IBM Research and University of Texas at Austin

Using on-line power modeling for server power capping

Madhu Saravana Sibi Govindan,
Charles Lefurgy, and Ajay Dholakia

Workshop on Energy-Efficient Design 2009

June 20, 2009

Power capping for servers

- **Control power consumption of server**
  - Constrained by power supply
  - Constrained by data center

- **Example: redundant power supply failure**
  - Redundant supply fails – load shifts to remaining supply
  - Power supply sees 125% load
  - Must reduce supply load from 125% to 100% in 1 second

Prior work and opportunities

- **Capping in literature (2006-2008)**
  - Use processor frequency to control system power
  - Use static, off-line power models
  - No direct measurement of settling time on real workloads

- **What happens when conditions dynamically change?**
  - Is an off-line power model enough?
- **Can power capping controller learn to adapt?**
  - Is on-line power model as good as off-line model?
  - Goal: reduce development time searching for good power model

Contributions

- **Develop on-line power models**
  - Continuously measure behavior of system at run-time
  - Enable self-tuning controller
- **Self-tuning controller performance is acceptable**
  - Directly measure settling time of power controllers
- **No hand-tuning → saves development time**
  - Adapt to different server configurations

Measure power

- **Entire power of blade is measured**
- **Every 64 ms (1 control period)**

IBM HS21 blade server

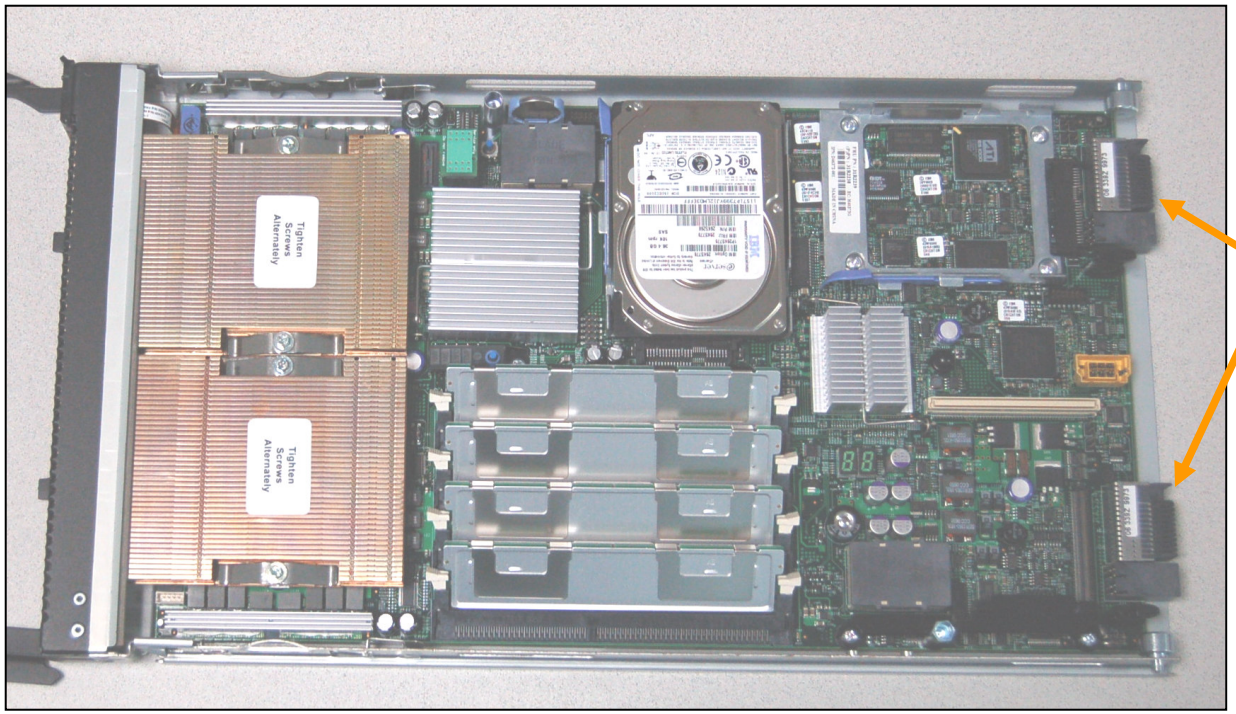

Measure 12V bulk power

0.1 W precision, 2% error

Measure settling time

- **Firmware tracks each overshoot of power cap**
- **Record maximum settling time**

