$X = 0$

Thread 2



wait for concurrent
ROTs
non-transactionally

ROTs

$X = 0$

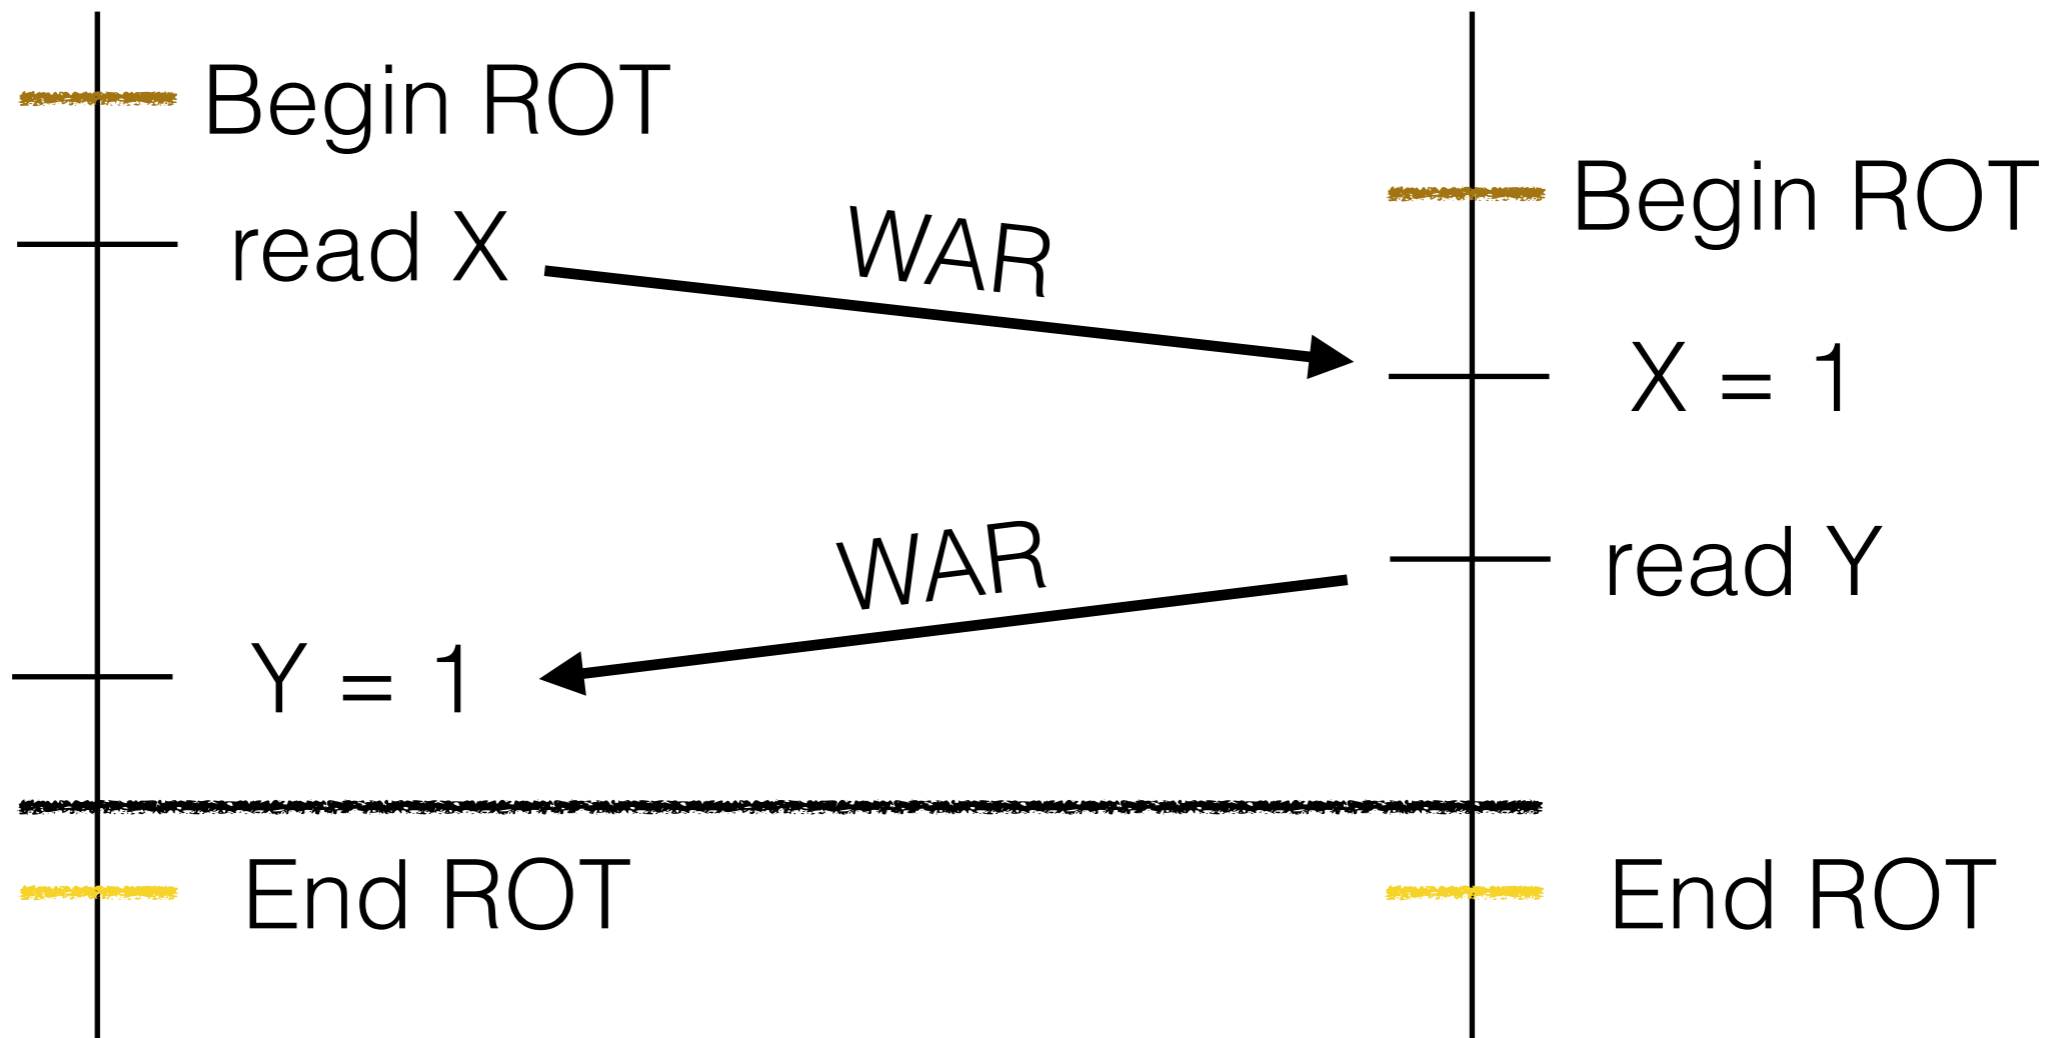
$Y = 0$

$X = 0$

$Y = 0$

Thread 1

Thread 2



ROTs

$X = 0$

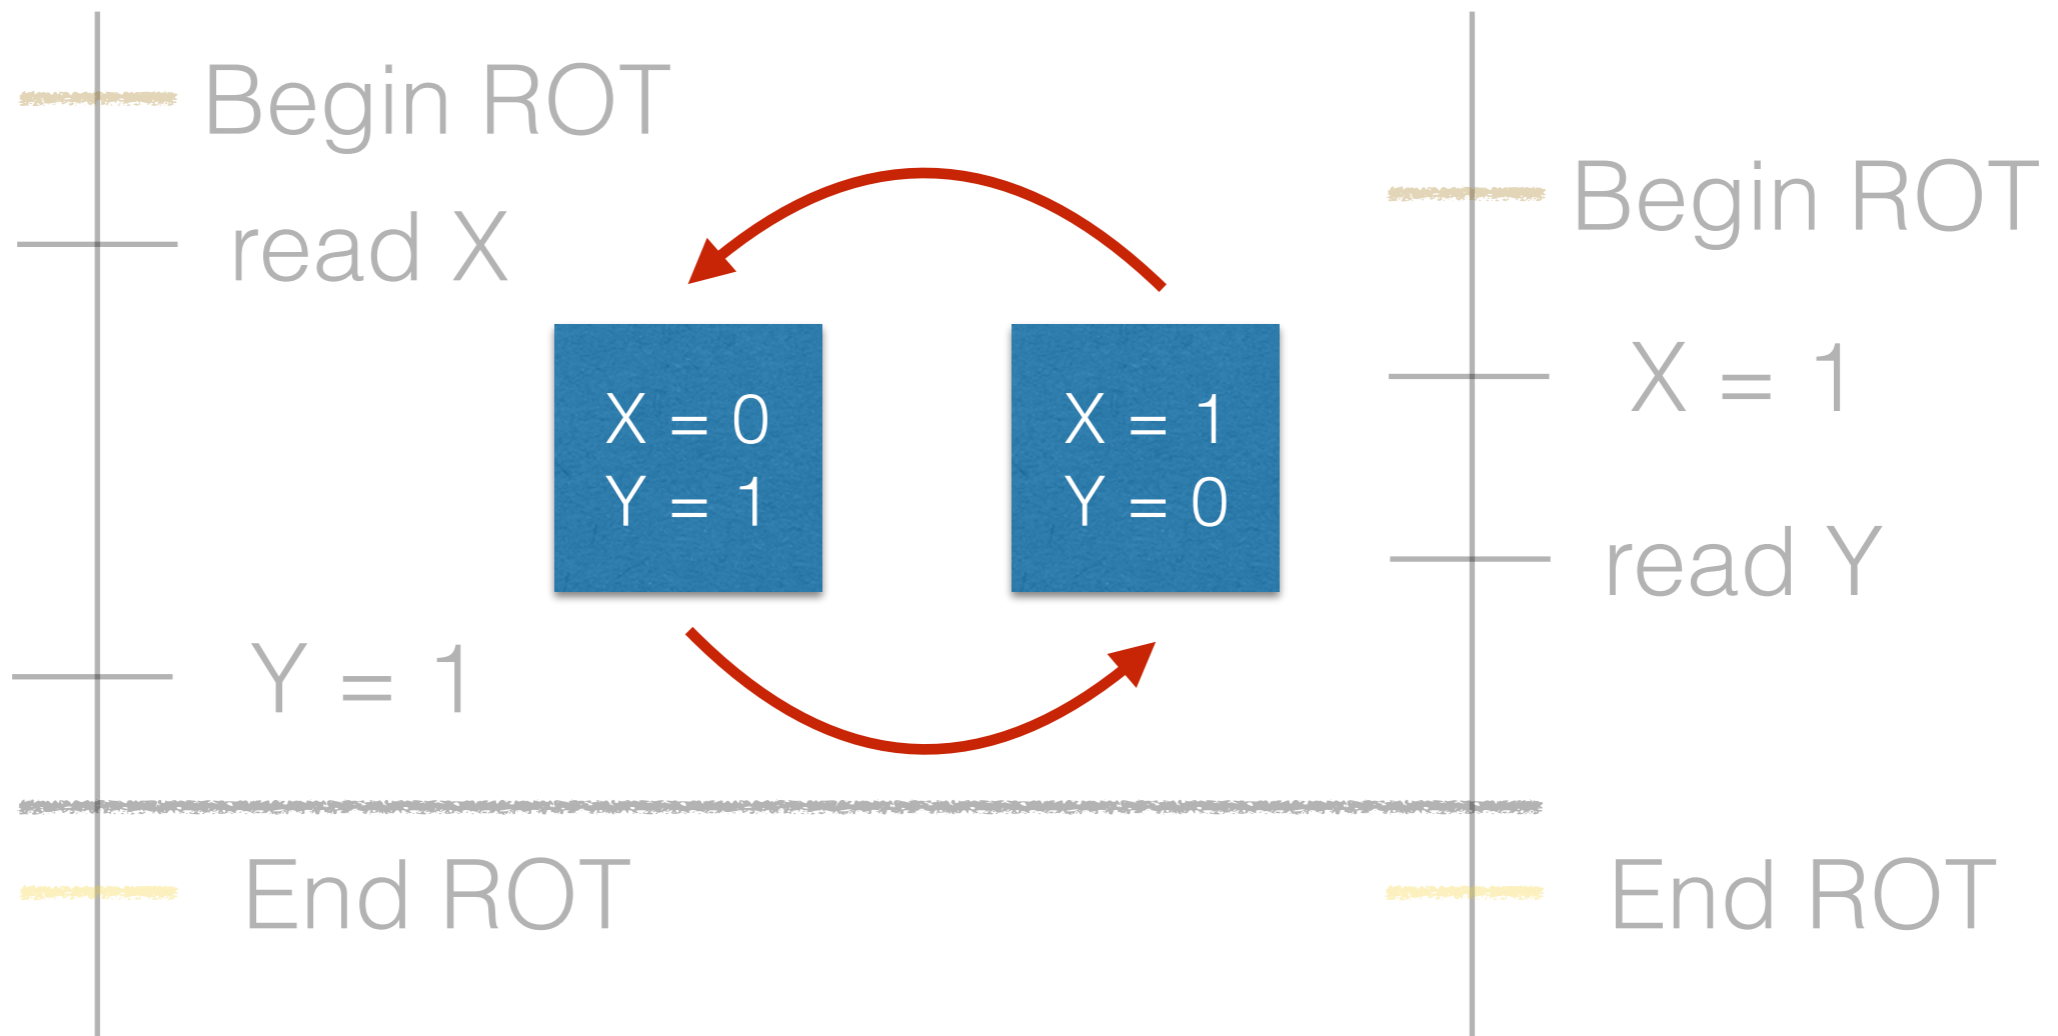
$Y = 0$

$X = 0$

$Y = 0$

Thread 1

Thread 2



Touch-to-Validate


- core algorithm of P8TM
- to make concurrent execution of ROTs safe and serialisable
- basic intuition: convert WAR to RAW

