Power-Aware Throughput Control for Database Management Systems

Zichen Xu, Xiaorui Wang, Yi-Cheng Tu*
The Ohio State University
* The University of South Florida

Problem Overview and Motivation

- Data centers are energy starving
- Database management system (DBMS) is one of the major services in data centers
- Control the DBMS throughput to save energy

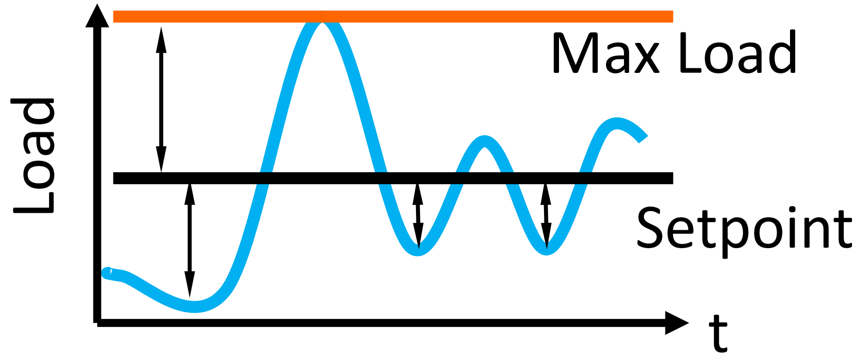

  - Power saving to the setpoint: low power hardware states
  - Balanced performance: performance control

Power-Aware Throughput Control (PAT)

- Goal: minimize power consumption of a database system while maintaining its desired performance
- Challenges
  - Energy and DBMS: power consumption minimization
  - Control and DBMS: performance guarantee

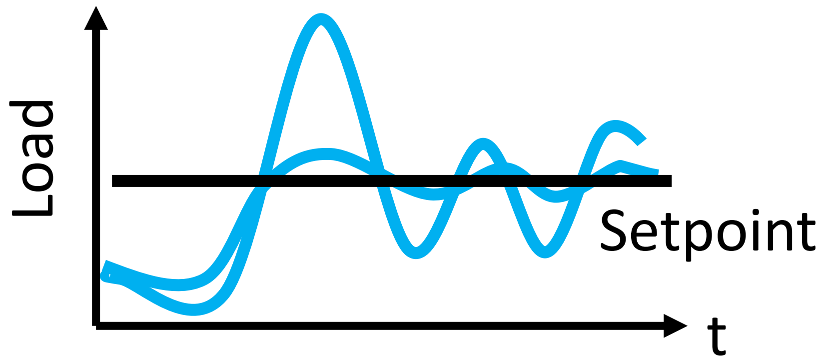

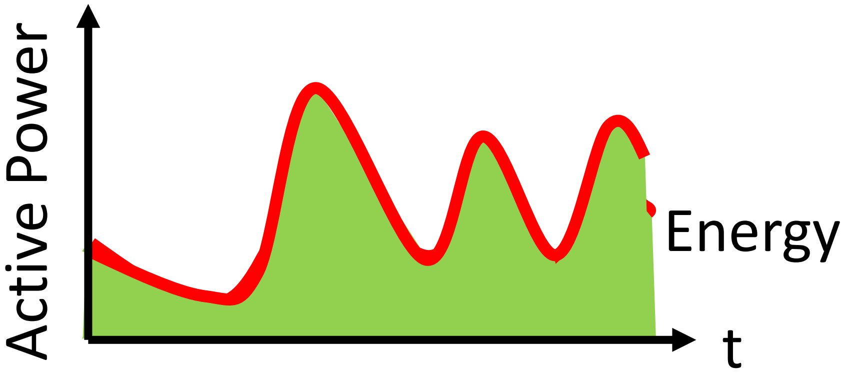

- Contributions:
  - Energy profiling in DBMS
  - Feedback control design in DBMS

Power-Aware Throughput Control (PAT)

- Goal: minimize power consumption of a database system while maintaining its desired performance
- Challenges
  - Energy and DBMS: power consumption minimization
  - Control and DBMS: performance guarantee

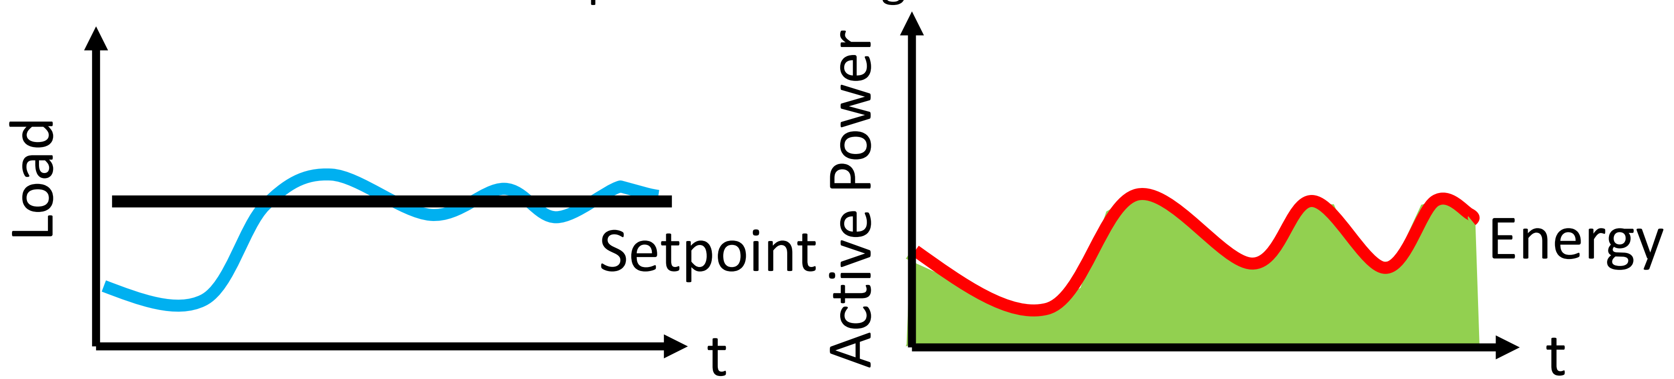

- Contributions:
  - Energy profiling in DBMS
  - Feedback control design in DBMS