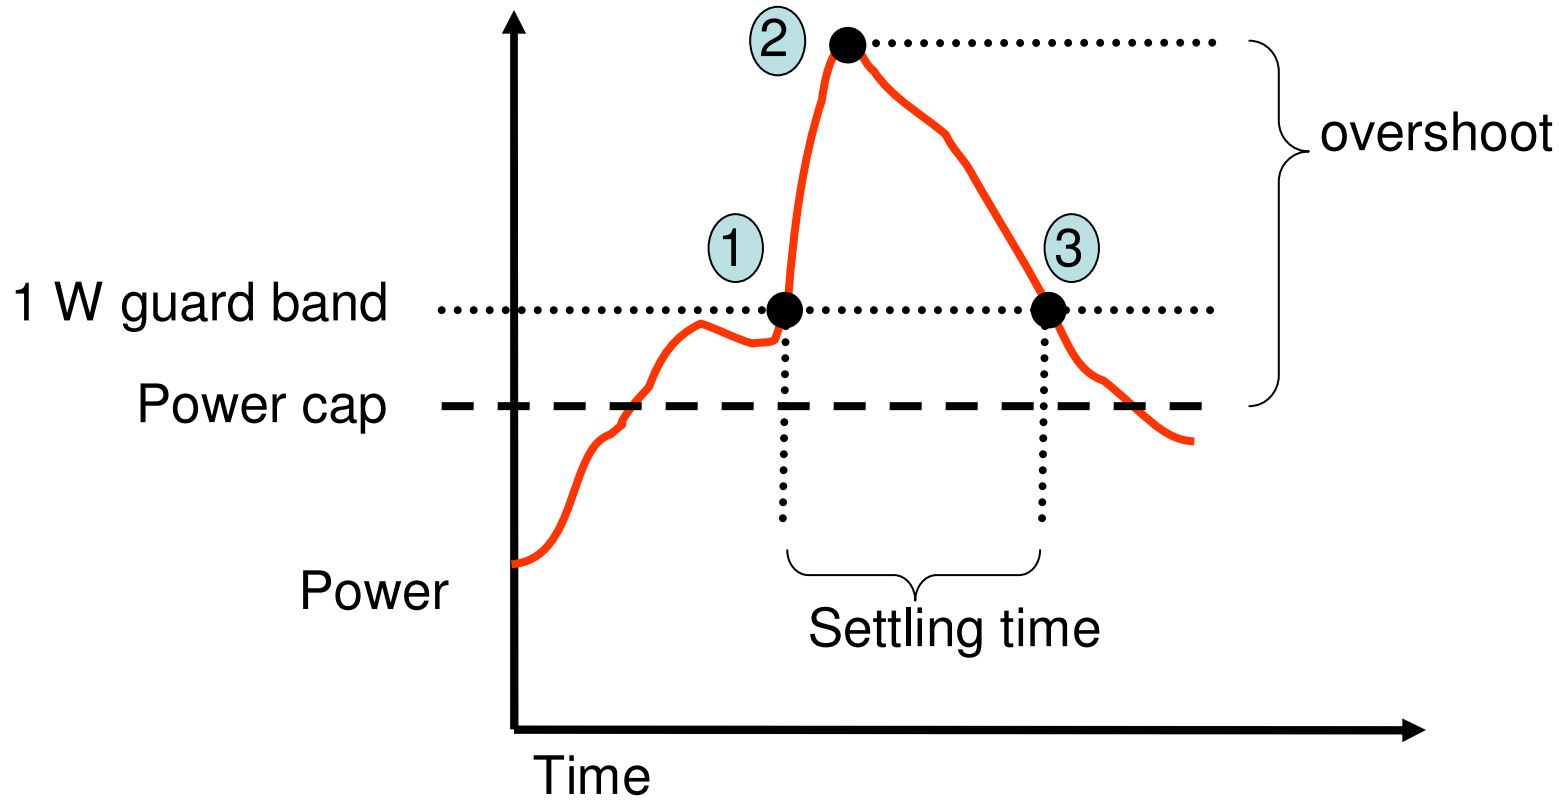

HS21 blade (2 - 5160 x86, 4 GB, 1 disk)

