


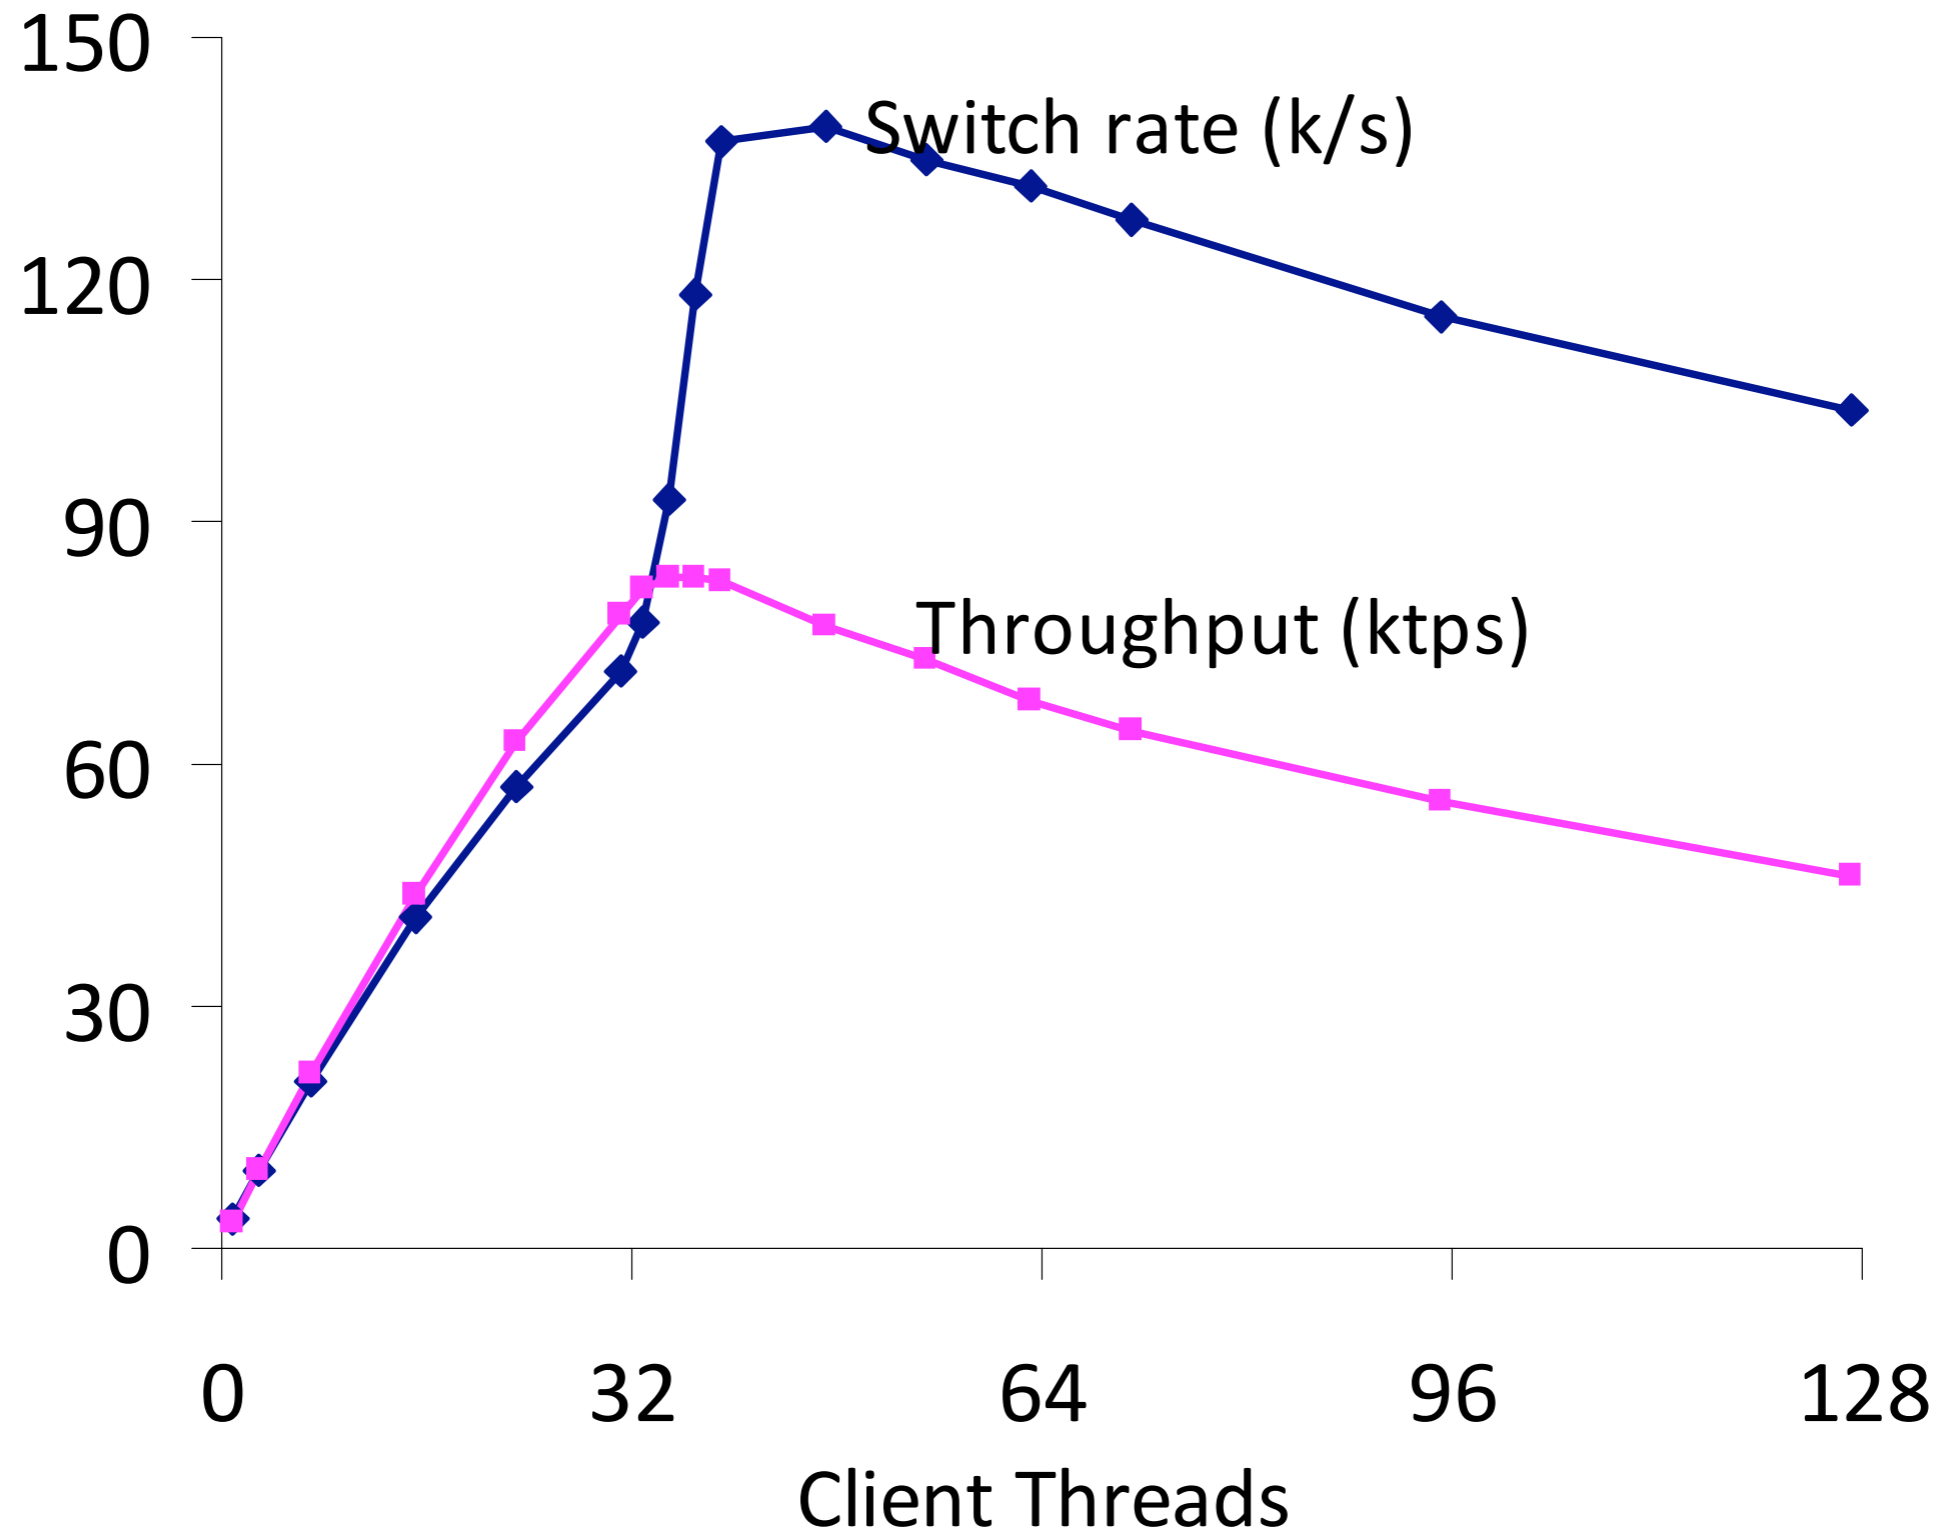
DECOUPLING CONTENTION MANAGEMENT