Characterization Study – Memory

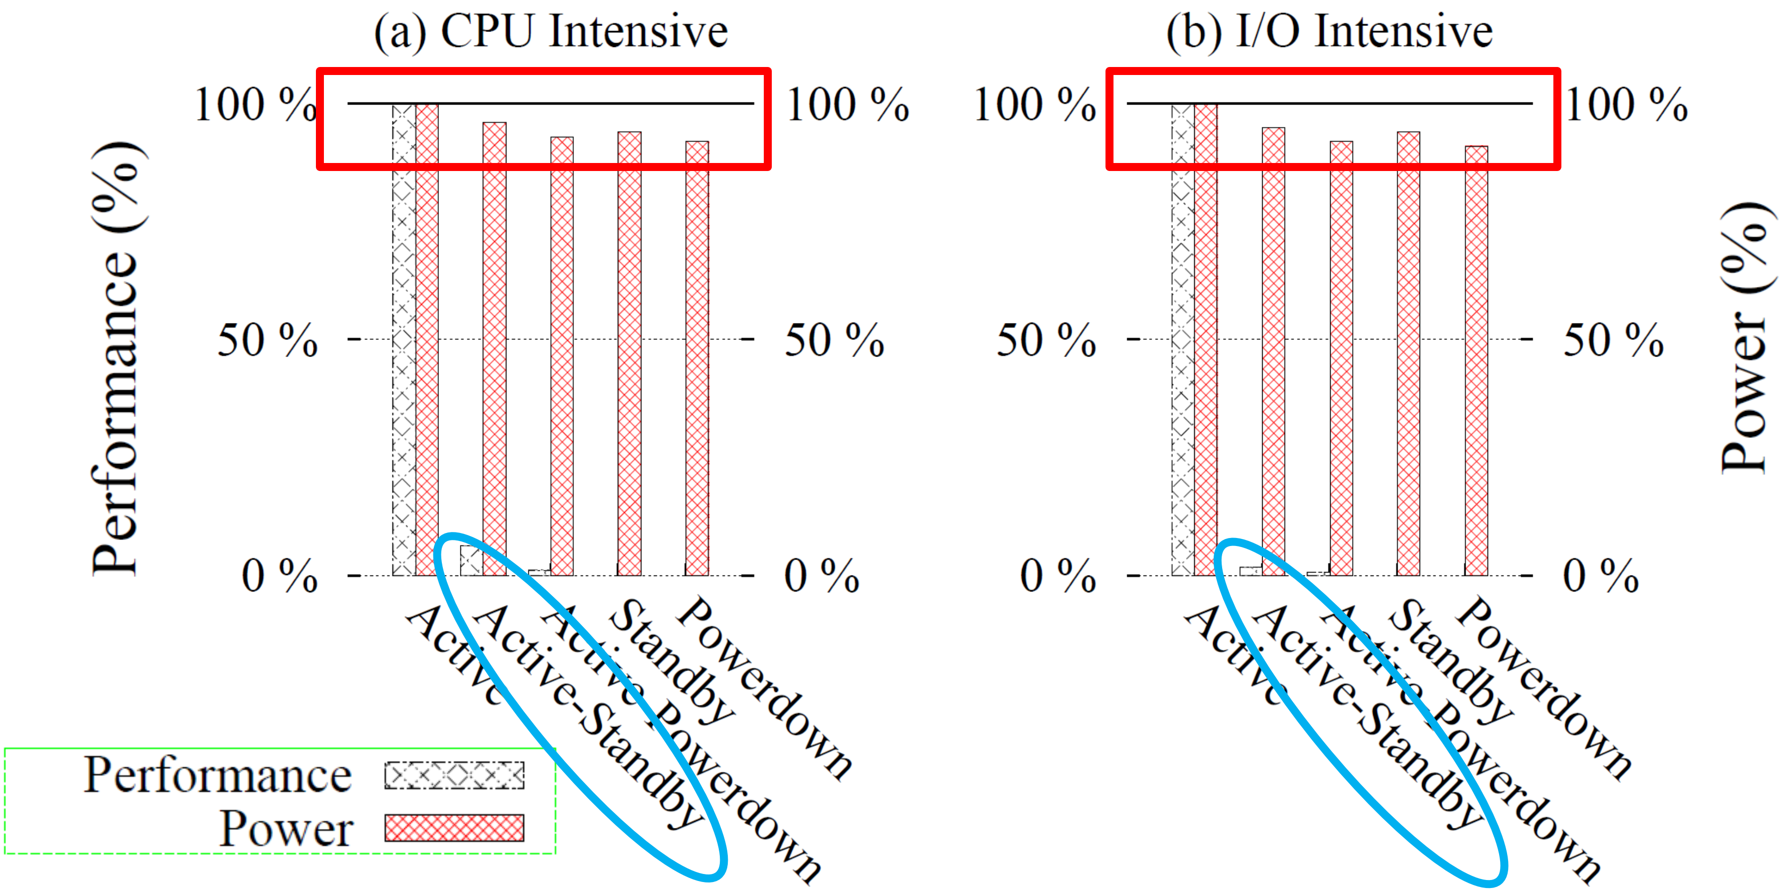

- Low power state in memory may not be a good choice for power saving in DBMS workloads

Charac

Less than 20% power savings

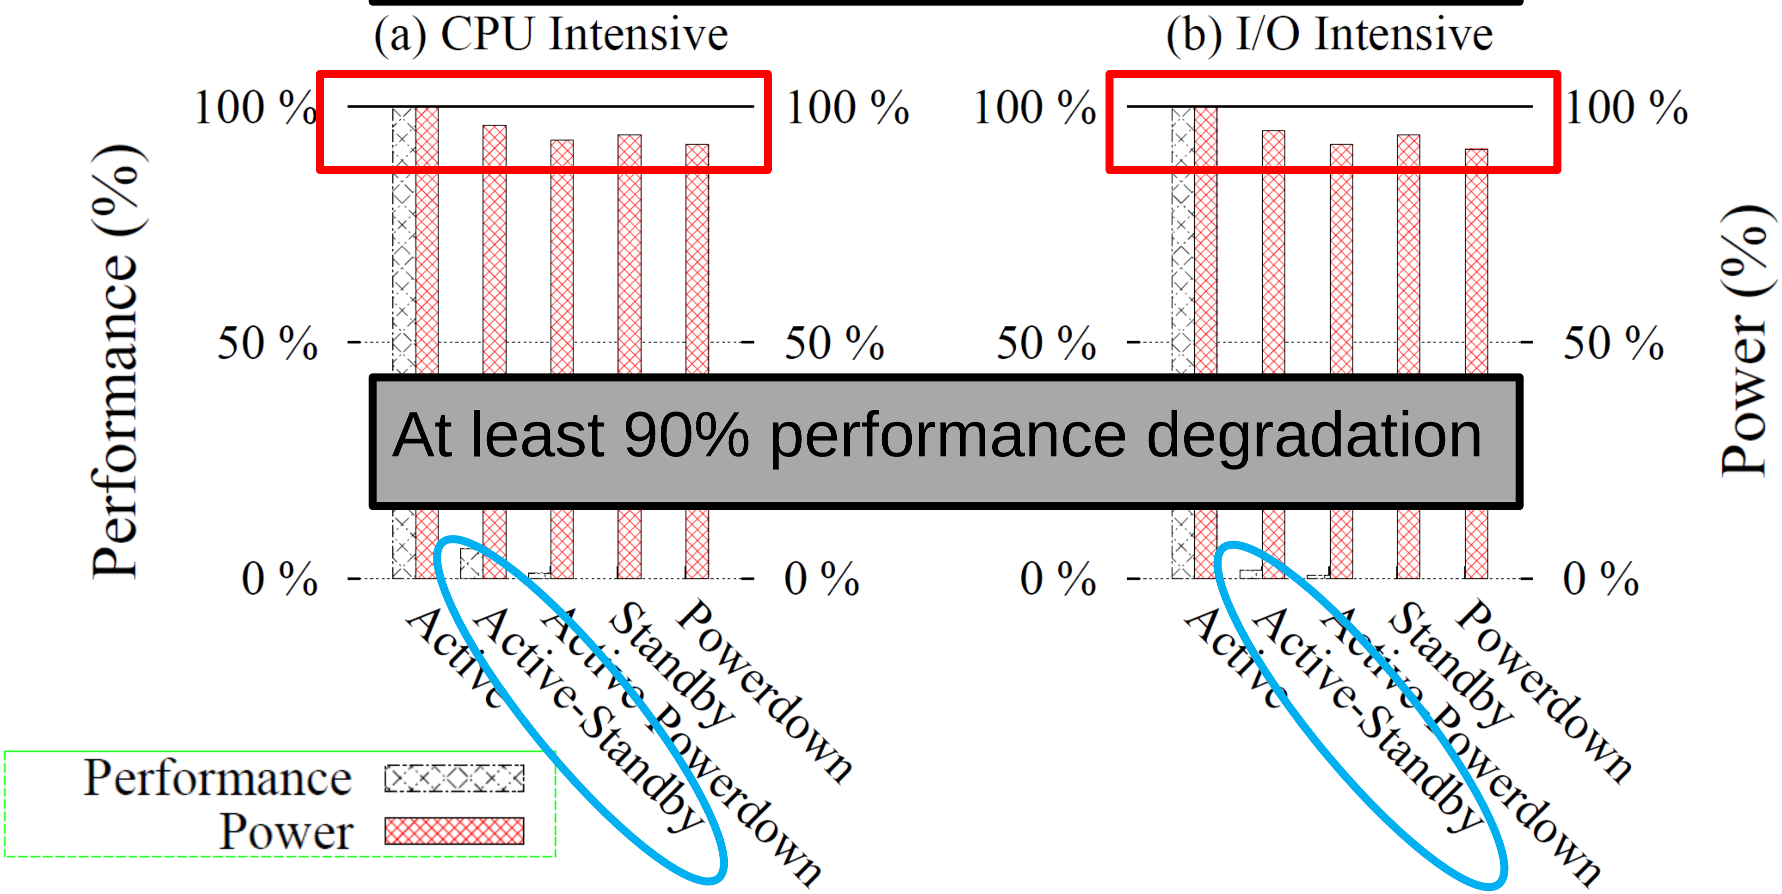

At least 90% performance degradation

- Low power state in memory may not be a good choice for power saving in DBMS workloads