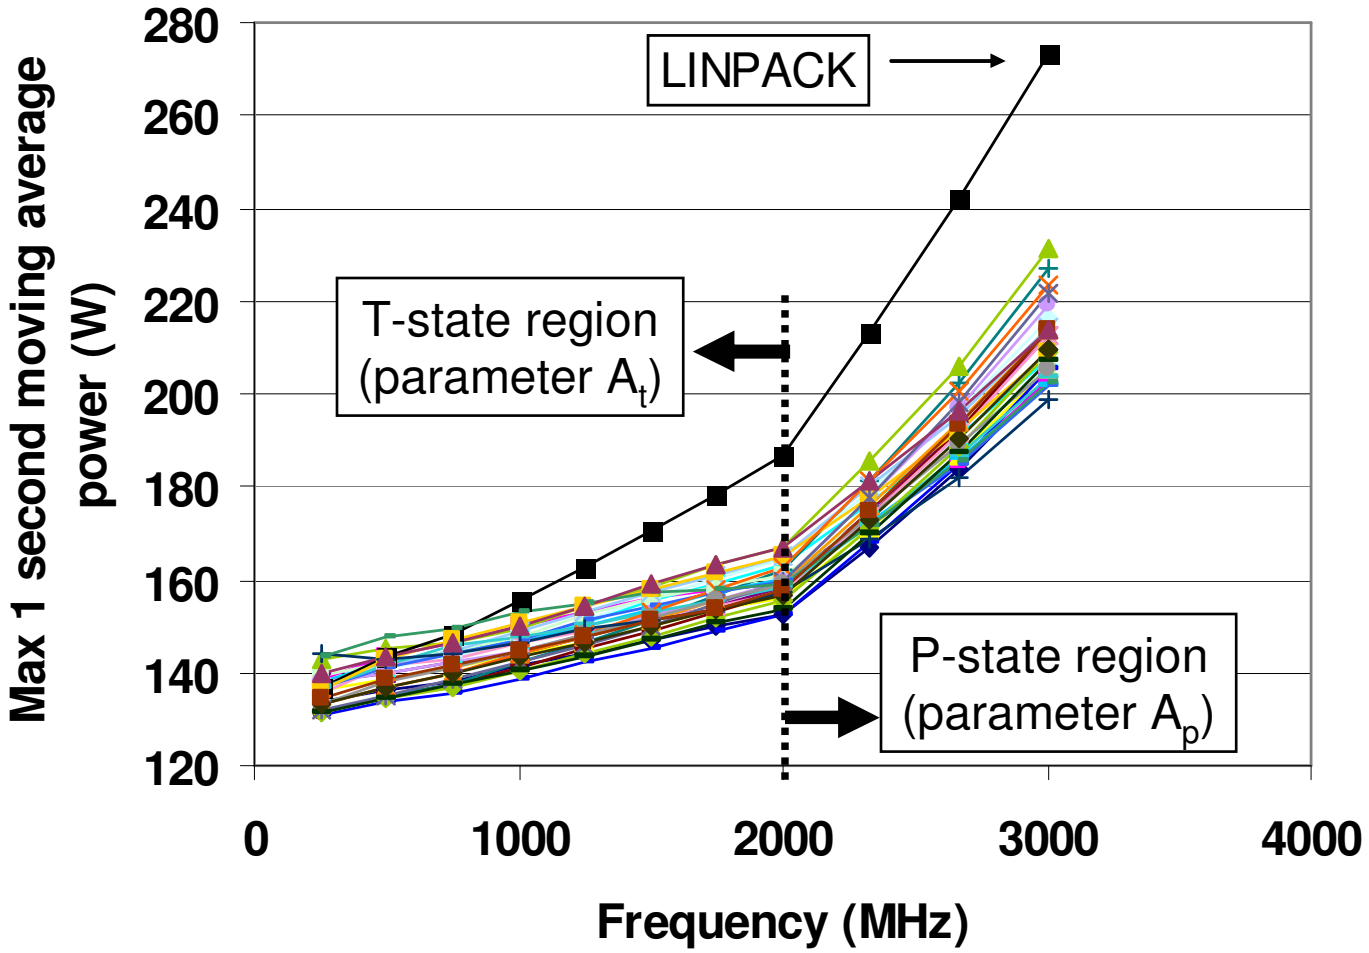

Two types of controllers

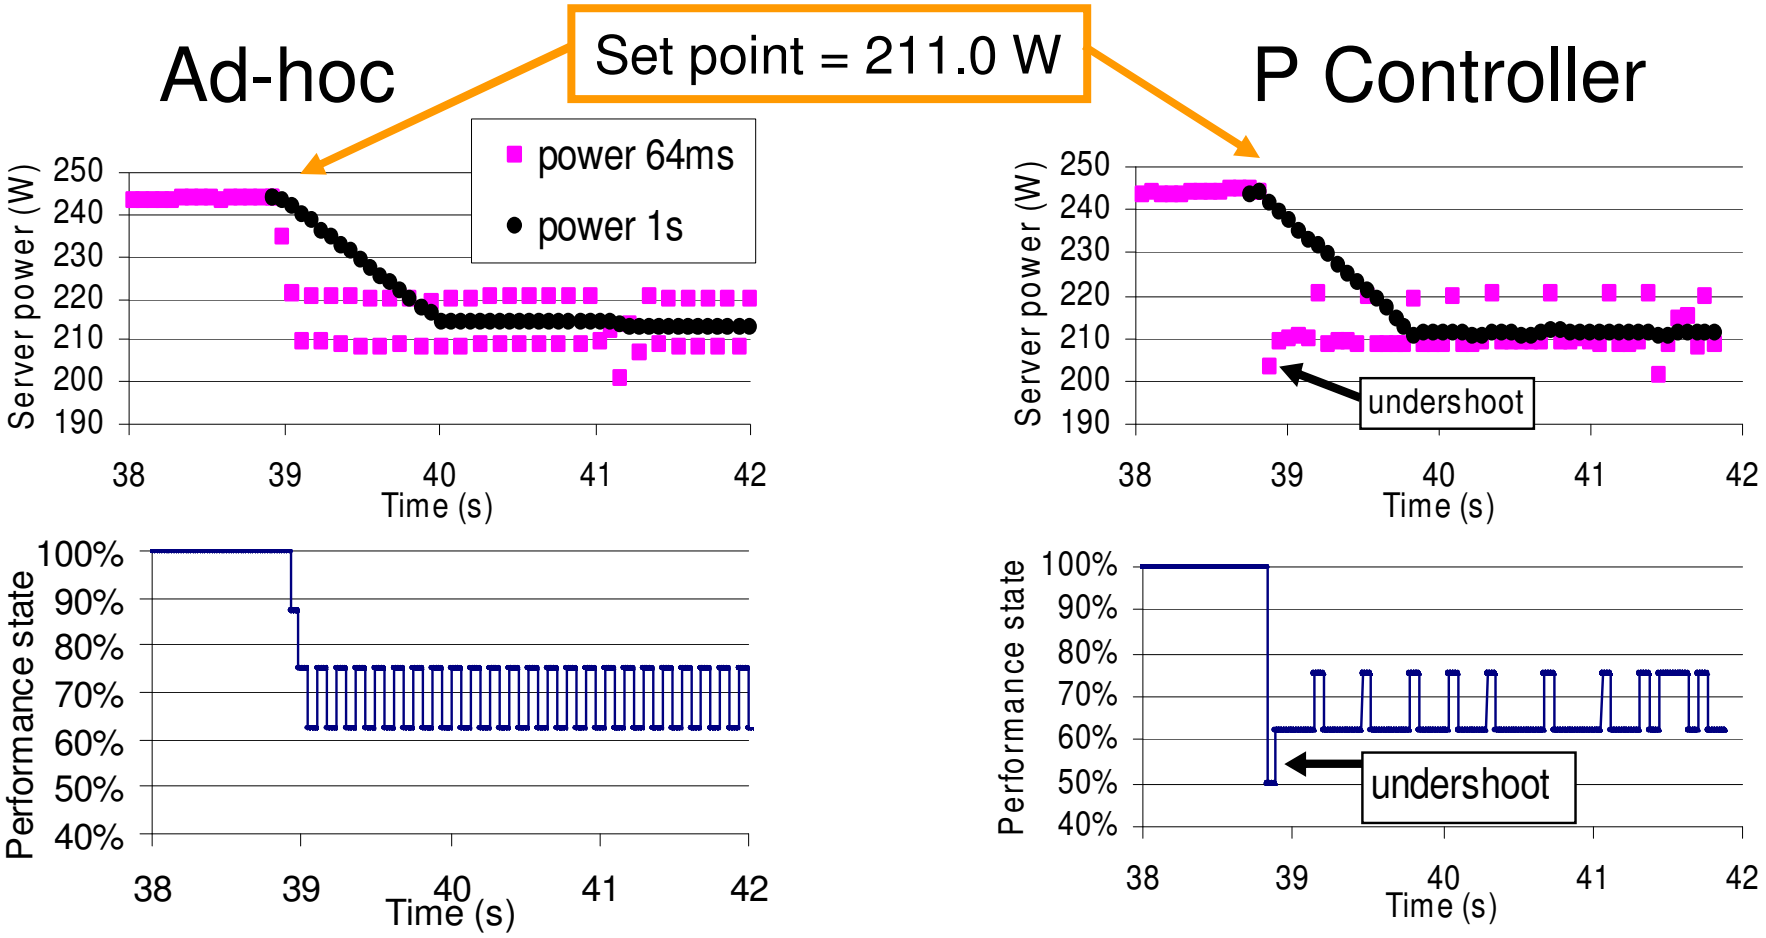

Settles to 216.0 W 5 W Violation
CPU speed: 68.8%

Settles to 211.0 W No violation
CPU speed: 65.8%

How proportional capping works

1. Measure Δ power = ($P_{cap} - P_{current}$)