FROM SCHEDULING

F. Ryan Johnson
Anastasia Ailamaki

Radu Stoica
Todd C. Mowry

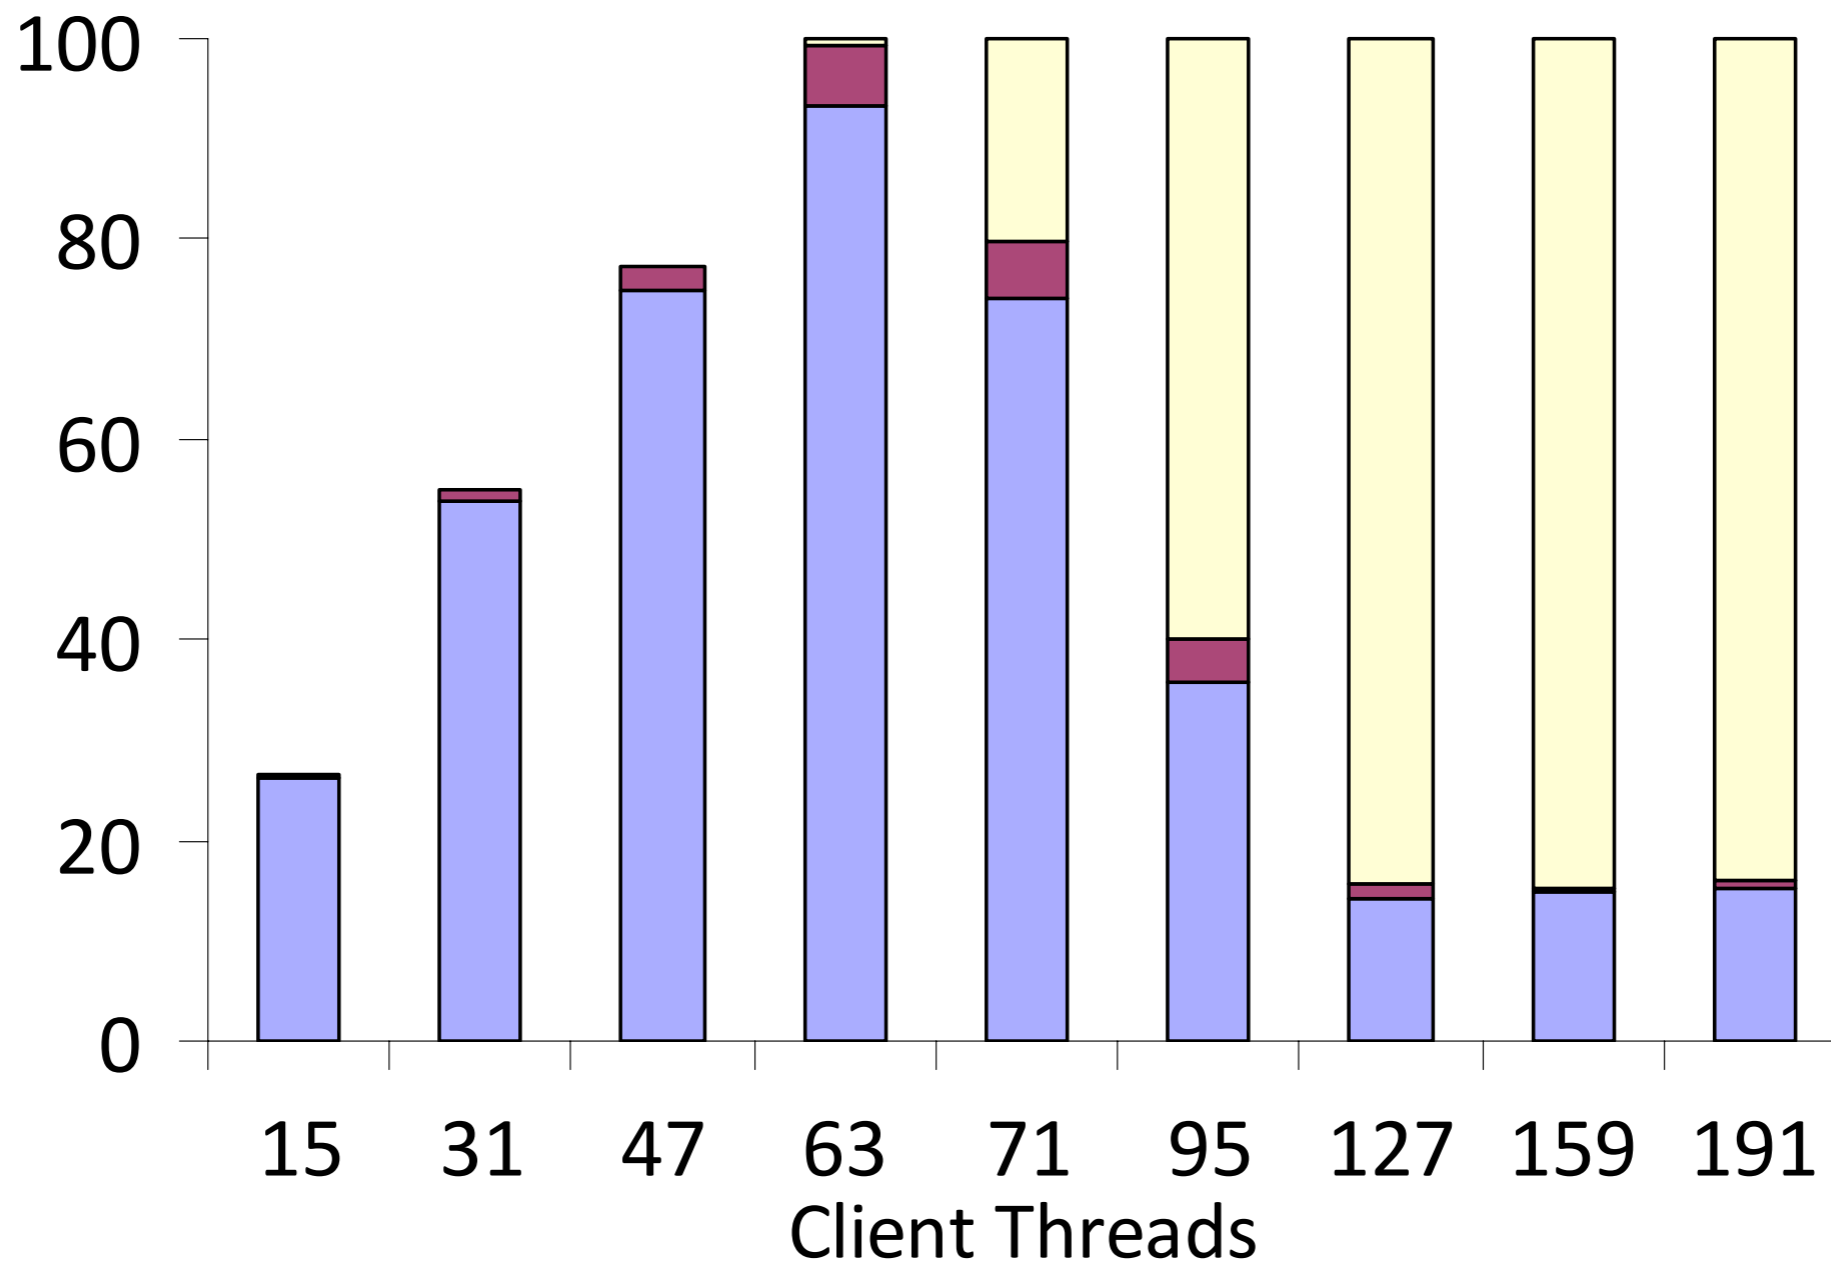
Throughput (ktps)