Characterization Study – CPU

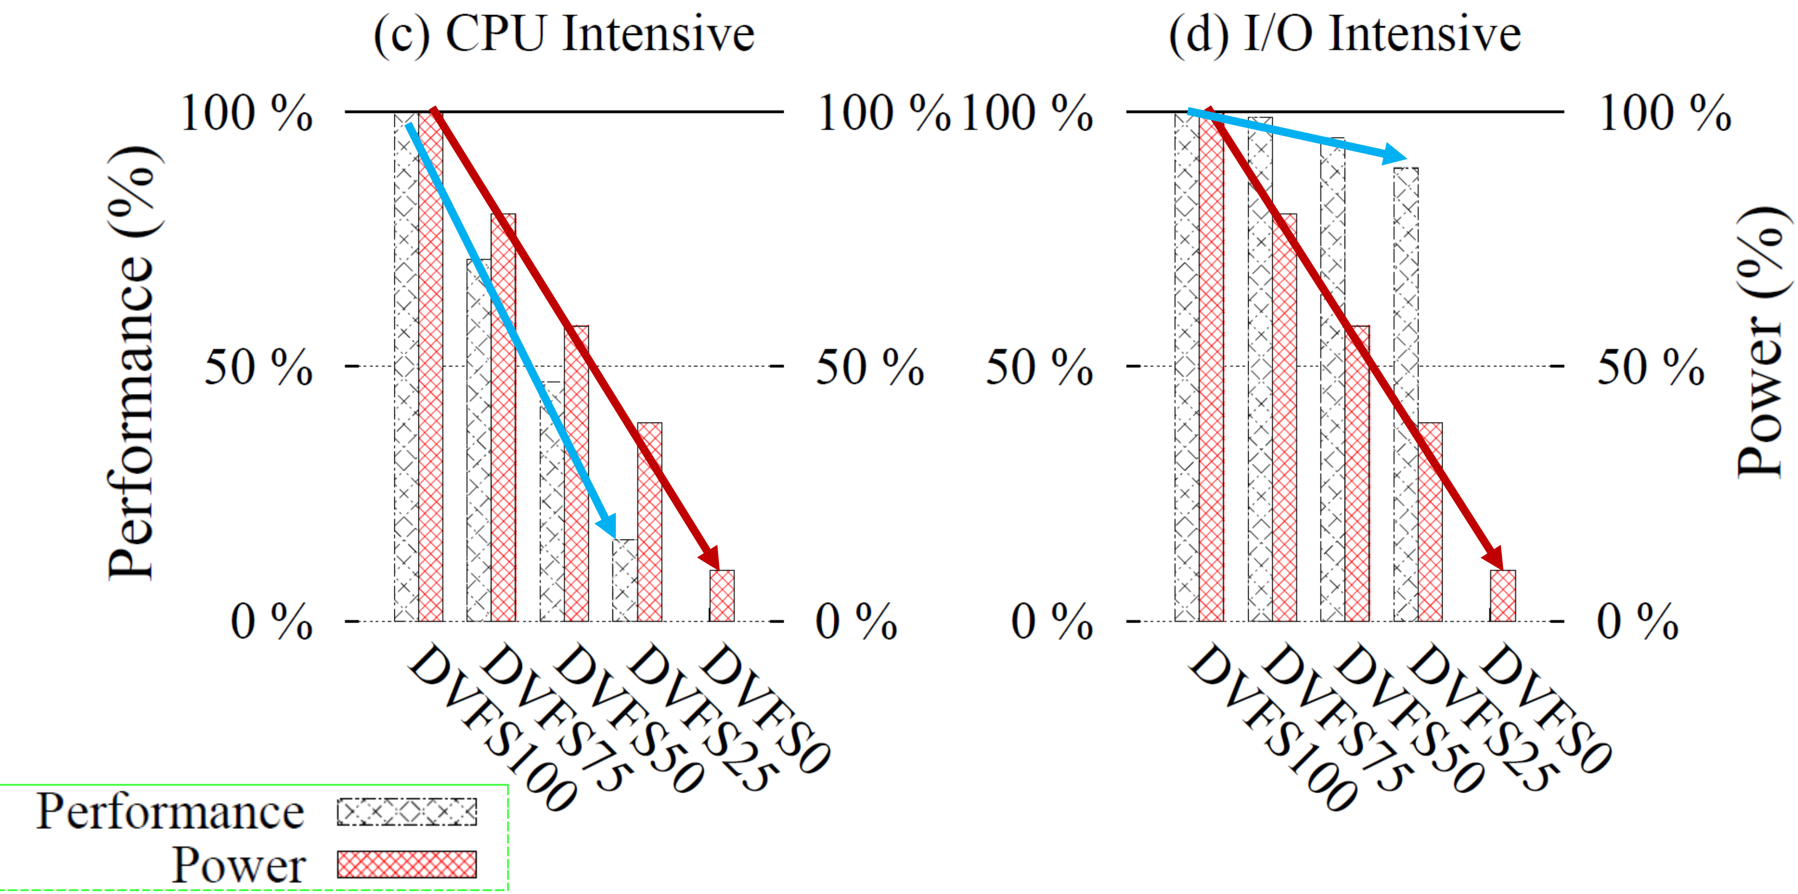

- CPU provides great power-saving and the relationship can be approximated as a linear model