2. Use power model to find Δ frequency
   - A = expected slope
   - $F_{new} = F_{current} + \Delta P / A$

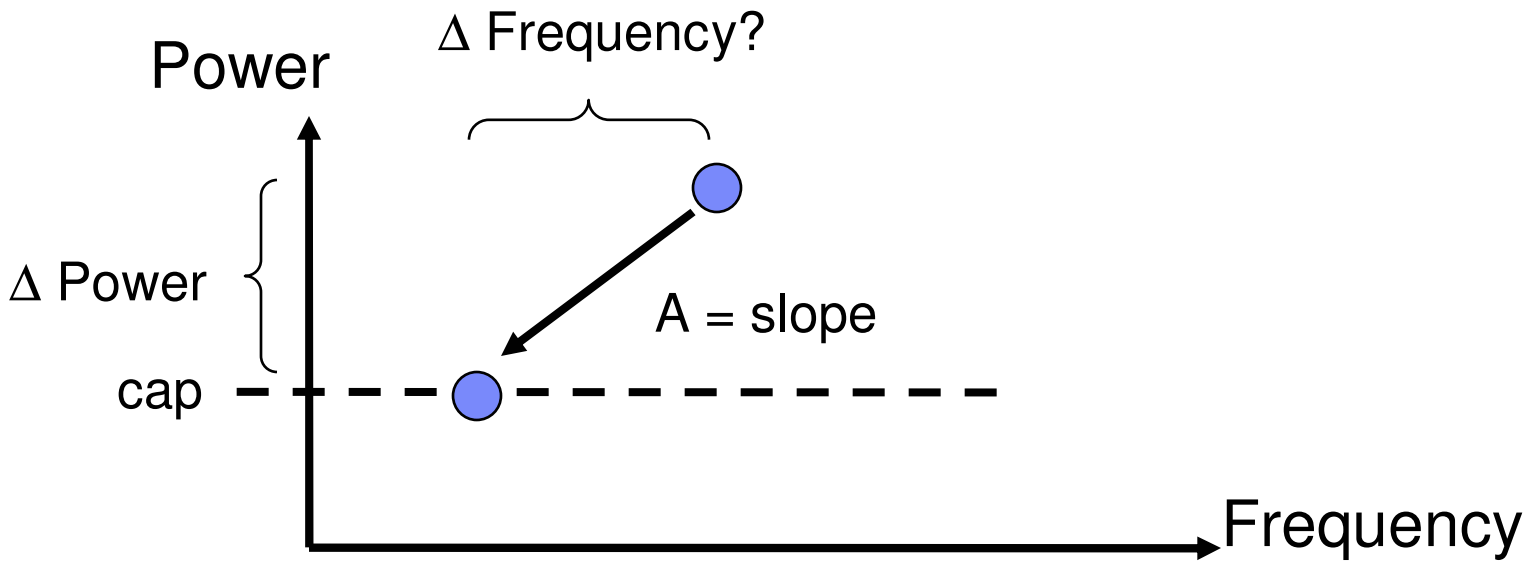

Using two actuators

- **When crossing actuator domains:**
  - Estimate power at crossover frequency (2 GHz)
  - Estimate frequency change from crossover frequency

