Impact of DVFS on n-Tier Application Performance

Qingyang Wang, Yasuhiko Kanemasa, Jack Li,
Chien-An Lai, Masazumi Matsubara, Calton Pu

Resource Utilization Paradox

- Data centers are supposed to run at **high utilization** (for high return on investment)

- But servers in typical data centers are only utilized 18% time on average.

  Gartner *[Dec 2010]*

- Why? SLA sensitive applications suffer from wide response time fluctuations at relatively low utilization (e.g., 60%) *–[Wang et al. ICDCS'13]*
  - One important reason: **Bursty workload** in web-facing applications

Bursty Workload

- Web-facing n-tier applications
  - Workload fluctuates between high and low CPU requirement *-[Mi et al. Middleware08]*

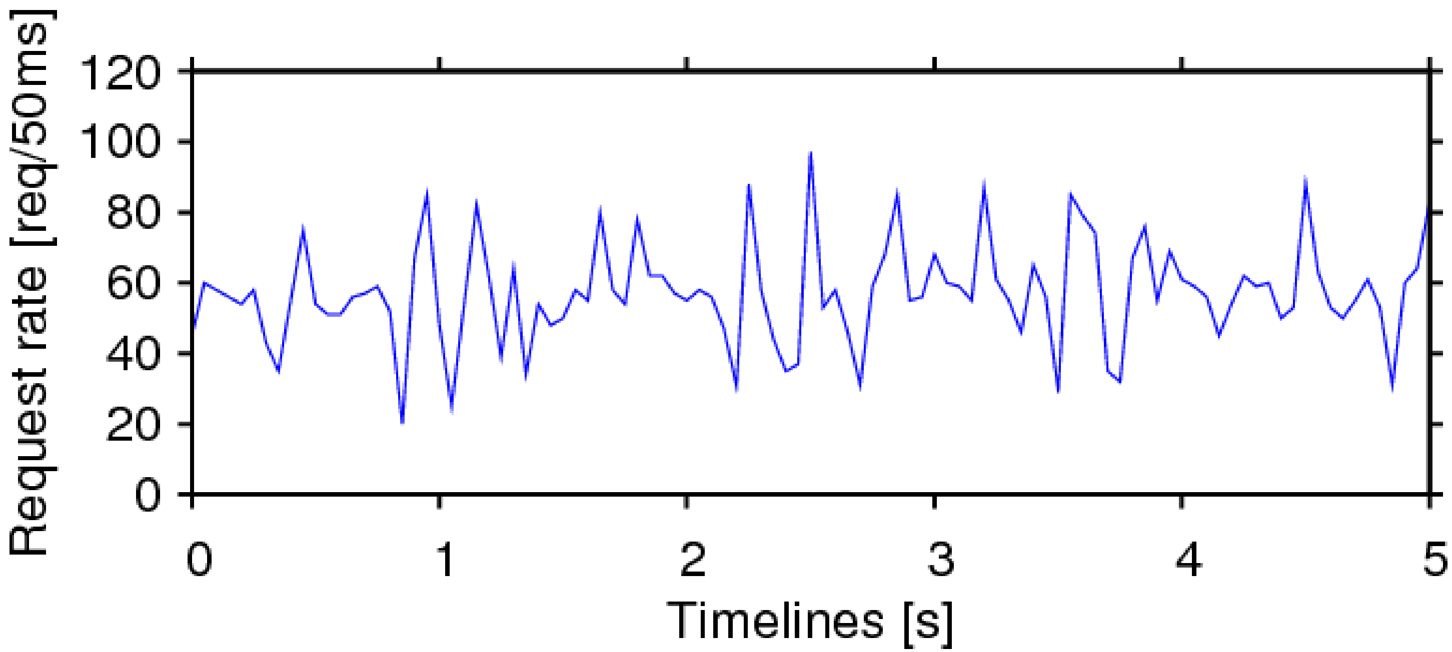

What's the Impact of DVFS on Bursty Workload?

- Dynamic Voltage and Frequency Scaling (DVFS) can help handle bursty workload.
  - Adjusting CPU voltage/frequency on-demand

- Potential mismatch between CPU requirement and DVFS adjustment
  - Workload burst length can be close to DVFS control period (e.g., 500ms)

- Evaluation using measurements and exploration using simulation

Experimental Setup

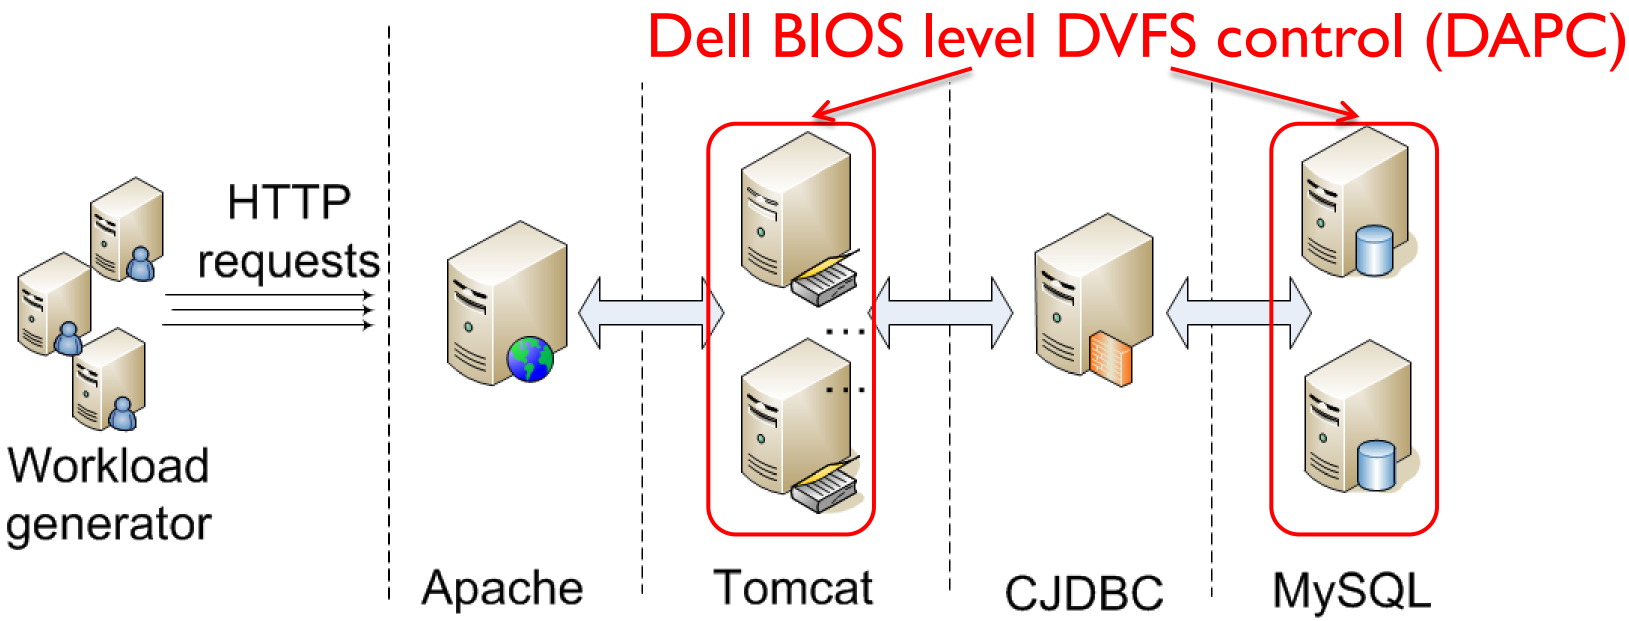

- ❑ RUBBoS benchmark: a bulletin board system modeled after Slashdot
- ❑ 24 web interactions
  CPU intensive
- ❑ Default workload generator
  naturally bursty

- ❑ Intel Xeon E5607
  2 quad-core 2.26 GHz
  16 GB memory
- ❑ **Support P0~P8**
  P0: (2.26GHz/1.35v)
  P8: (1.12 GHz/0.75v)