Characterization Study – DBMS Workloads

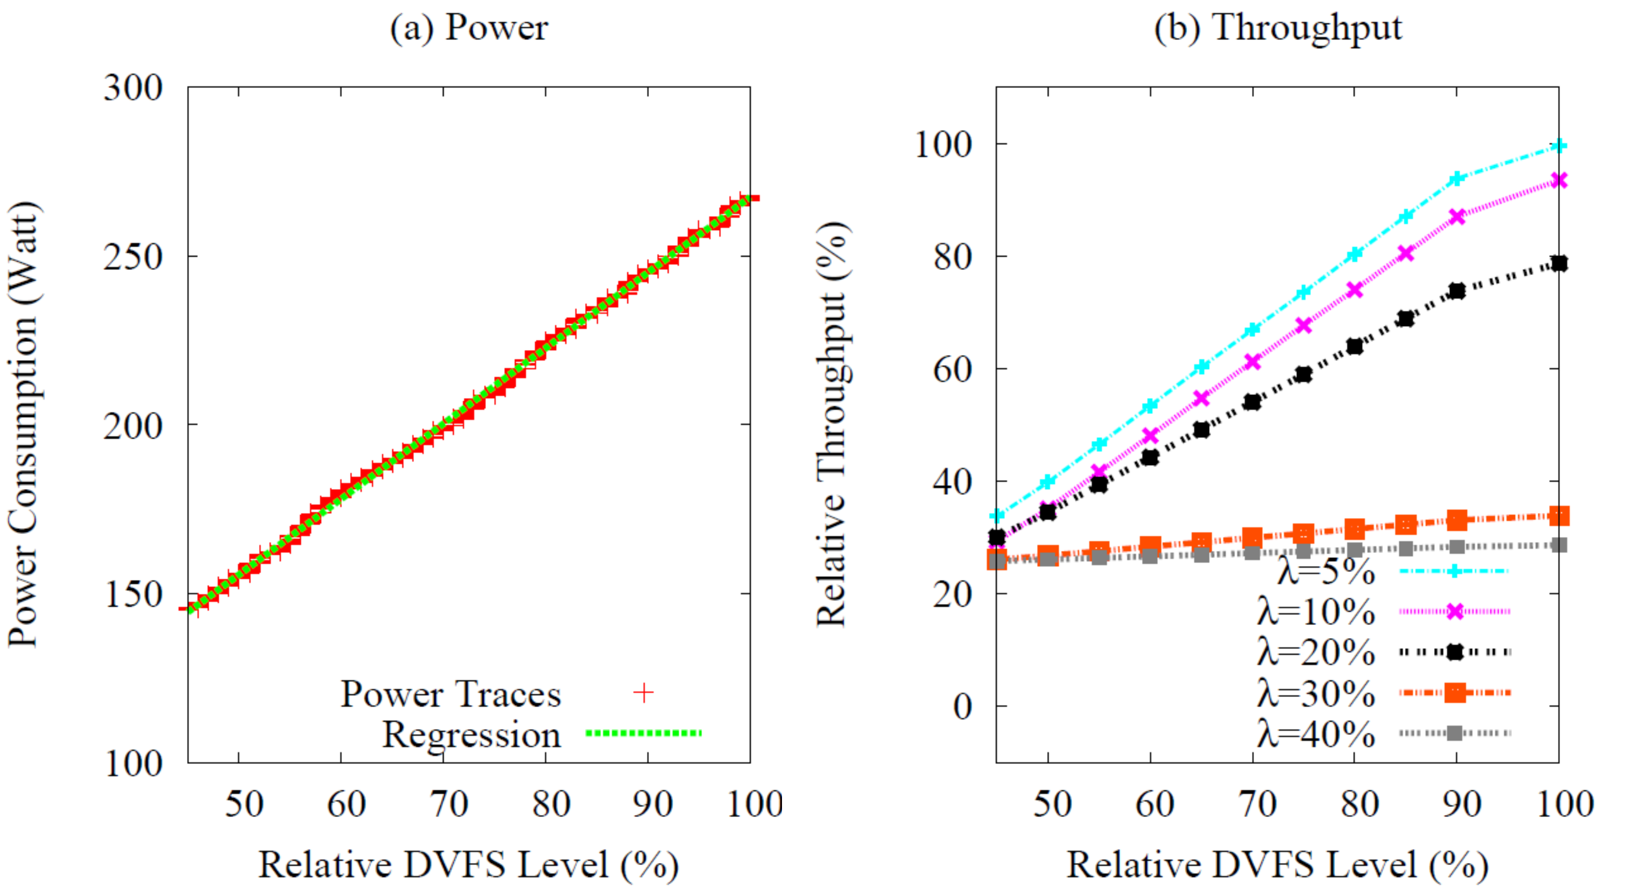

- The ratio of I/O intensive queries (**λ**) is the key in the workload statistics for power control

Characterization Study – DBMS Workloads

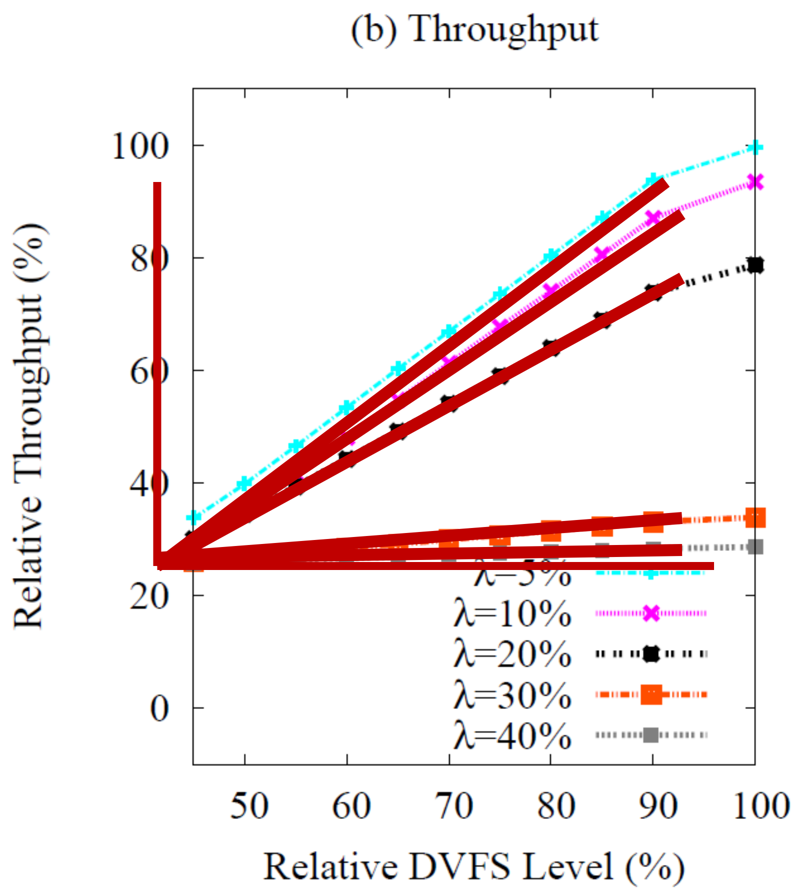

- The ratio of I/O intensive queries (**λ**) is the key in the workload statistics for power control

Characterization Study – DBMS Workloads

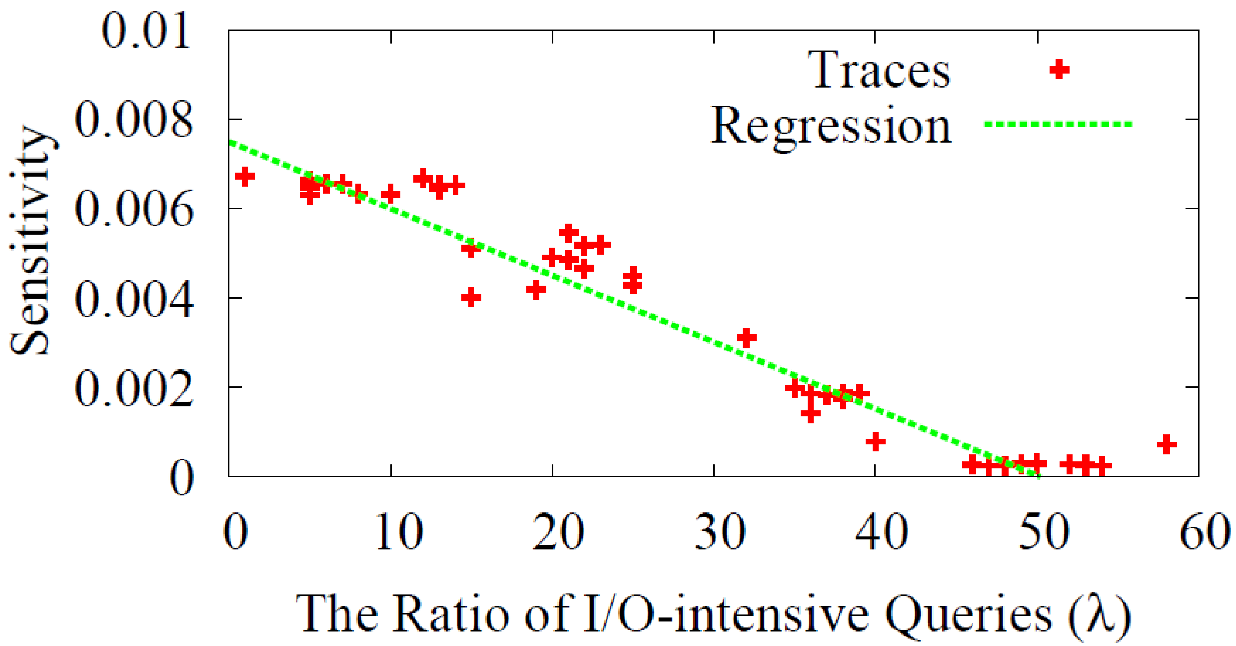

- The ratio of I/O intensive queries (**λ**) is the key in the workload statistics for power control

Characterization Study – Insights

- Low power state in CPU provides great opportunities for power-saving
- DBMS throughput and CPU frequency can be approximated as a linear model
  - $r = A * f + B$
- The ratio of I/O intensive queries (**λ**) is the key in the workload statistics for power control
  - $r = \lambda A * f + B$