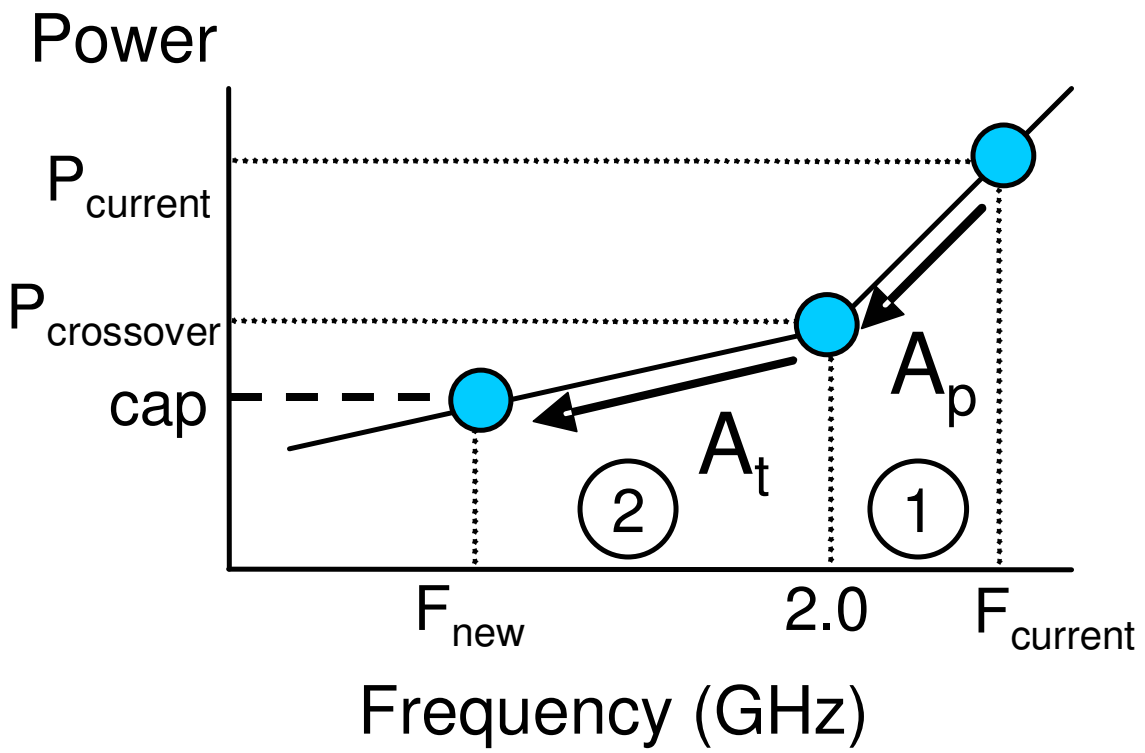

Affect of control parameter A_p on settling time

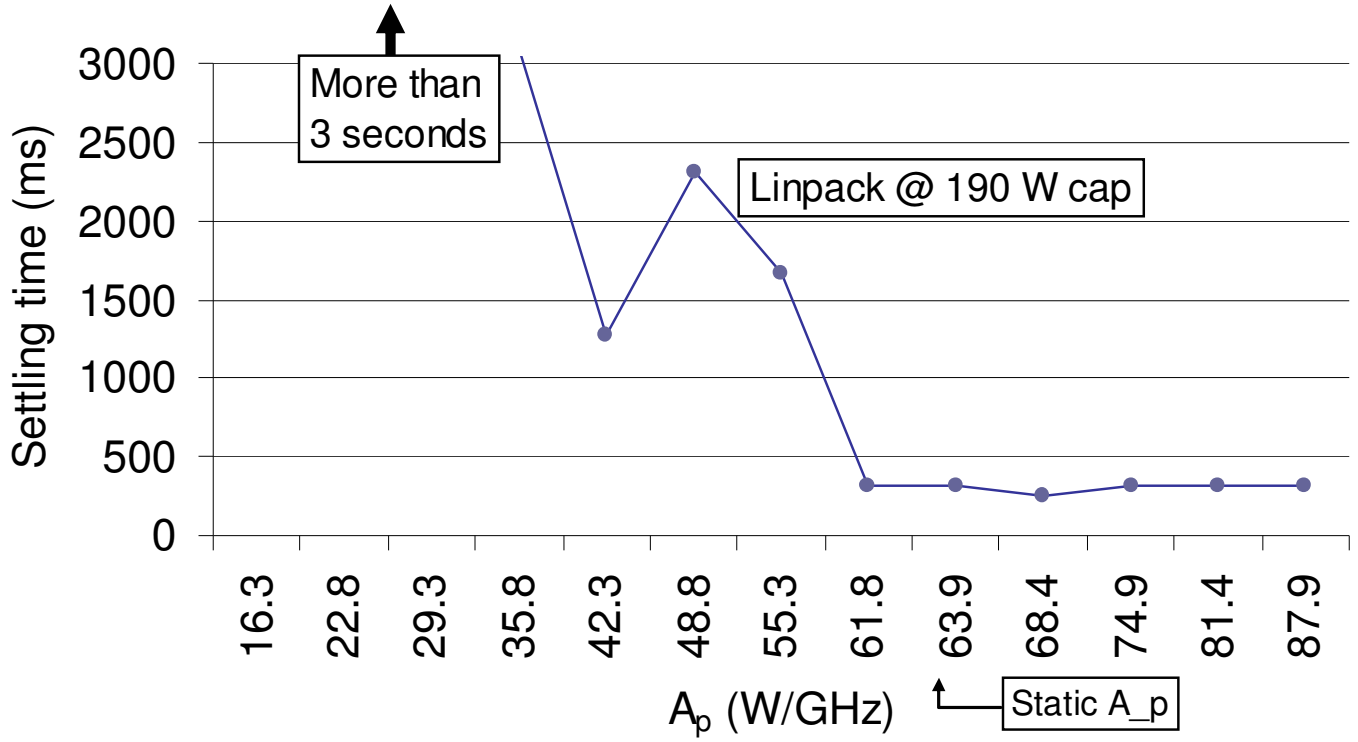

Affect of control parameter A_p on settling time