Simulation Setup

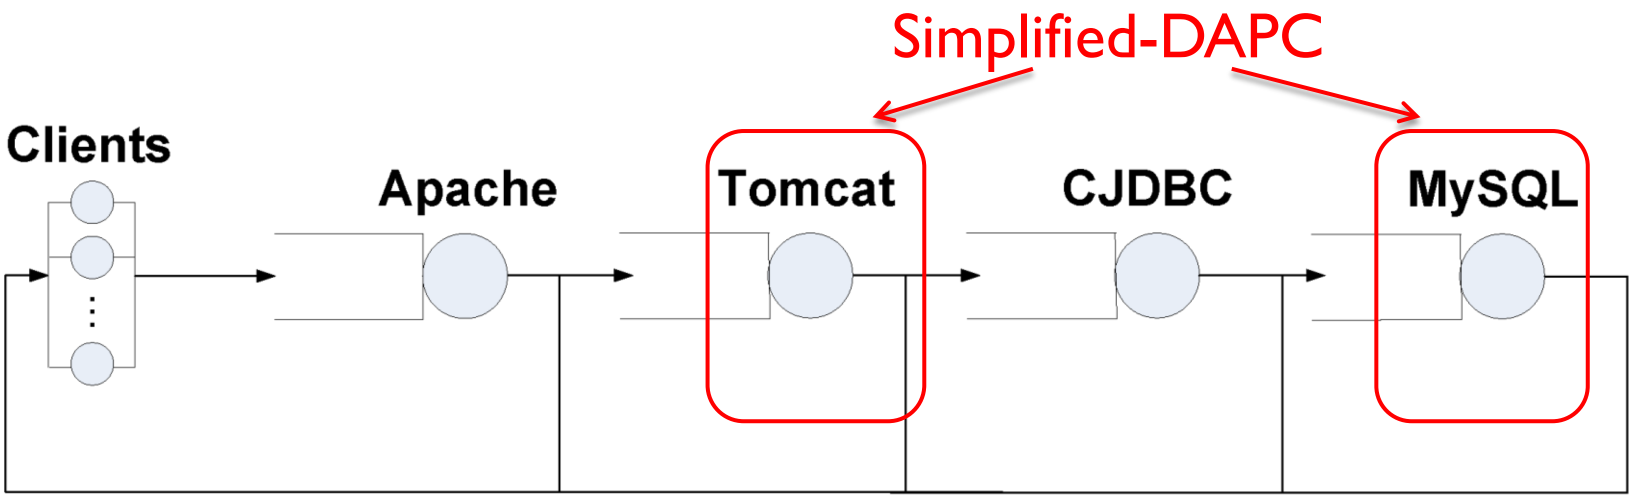

- Benefits of simulation:
  1. Extension of experimental study
  2. Abstraction from different levels of DVFS controller
- Parameterized from real experimental measurements
- Default workload generator naturally bursty