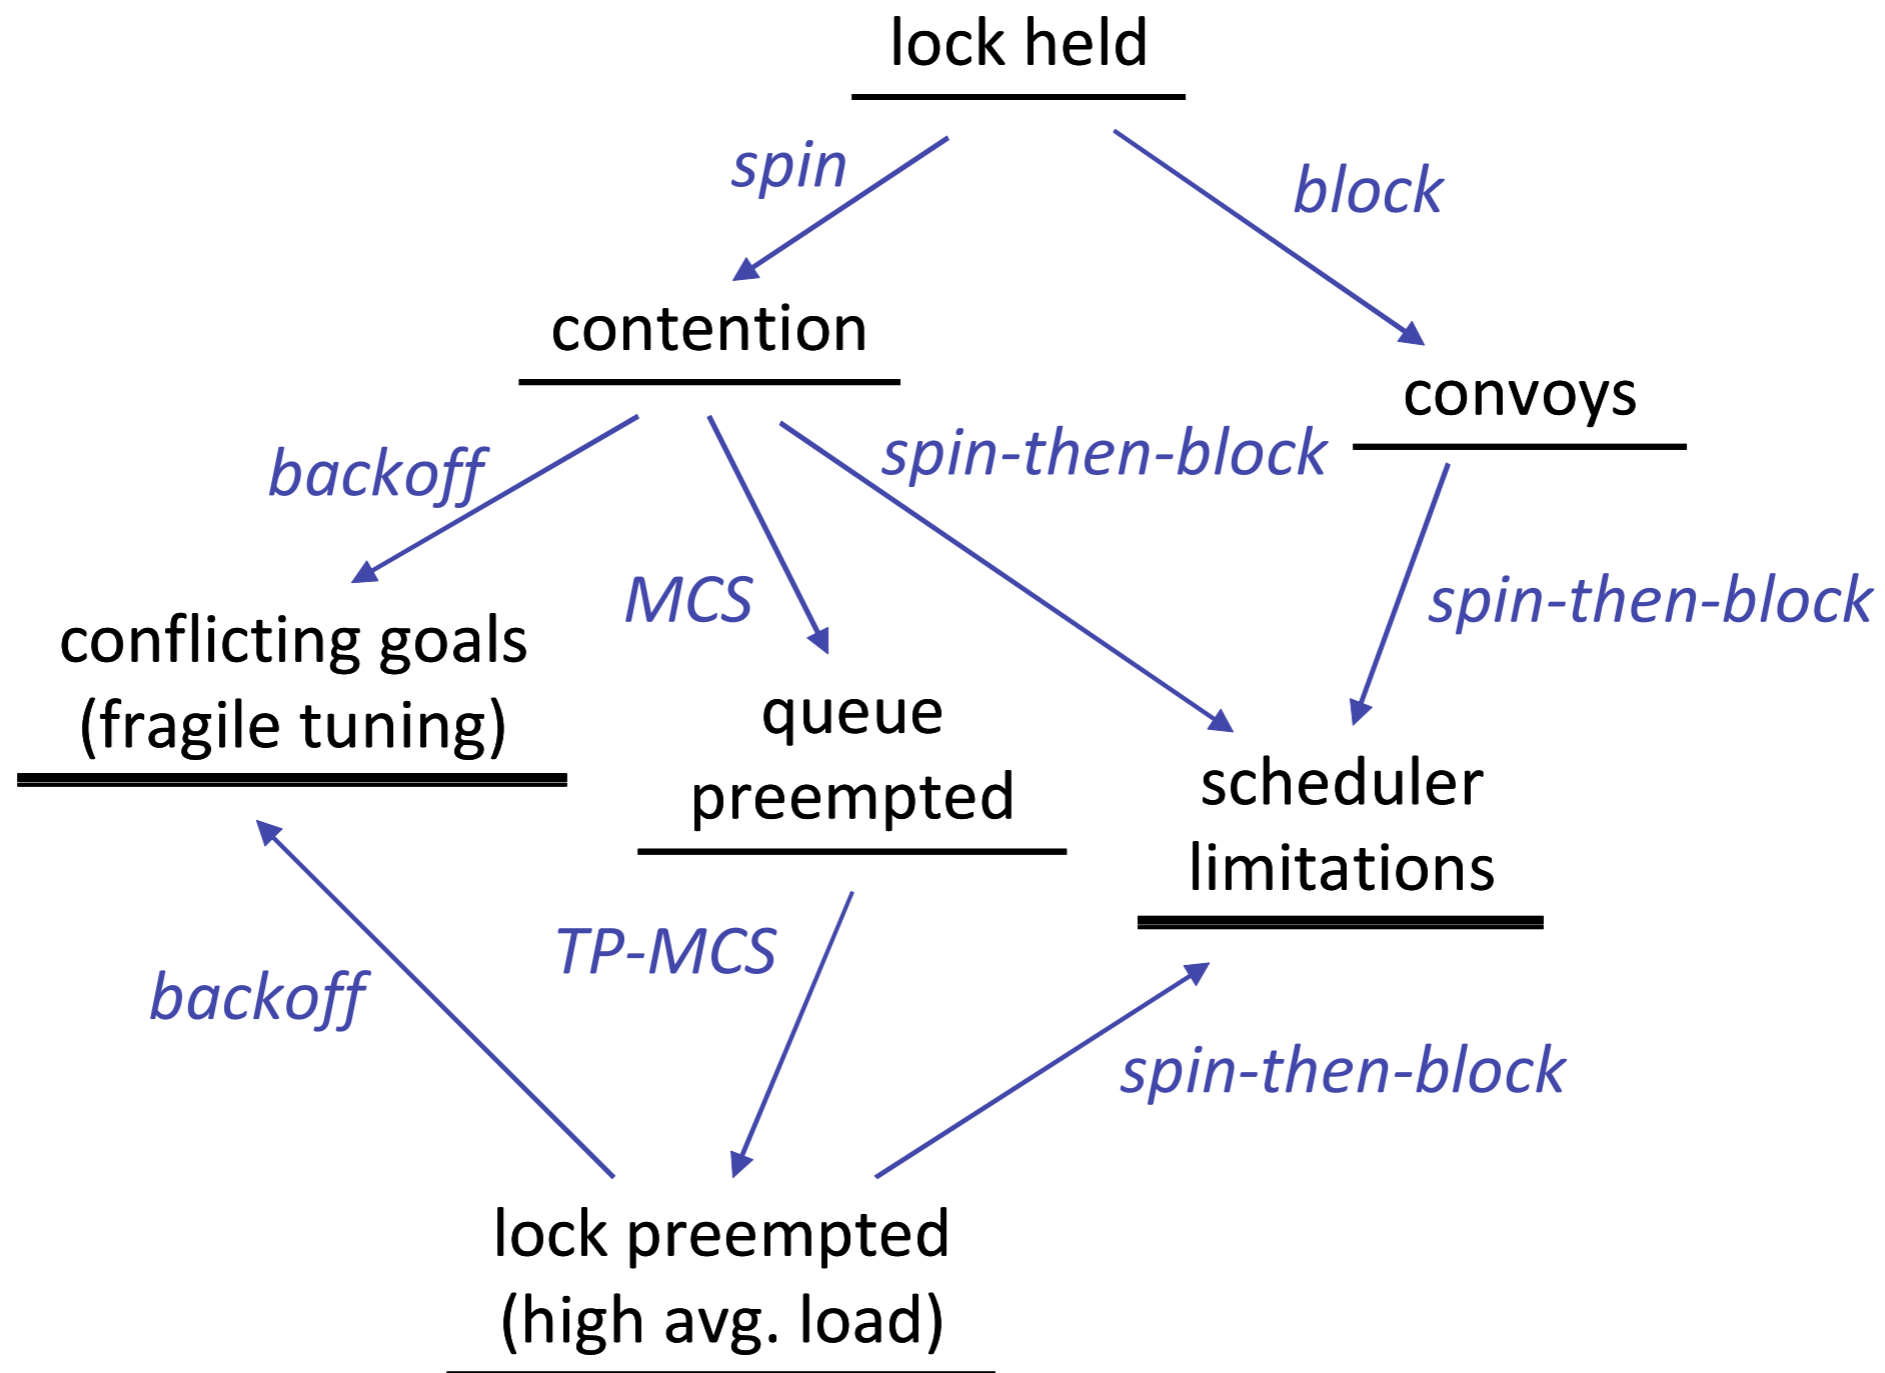




Machine Util (%)