T2V

X = 0

Y = 0

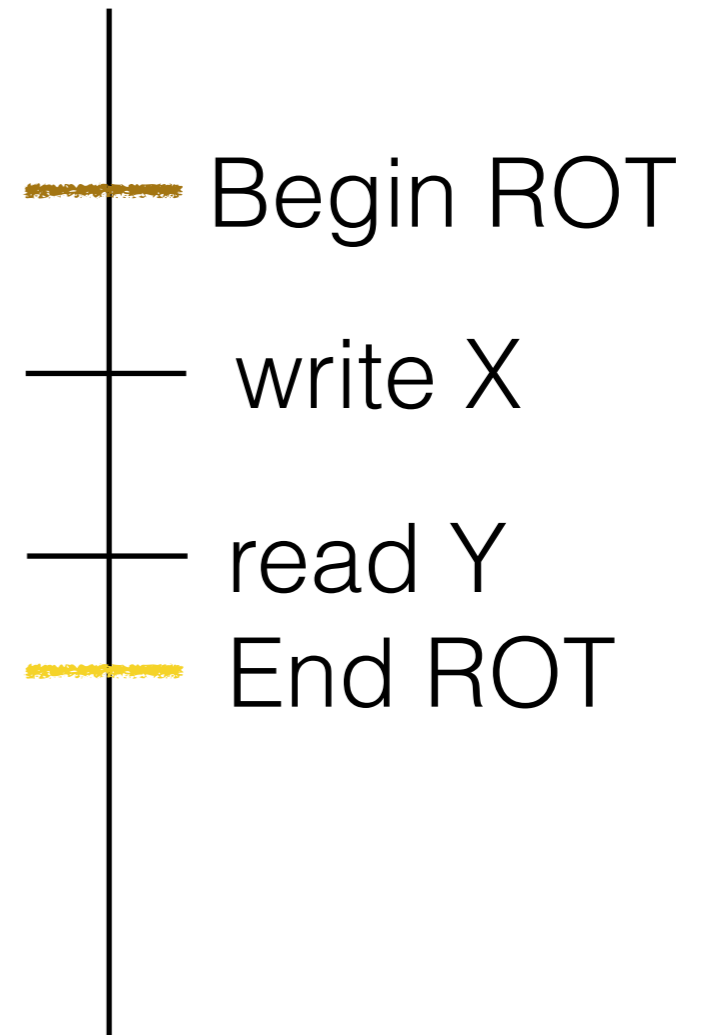
Thread 1



X = 0

Y = 0

Thread 2




T2V

X = 0

Y = 0

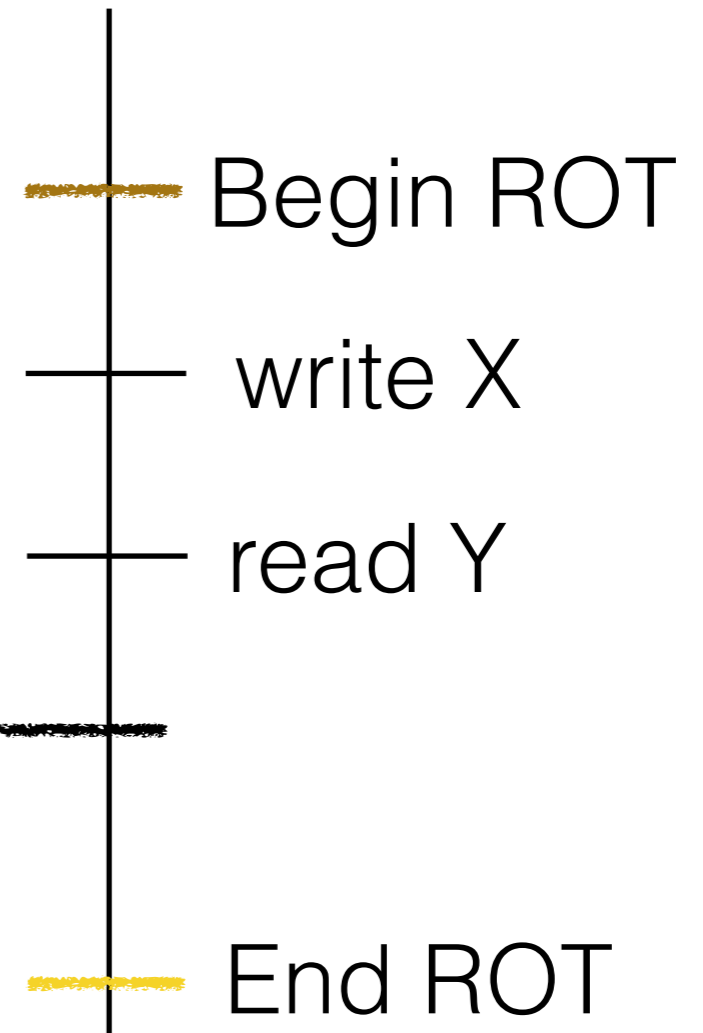
Thread 1



X = 0

Y = 0

Thread 2

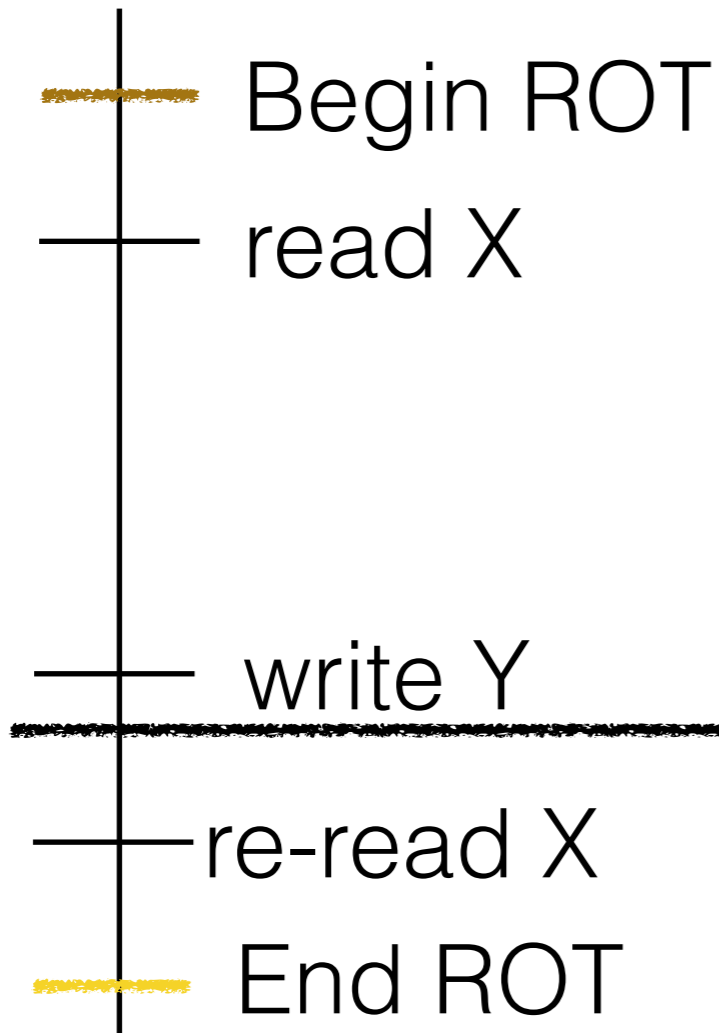


T2V

X = 0

Y = 0

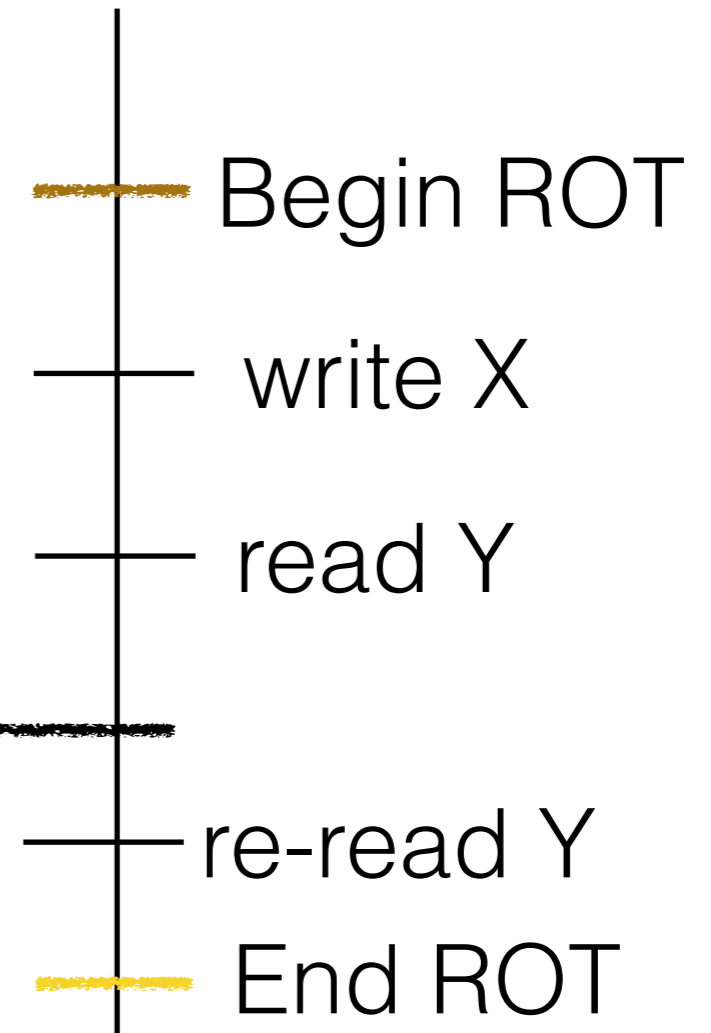
Thread 1



X = 0

Y = 0

Thread 2

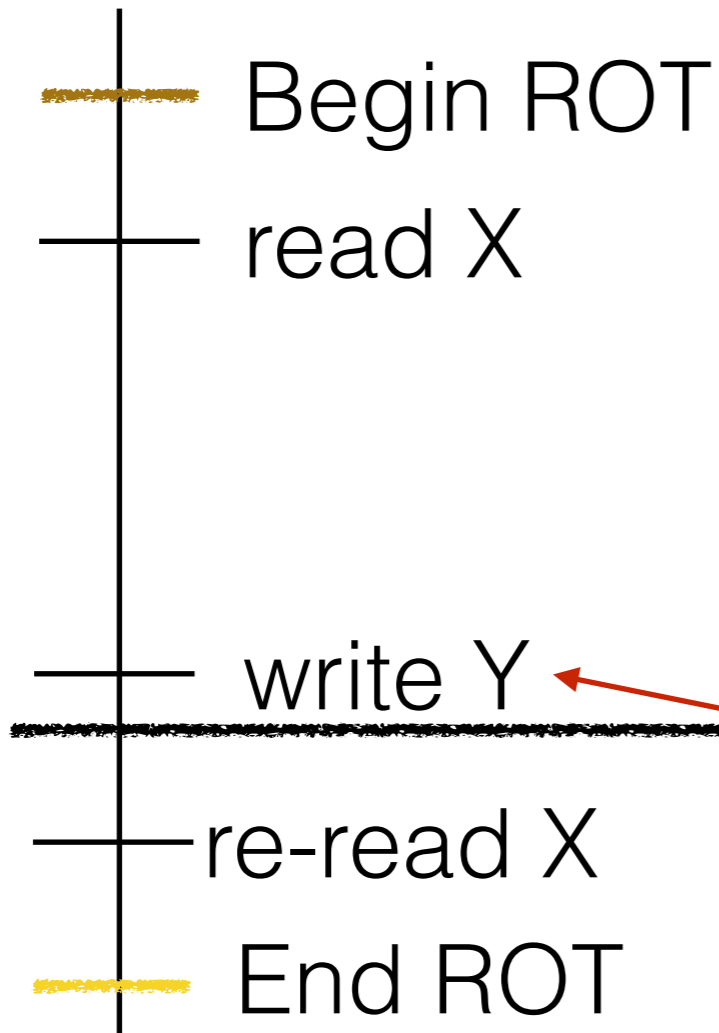


T2V

X = 0

Y = 0

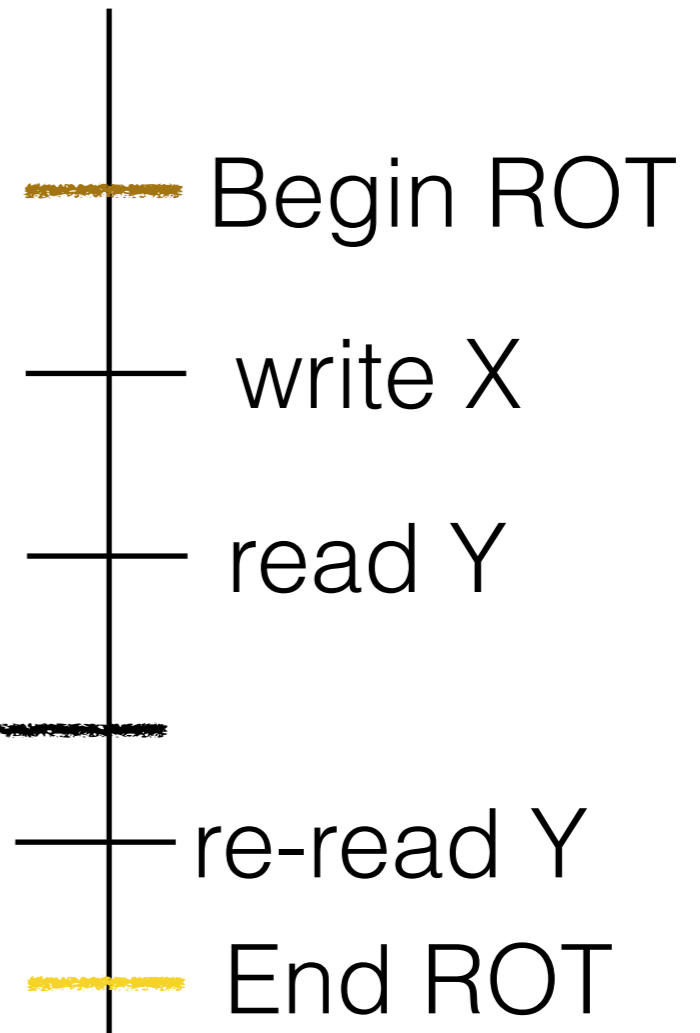
Thread 1



X = 0

Y = 0

Thread 2



T2V

X = 0

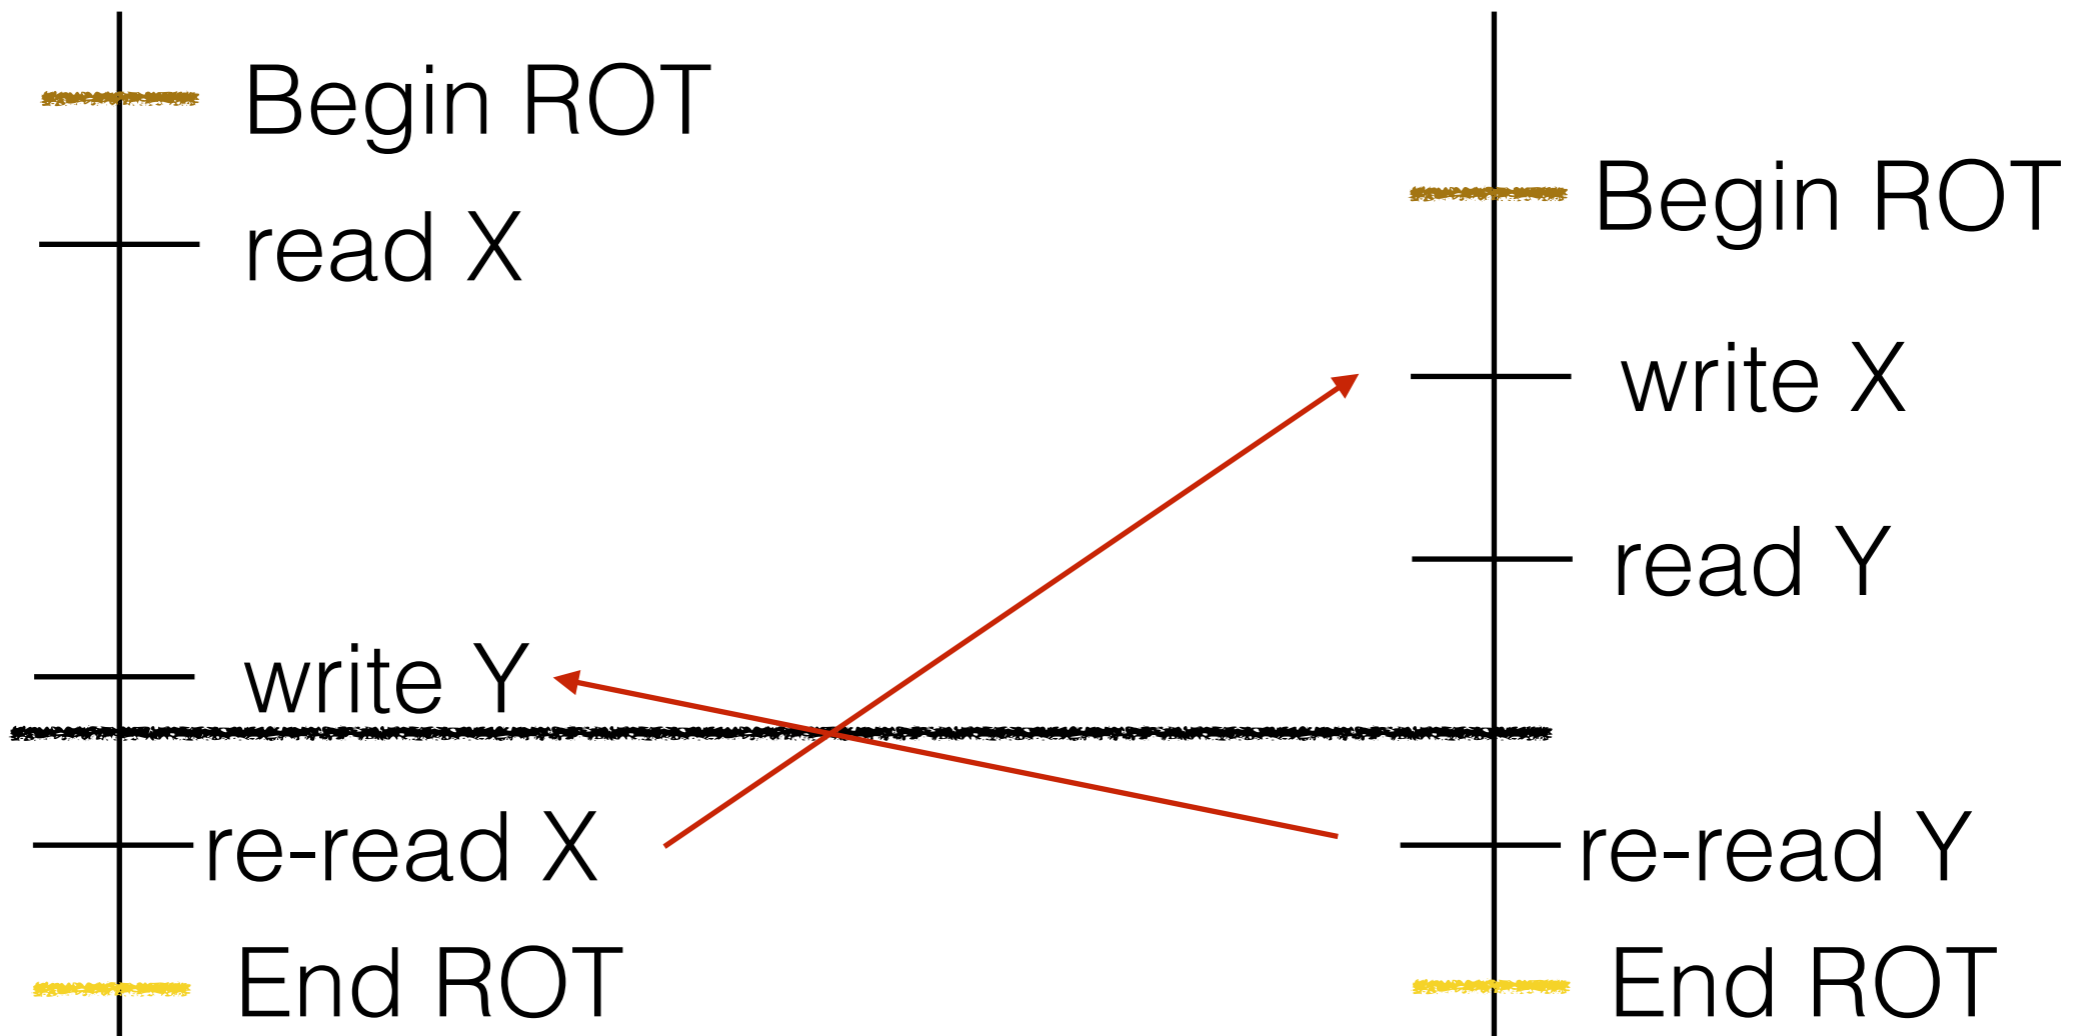
Y = 0

X = 0

Y = 0

Thread 1

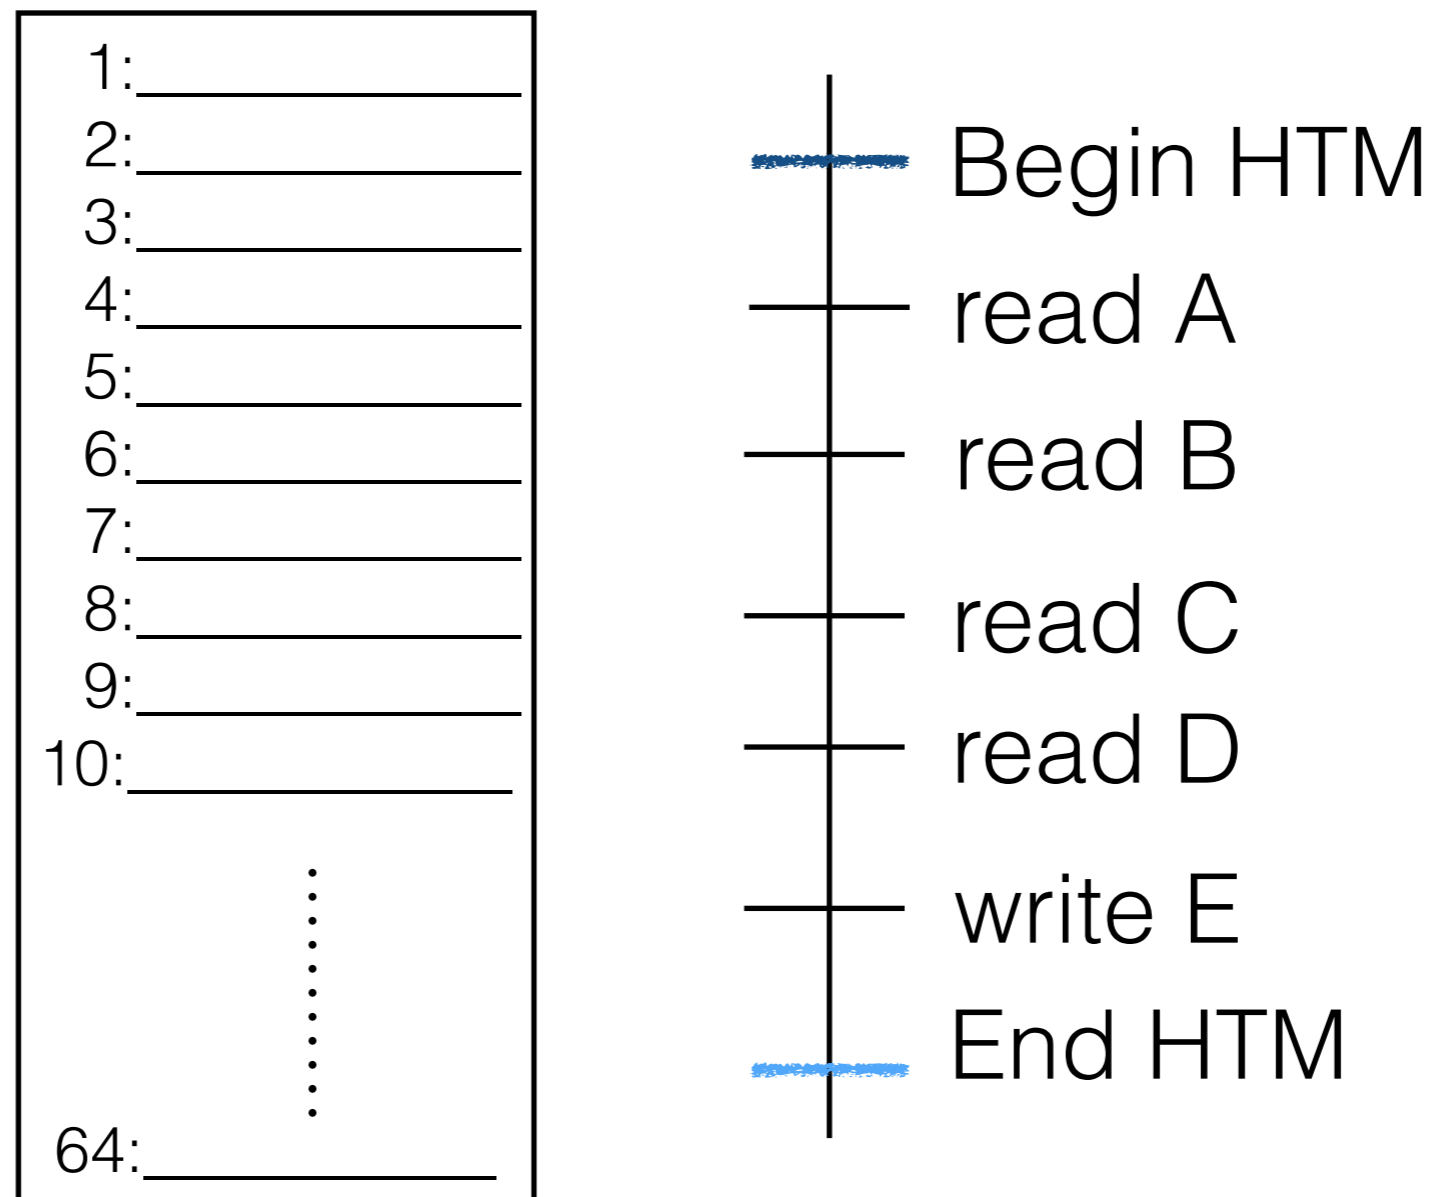
Thread 2



T2V

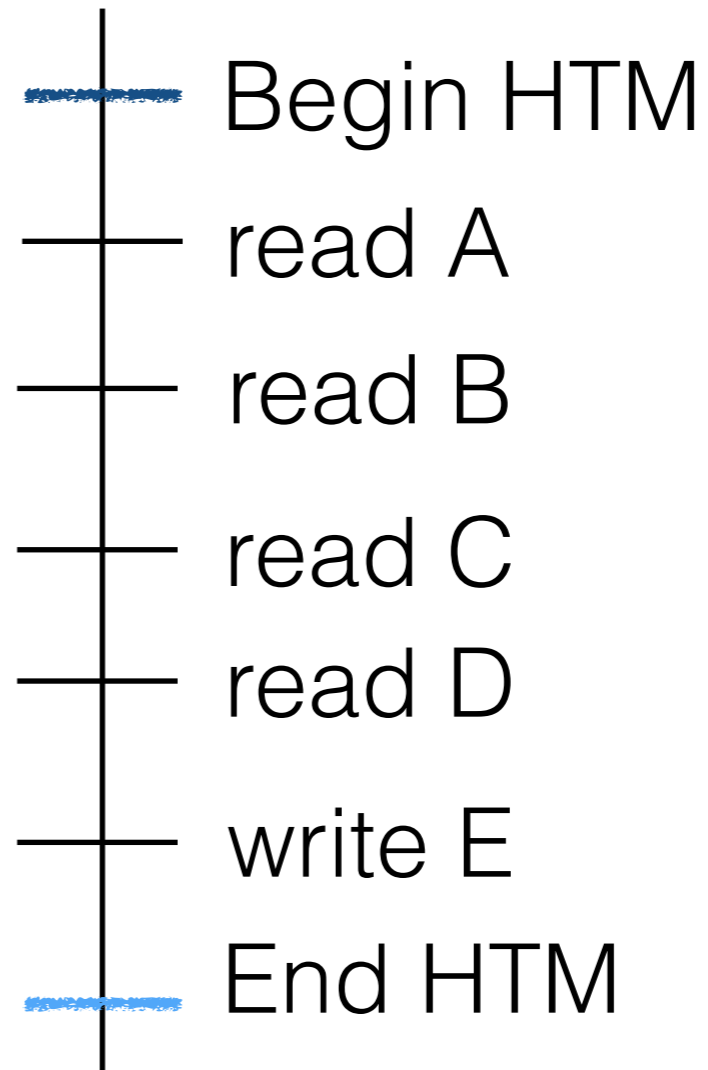

- needs to track only the addresses
- this must be done in software
- how can software outperform hardware?

TMCAM



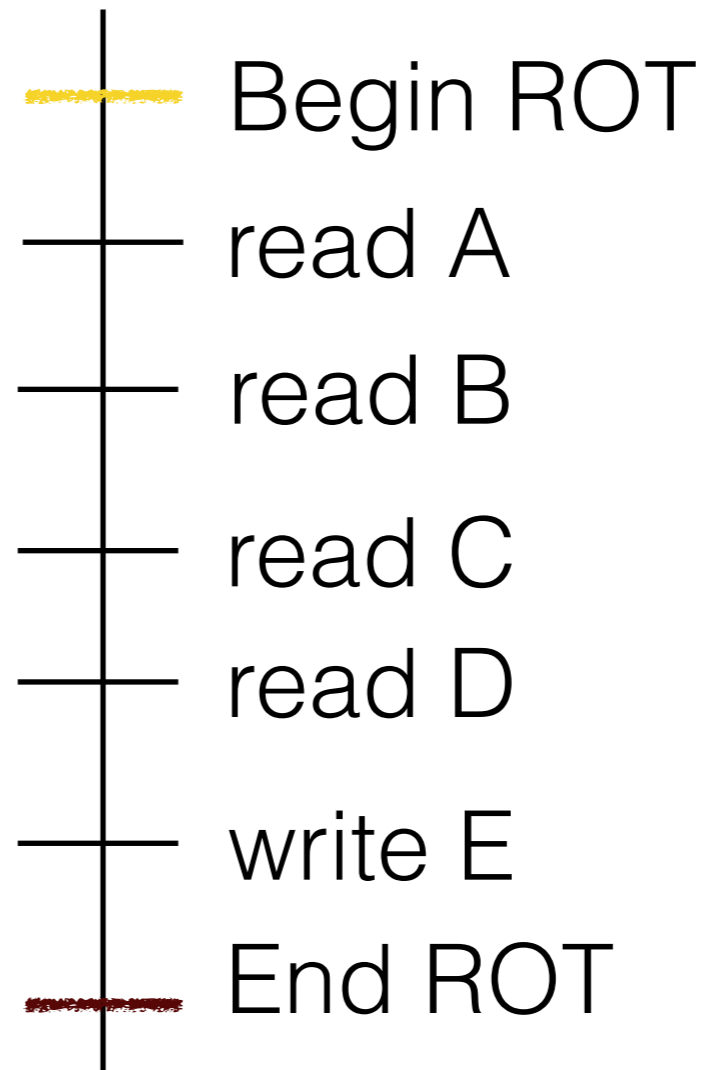

TMCAM

TMCAM





TMCAM



Read-set Tracking



Read-set Tracking




Read-set Tracking




Read-set Tracking

8 bytes




128bytes




Read-set Tracking

8 bytes



up to 16x
larger read-set

128bytes



HTM

- transactions may fit in HTM
- we need to avoid extra overheads of using ROTs
- try first in HTM, if it overflows, fallback to ROT
- how can HTMs and ROTs run concurrently?