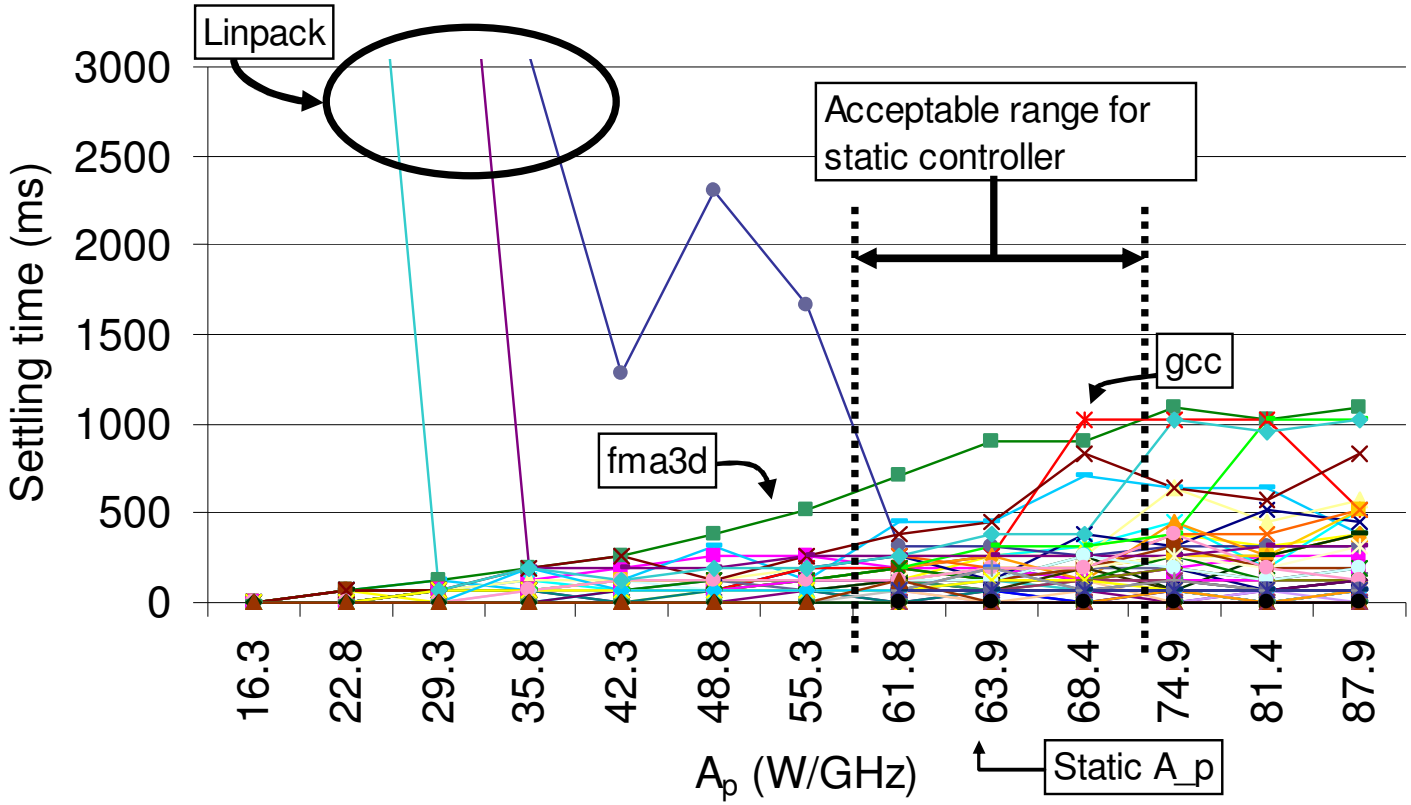

Our capping controller

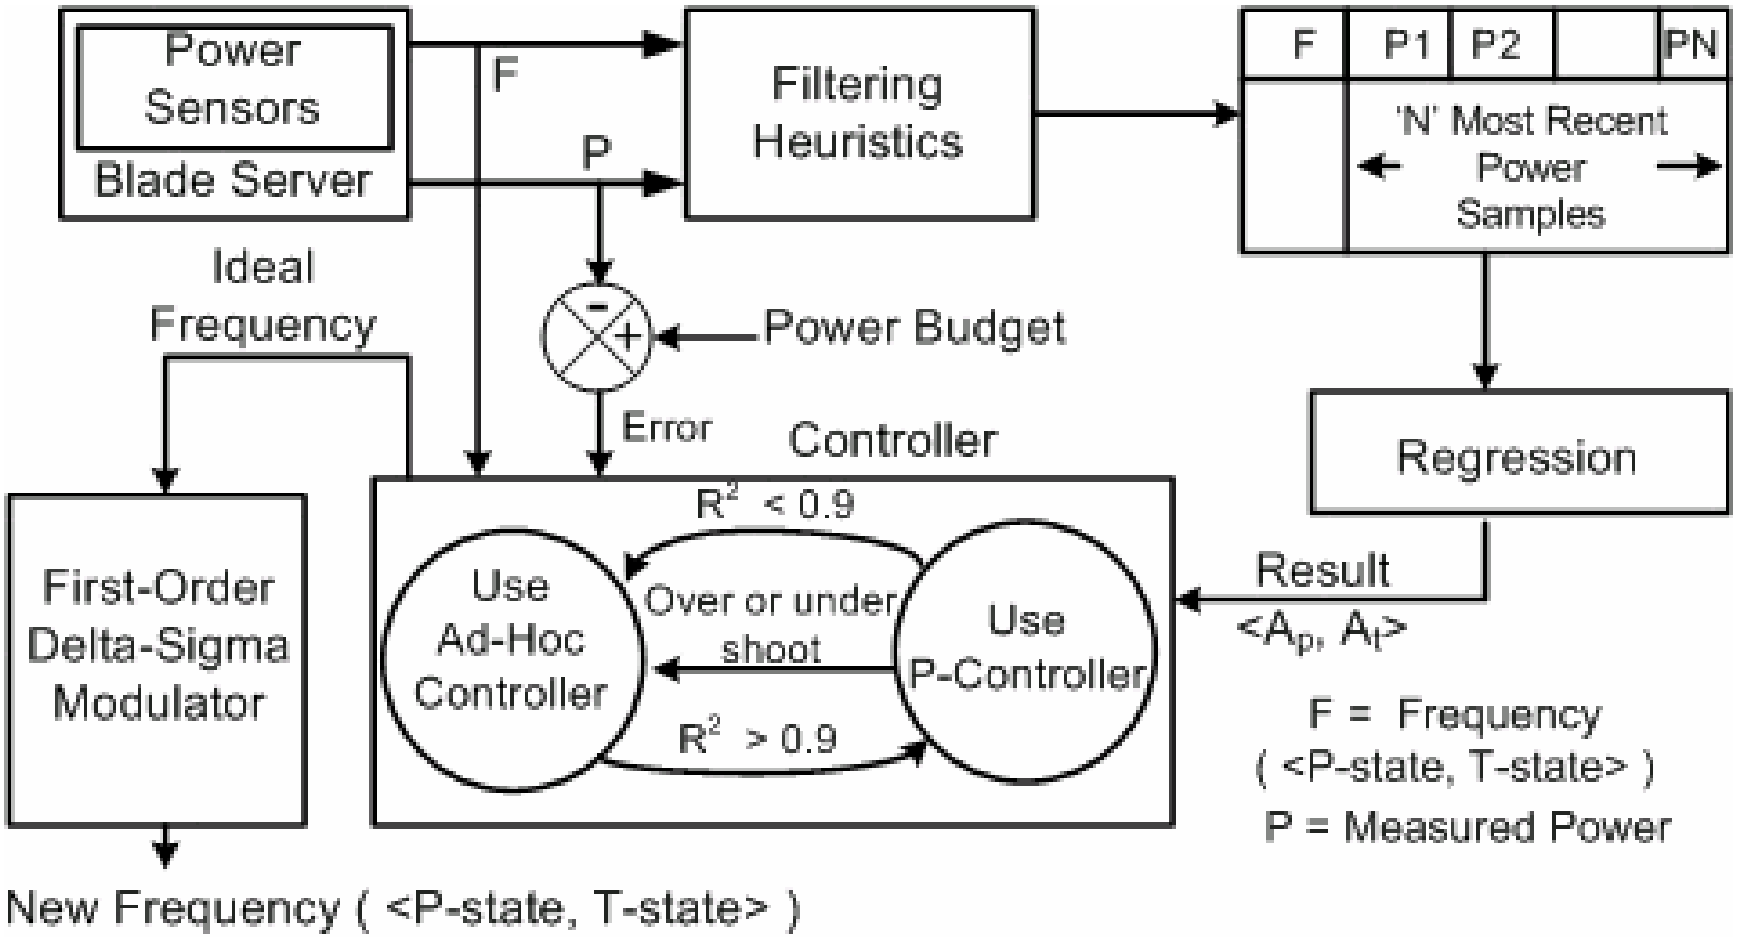

Results: settling time

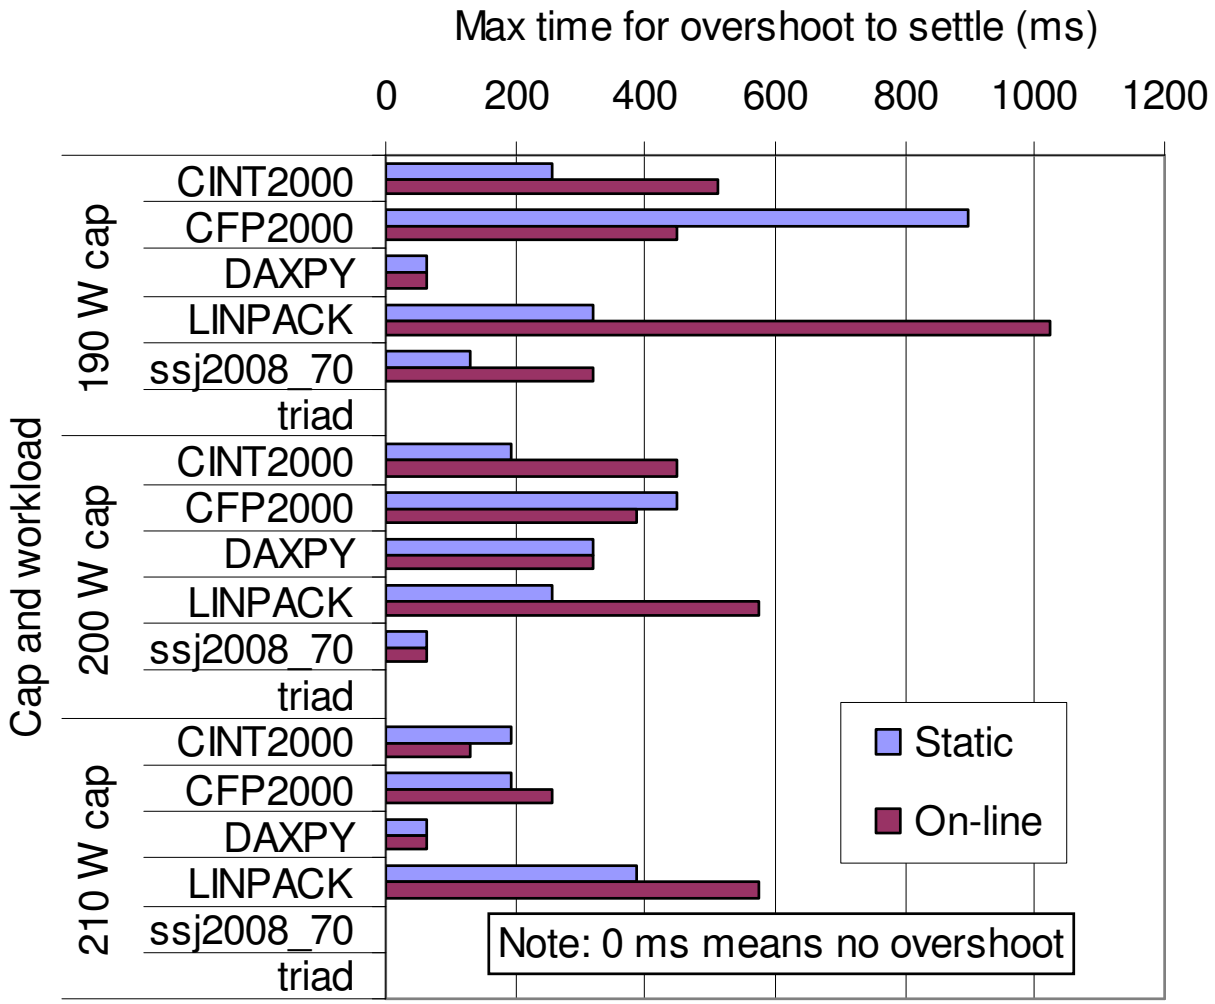

Results: workload performance