Work Contention Prio-Invert





DECOUPLING

contention management

spinning

fast lock handoff

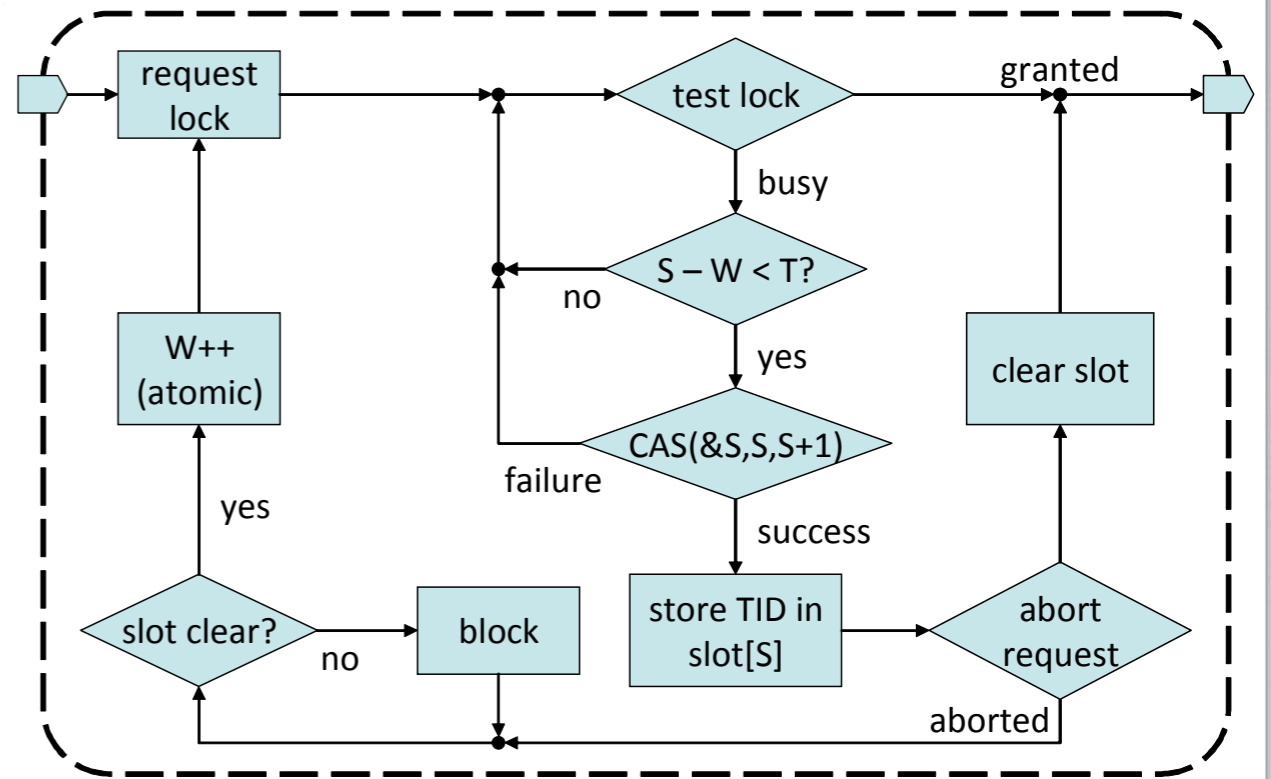

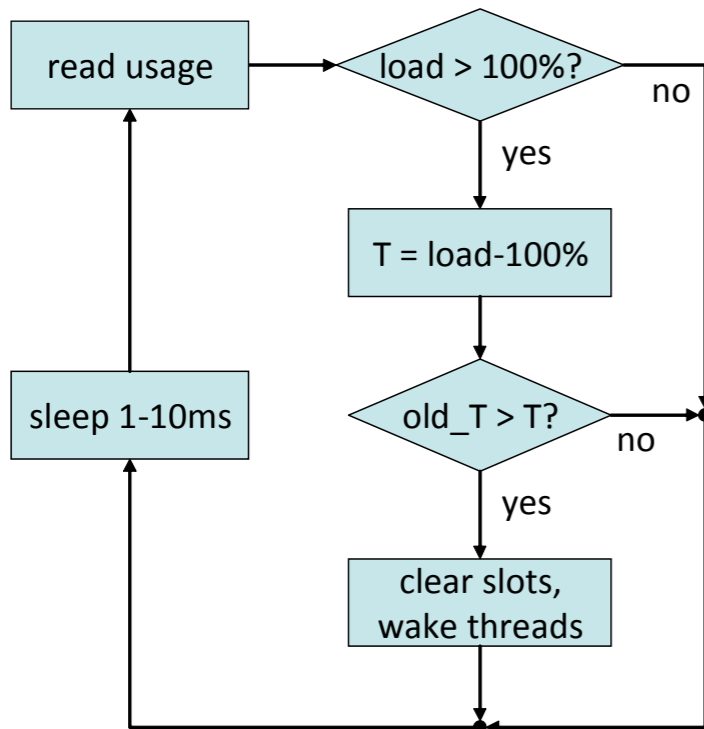
on critical path

load control

blocking

scheduler overhead

off critical path



OS SUPPORT

- schedule daemon periodically & independent from system tick

high-res timer

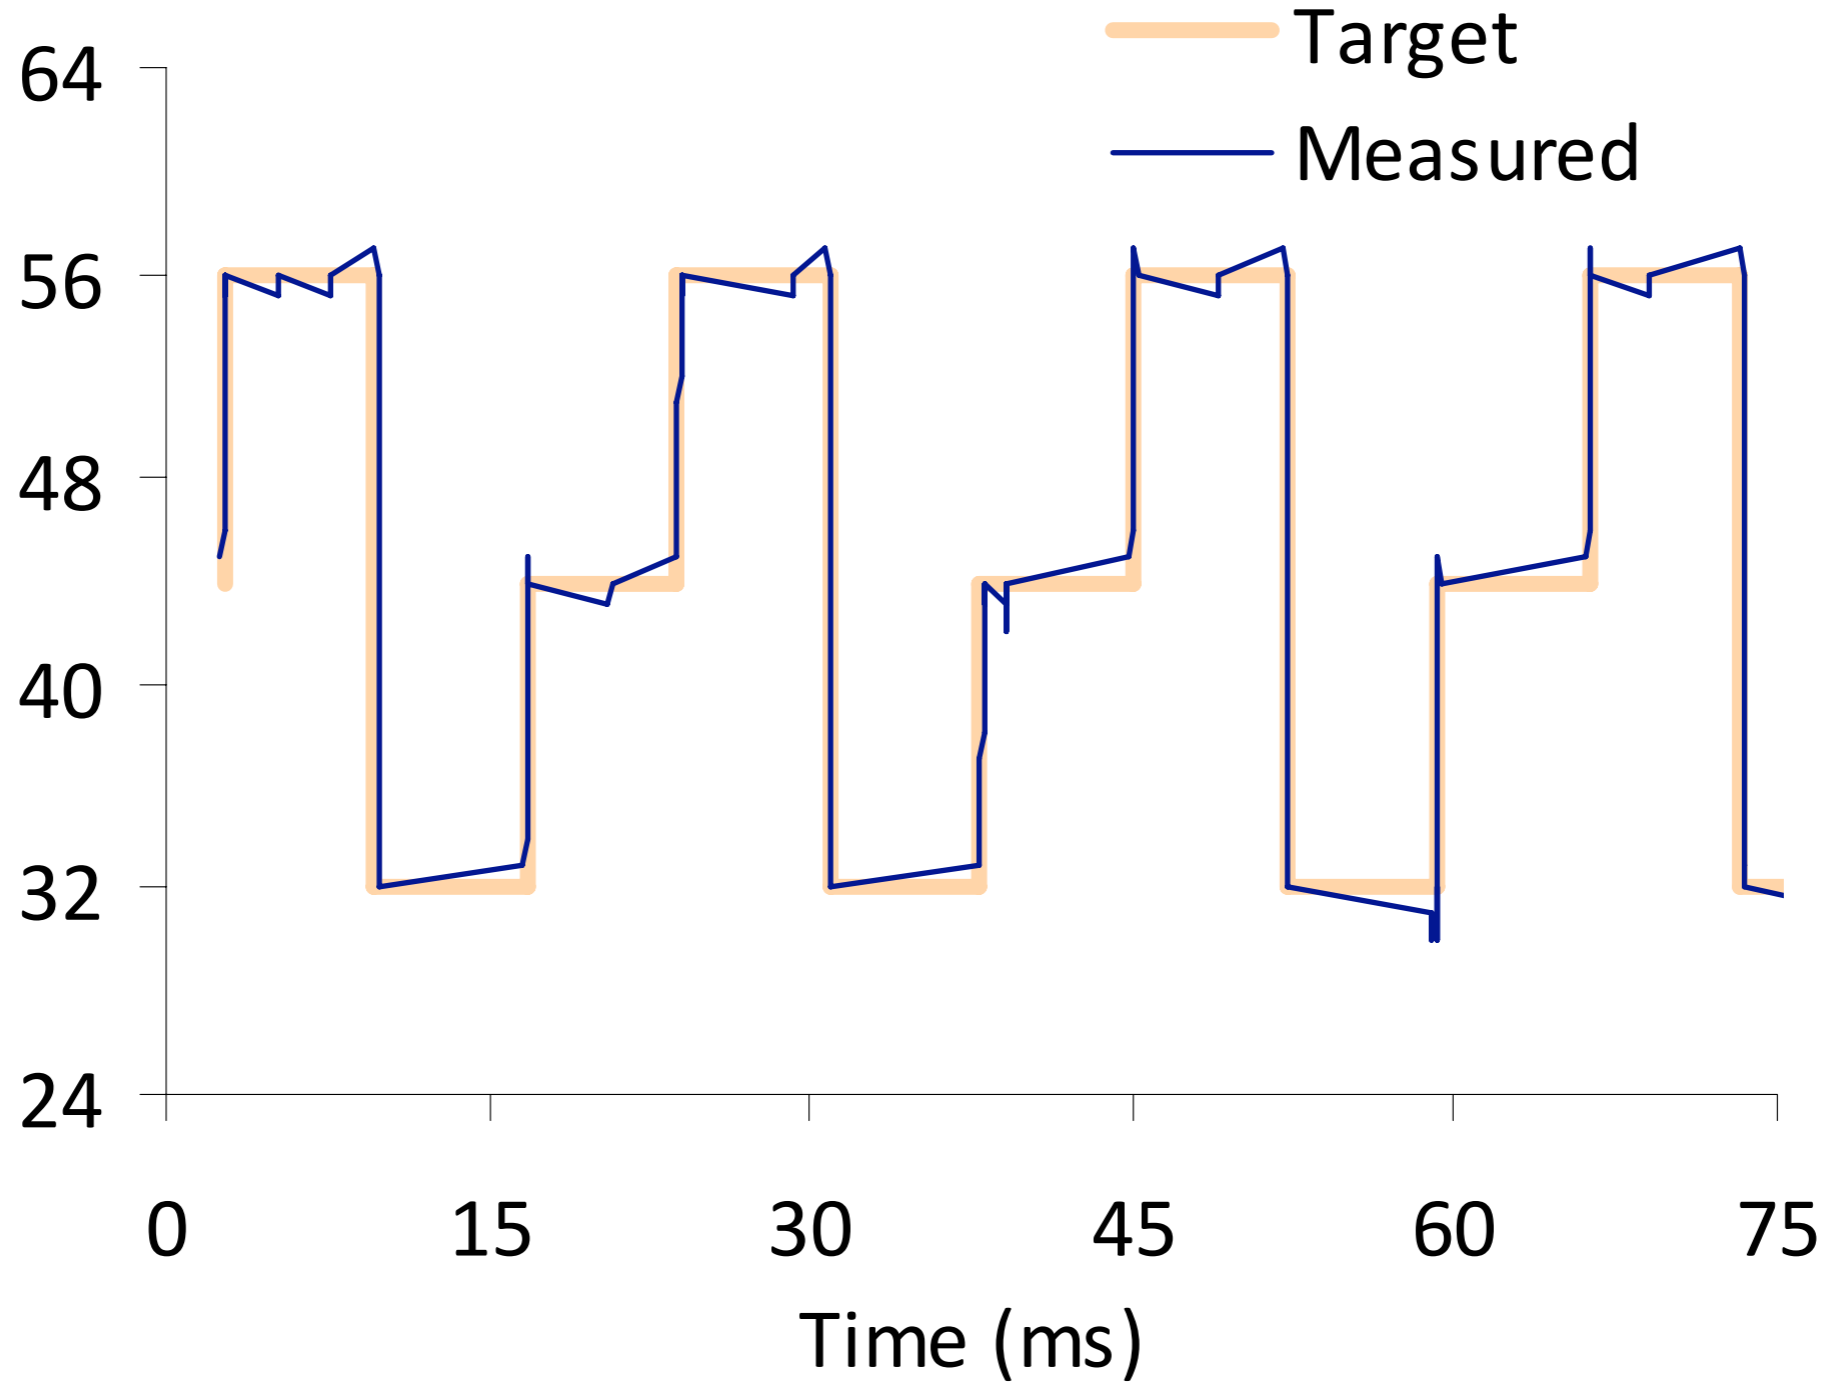
- measure load accurately and with high resolution


Solaris microstate

- efficient deschedule and wake


futex, lwp_park

Running Threads





Normalized Throughput



LIMITATIONS

- large, transient changes in load
 - apply some control theory
- nested critical sections
- load measurement cost linear in thread count
 - load-change notification instead of polling

CONCLUSION

- unwanted interaction between scheduling and contention management lead to poor performance
- decouple the two
 - use spinning in response to contention
 - use blocking to control the number of runnable threads
- implementation is transparent to the application

DISCUSSION

- Did you fully understand the presented algorithm?
- What about the increased energy consumption due to spinning?
- Do state-of-the-art spinlocks use MONITOR-MWAIT?
- Is this an argument for parallelism, where the number of threads is determined by the OS, not the application?