Framework

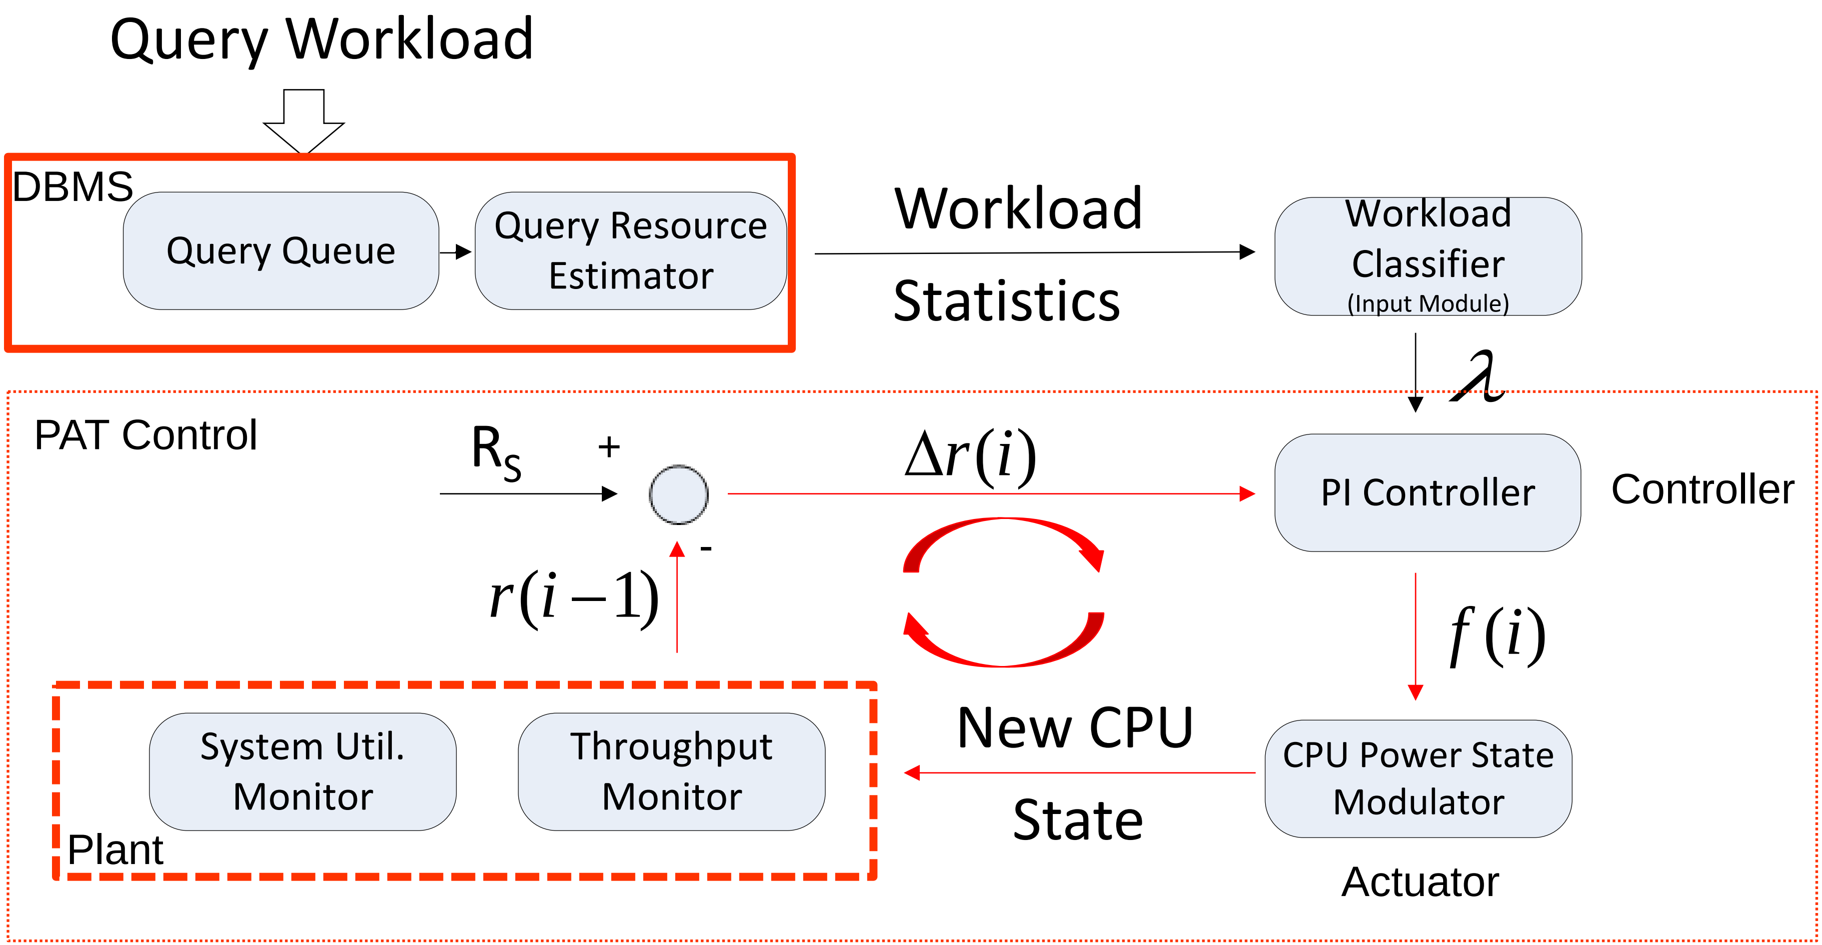
Query Workload
DBMS
Query Queue
Query Resource Estimator
Workload
Statistics
Workload Classifier
(Input Module)
λ
PAT Control
R_S
+
-
$\Delta r(i)$
PI Controller
Controller
$r(i-1)$
$f(i)$
System Util. Monitor
Throughput Monitor
Plant
New CPU
State
CPU Power State Modulator
Actuator

Fuzzy Workload Classifier

- One fuzzy rule base
  - Compute the membership of query to each learnt cluster
- The procedure of workload classification
  - Evaluation of antecedents
  - Implication calculation of consequents
  - Aggregation result of consequents

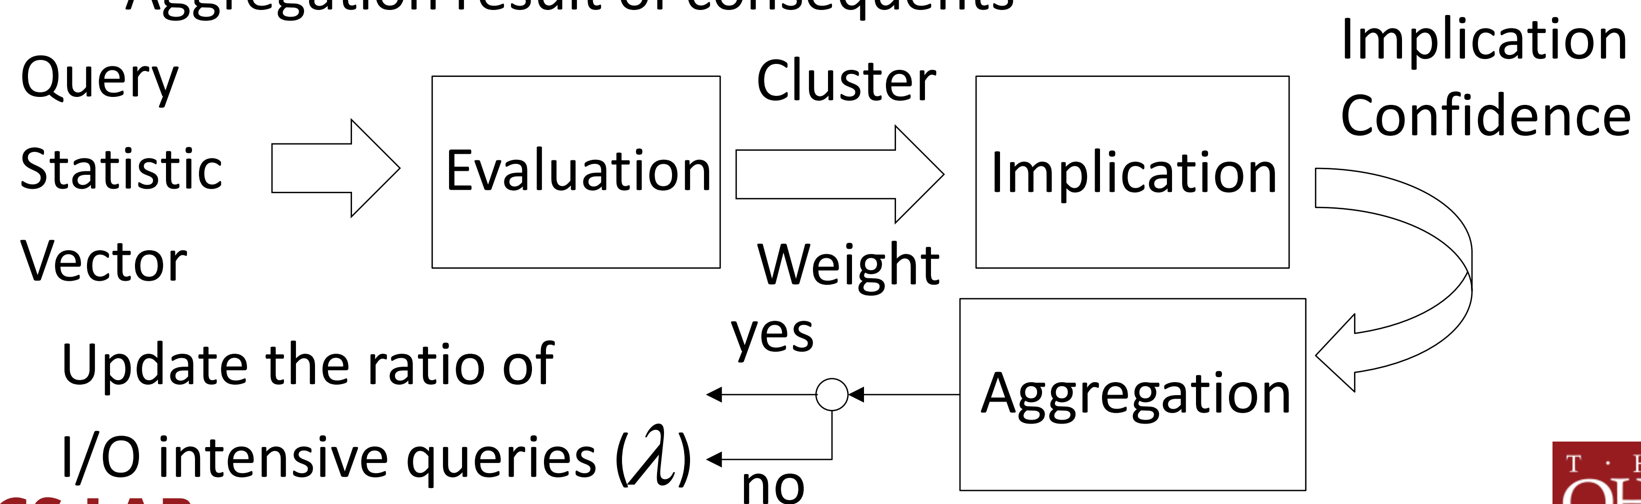

The Controller Design

- System model
  - $\Delta r(i) = \lambda A * f(i) + B$
- Control model (PI control)
  - Zero steady-state error
  - Short settling time
  - Stability