DAPC Power Saving vs. Performance Degradation

Power usage

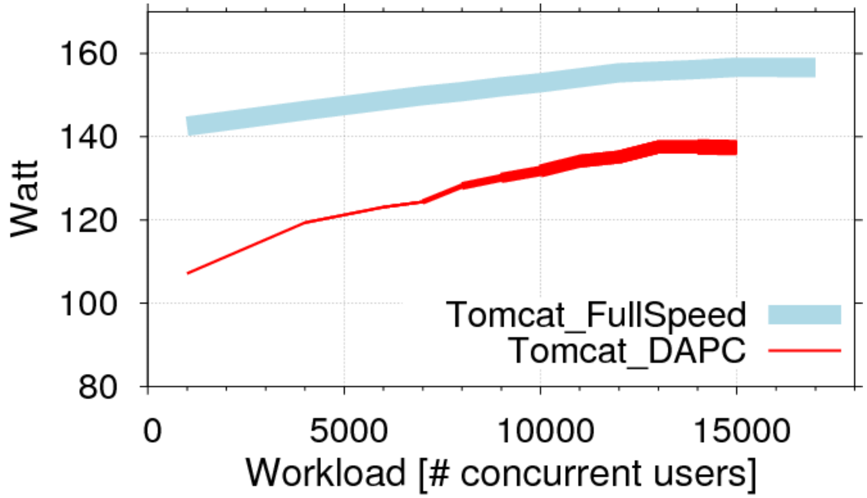

P0 state

P8 state

Throughput

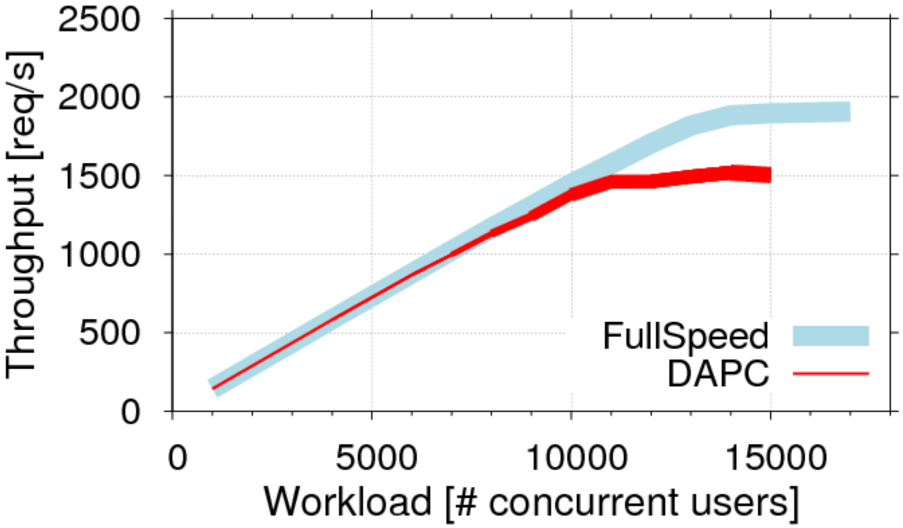

Response time

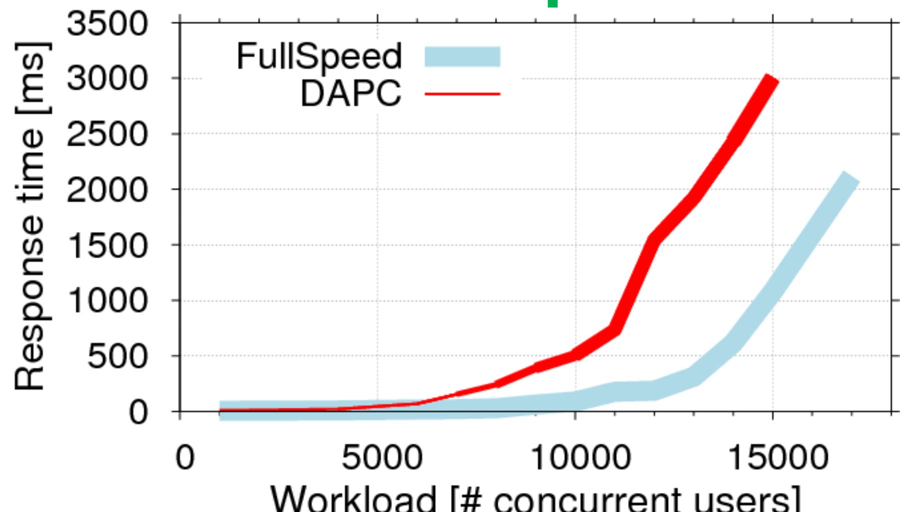

High Utilization Saves More

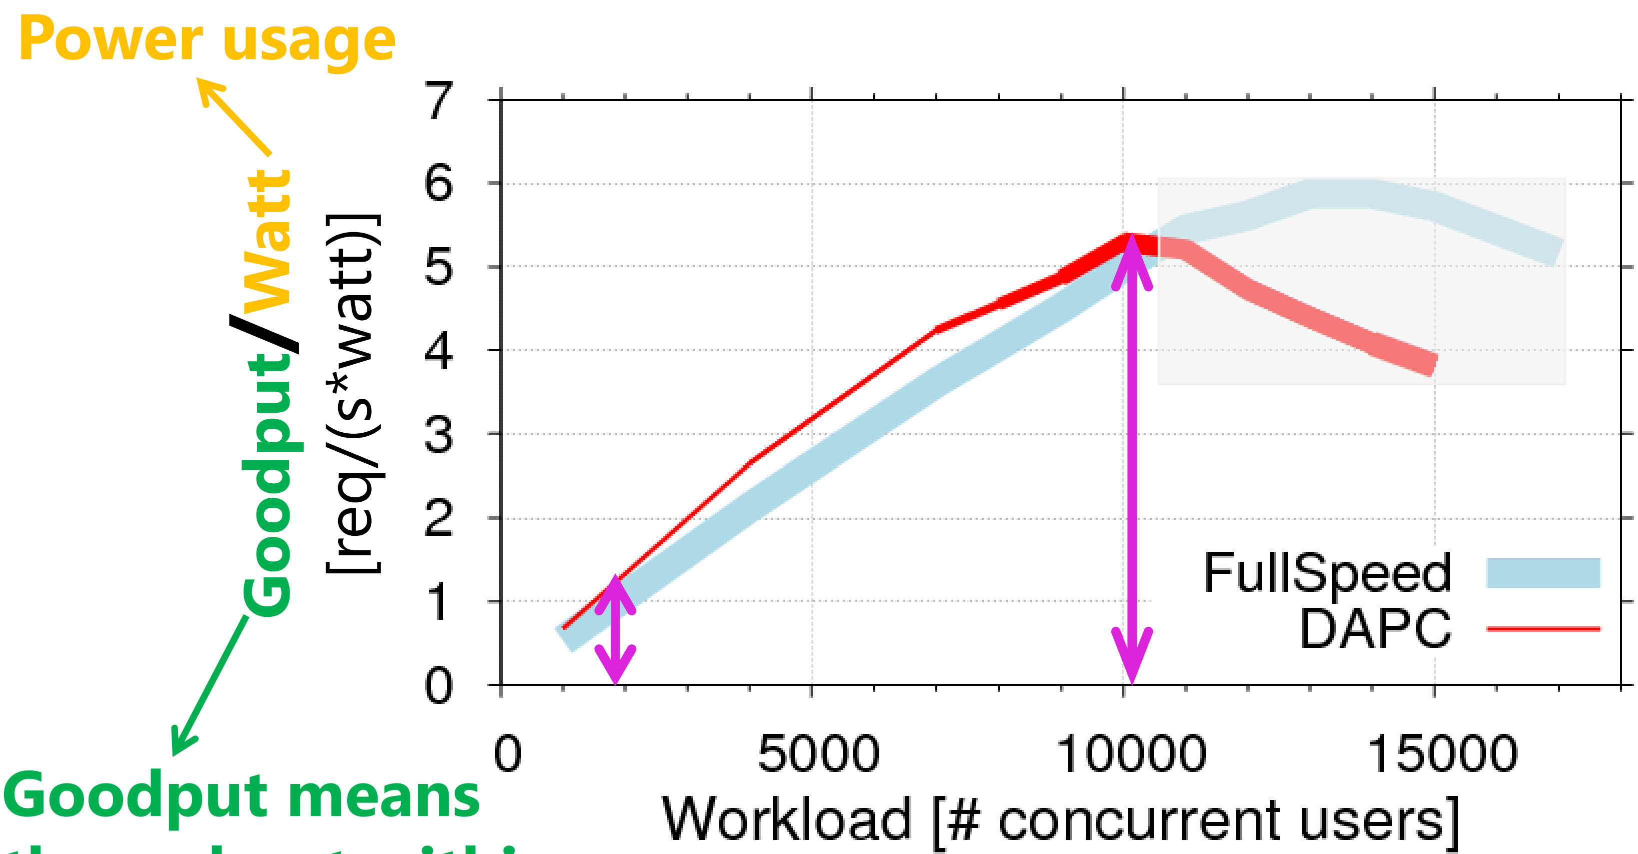

Goodput means throughput within SLA constraints

Georgia Tech College of Computing

DAPC at High Utilization

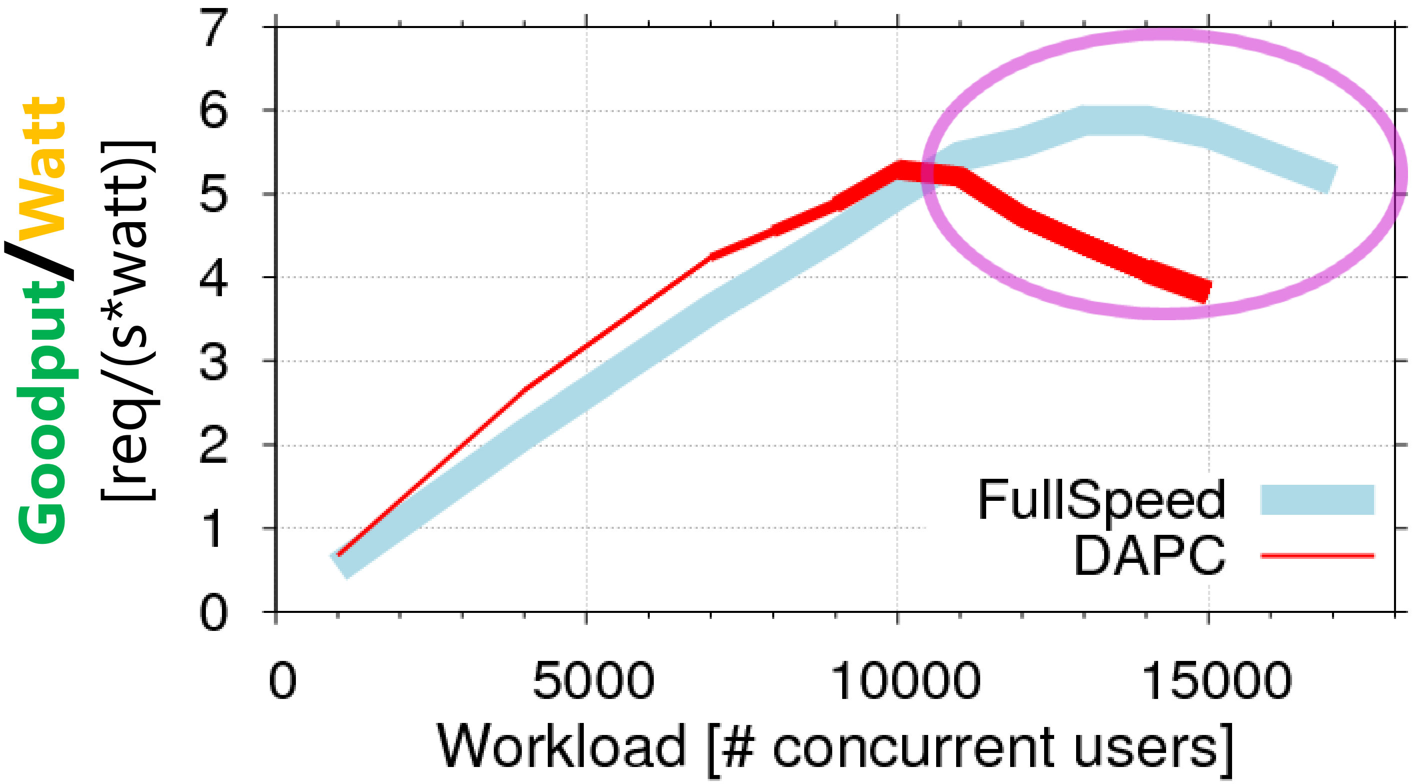

Georgia Tech College of Computing

Two Kinds of Performance Problems

- Large Response Time Fluctuations
  - Due to the delay of CPU P-state adaptation on bursty workload

- Throughput loss
  - Due to the rapidly oscillating bottlenecks between different tiers –*[Wang et al. Cloud'13]*