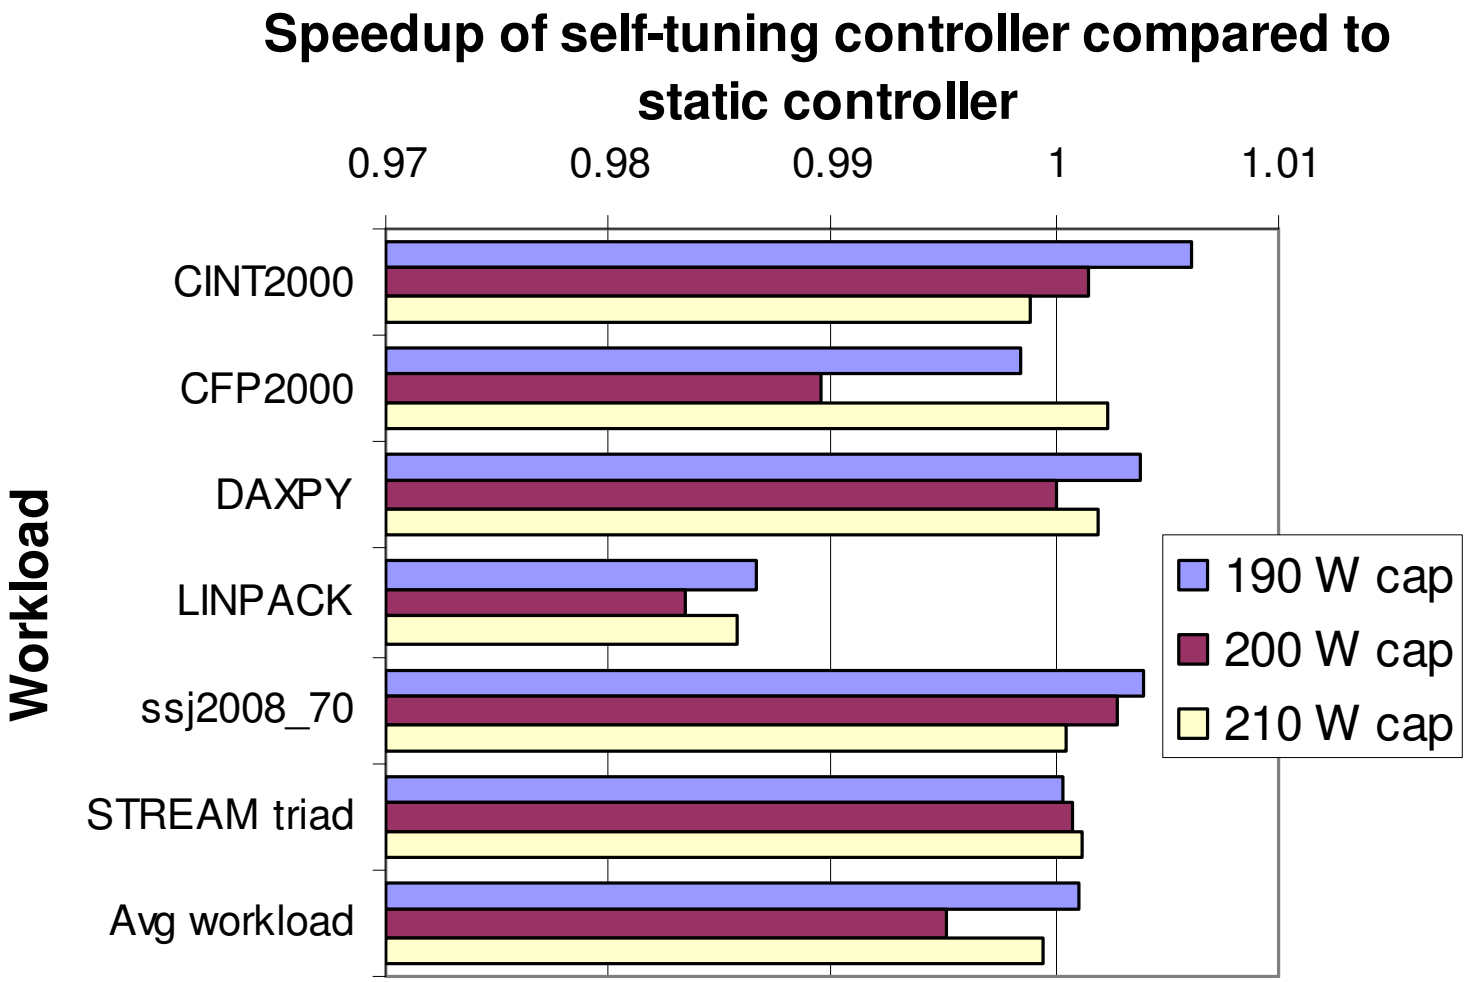

Conclusions

- **On-line power model achieved results close to best off-line power model**
- **More work required for Linpack**
  - Oscillating workloads can be difficult
- **Ad-hoc control**
  - Useful when power model is unknown
- **Future capping studies**
  - Would like to see more measurement of settling time and overshoot
  - Improve direct comparisons between capping methods