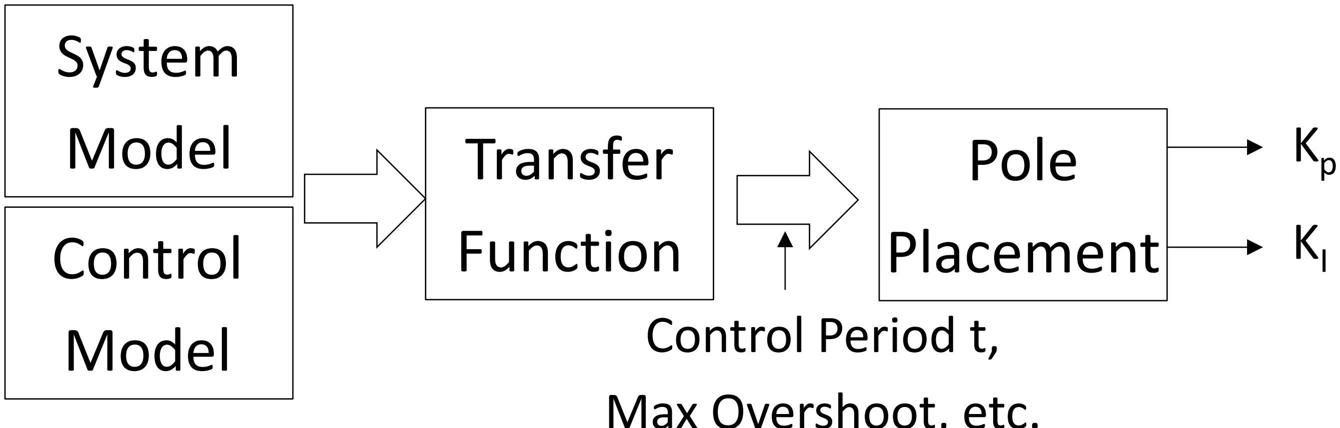

Methodology

- Testbed
  - Hardware: a DELL PowerEdge R710 with 12-core Intel Xeon E5645, 16GB memory and 1TB storage
  - Software: Redhat 5 (V3.0.0) + PostgreSQL (V8.3.18)
  - Workload: TPC and SDSS
  - Power measurement: WattsUp power meter
- System Contention
  - Ideal: running DBMS alone
  - Competing: running DBMS with other processes with equal priority.
  - Preemptive: running DBMS with other processes with higher priority.

Baselines

- State-of-the-Practice Baselines
  - NORMAL: no power management, performance first
  - SPEEDSTEP: Intel™ build-in power management

- State-of-the-Art Baselines
  - TRADITION: Open-loop control based on known workload statistics
  - HEURISTIC: Ad-hoc control based on performance-power model
  - SCTRL: System-level power control for DBMS performance