Large Response Time Fluctuations

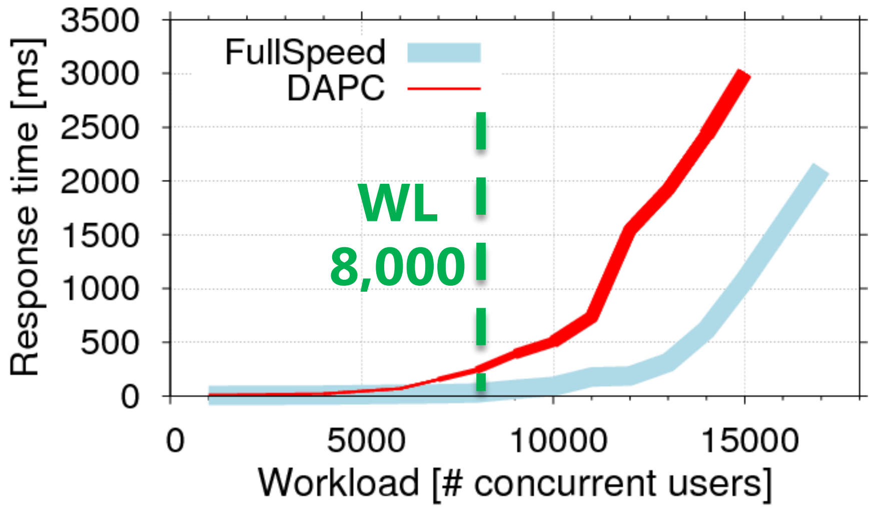

FullSpeed at WL 8,000

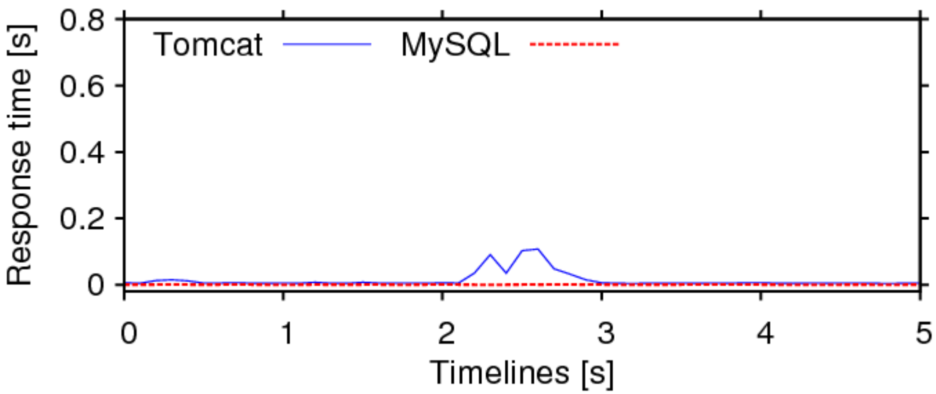

DAPC at WL 8,000

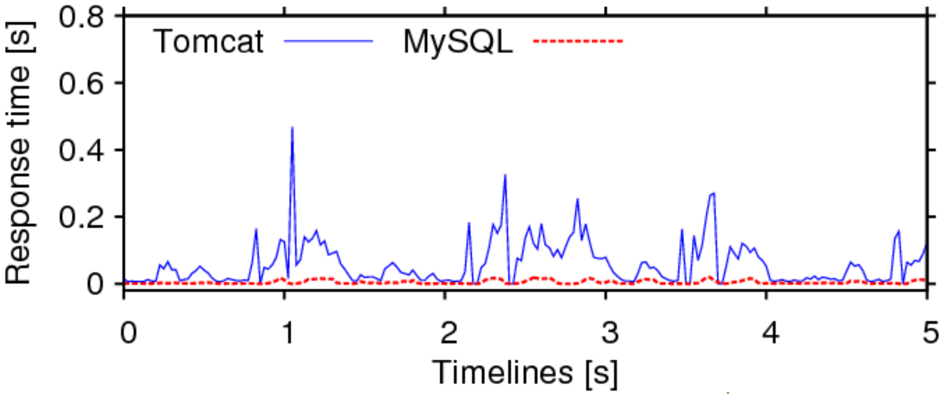

Georgia Tech College of Computing

Response Time Fluctuations Caused by Queued Requests

DAPC at WL 8,000

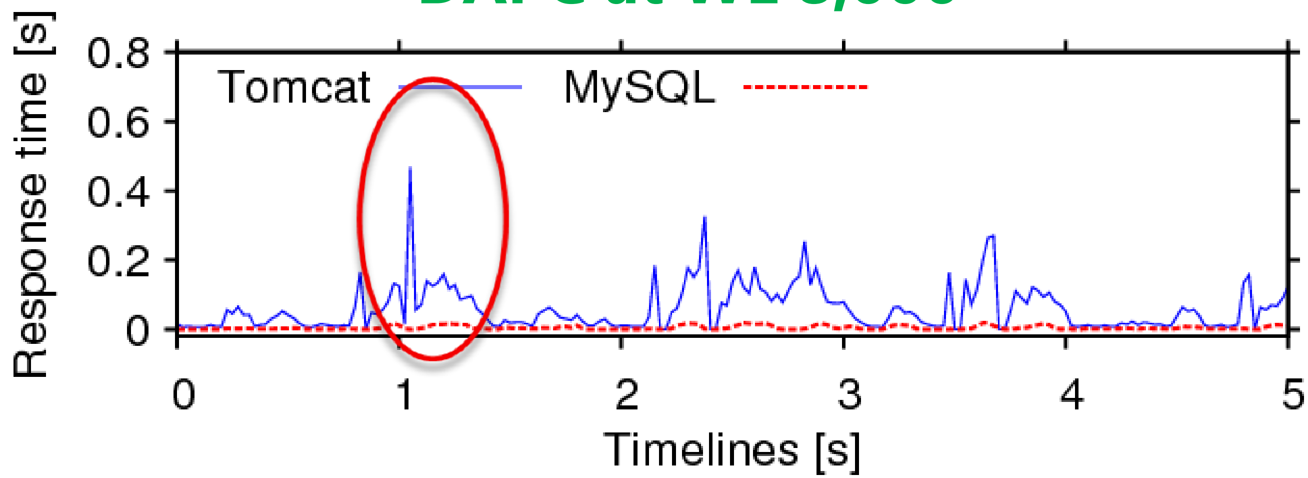

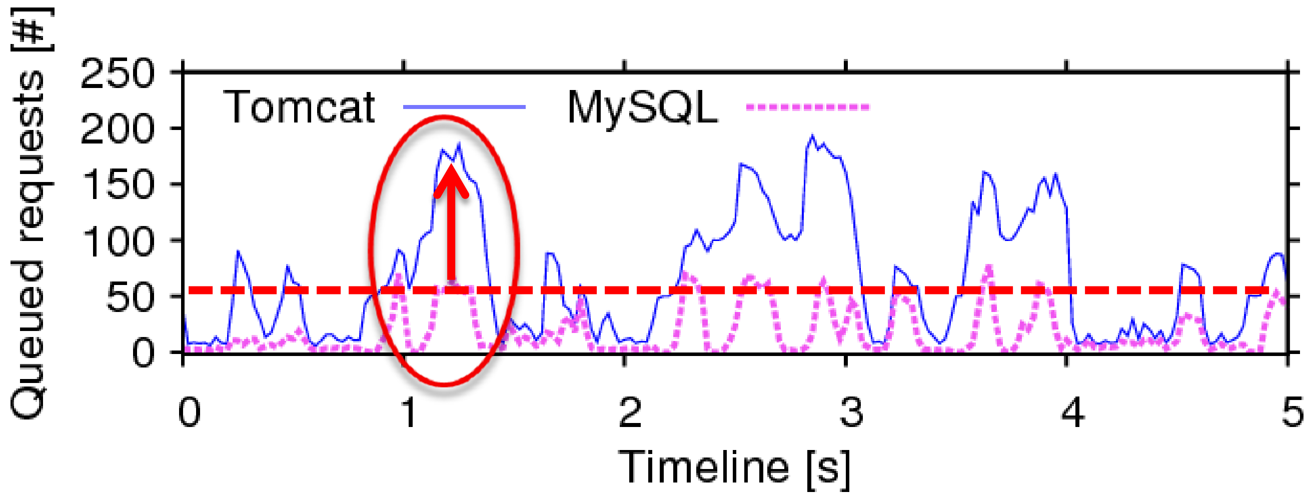

Push-Back: Upstream Queuing Amplification

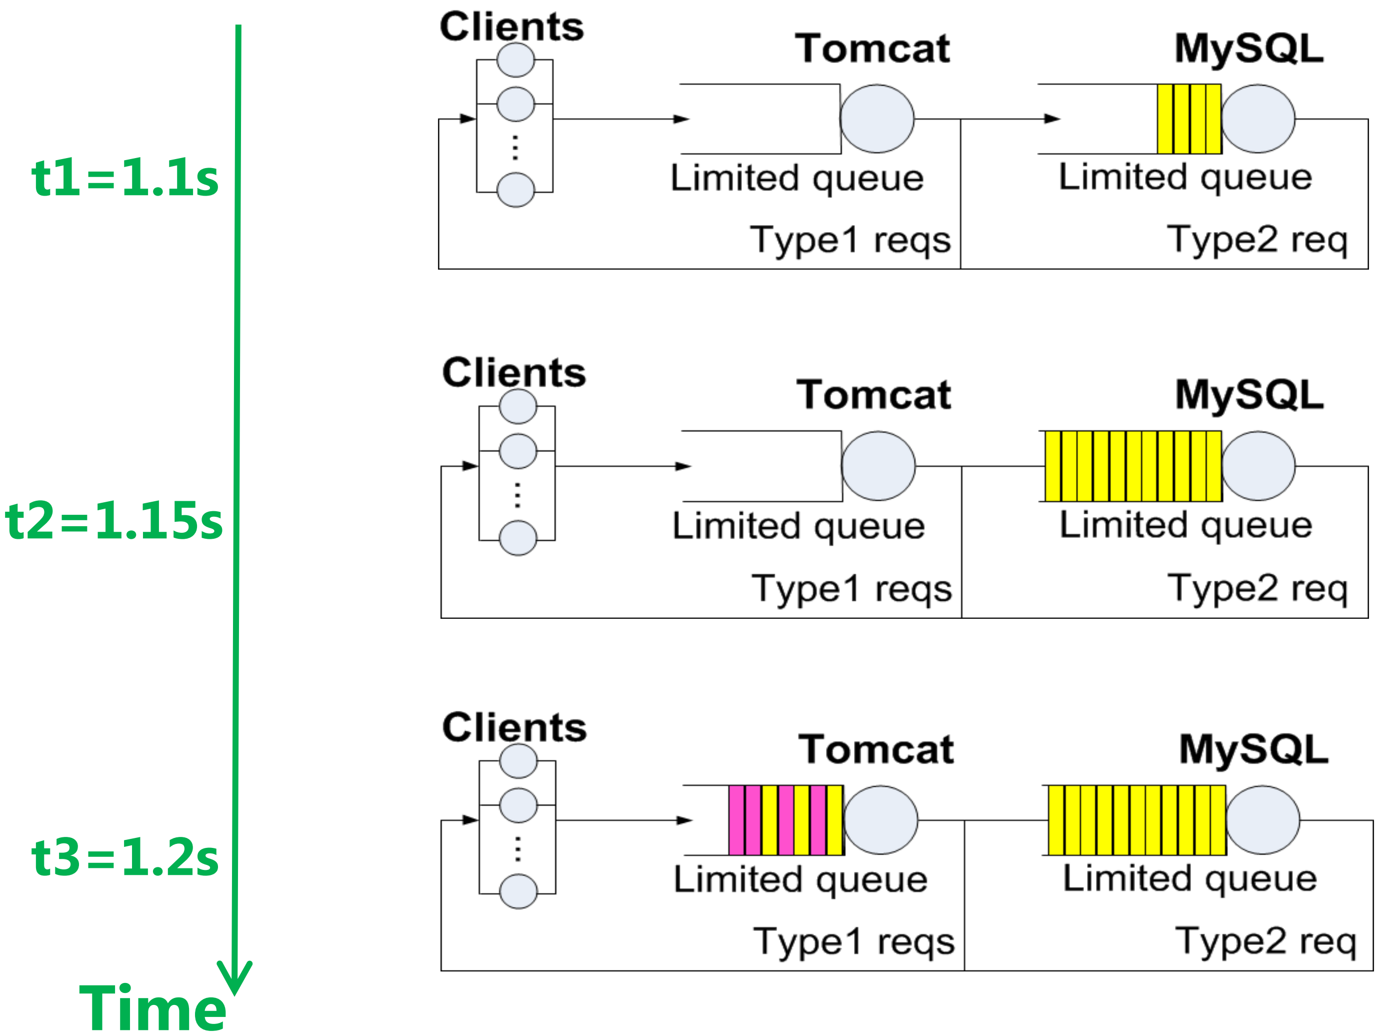

Type 1 req

Type 2 req