HTM + ROT


$X = 0$

$Y = 0$

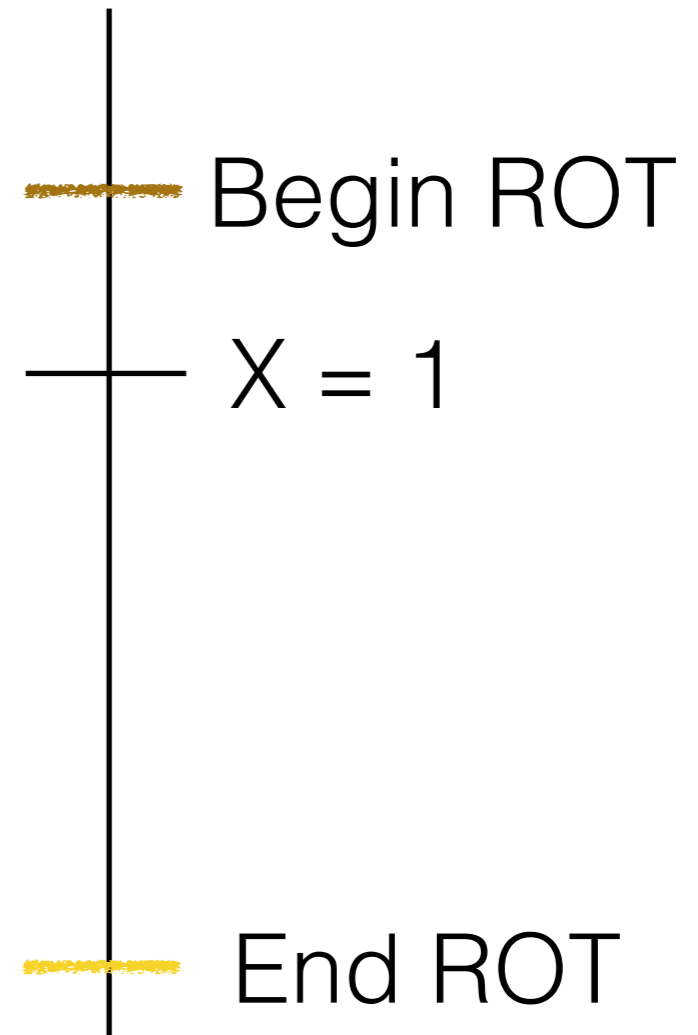
$X = 0$

$Y = 0$

Thread 1



Thread 2



HTM + ROT

$X = 0$



$Y = 0$

$X = 0$

$Y = 0$

Thread 1

Thread 2




HTM is protected
by H/W

HTM + ROT

$X = 0$

$Y = 0$

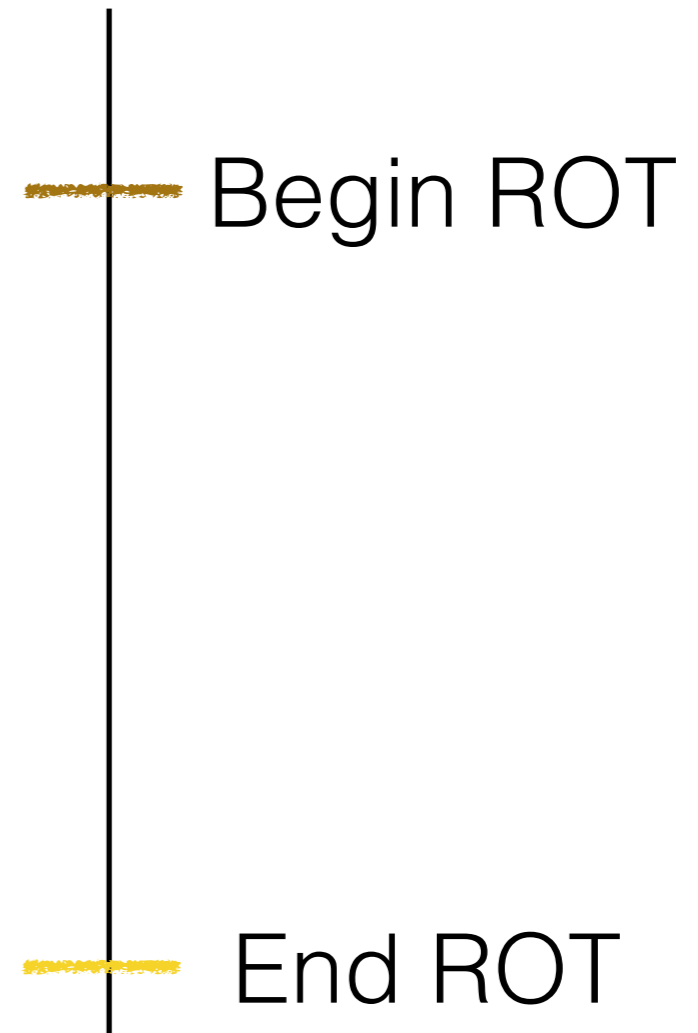
Thread 1



$X = 0$

$Y = 0$

Thread 2




HTM is protected
by H/W

HTM + ROT

$X = 0$

$Y = 0$

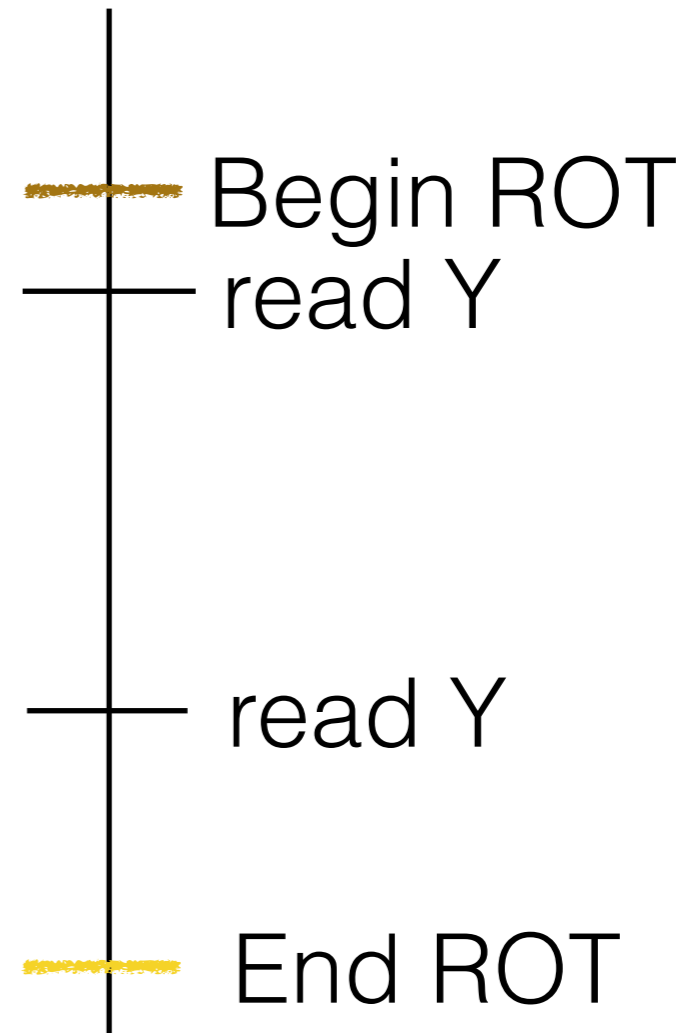
Thread 1



$X = 0$

$Y = 0$

Thread 2




HTM is protected
by H/W

HTM + ROT

X = 0

Y = 0

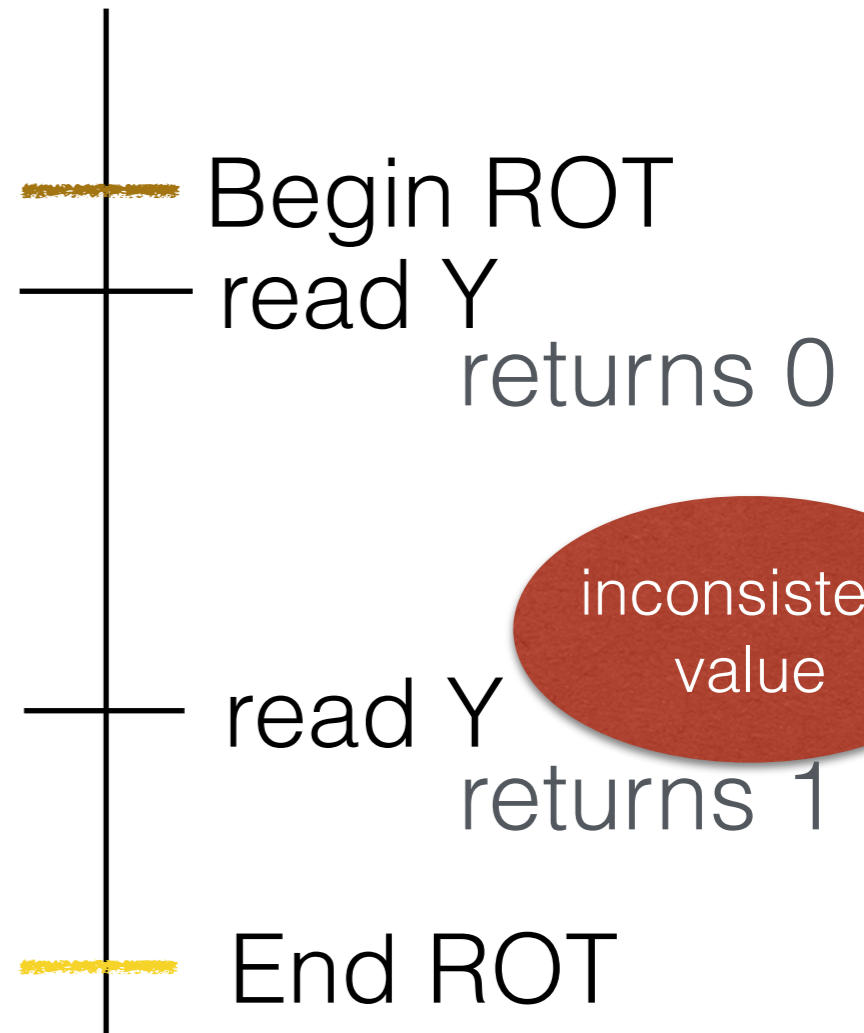
Thread 1



X = 0

Y = 0

Thread 2



HTM is protected
by H/W

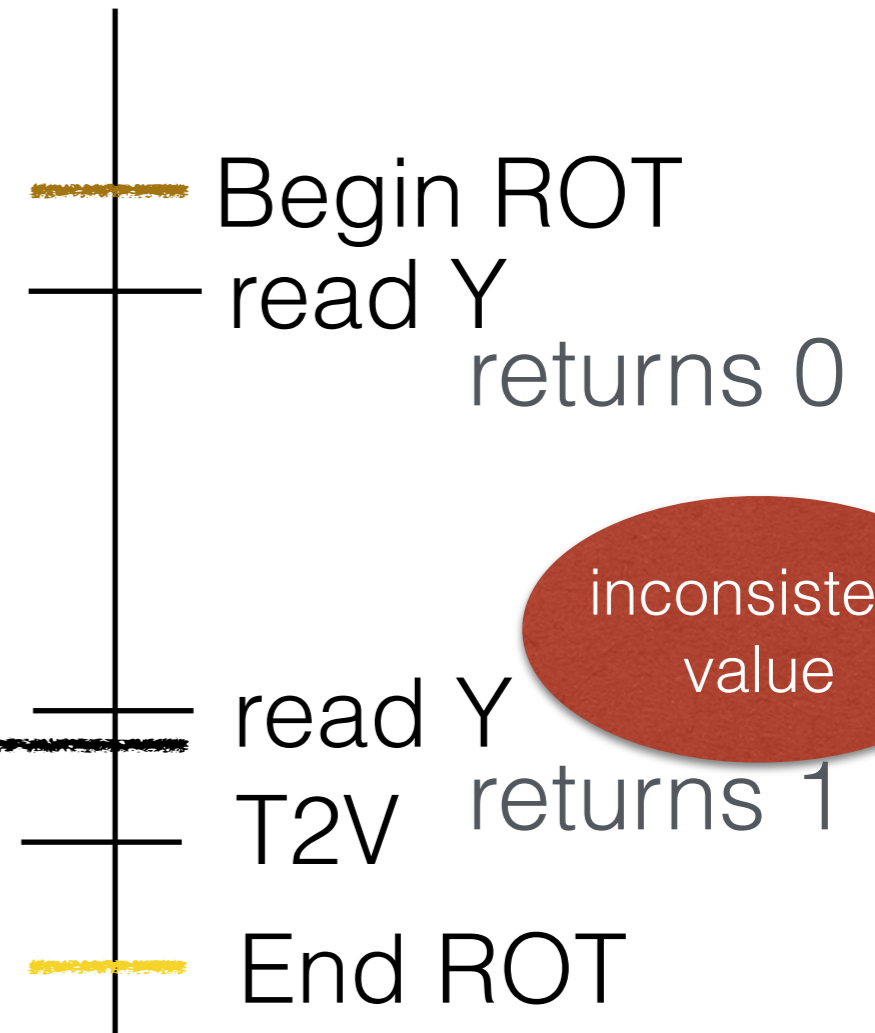

HTM + ROT

X = 0
Y = 0

X = 0
Y = 0

Thread 1

Thread 2



HTM is protected by H/W

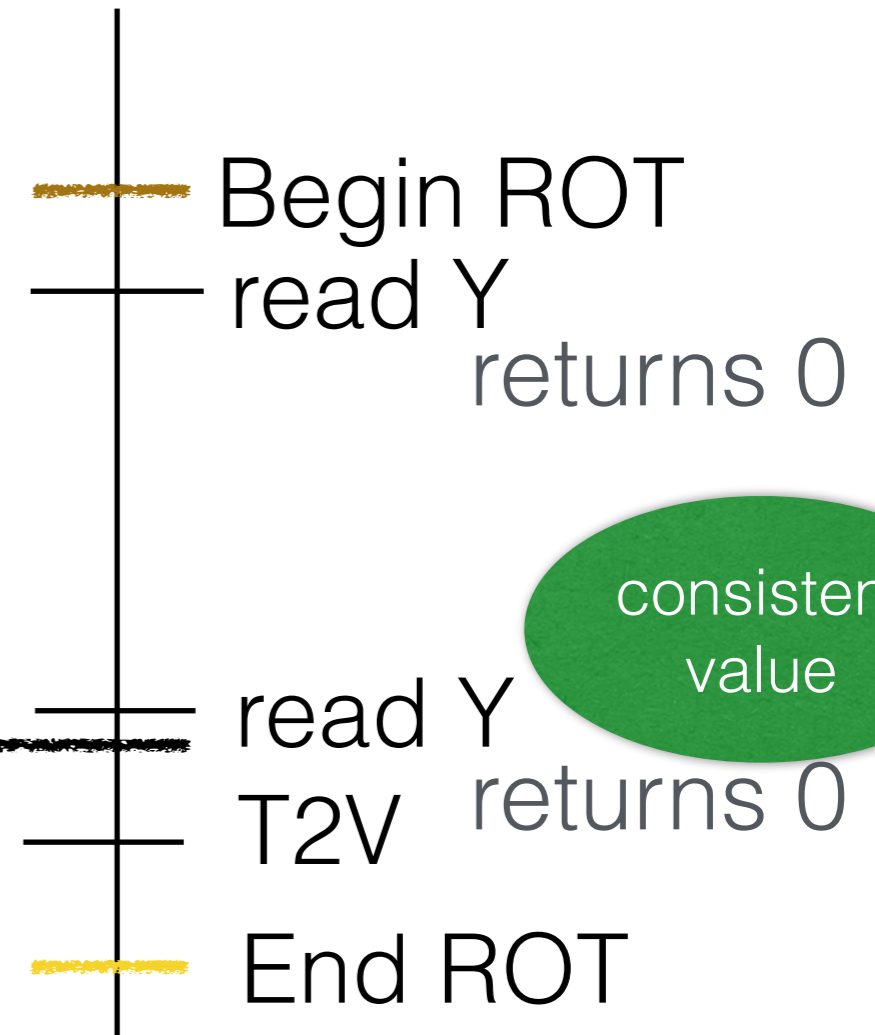
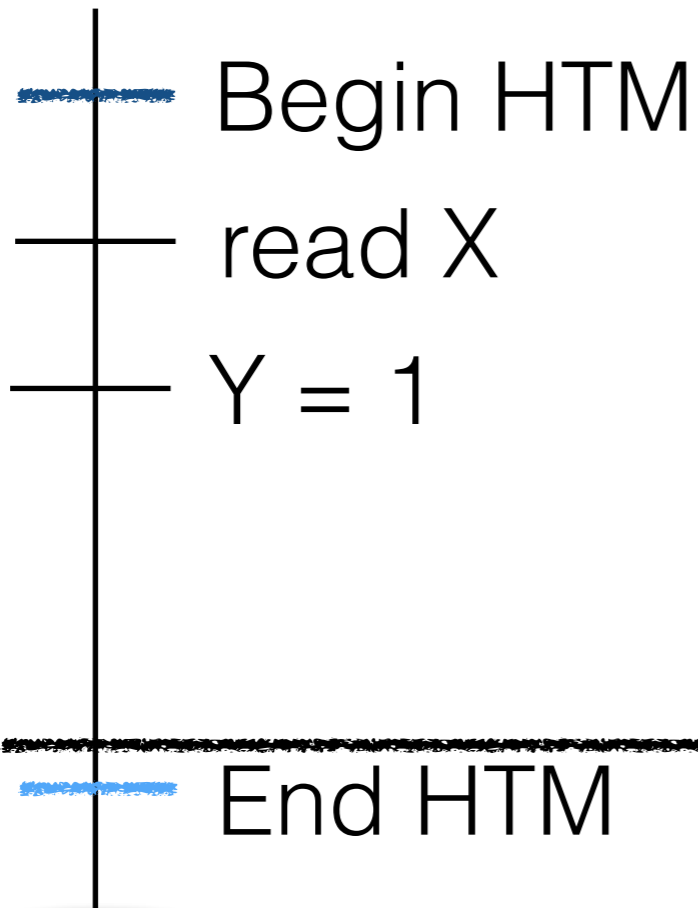
HTM + ROT