A Snapshot of PAT in Competing Setup

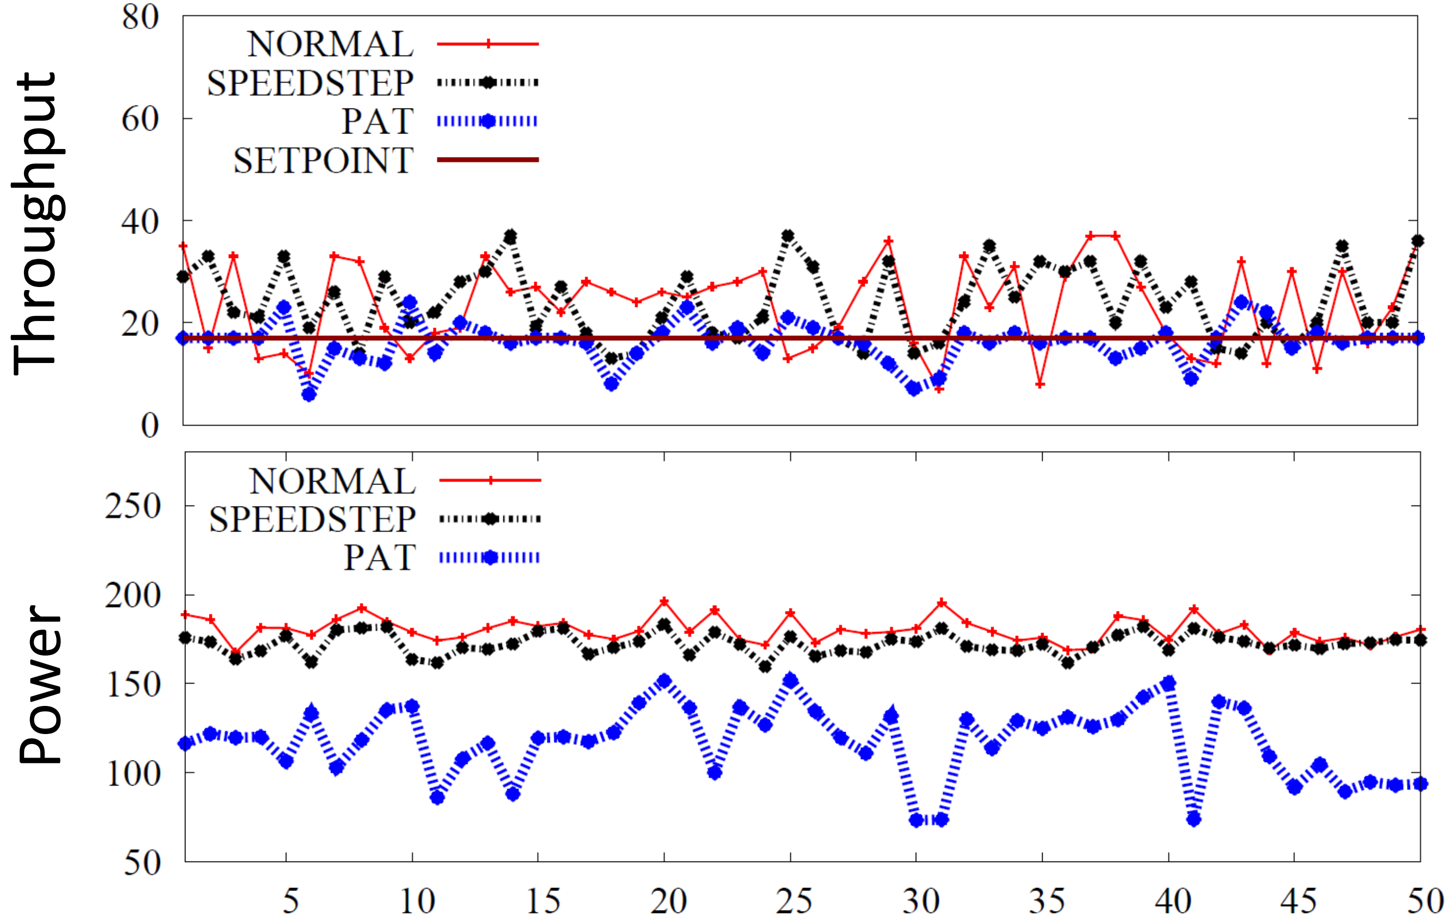

How PAT Performs in All Scenarios

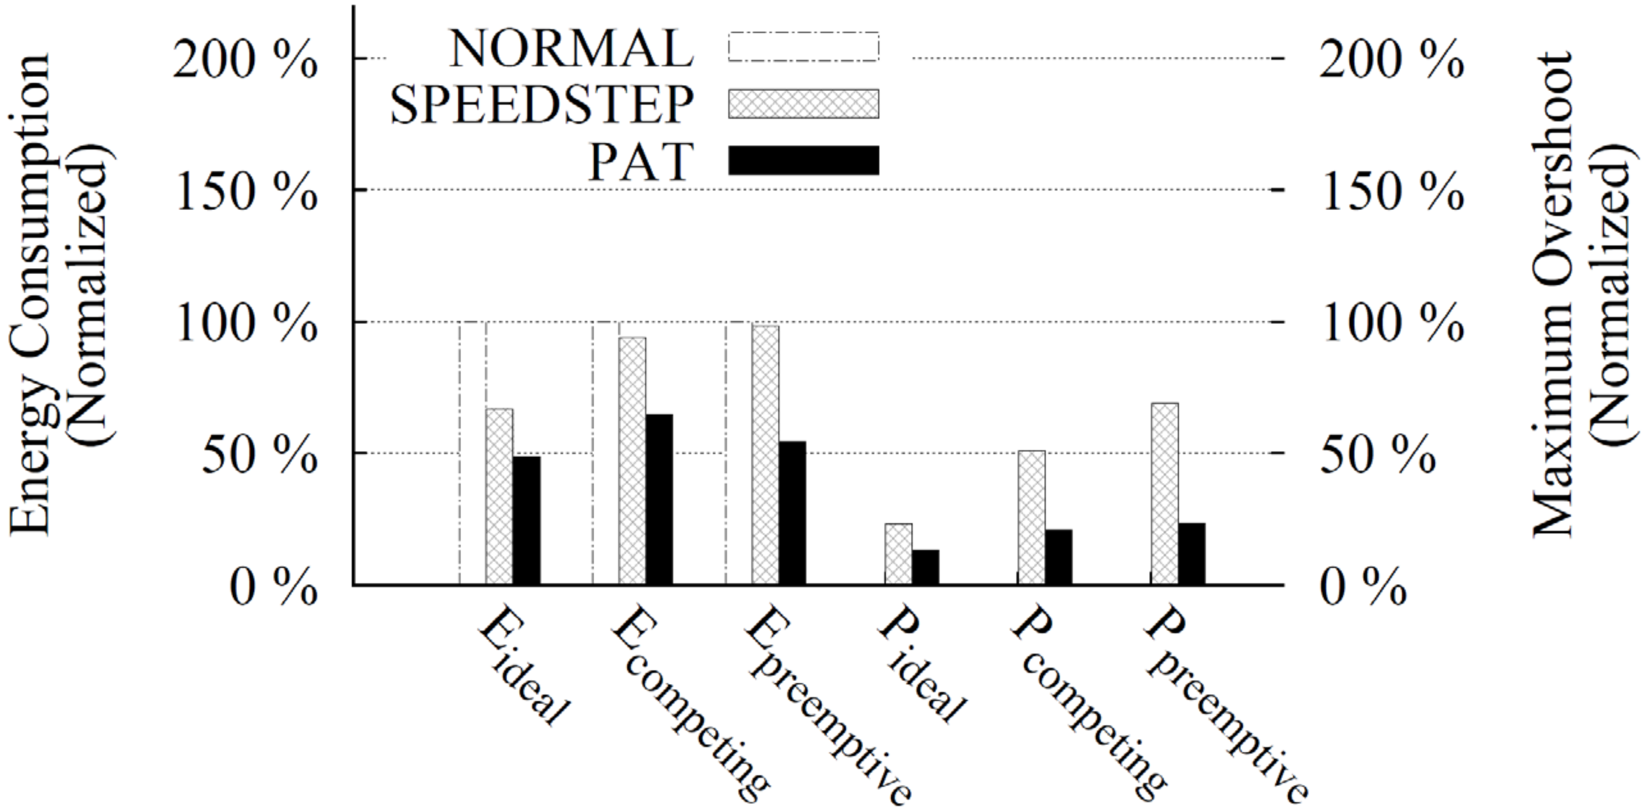

- PAT saves energy cost (15% more than the SPEEDSTEP, 51.3% more than NORMAL)