Type2 requests send SQL queries to MySQL while Type 1 requests don’t

Queuing Amplification Can Happen in 500ms

Simulation of DAPC at WL 8,000

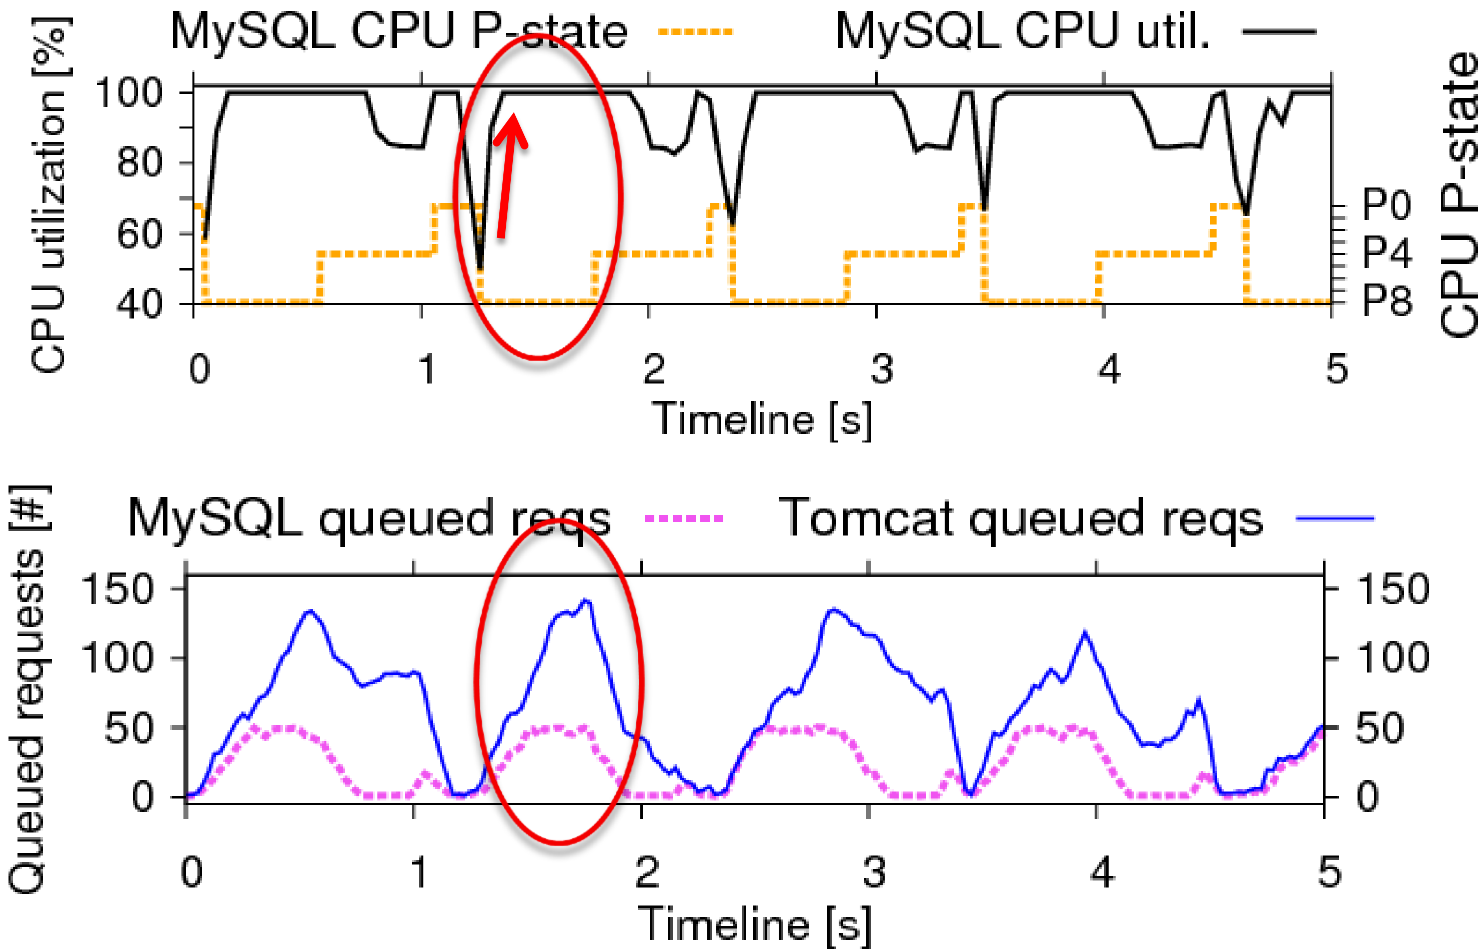

Georgia Tech College of Computing

Workload Burst Length: An Important Parameter

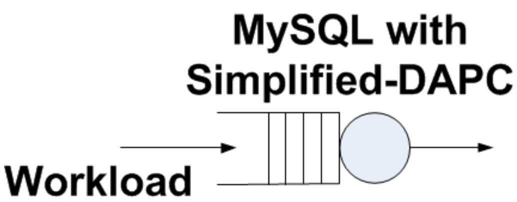

Simulation configuration

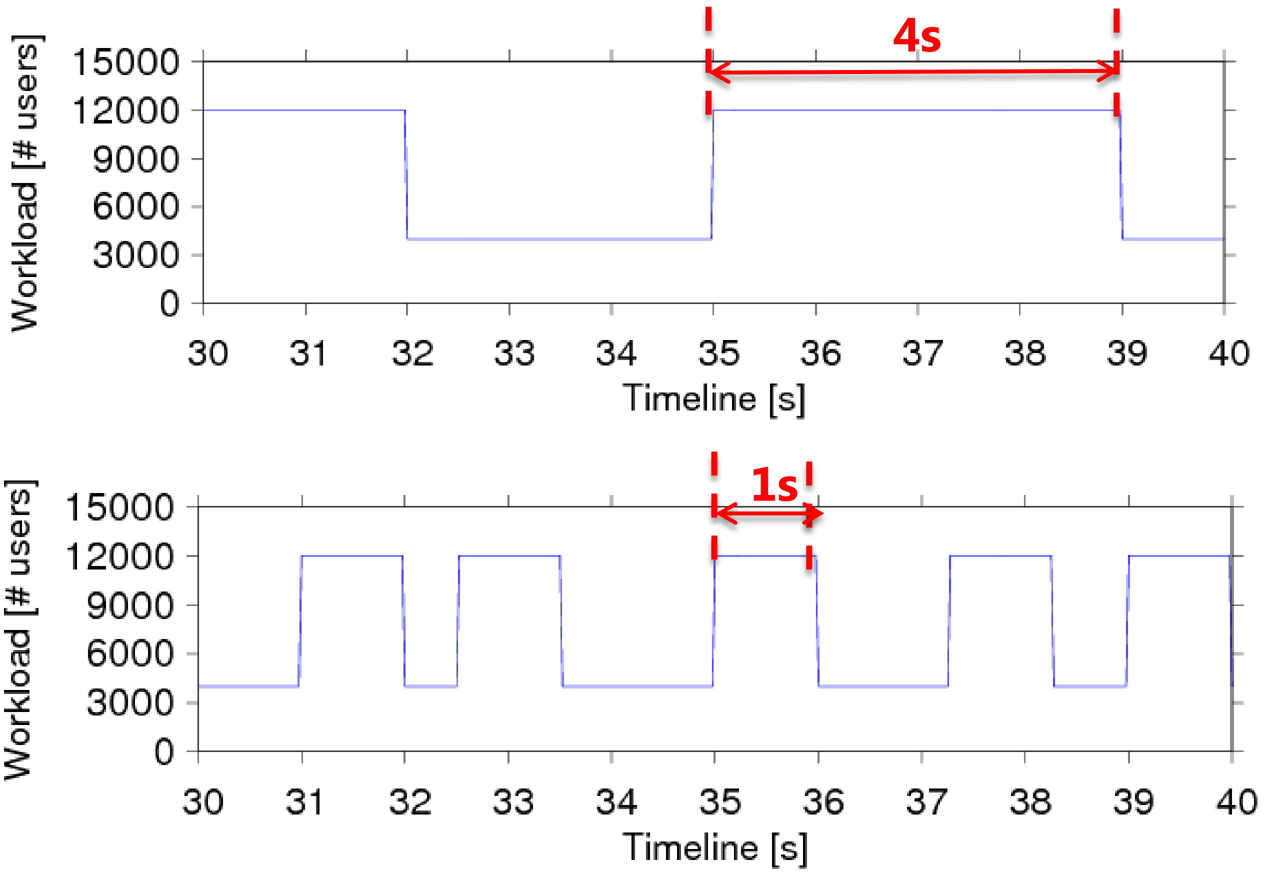

Workload Burst Length Sensitivity Analysis

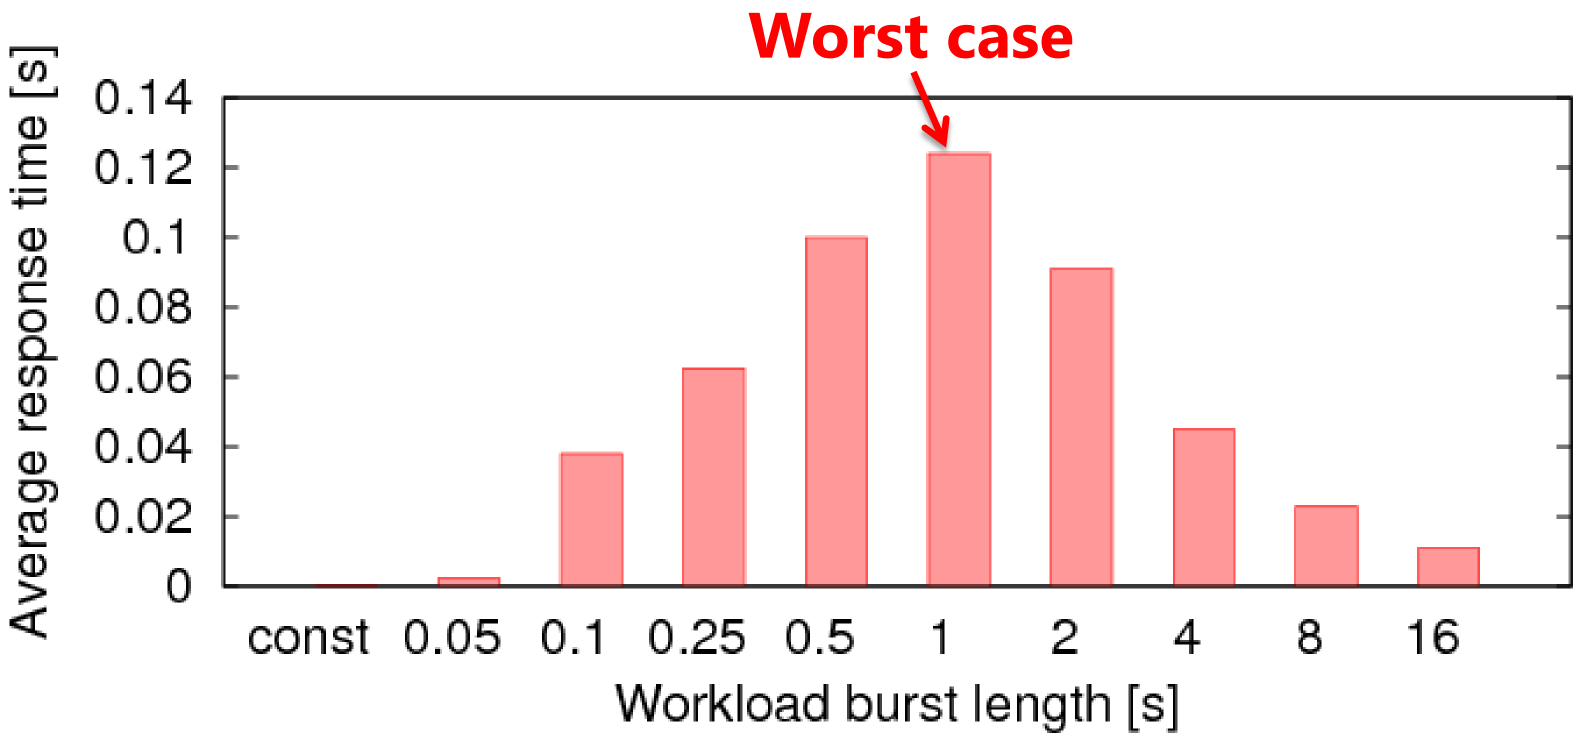

Georgia Tech College of Computing

Anti-Synchrony: Workload Burst Length vs. DAPC Adaptation Period