X = 0
Y = 0

X = 0
Y = 0

Thread 1

Thread 2



using S/R

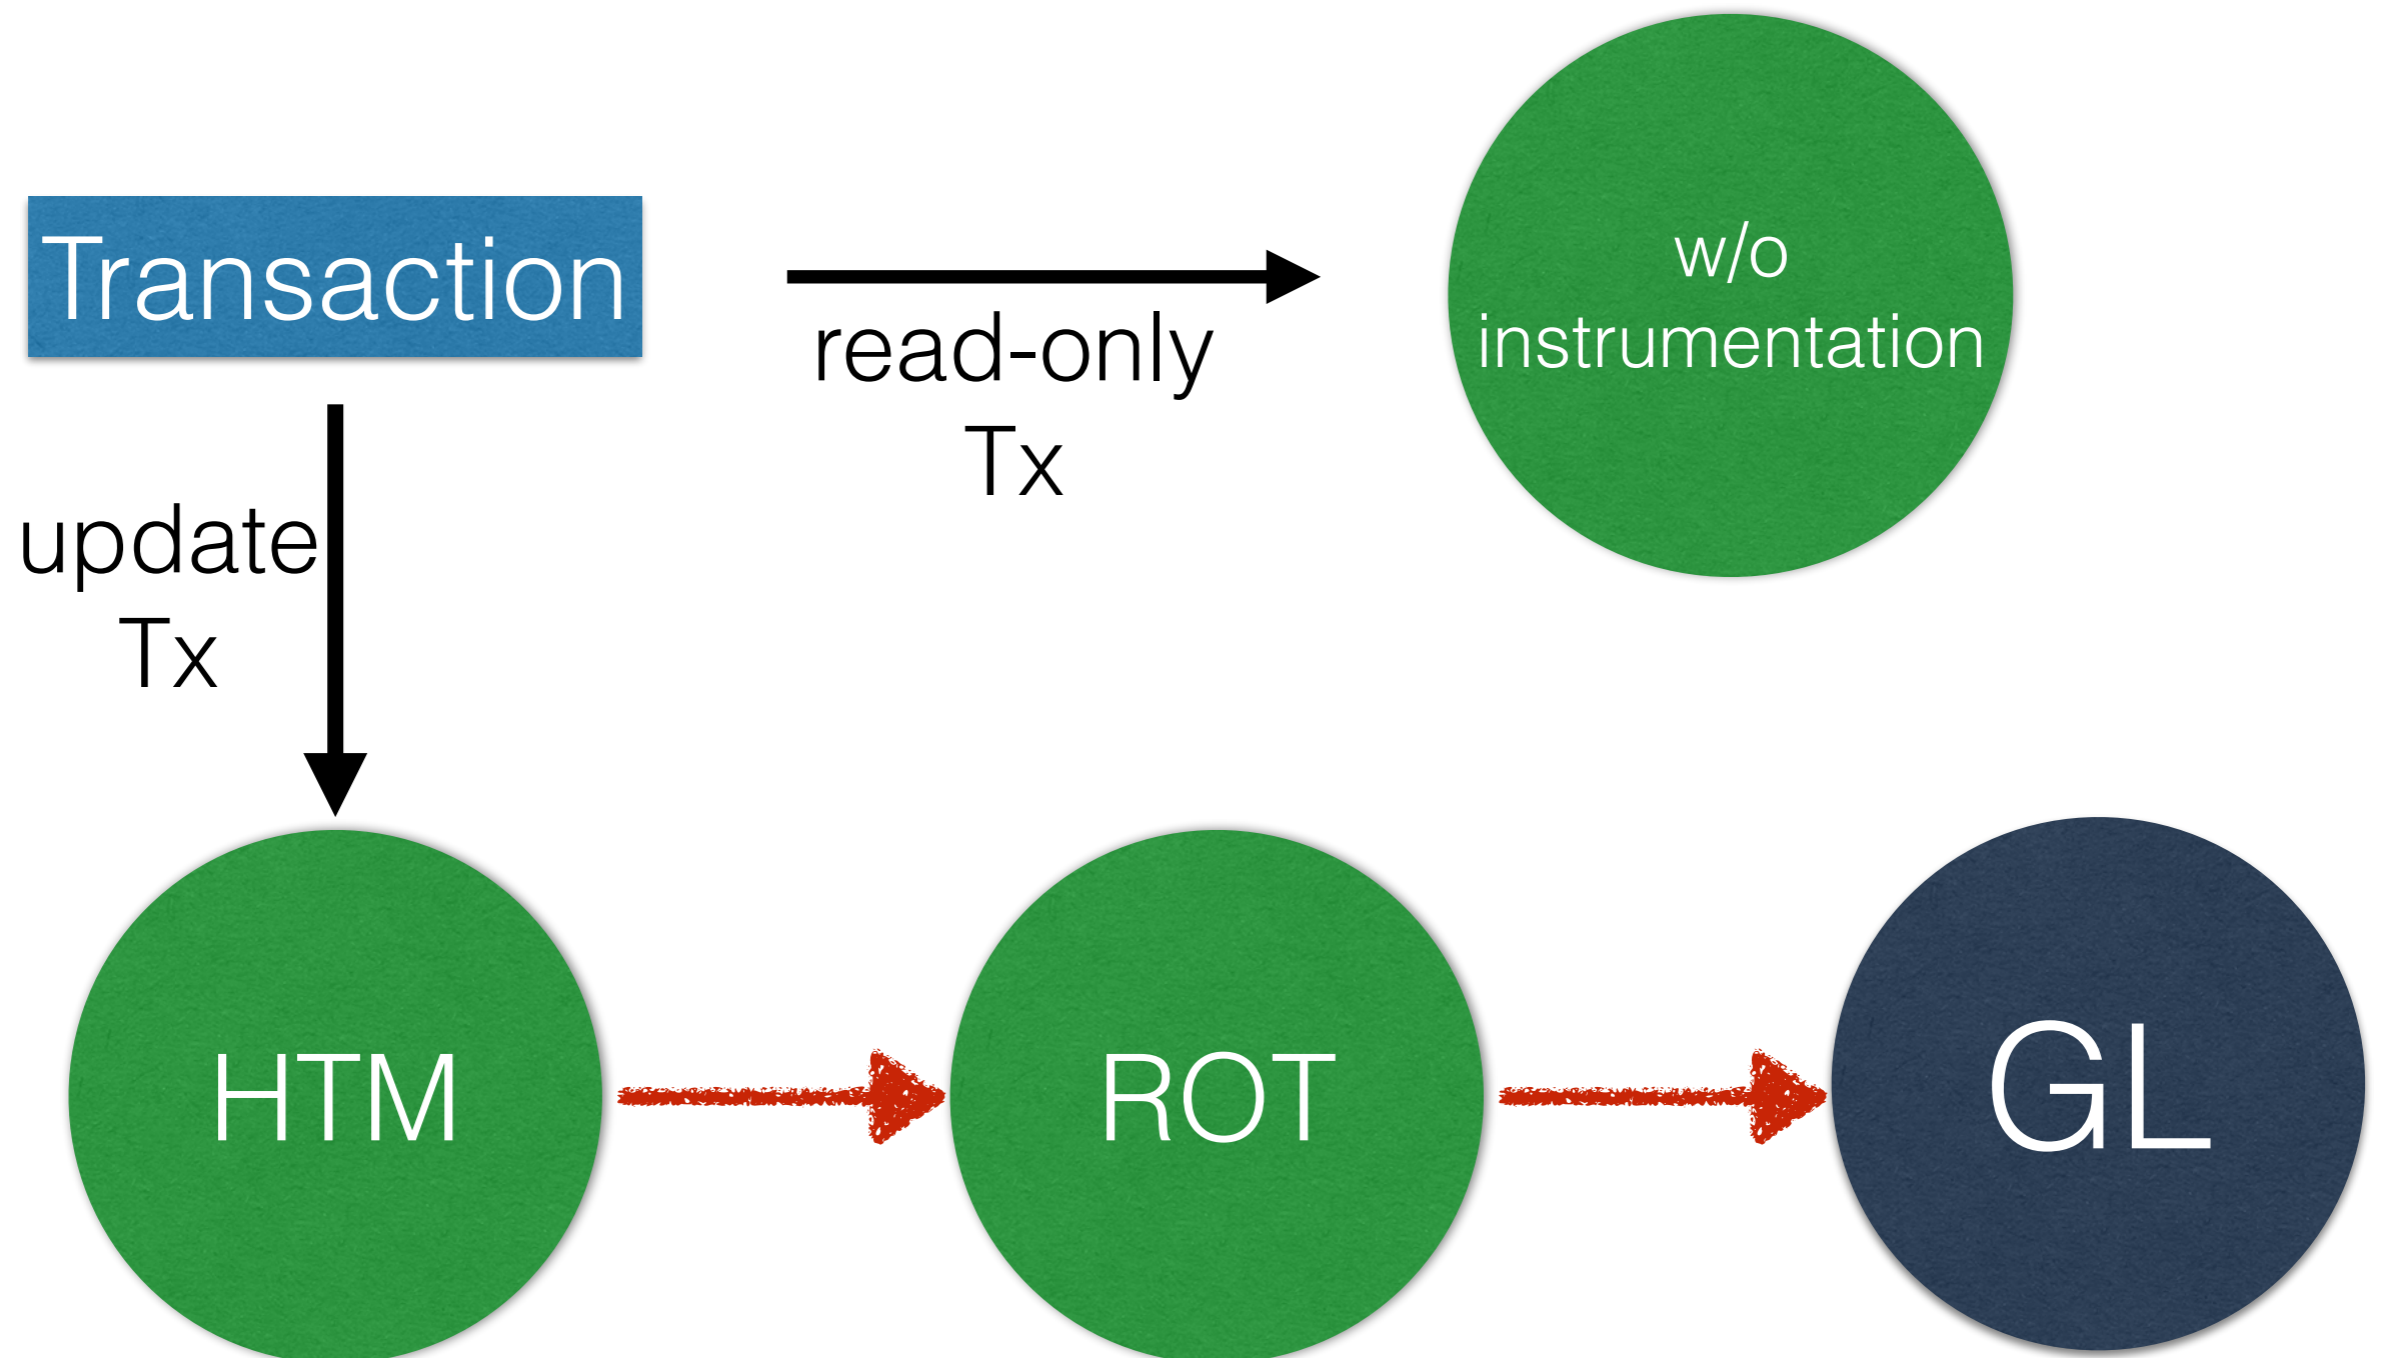
consistent value

HTM is protected by H/W

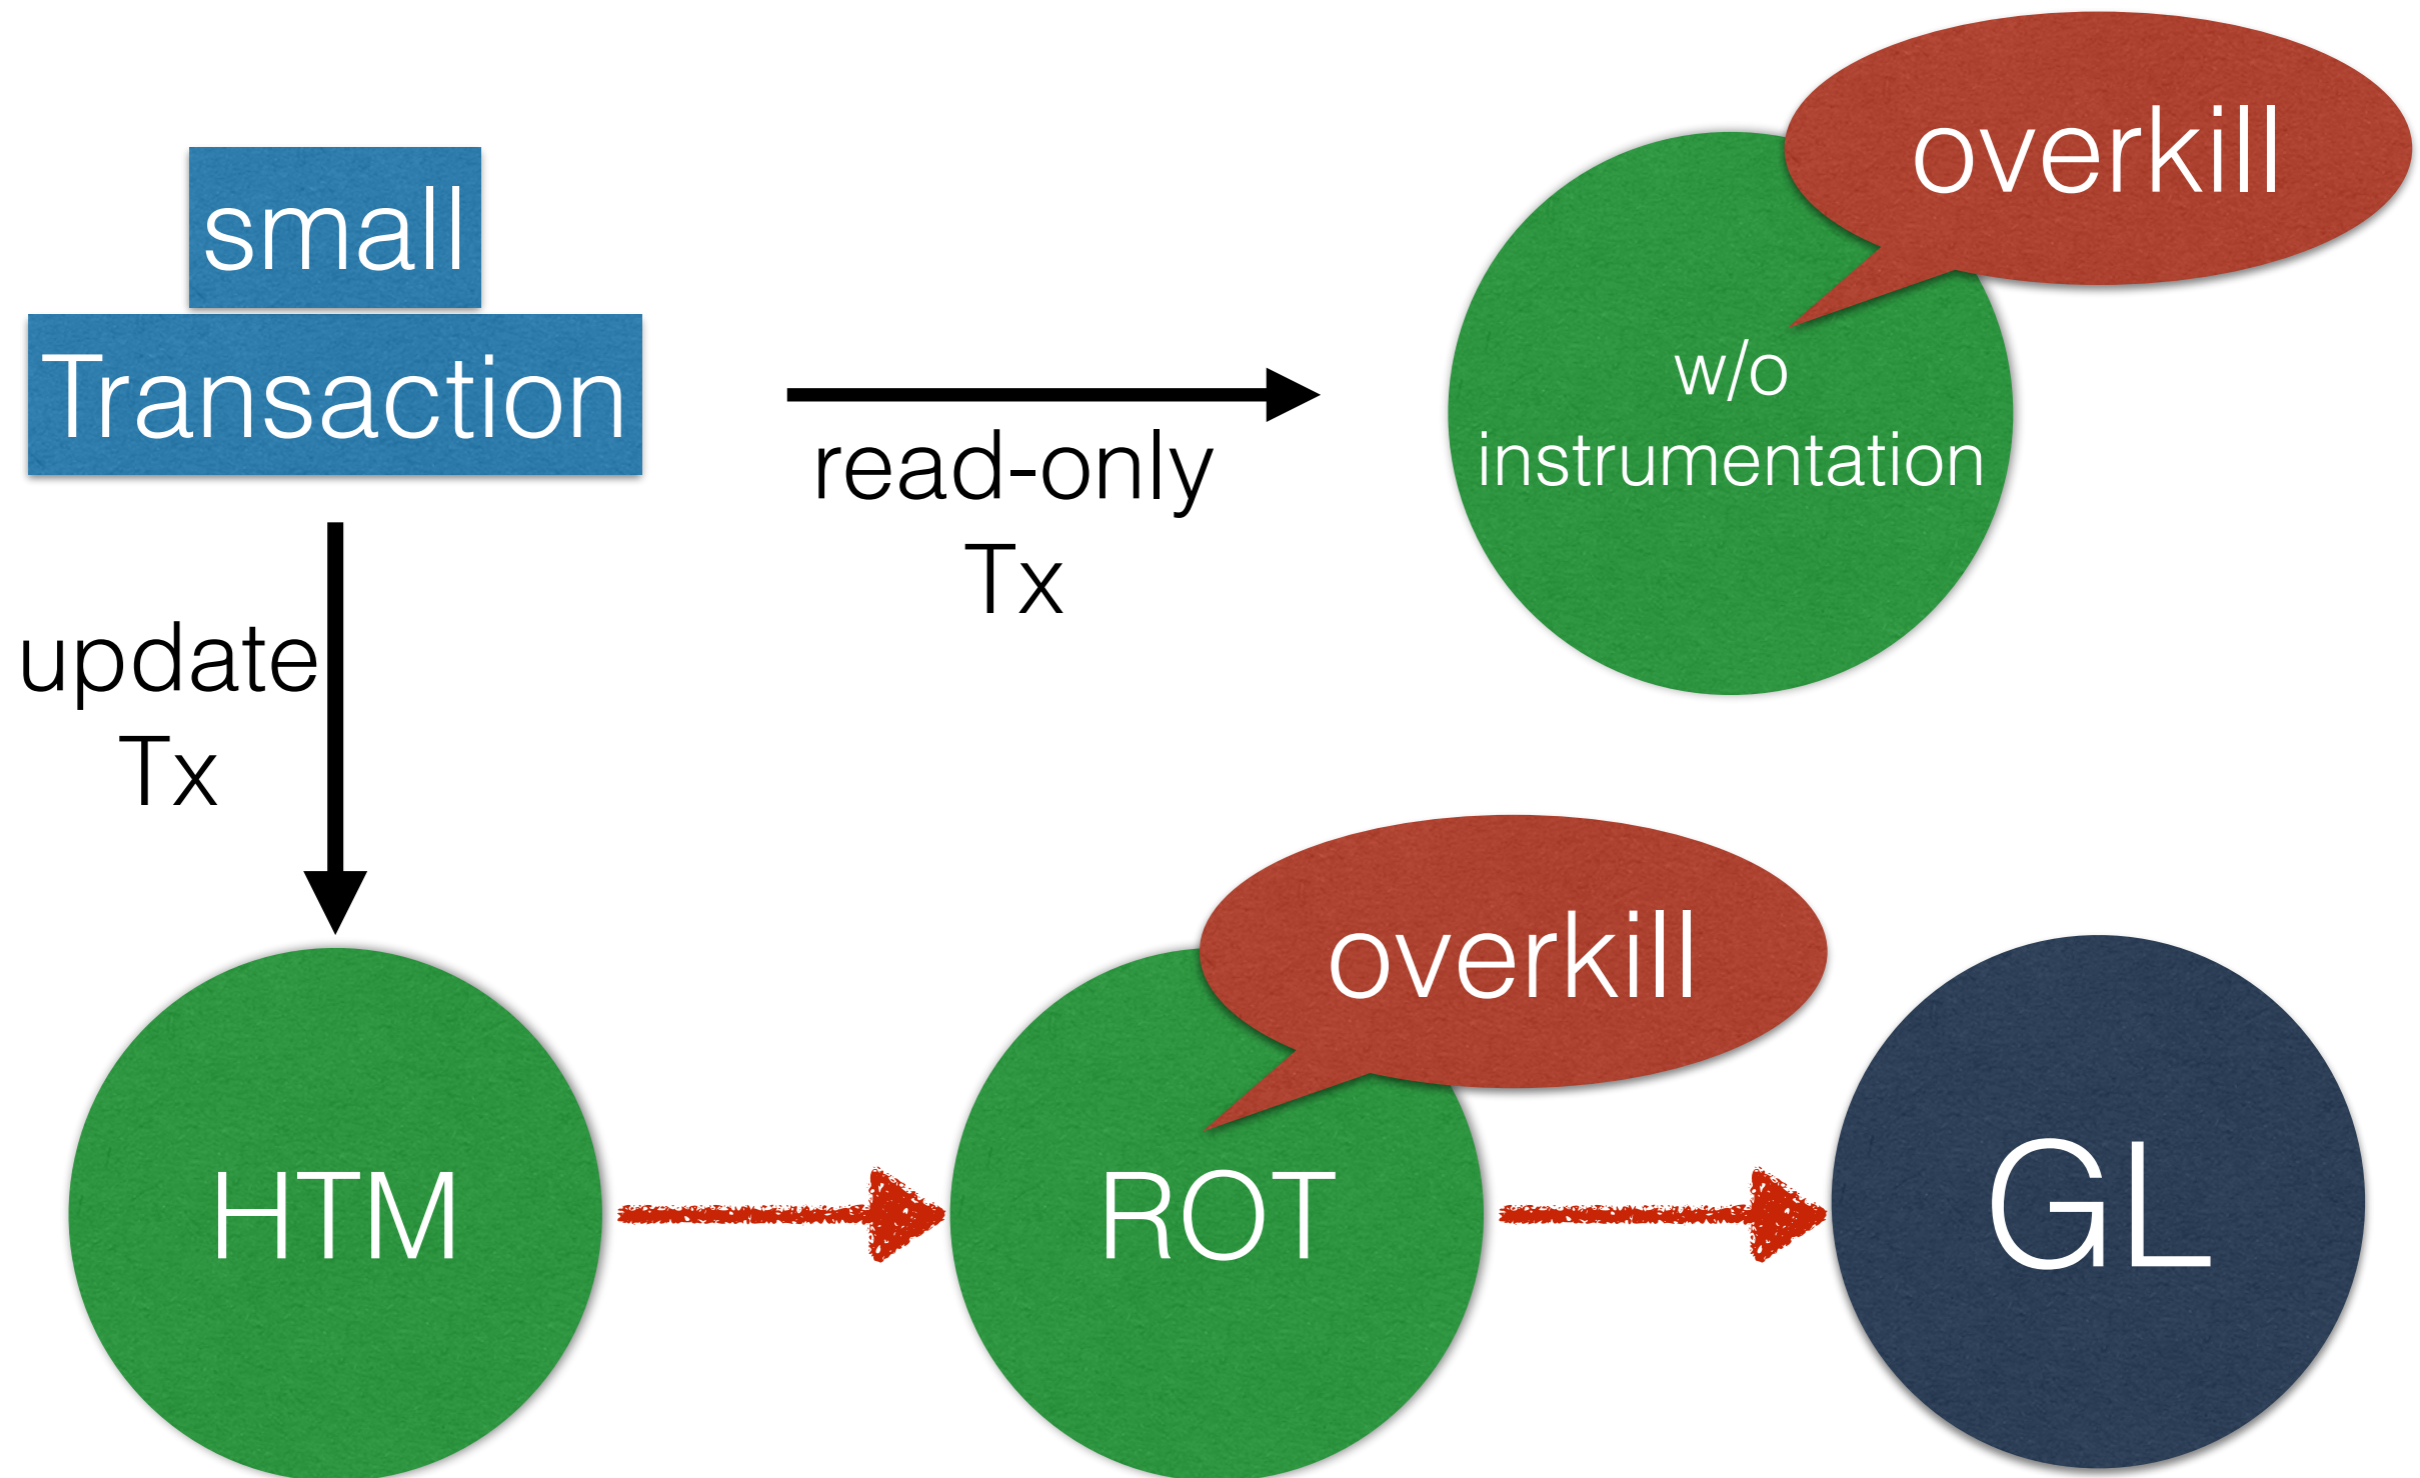
Uninstrumented Read-only

- read only transactions without any instrumentation
- outside the context of HTM or ROT
- no bounds on Tx size
- HTMs and ROTs must wait for UROs