MySQL Workload burst length 1s

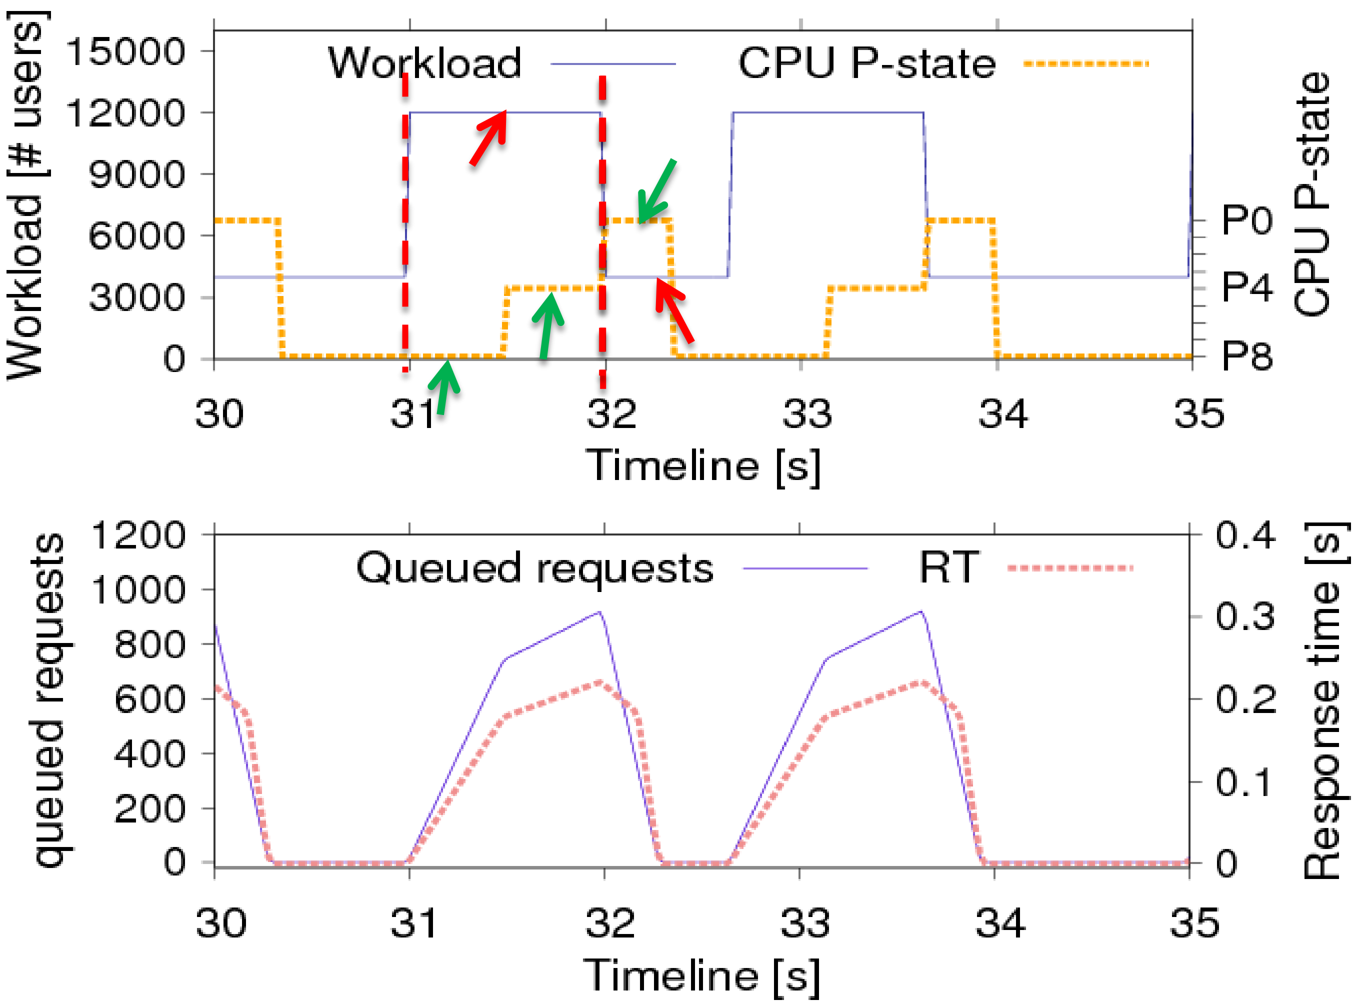

Georgia Tech College of Computing

Solutions

- Several candidate solutions
  - Assume constant workload or batch workload
    - Not for web-facing applications
  - Stay below the crossing point
    - Need measurements to find it
    - Significant research challenge
  - ***Short*** DAPC adjustment period
    - High overhead
    - Throughput loss may still exist with a fixed adjustment period

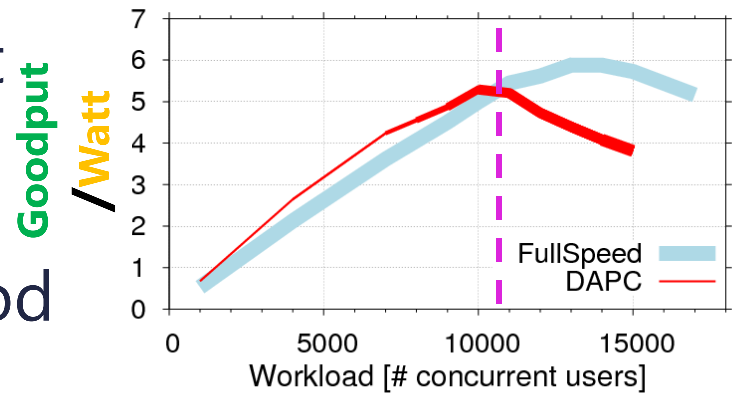

- Proposed solution: workload-sensitive adaptive control

Workload-Sensitive Adaptive Control

- Disrupt the anti-synchrony between workload burst length and DVFS adaptation period
  1. "Learn" the average interval between workload bursts
  2. Keep appropriate DVFS control adaptation period (e.g., 5 times smaller than the estimated workload burst interval)

How Does It Work

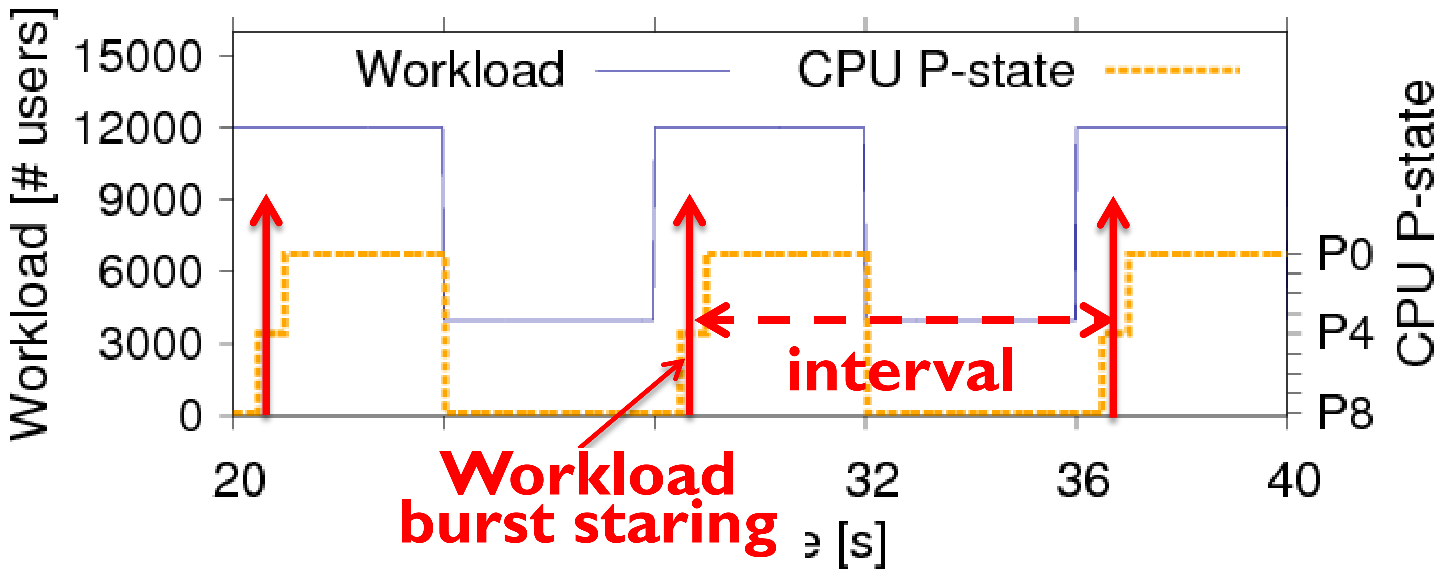

"Learn" the interval between workload bursts using observed workload bursts and a simple moving-average model