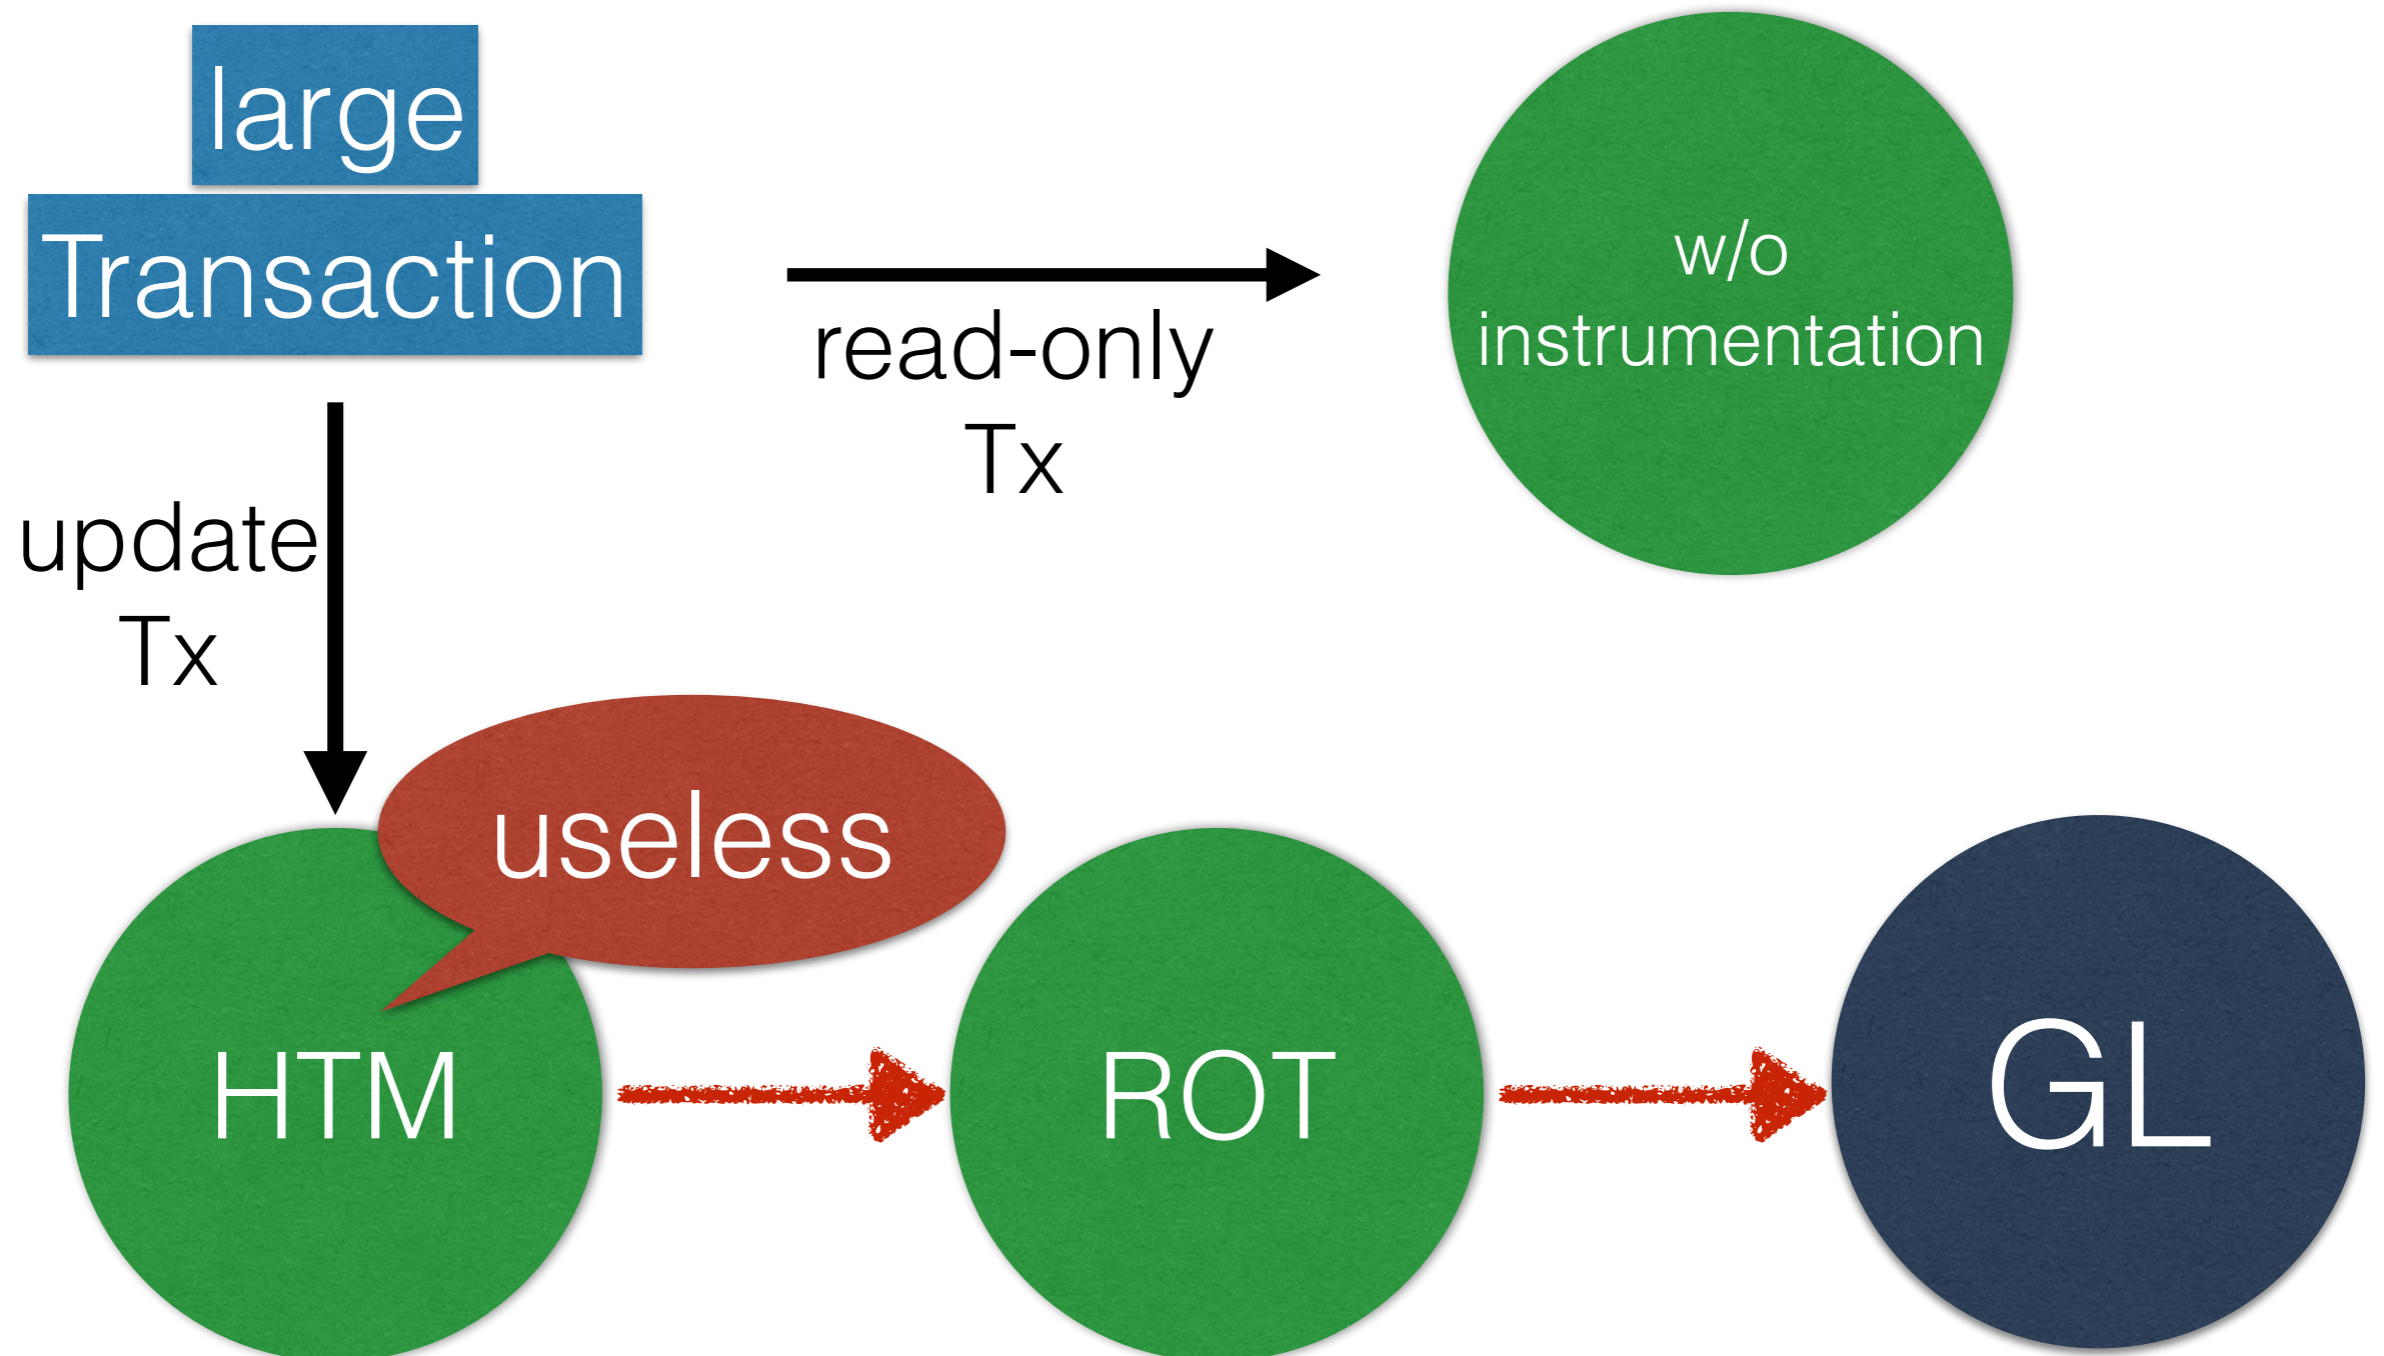
POWER8-TM



POWER8-TM



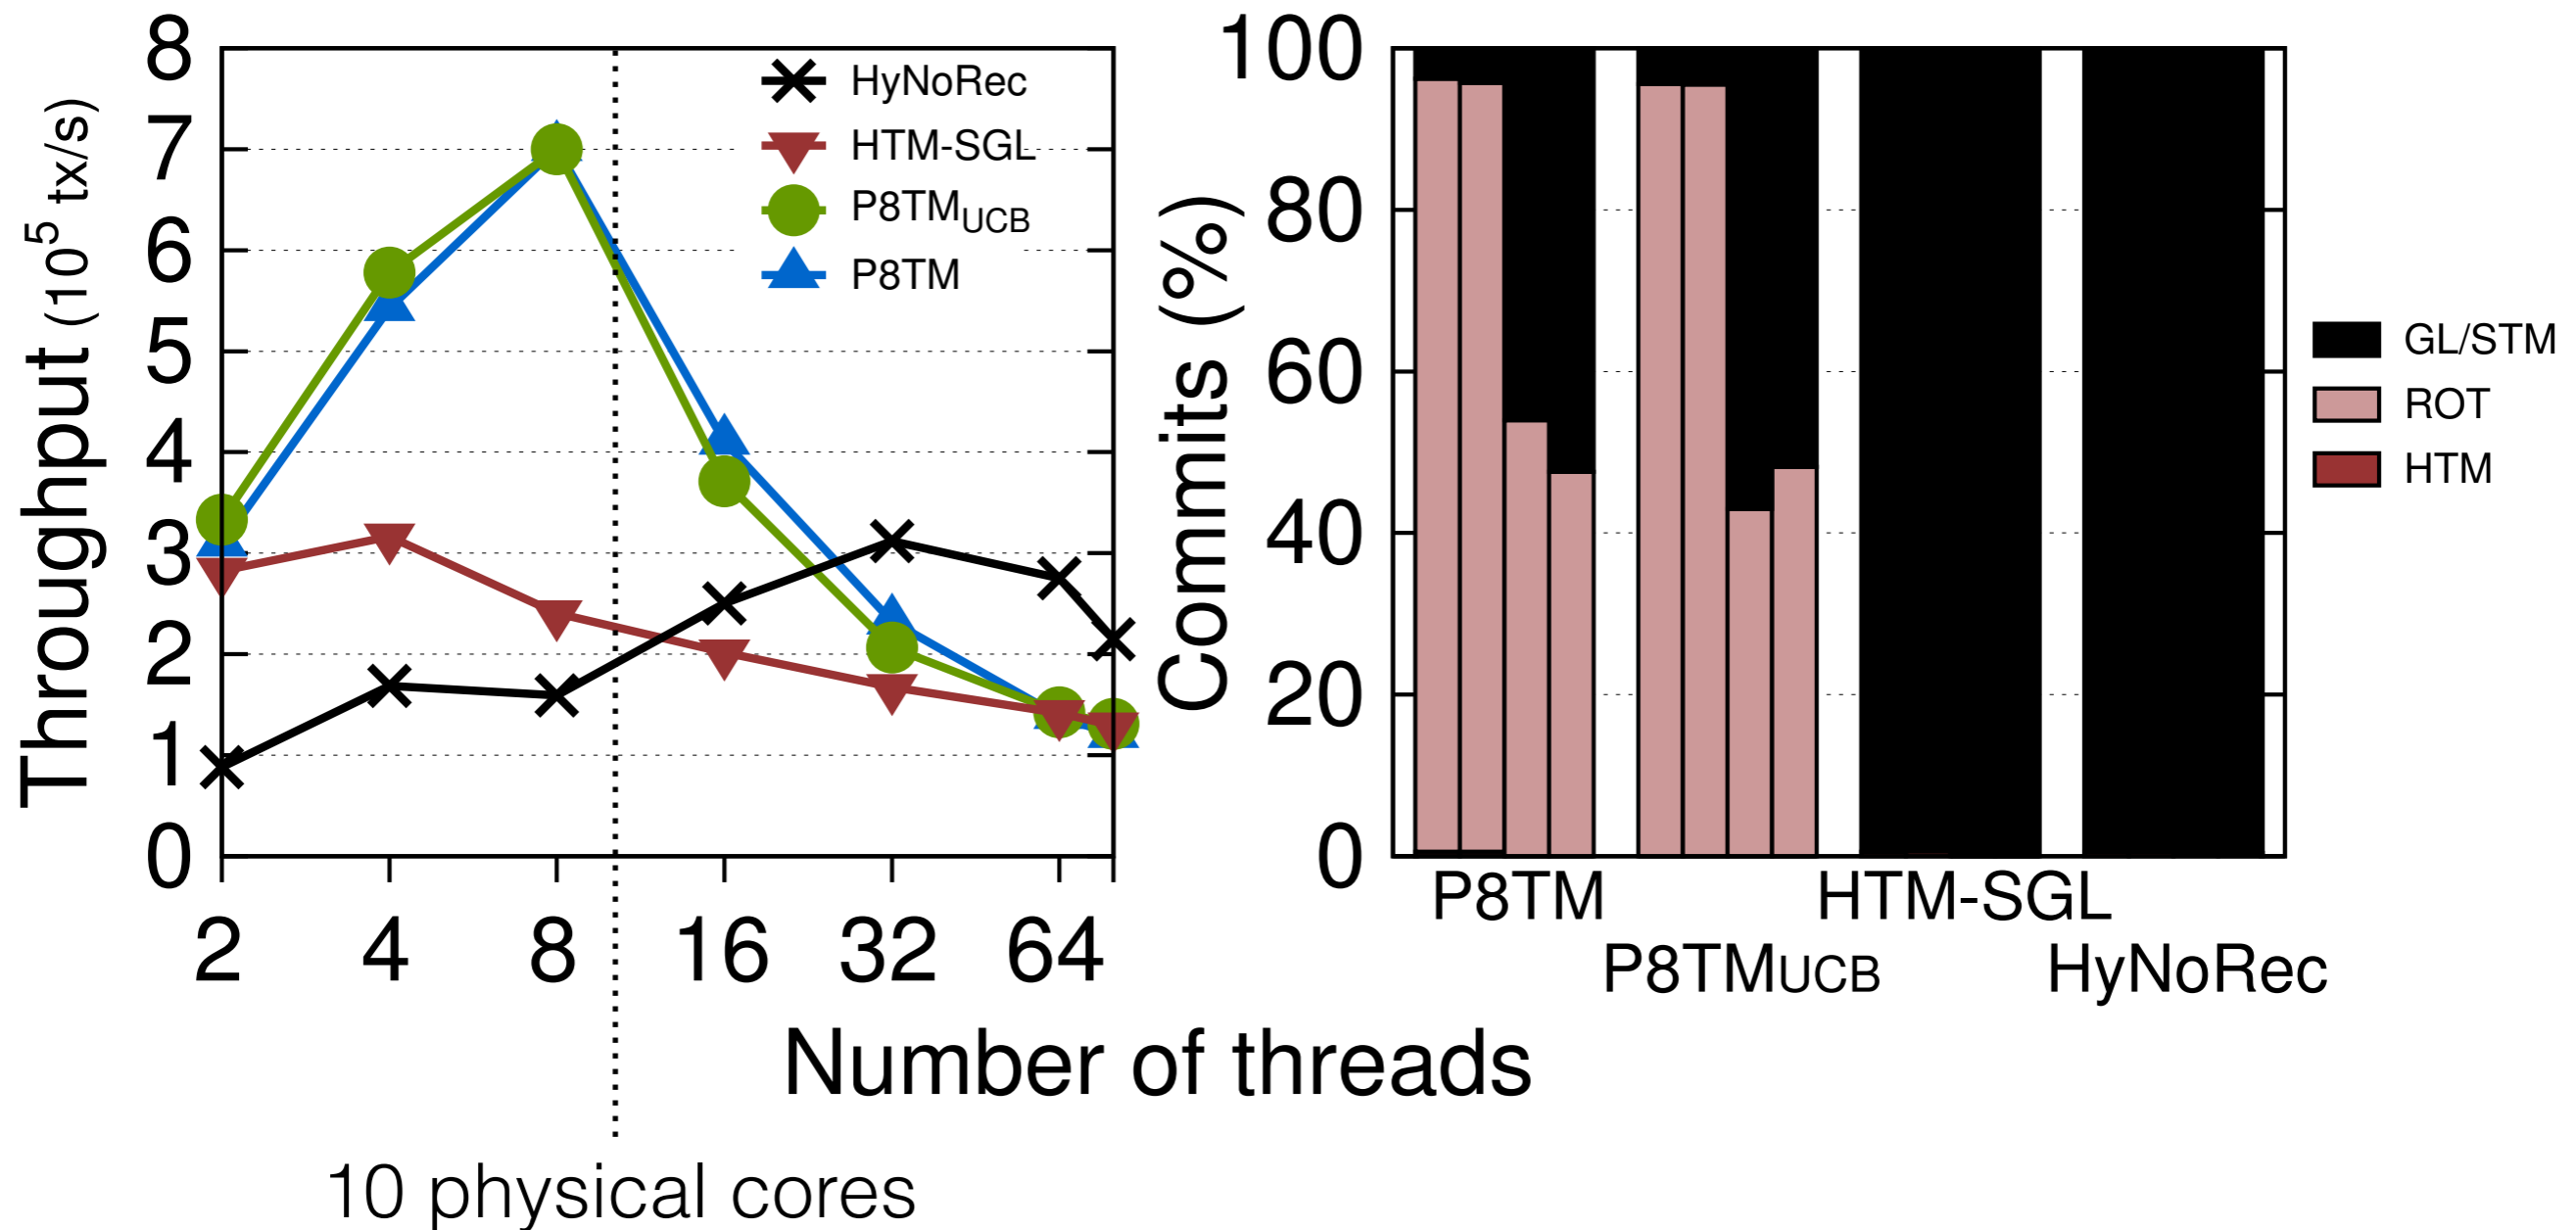
POWER8-TM



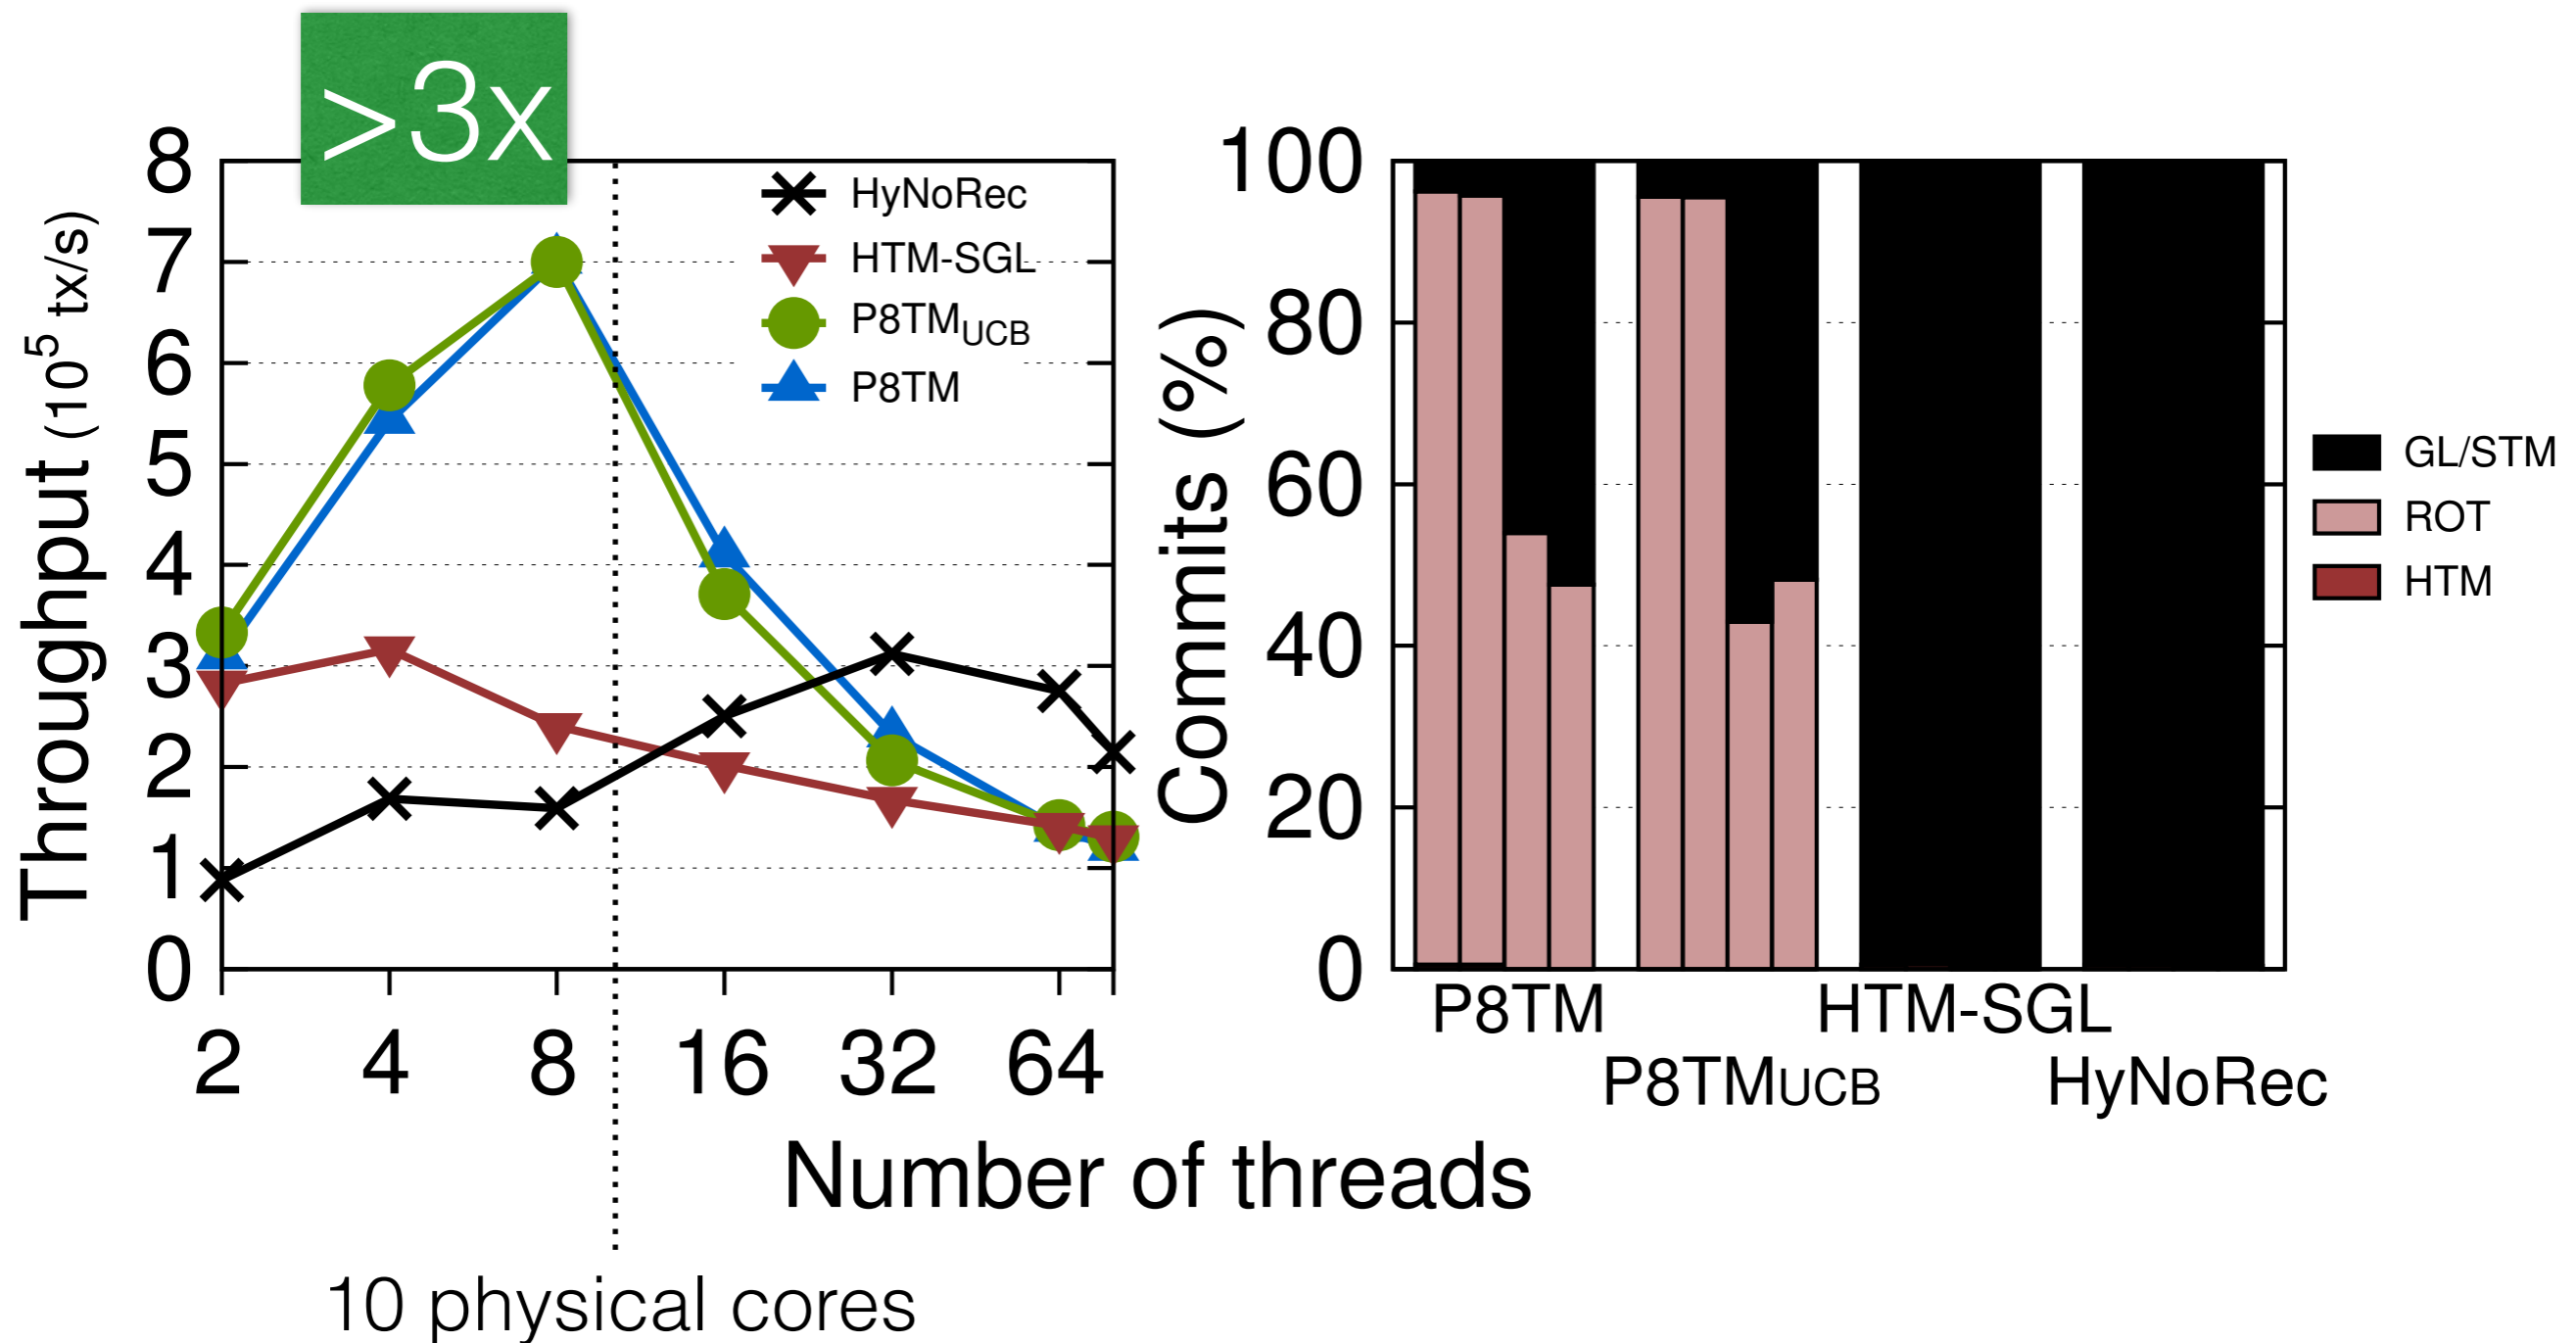
Self-tuning

- lightweight, online reinforcement learning
- determine execution path:
 - HTM \longrightarrow GL : small Txs
 - ROT \longrightarrow GL : large Txs
 - HTM \longrightarrow ROT \longrightarrow GL : mixed workload

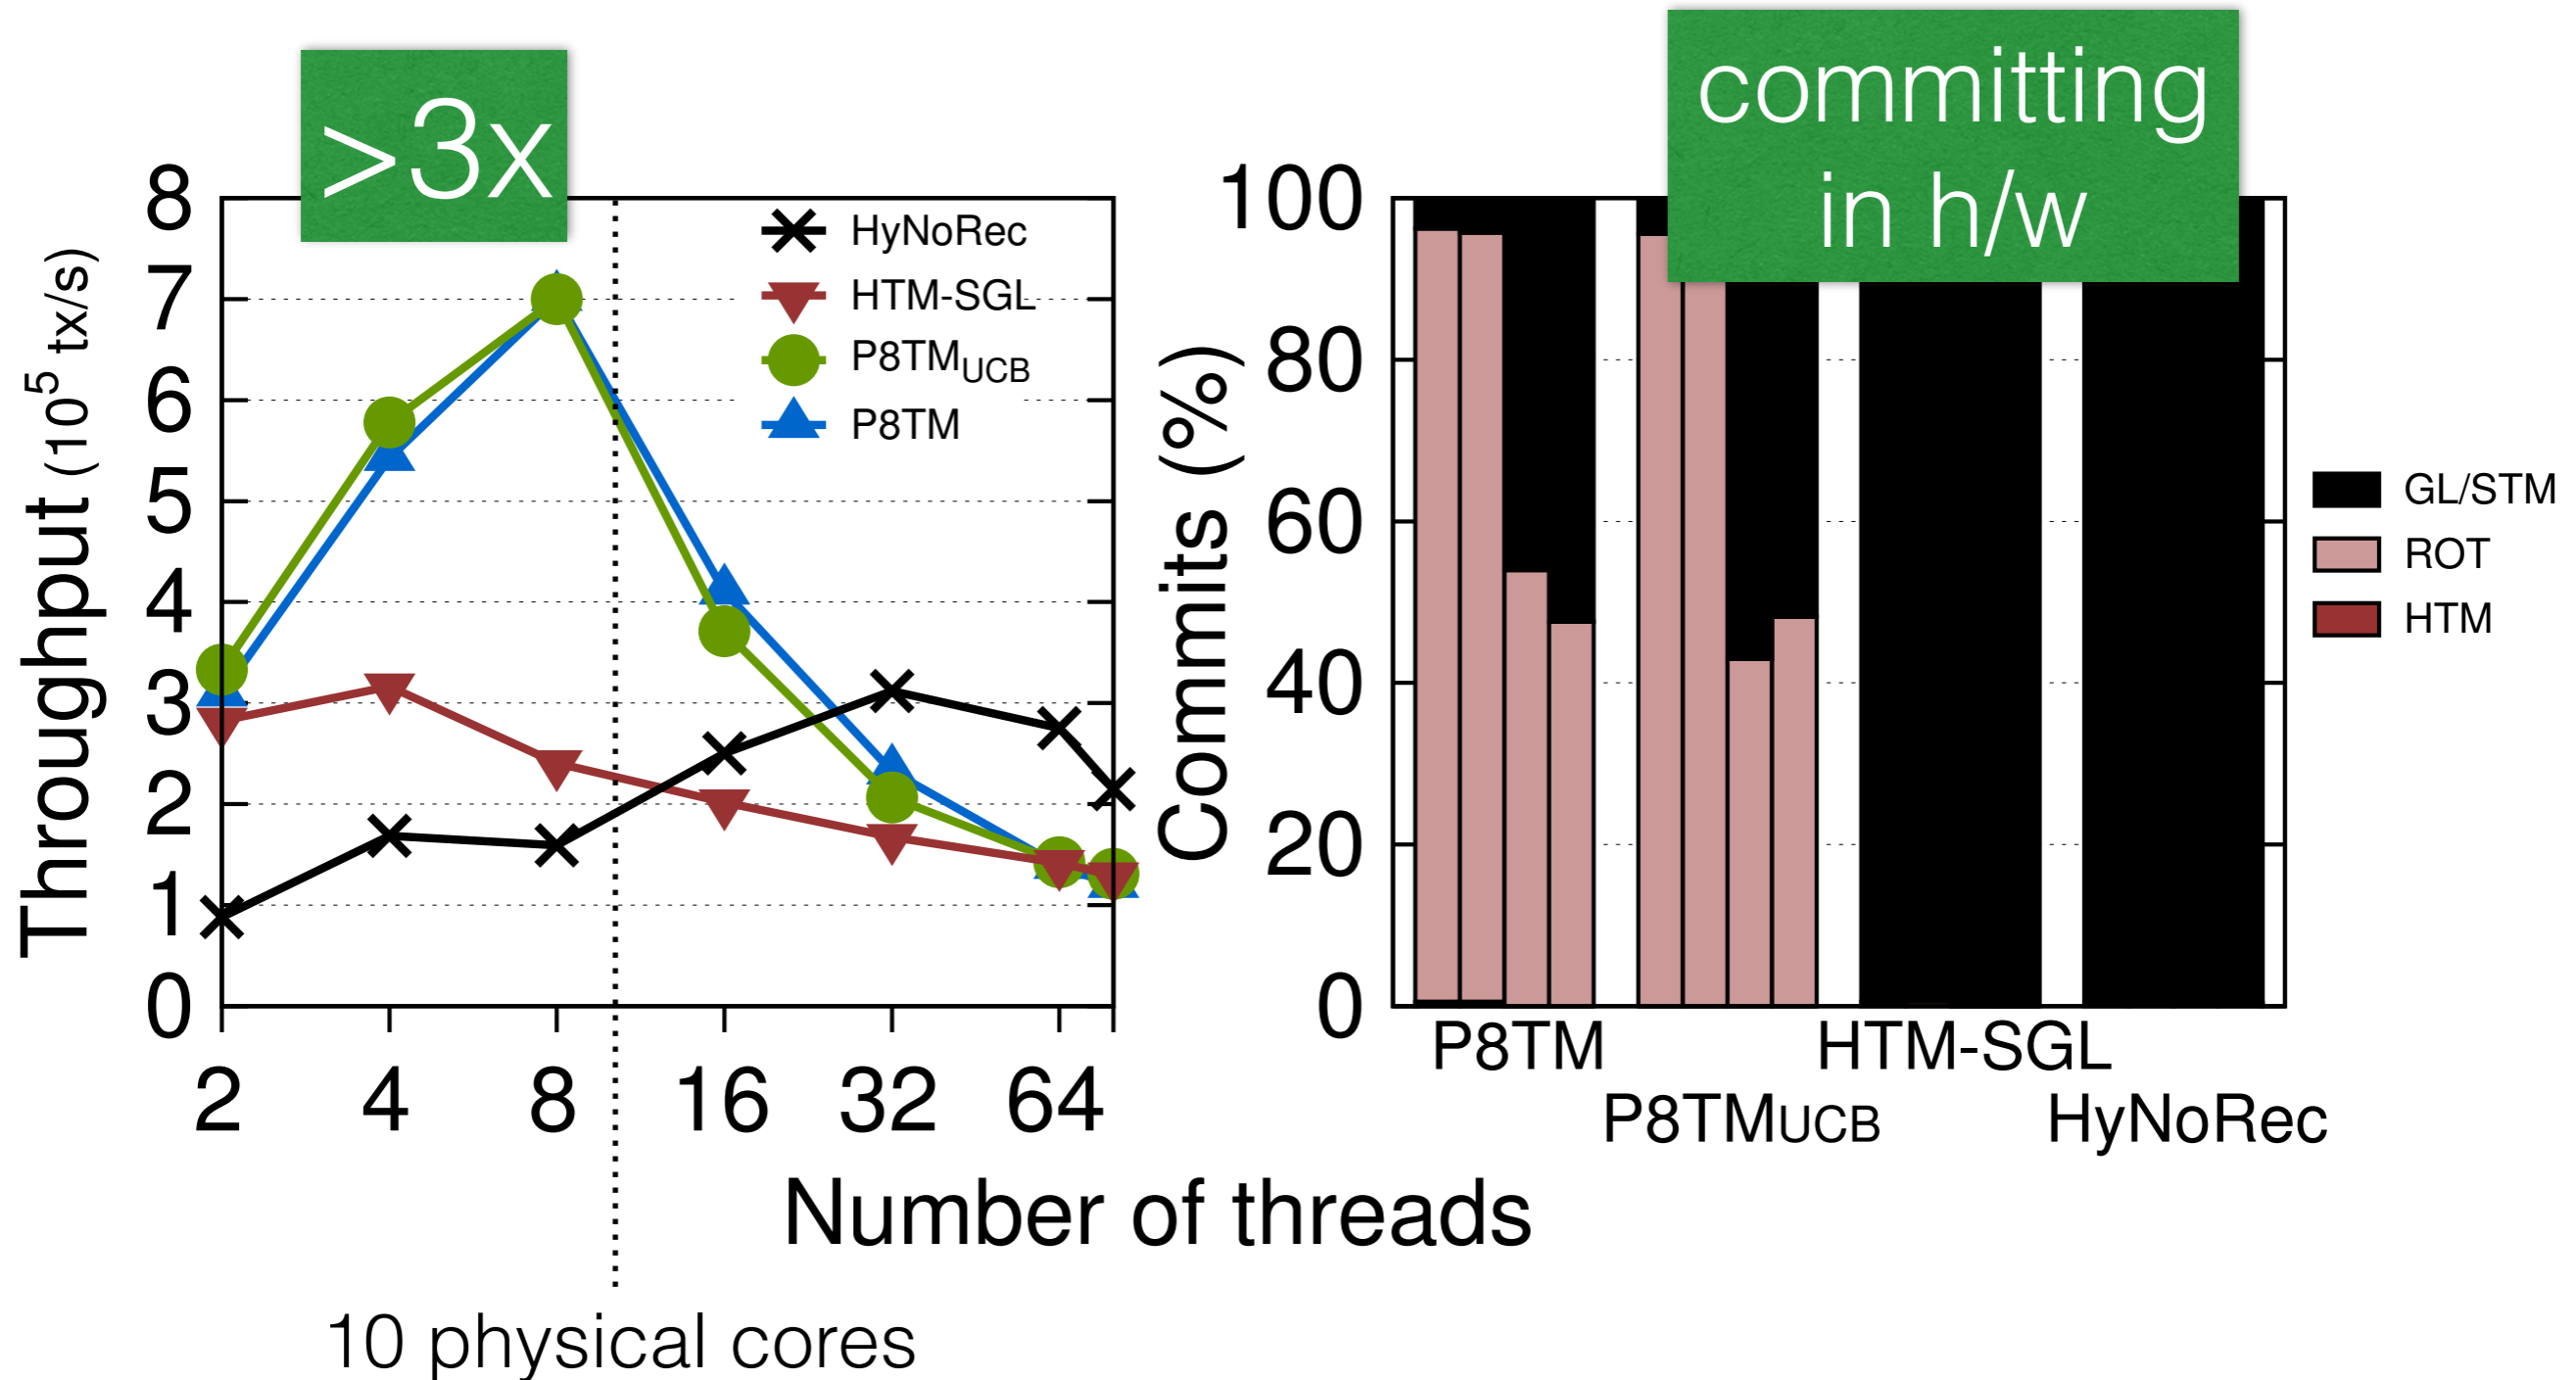
Evaluation: Vacation



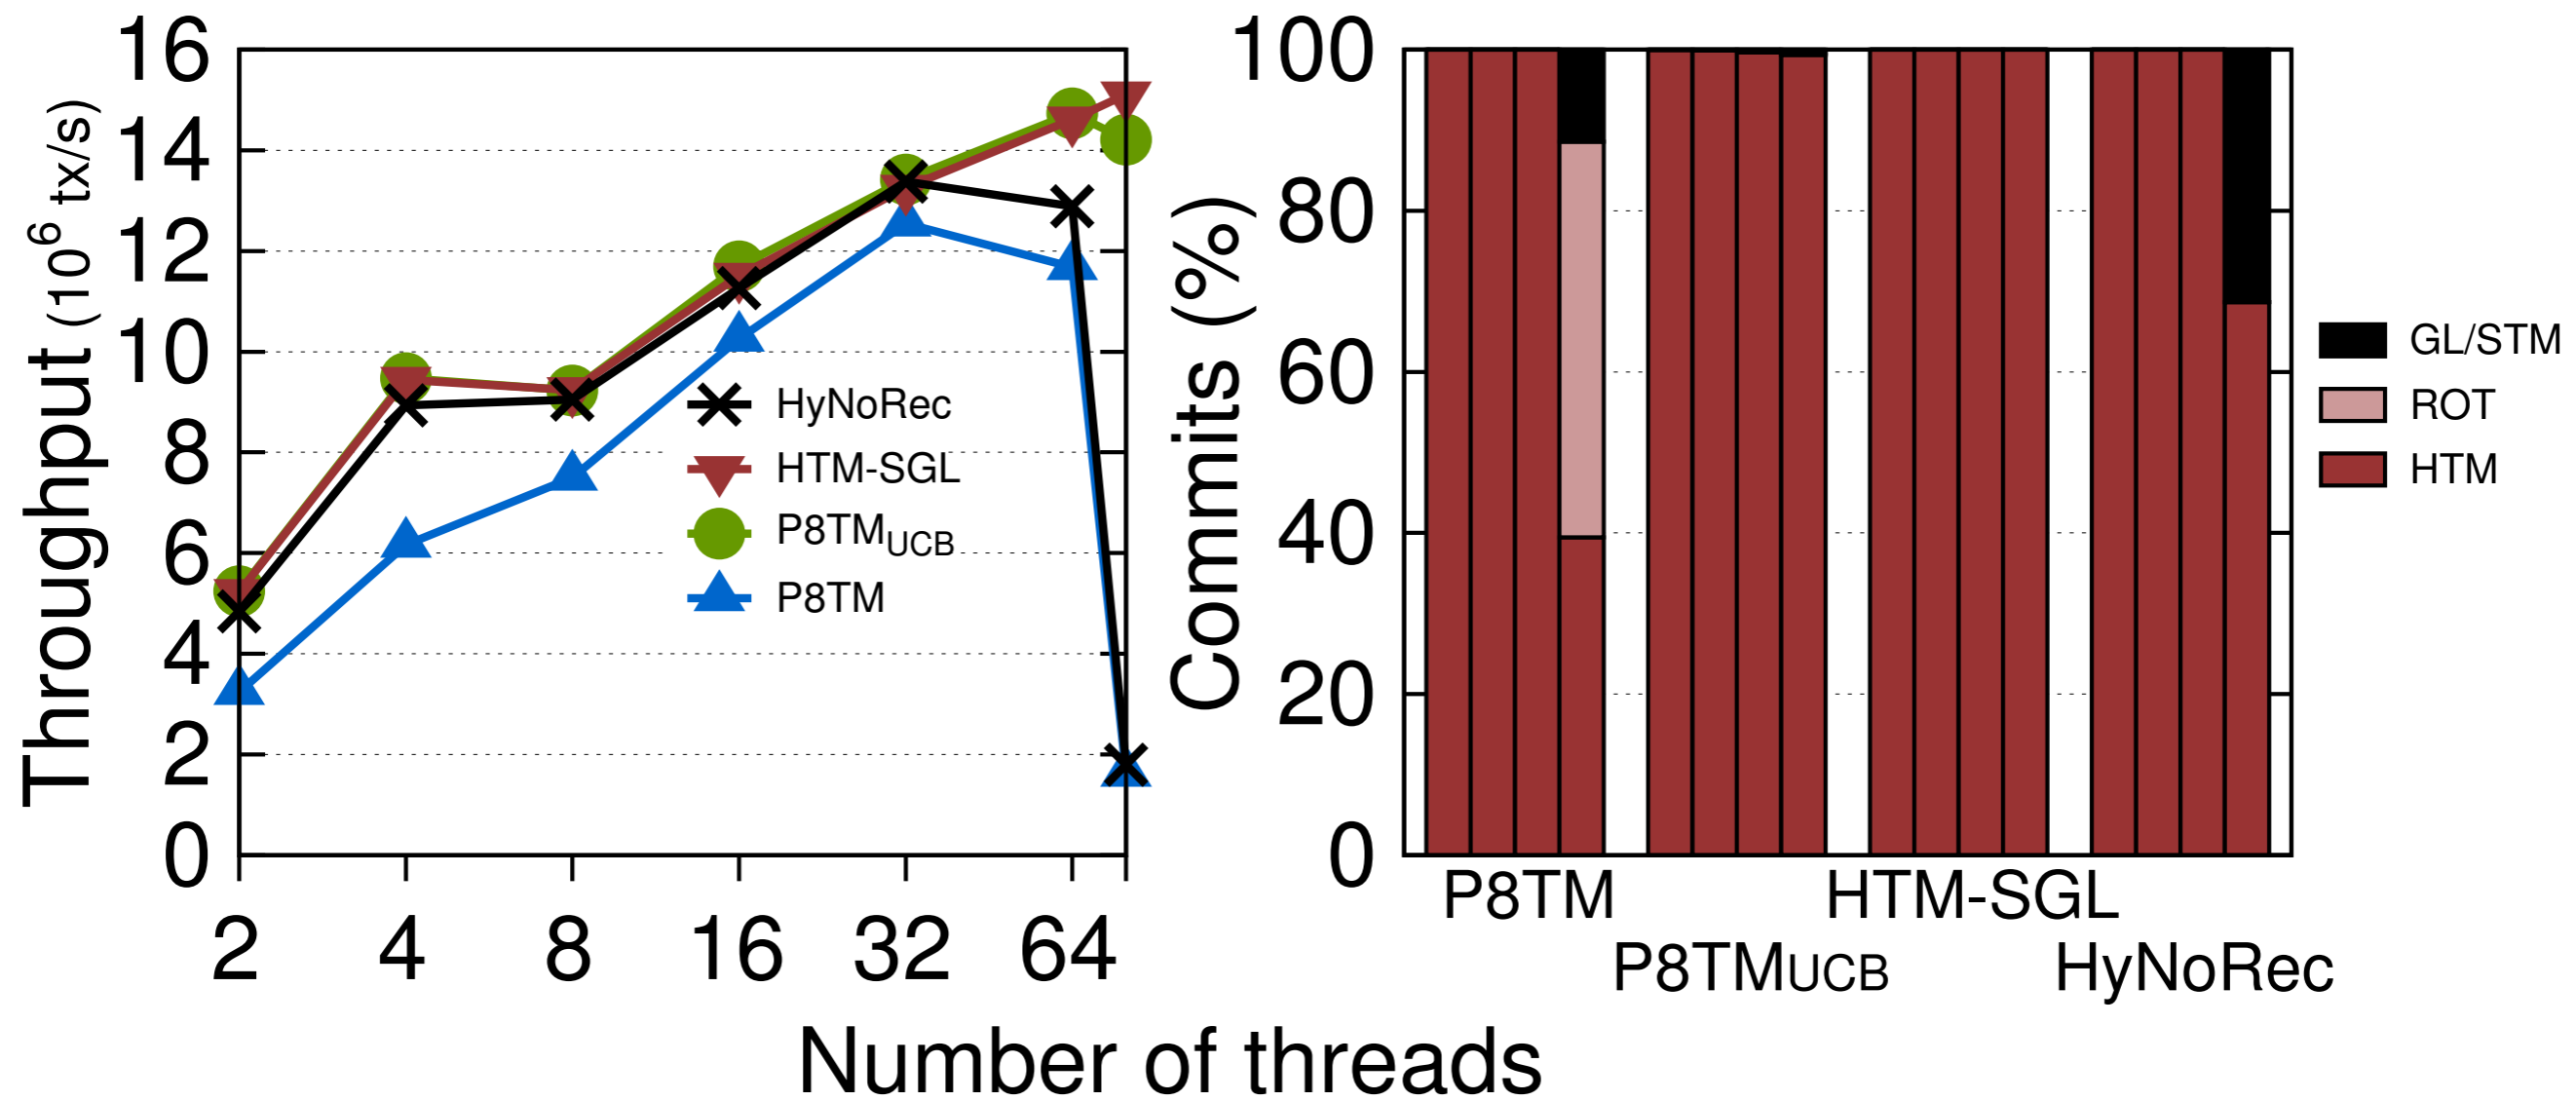
Evaluation: Vacation



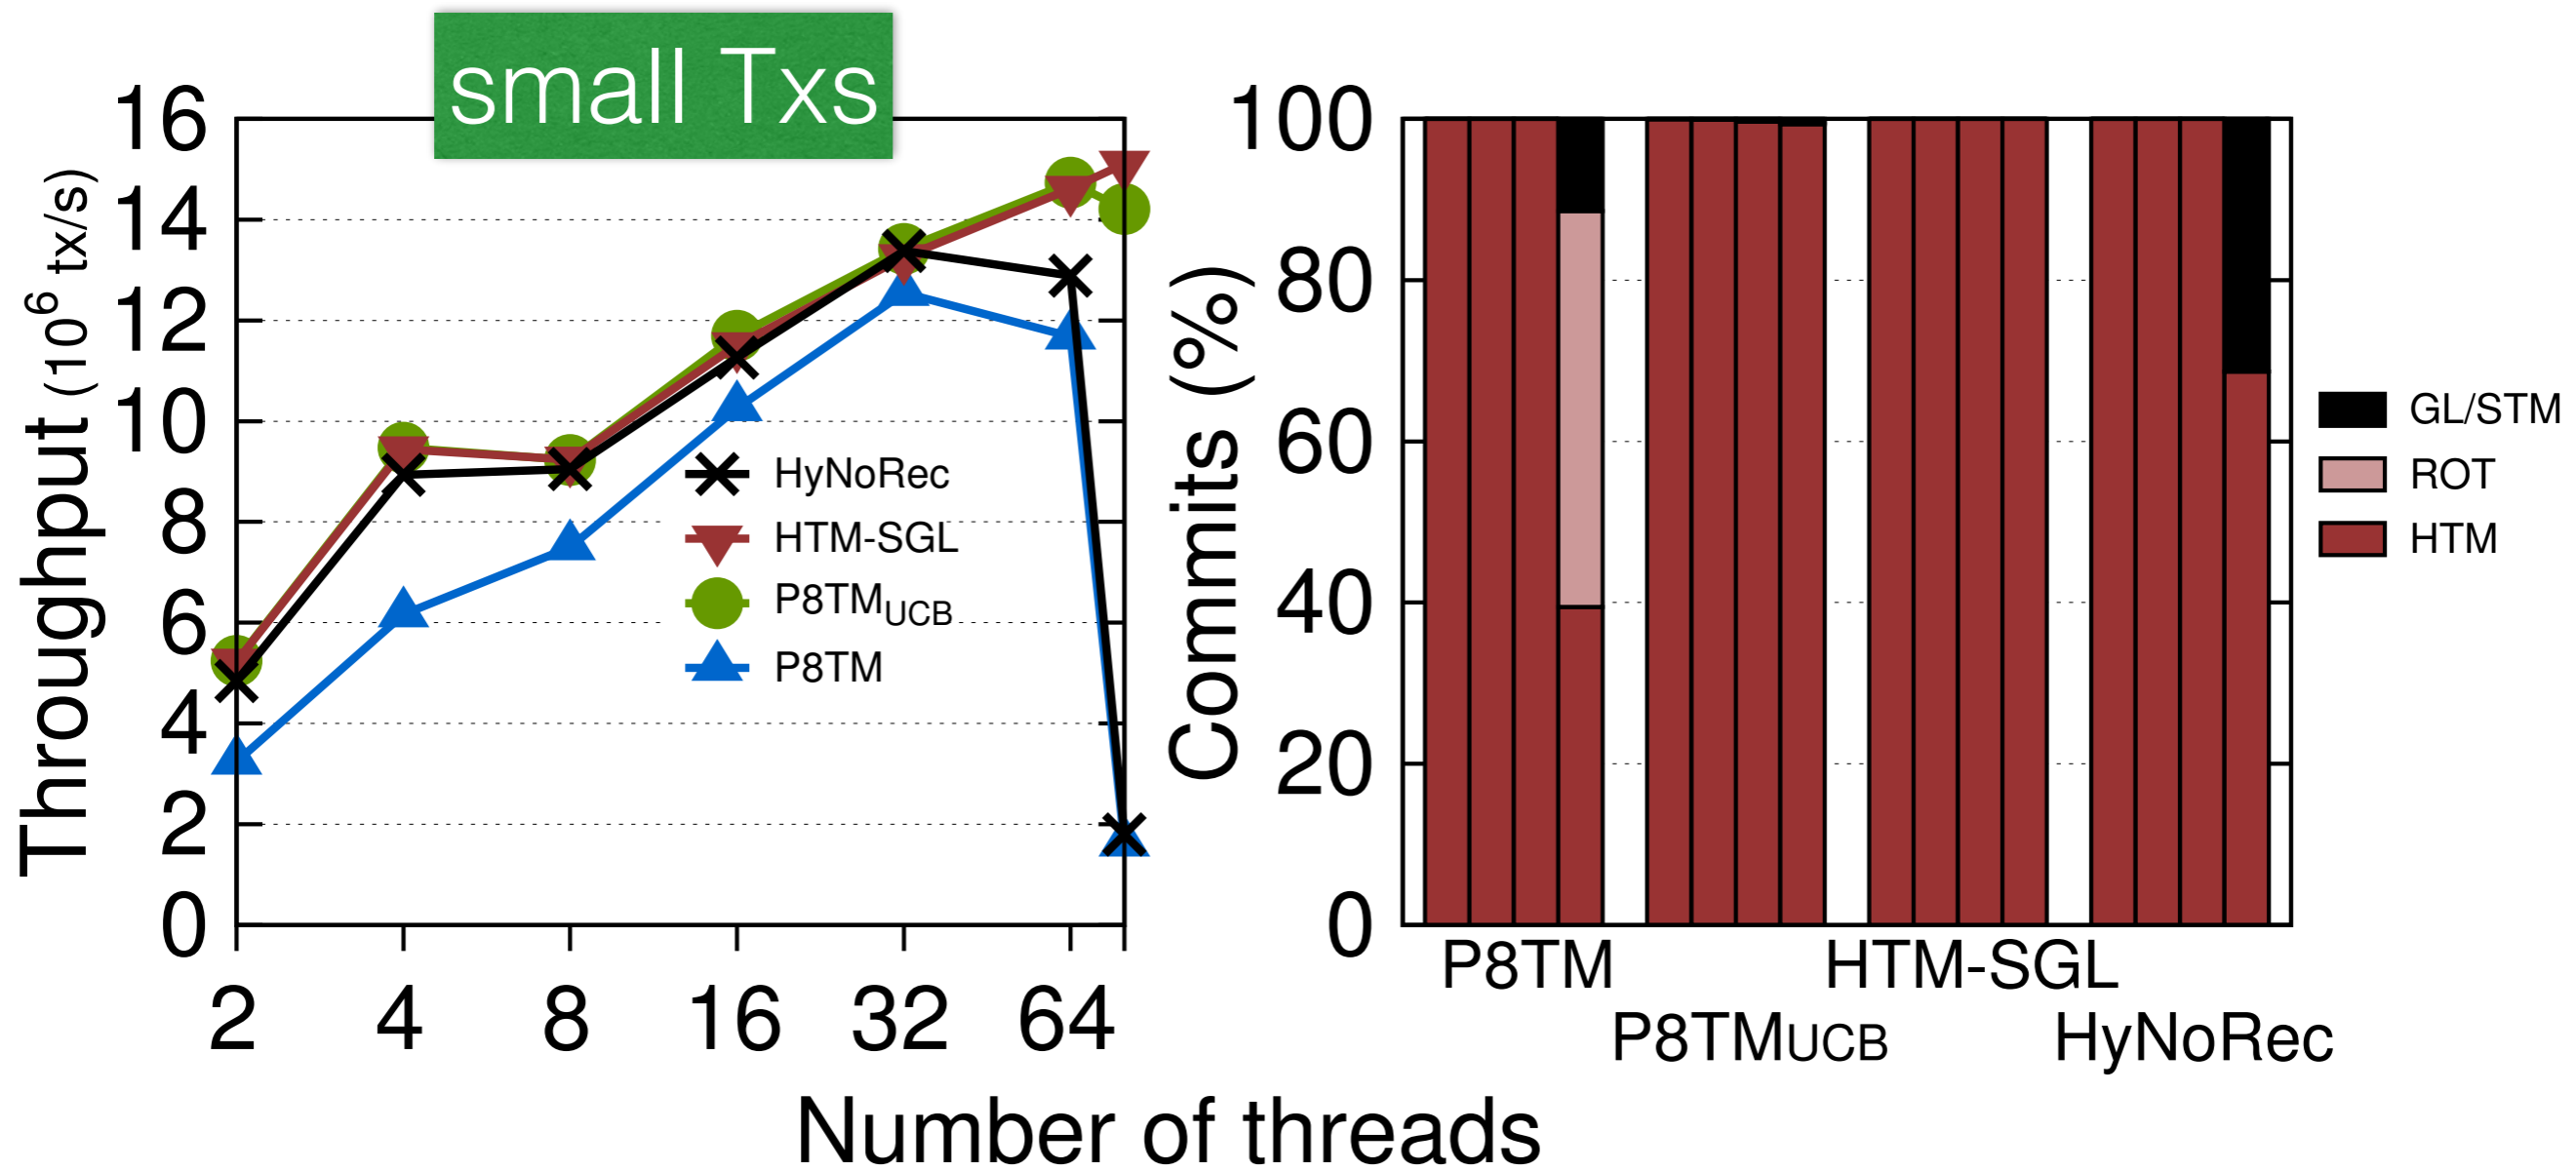
Evaluation: Vacation



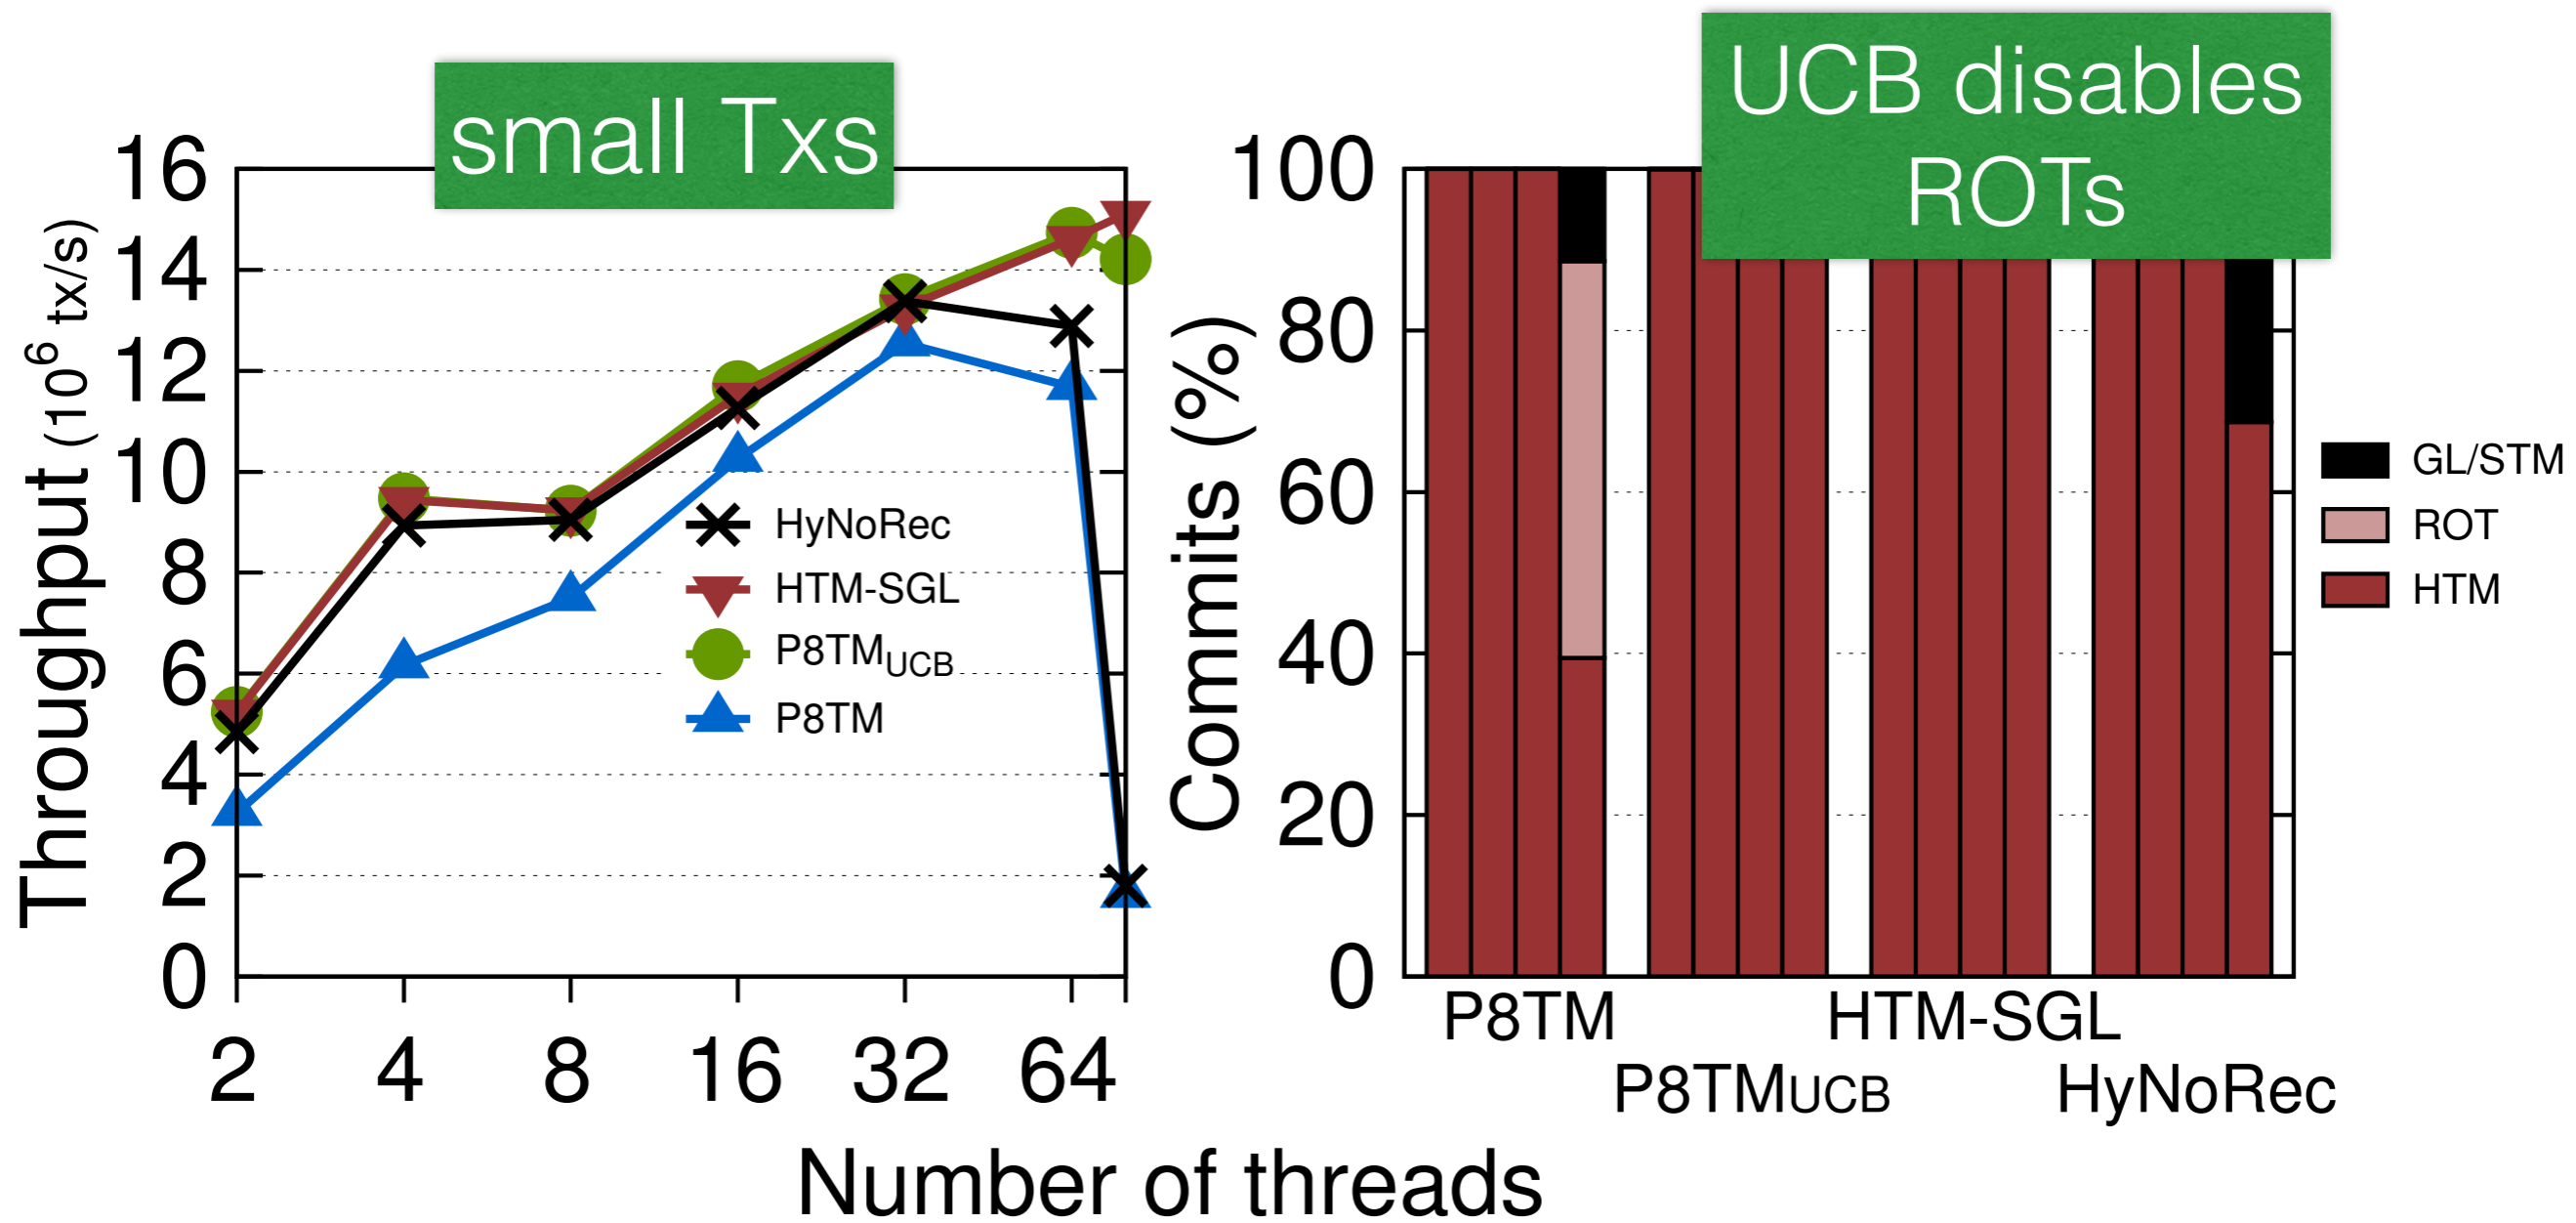
Evaluation: SSCA2



Evaluation: SSCA2



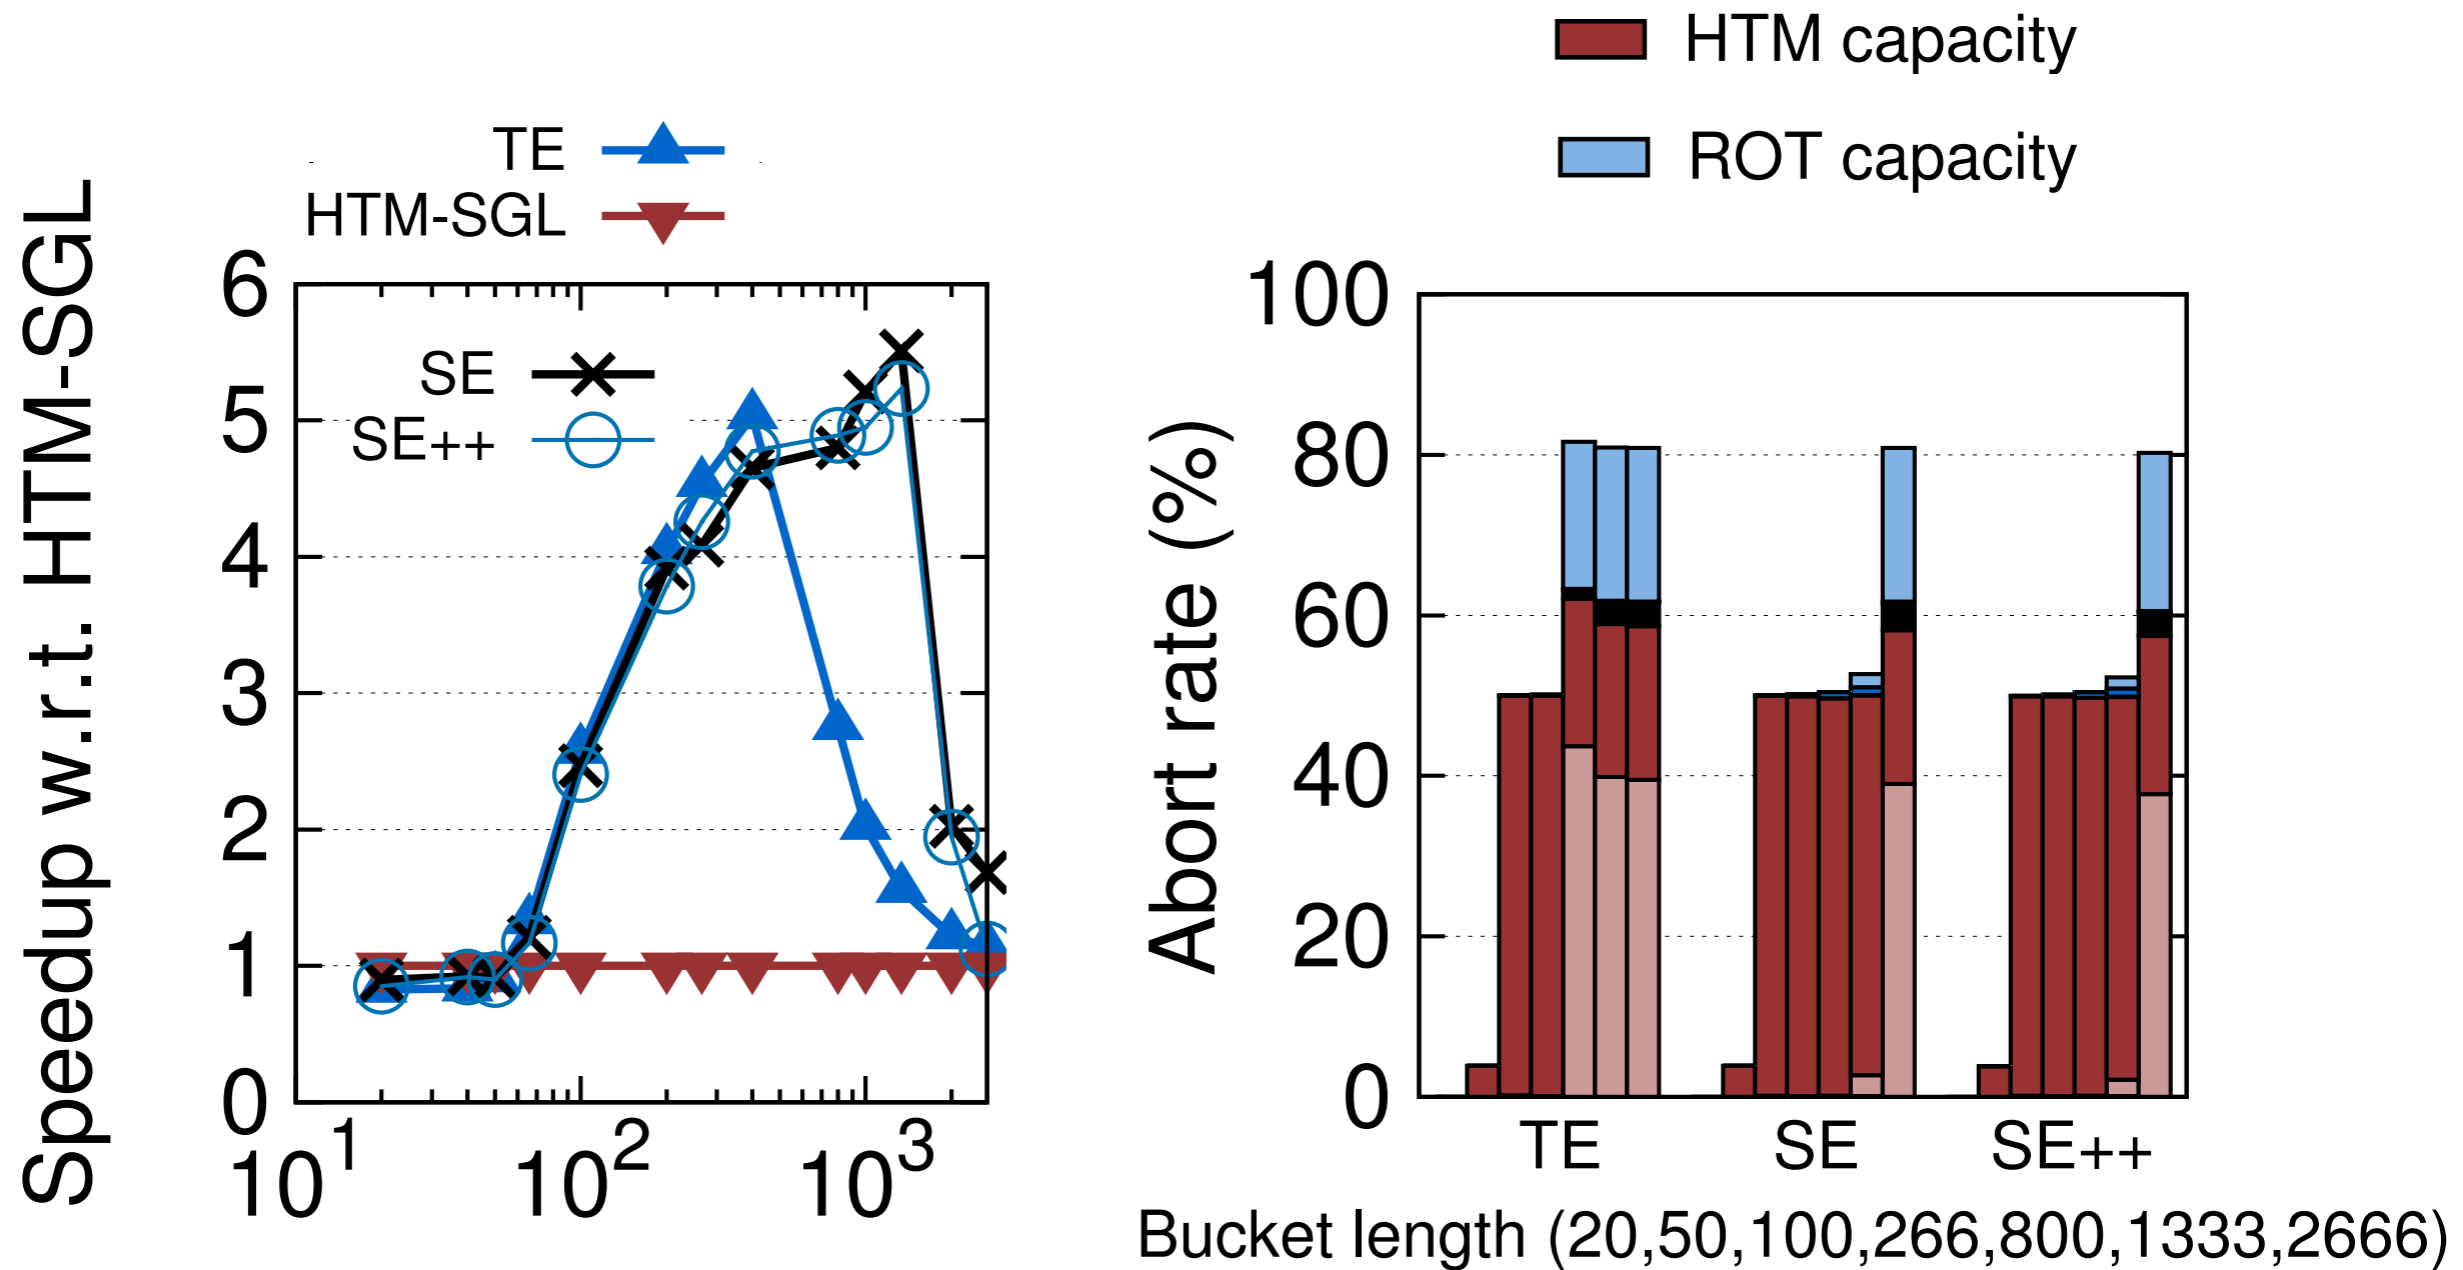
Evaluation: SSCA2



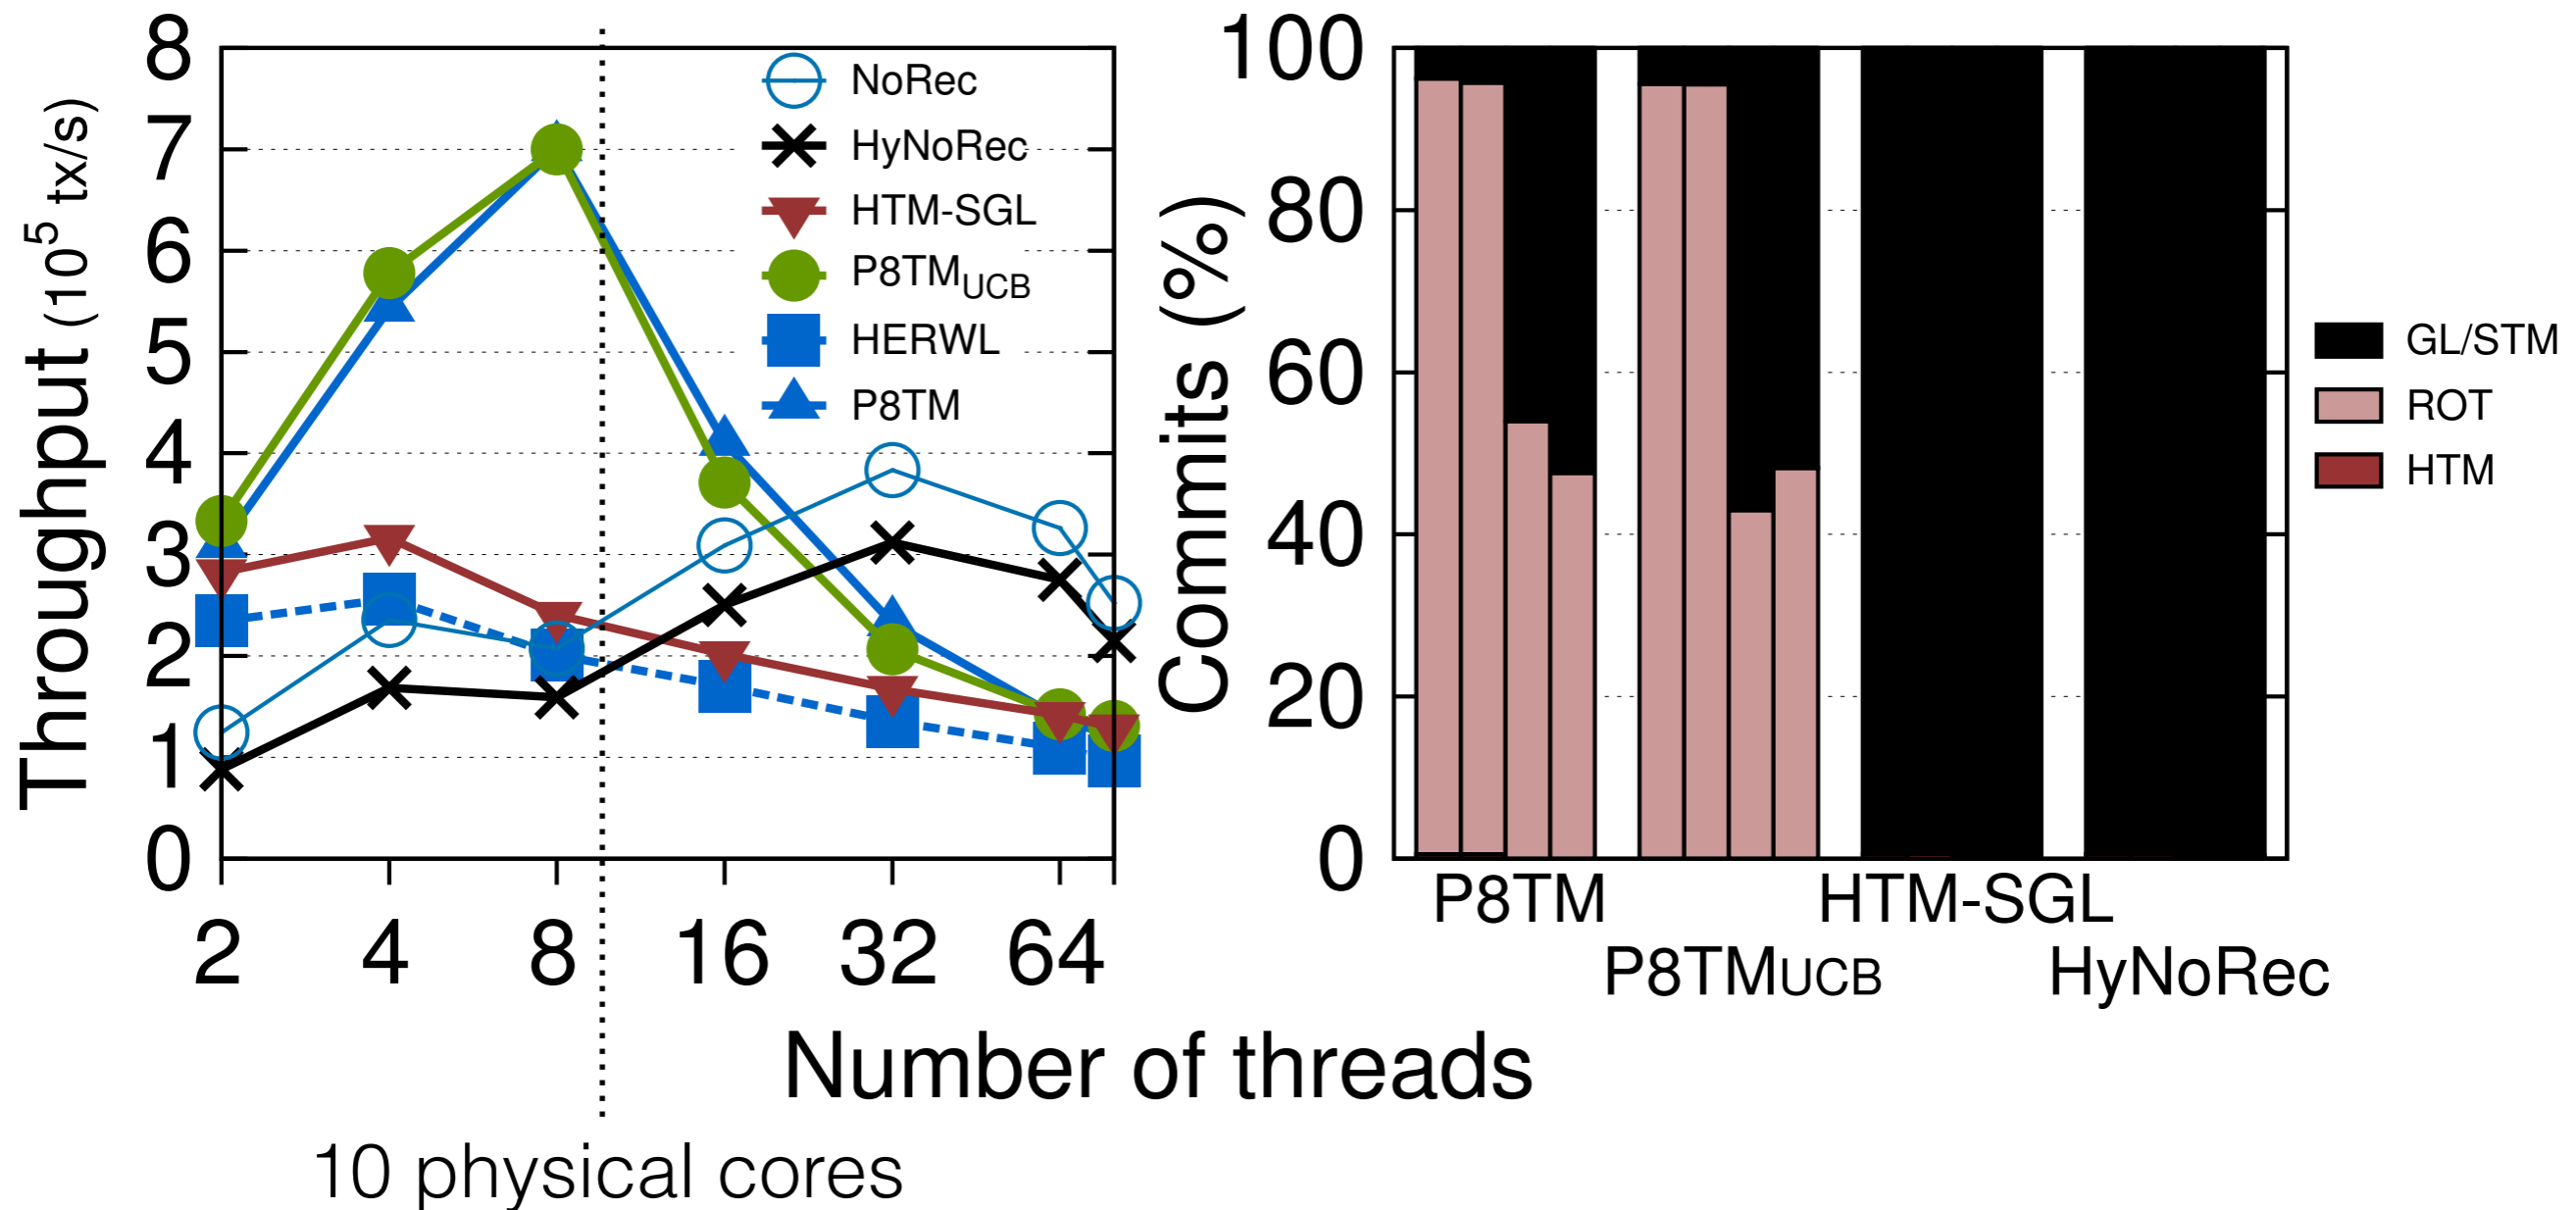
Conclusion

- POWER8-TM was able to exploit ROTs and suspend/resume to expand the capacity limitations
- TMCAM aware read-set tracking was necessary
- Self-tuning was effective in adapting to different workloads
- POWER8-TM promotes the importance of such features that can be used in innovative techniques to mitigate hardware limitations