Adaptive Controller Reduces Queue Length

Simulation at WL 11,000 for MySQL

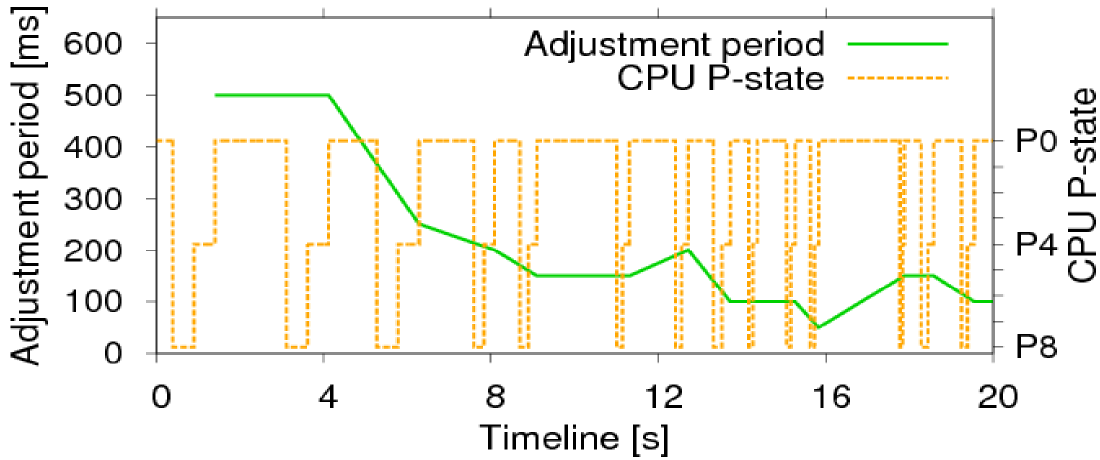

Original DAPC controller

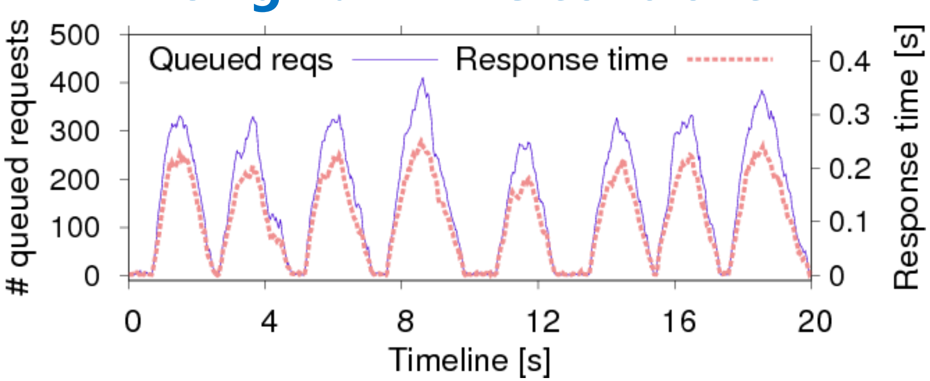

Adaptive controller

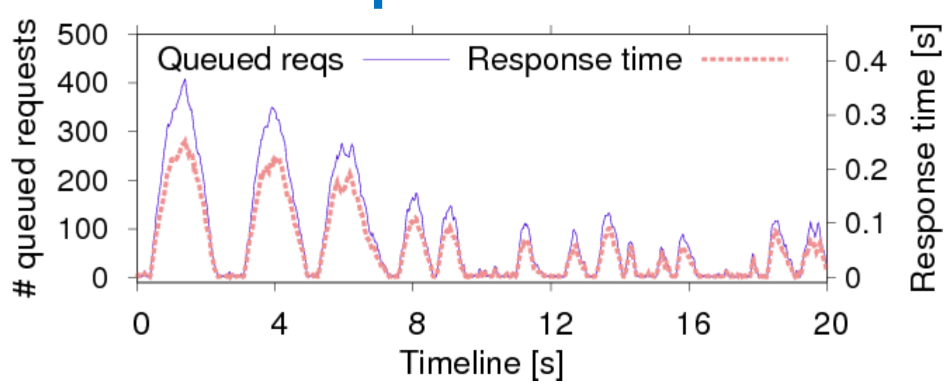

Adaptive Control Achieves QoS and Energy Savings

Power usage

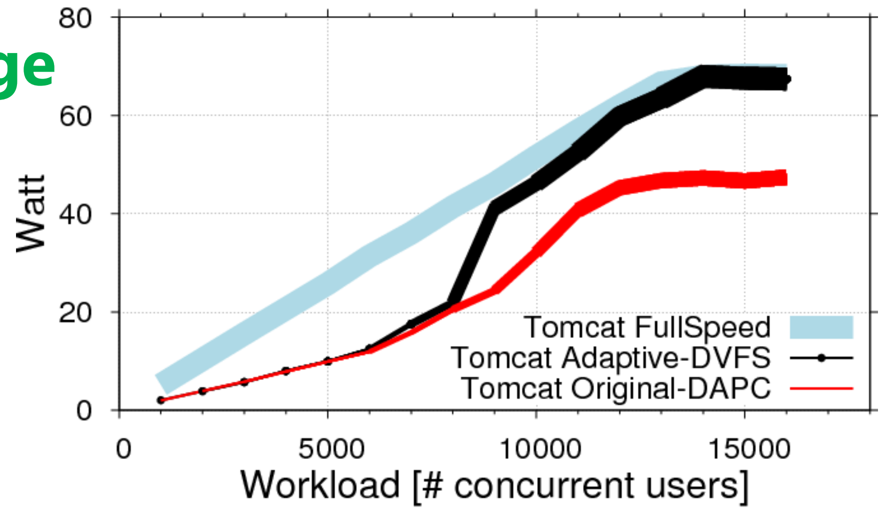

Throughput

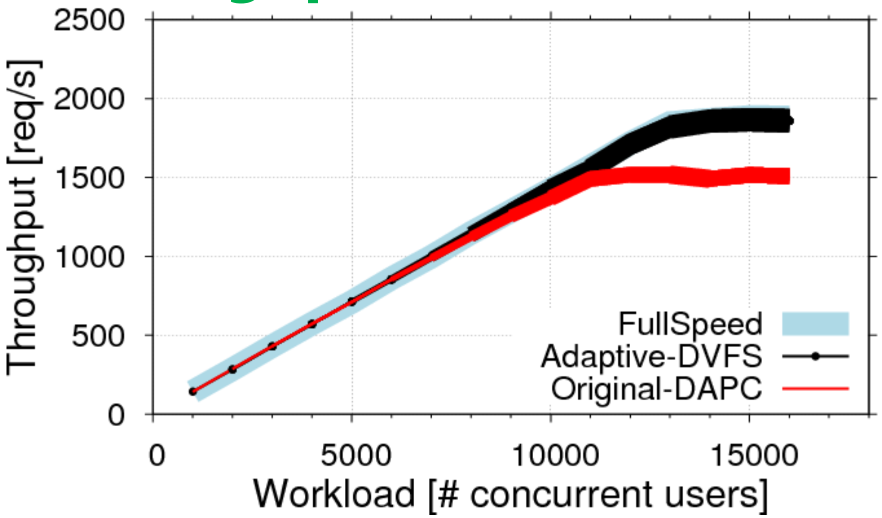

Response time

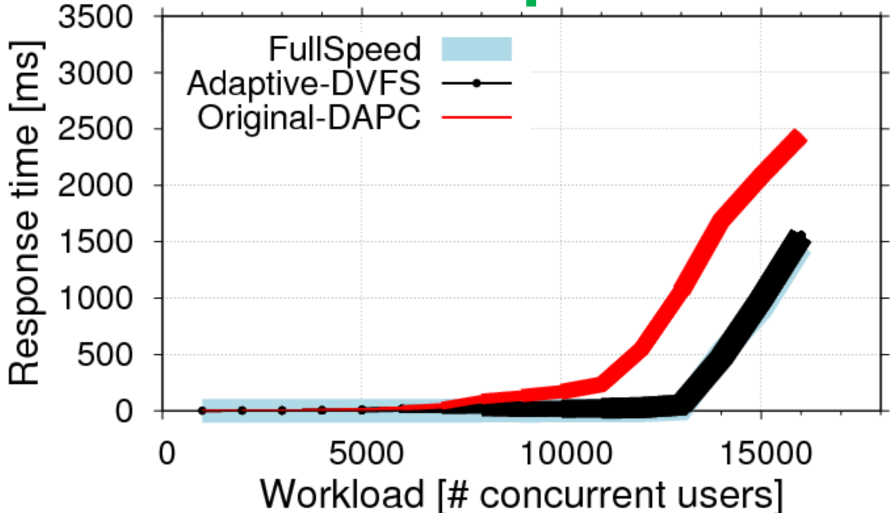

Conclusion

- Significant performance degradation of n-tier web applications at high utilization
  - Large response time fluctuations due to the push-back phenomenon
  - Throughput loss due to rapidly alternating bottlenecks
- The cause is the anti-synchrony between DVFS adaptation period and workload burst length
- Workload-sensitive adaptive control is able to mitigate performance impact while saving power

Thank You. Any Questions?

Qingyang Wang
qywang@cc.gatech.edu

Georgia Tech College of Computing

cercs

FUJITSU