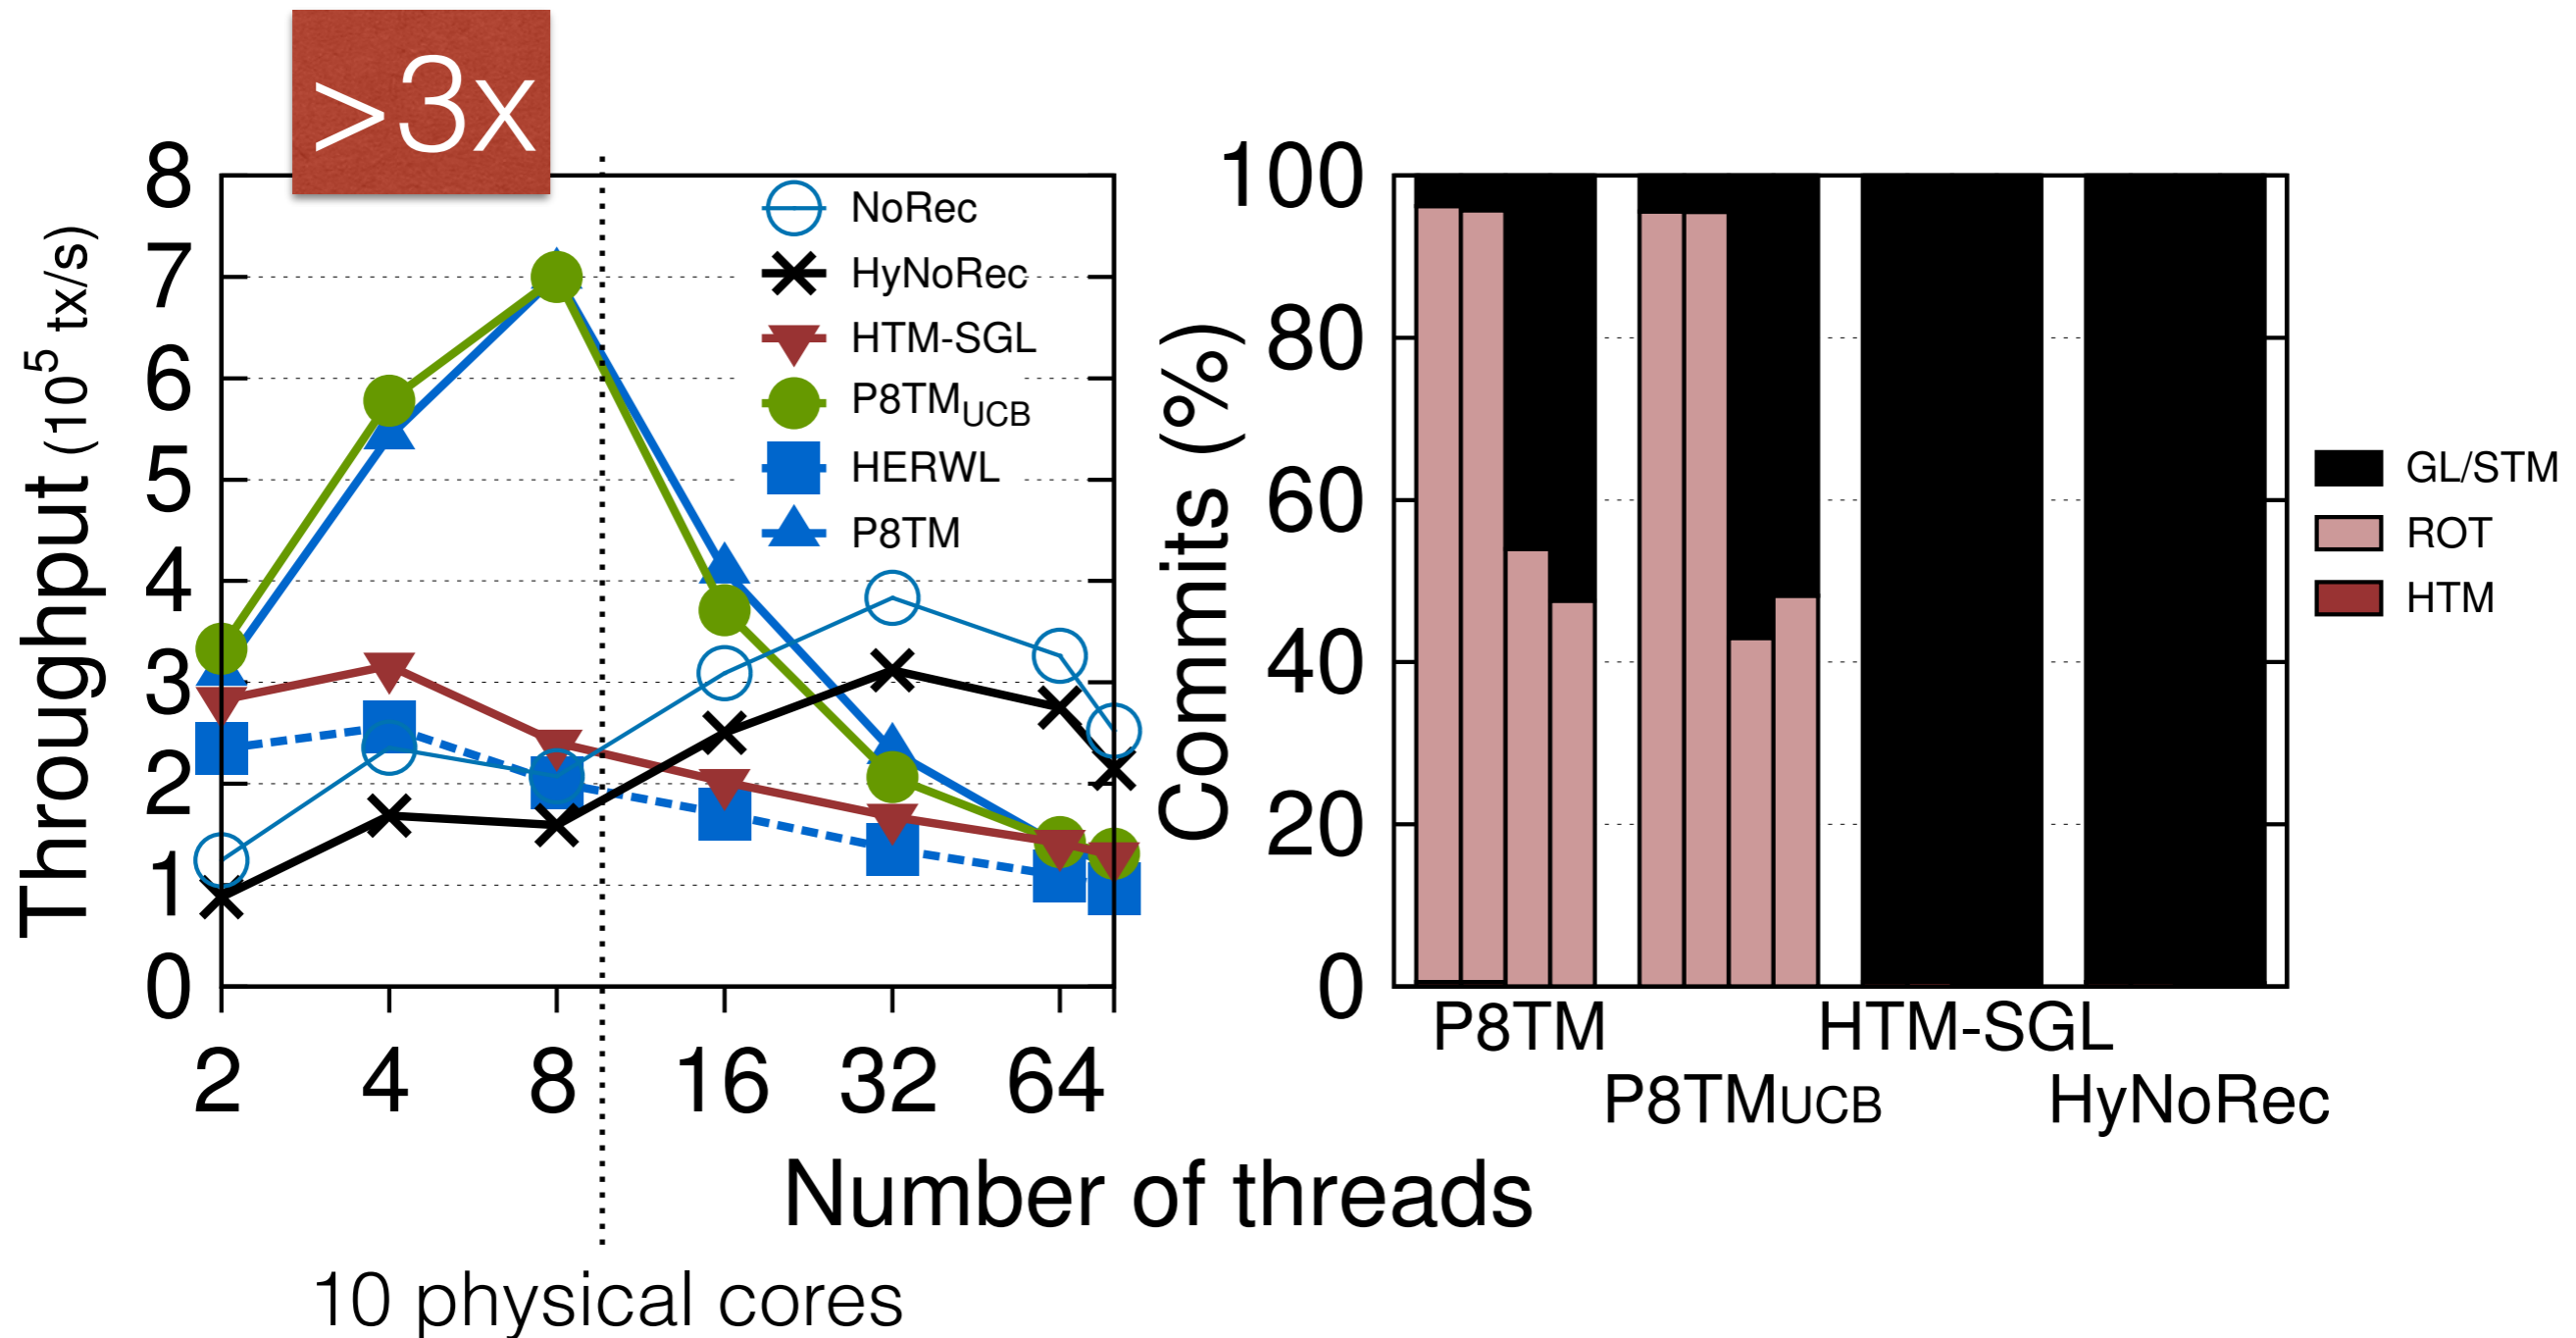
Results: read-set tracking



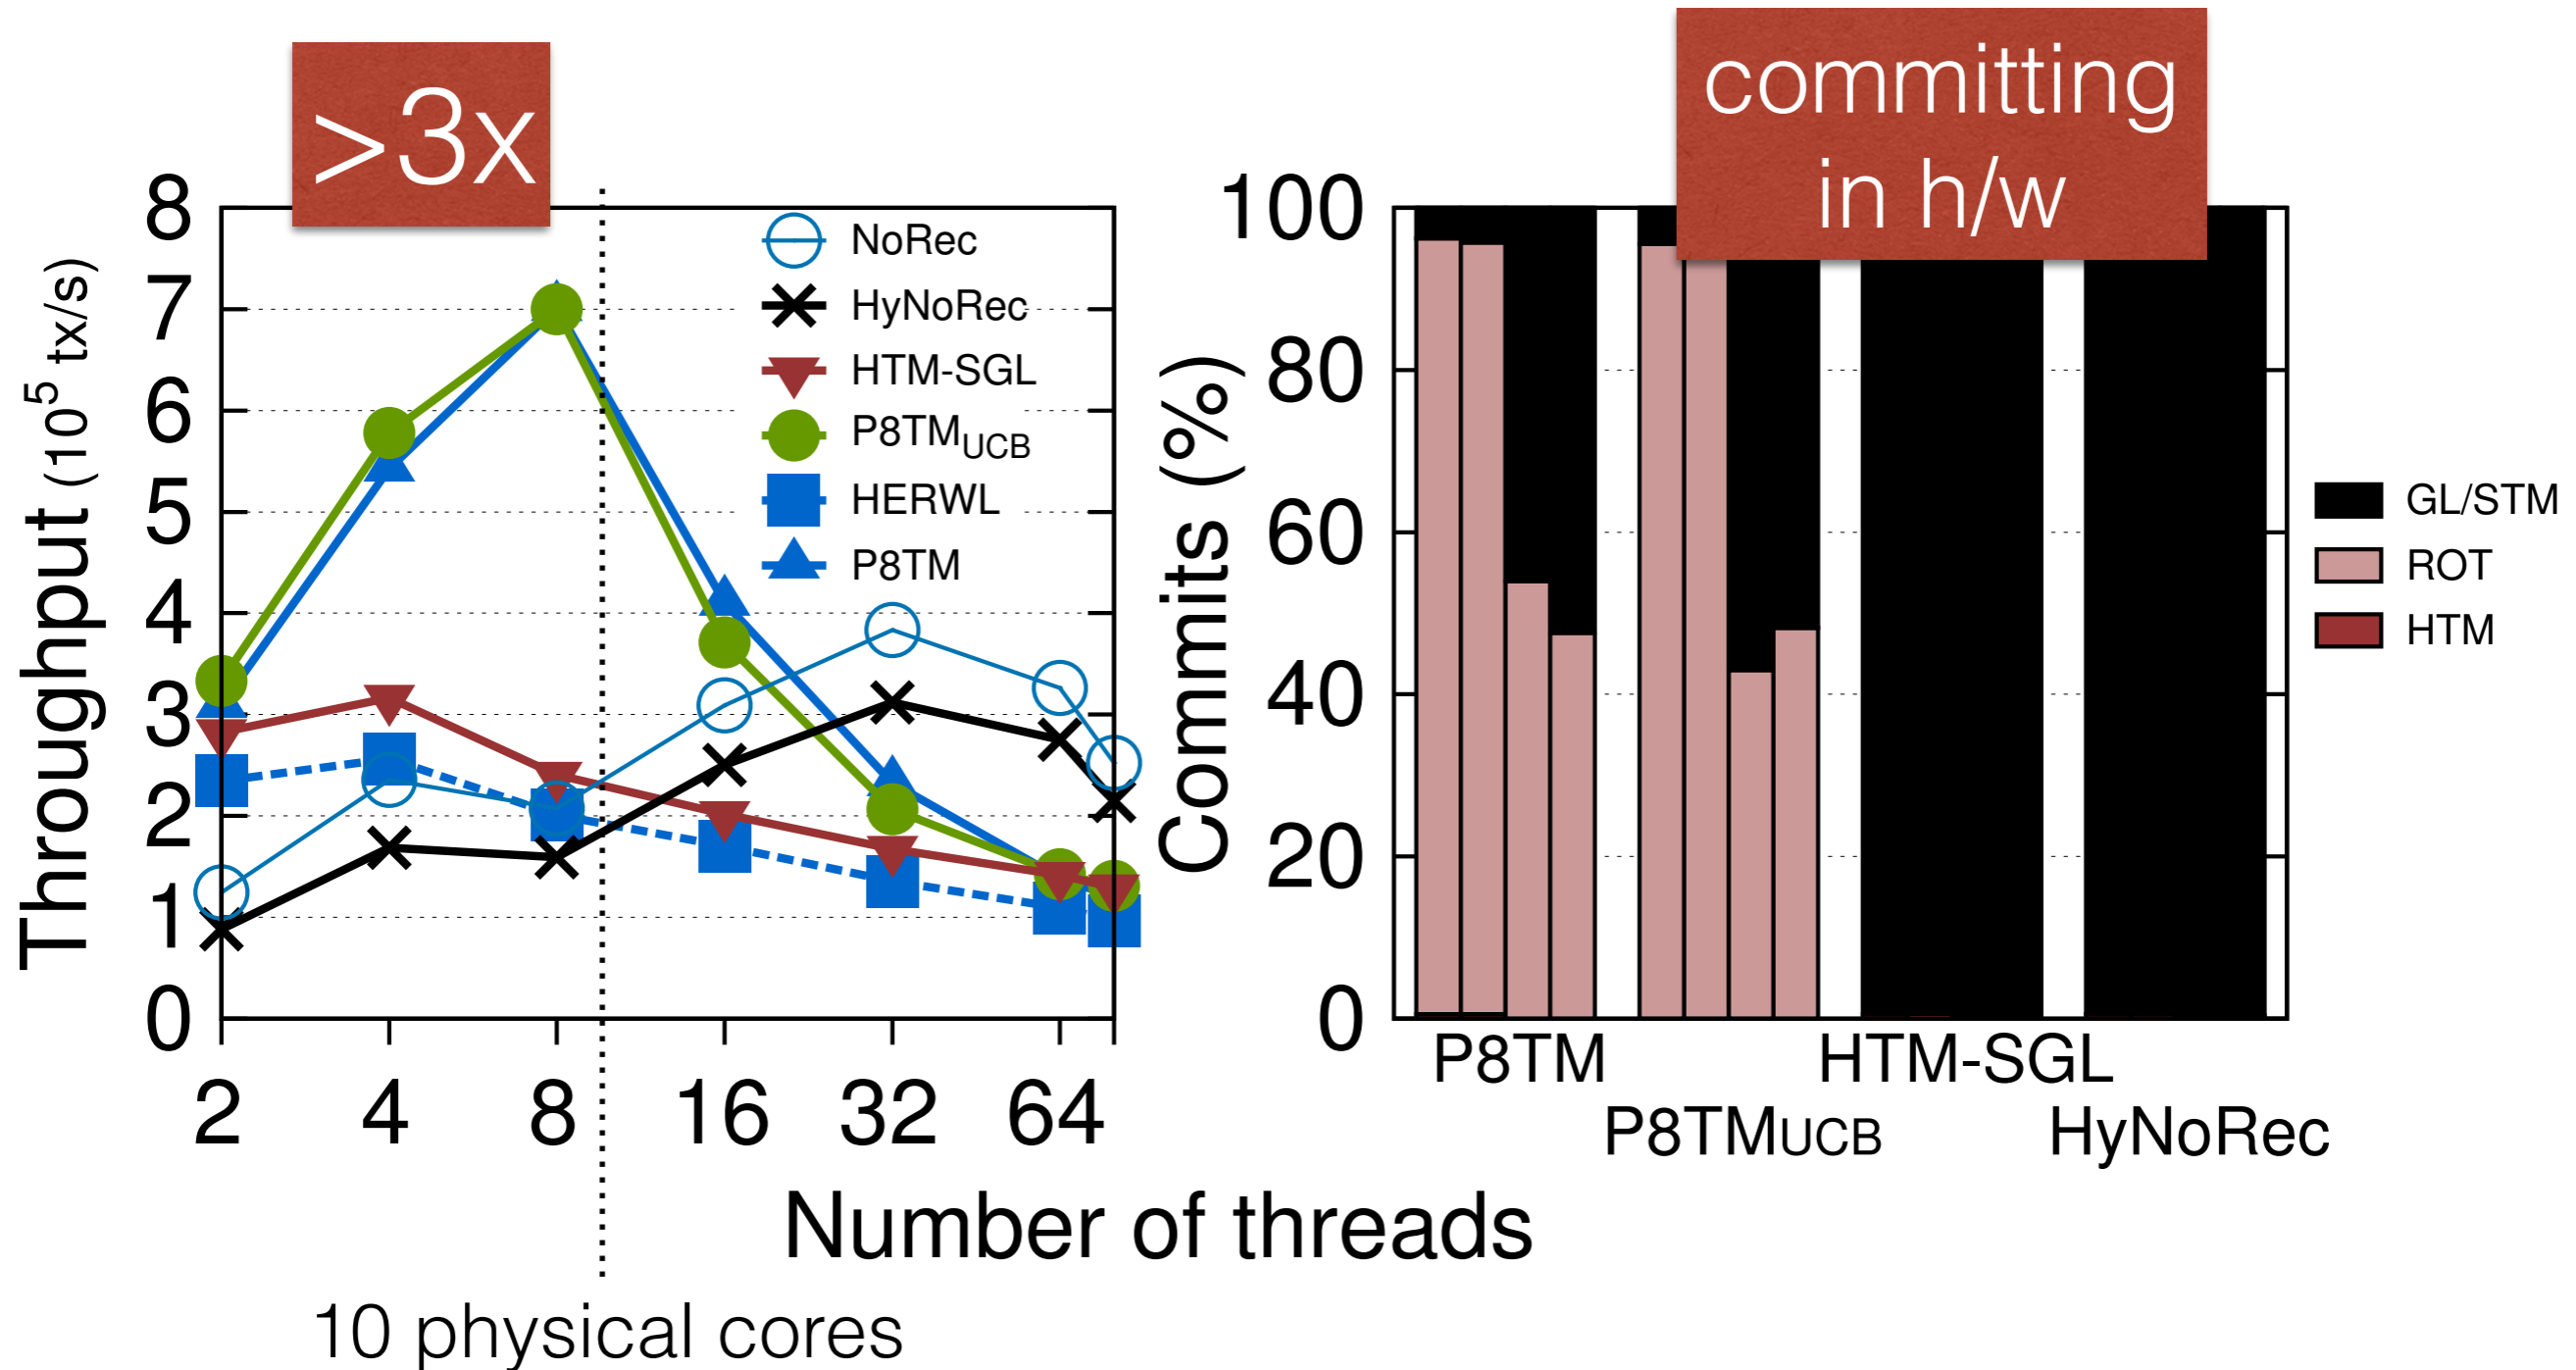
Evaluation: Vacation



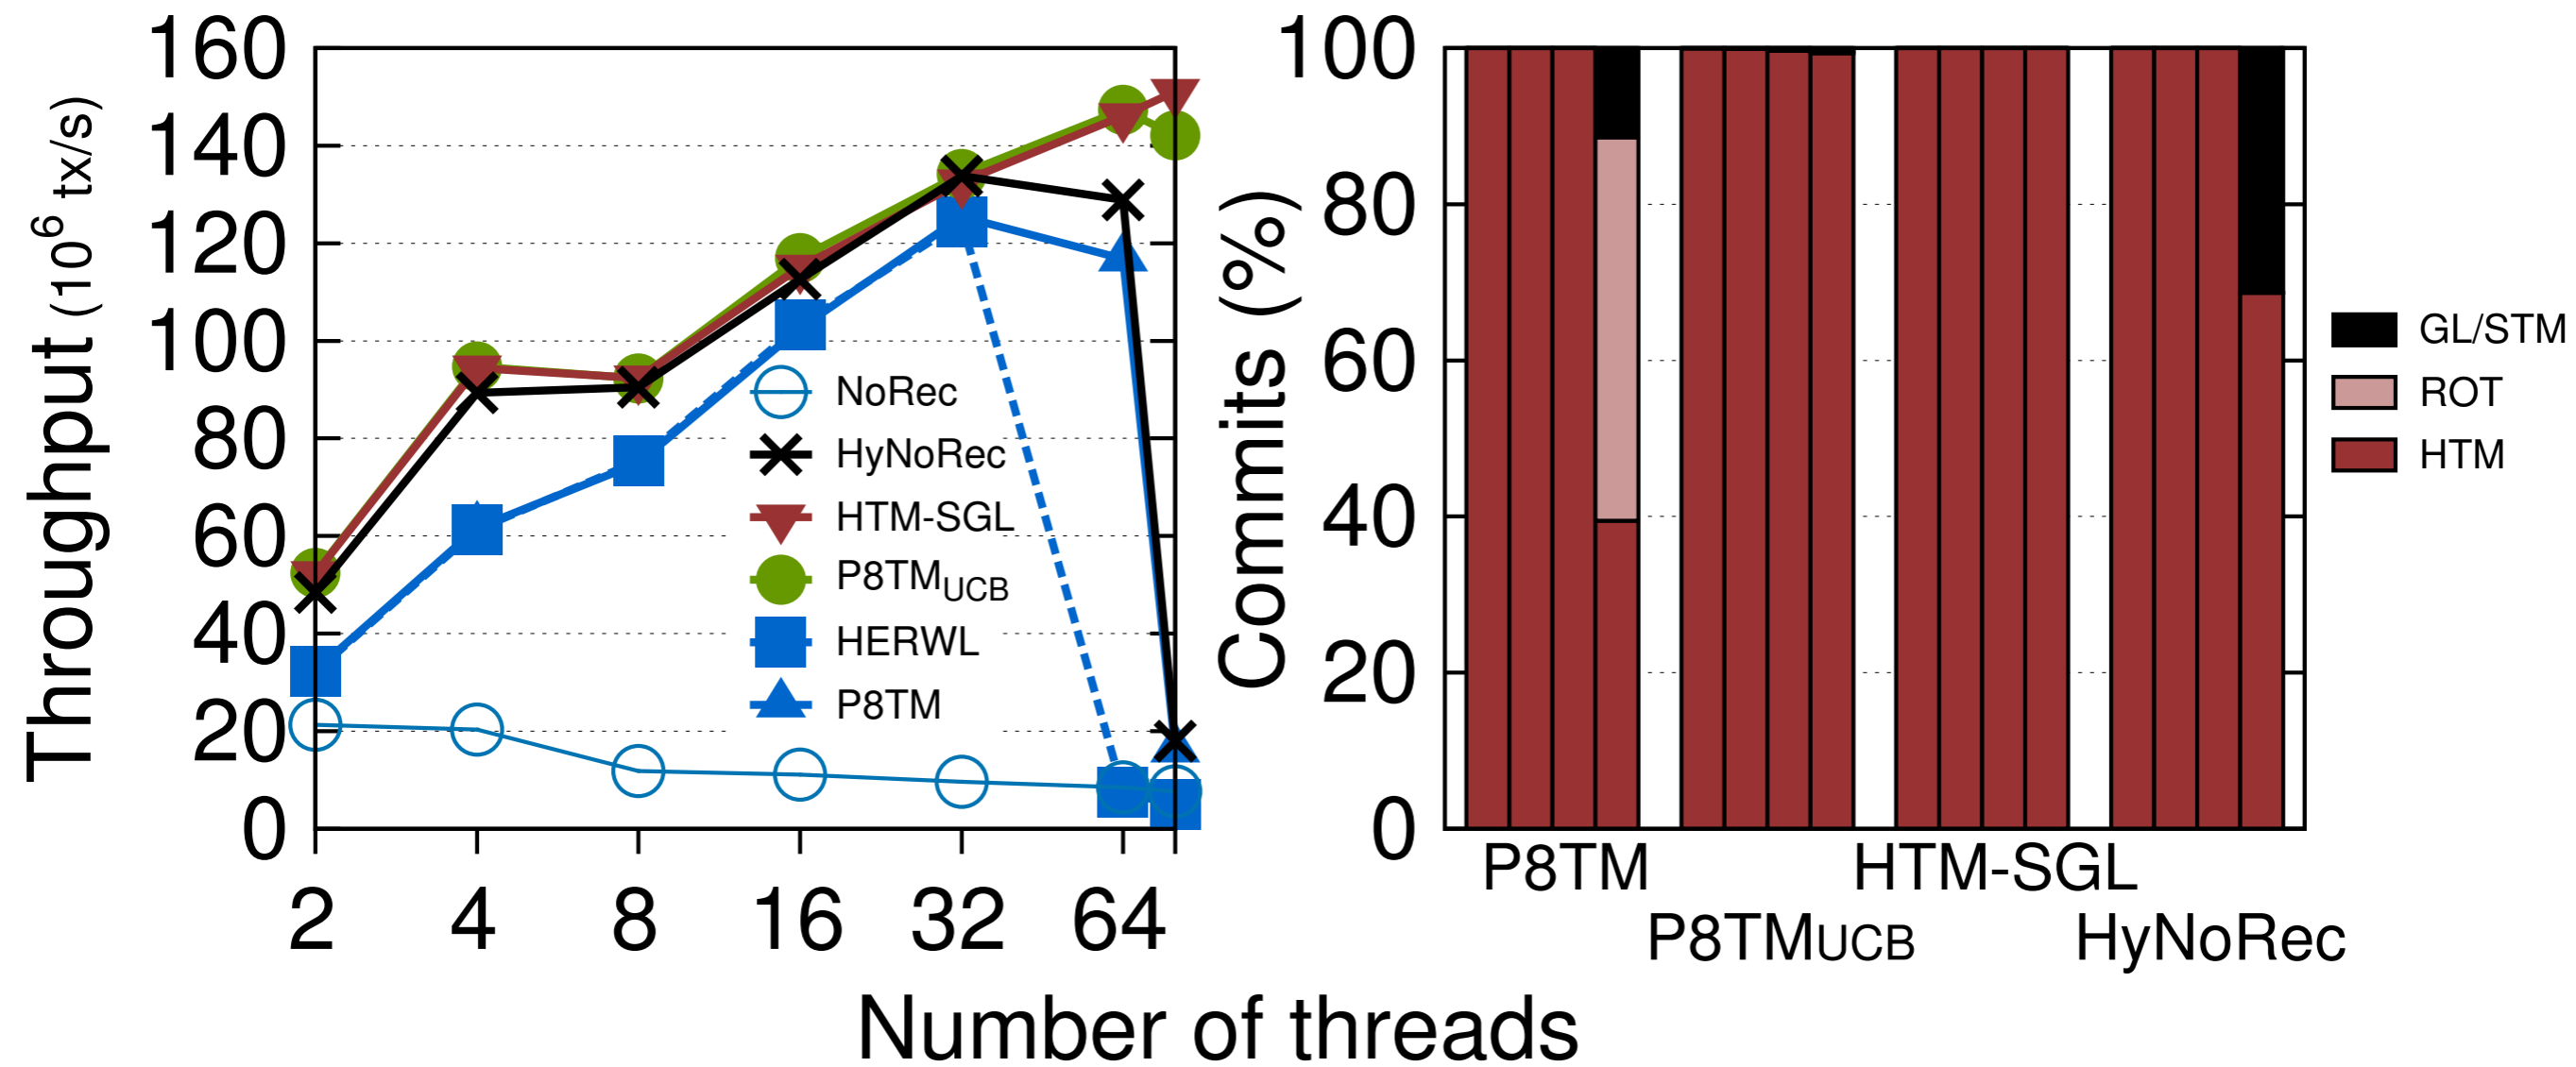
Evaluation: Vacation



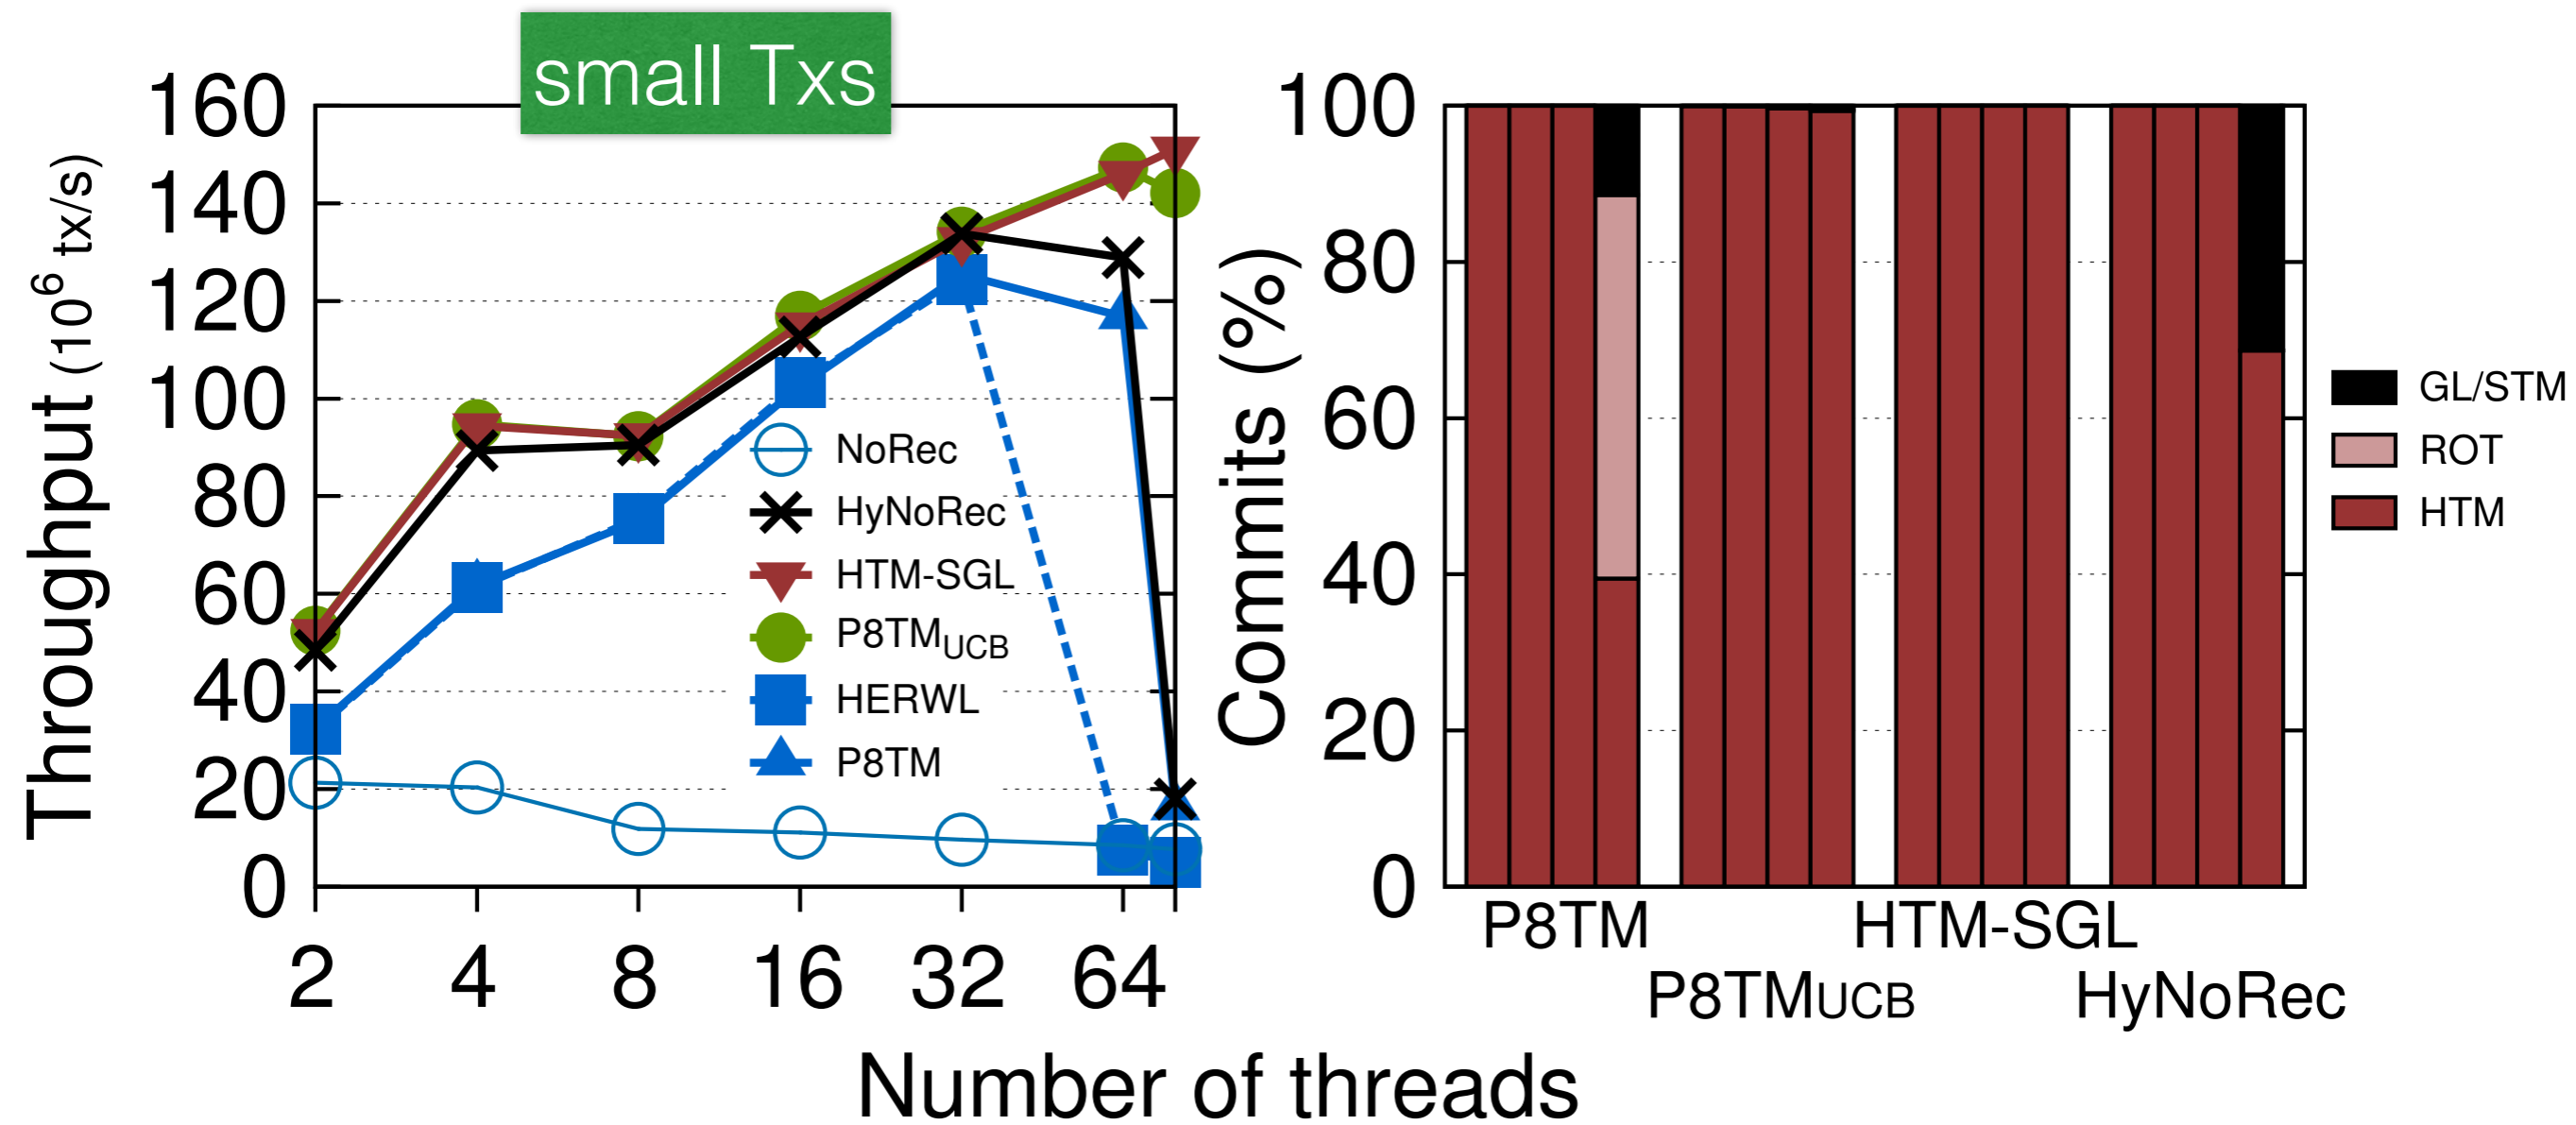
Evaluation: Vacation



Evaluation: SSCA2



Evaluation: SSCA2



Evaluation: SSCA2