A Snapshot of Comparison

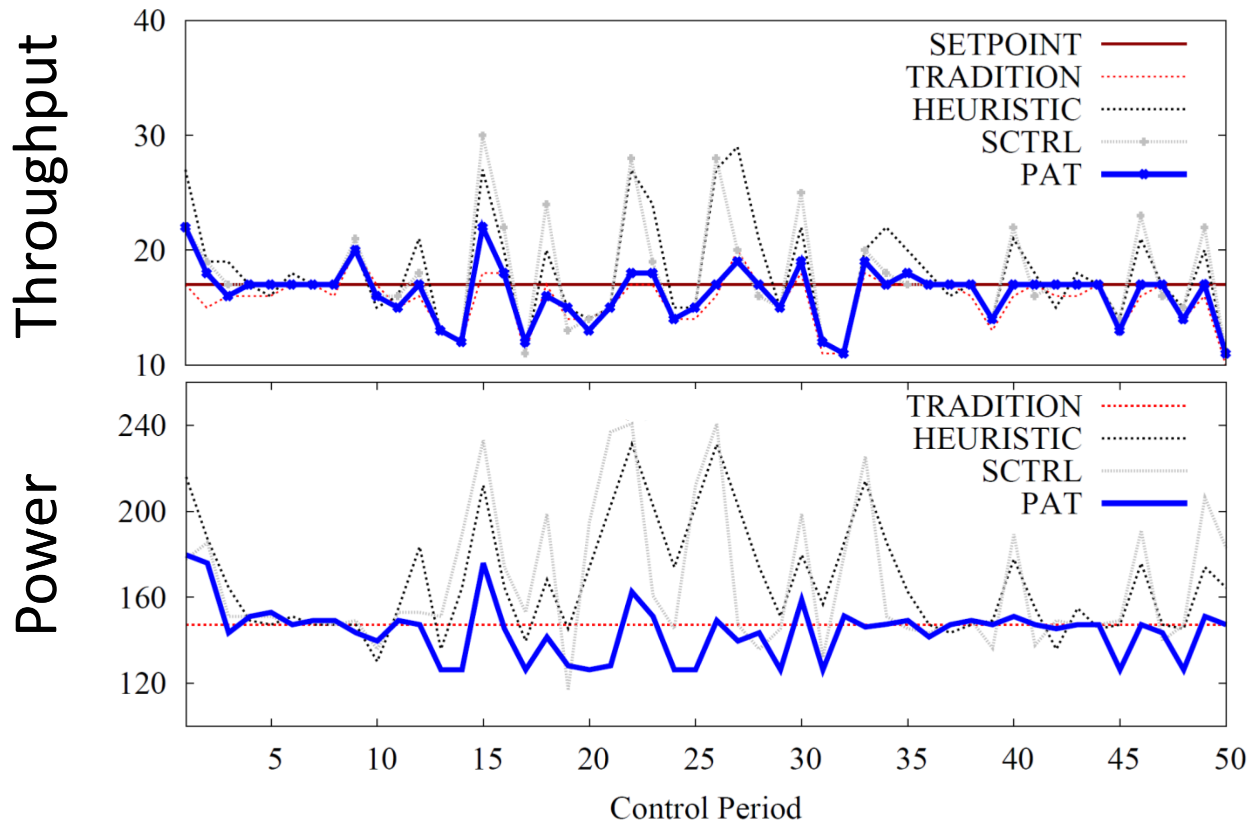

Comparison with Other Control Methods

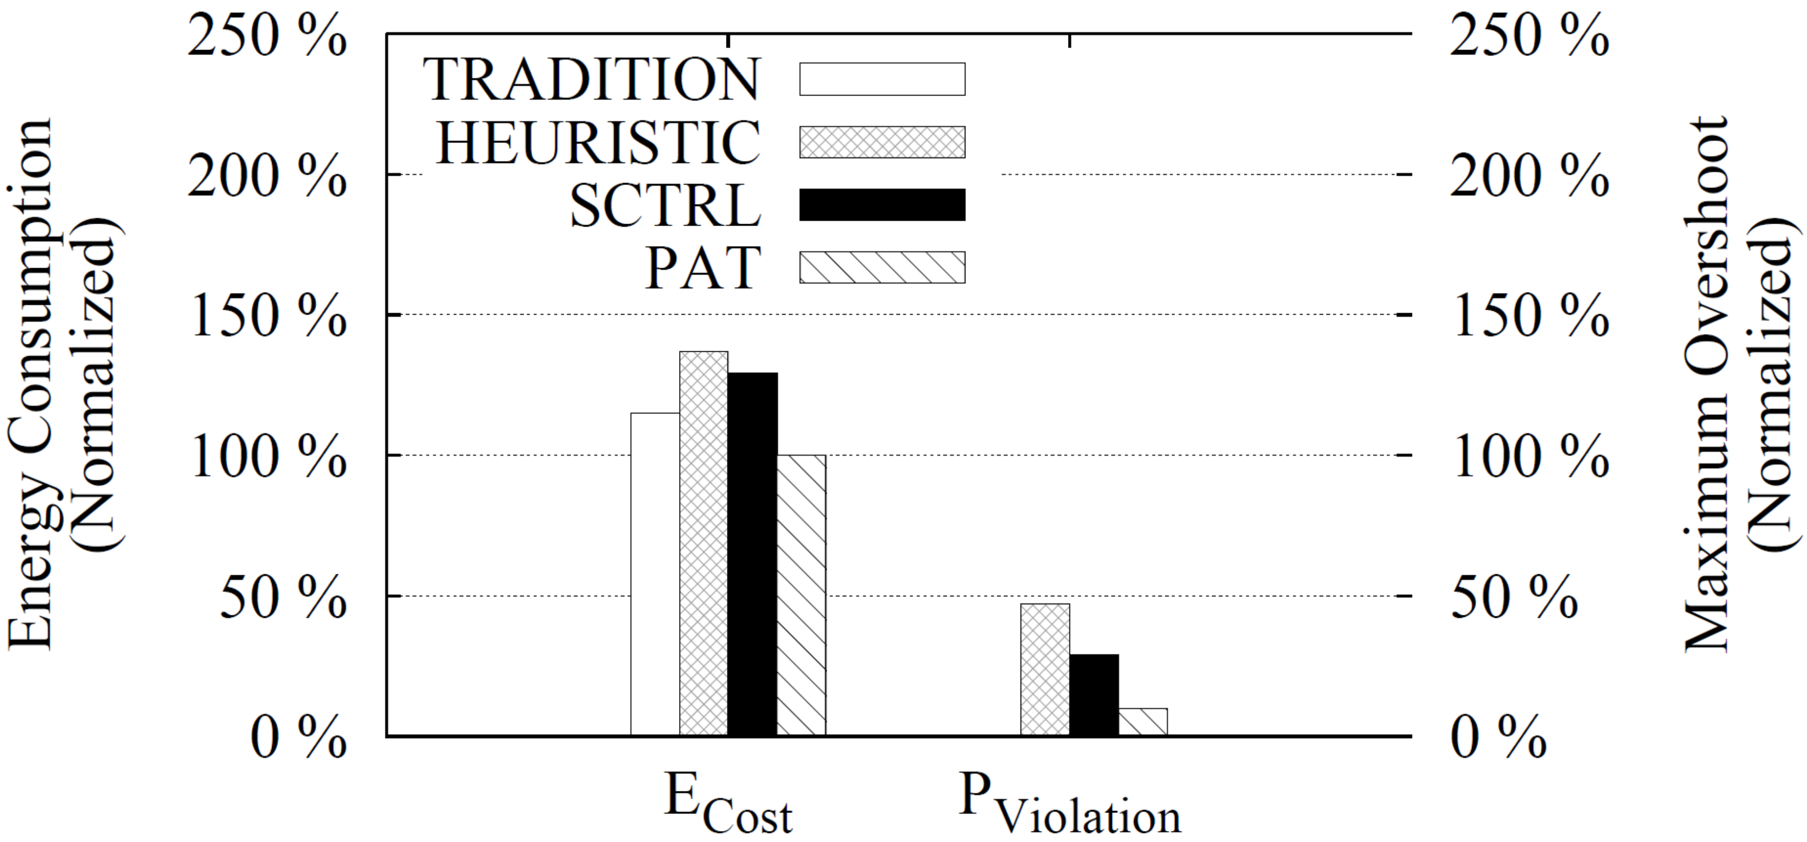

- PAT has the smallest overshoot and the best energy saving

Conclusion

- Minimizing power consumption of DBMS under a user-specified performance bound
- Controller design on the system characteristics is the solution
- Empirically study the relationship between energy and DBMS processing
- PAT: a feedback control framework with system characteristics

- Thanks

Acknowledge

- This research was supported by the US National Science Foundation (NSF) grants IIS-1117699, IIS-1156435, and CSR-1143607

Related Work

- Control in DBMS
  - Feedback Control in DBMS [Tu et al. VLDB 2007]
- Other Applications
  - Power Shaving
  - Power Over-commitment
  - Proportional Energy Consumption

A Snapshot of PAT in Ideal Setup

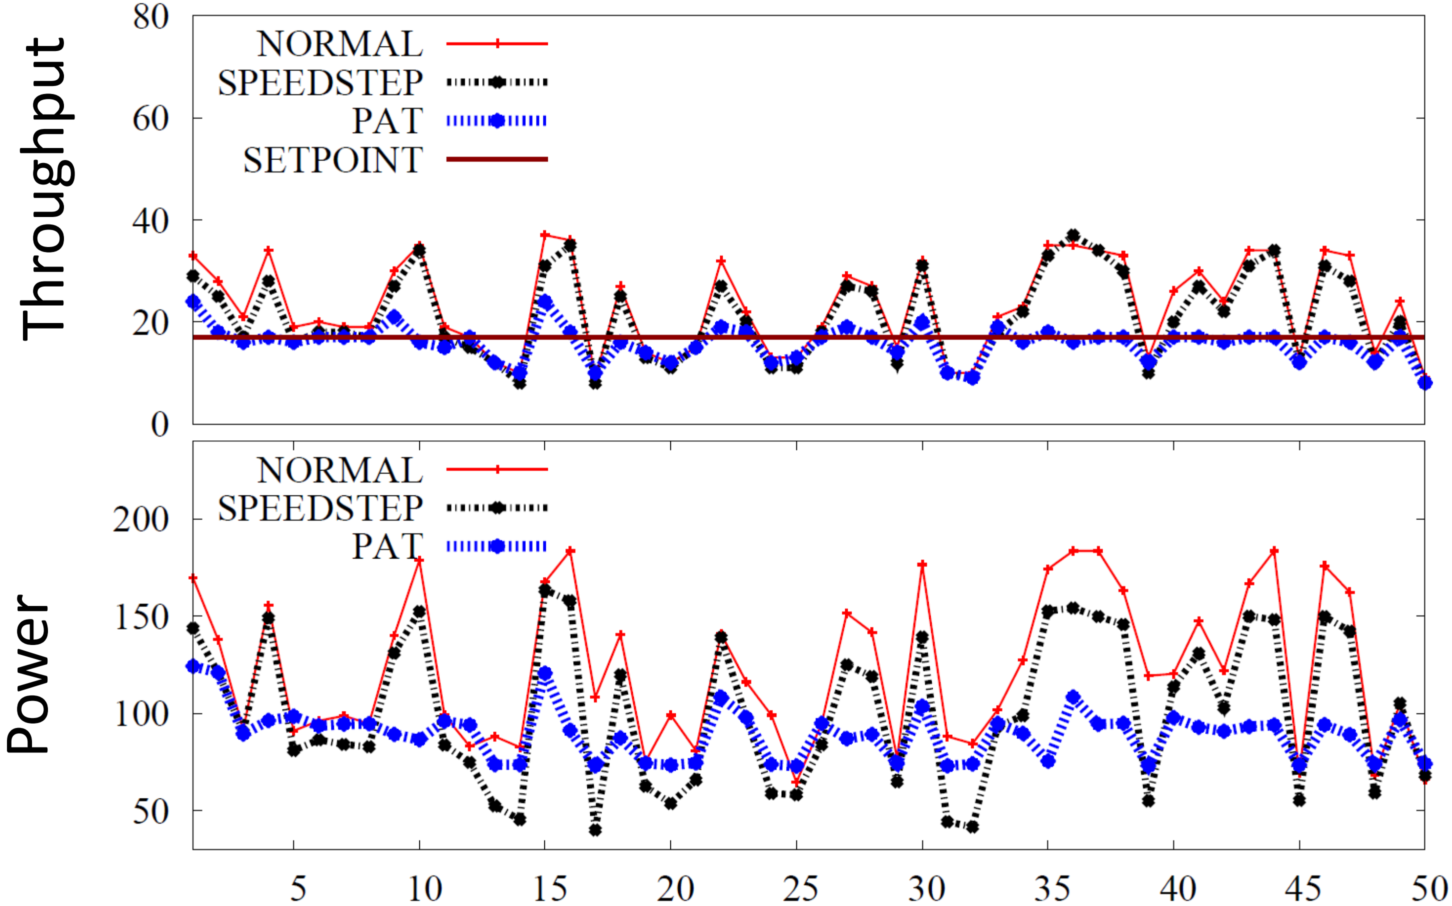

A Snapshot of PAT in Competing Setup

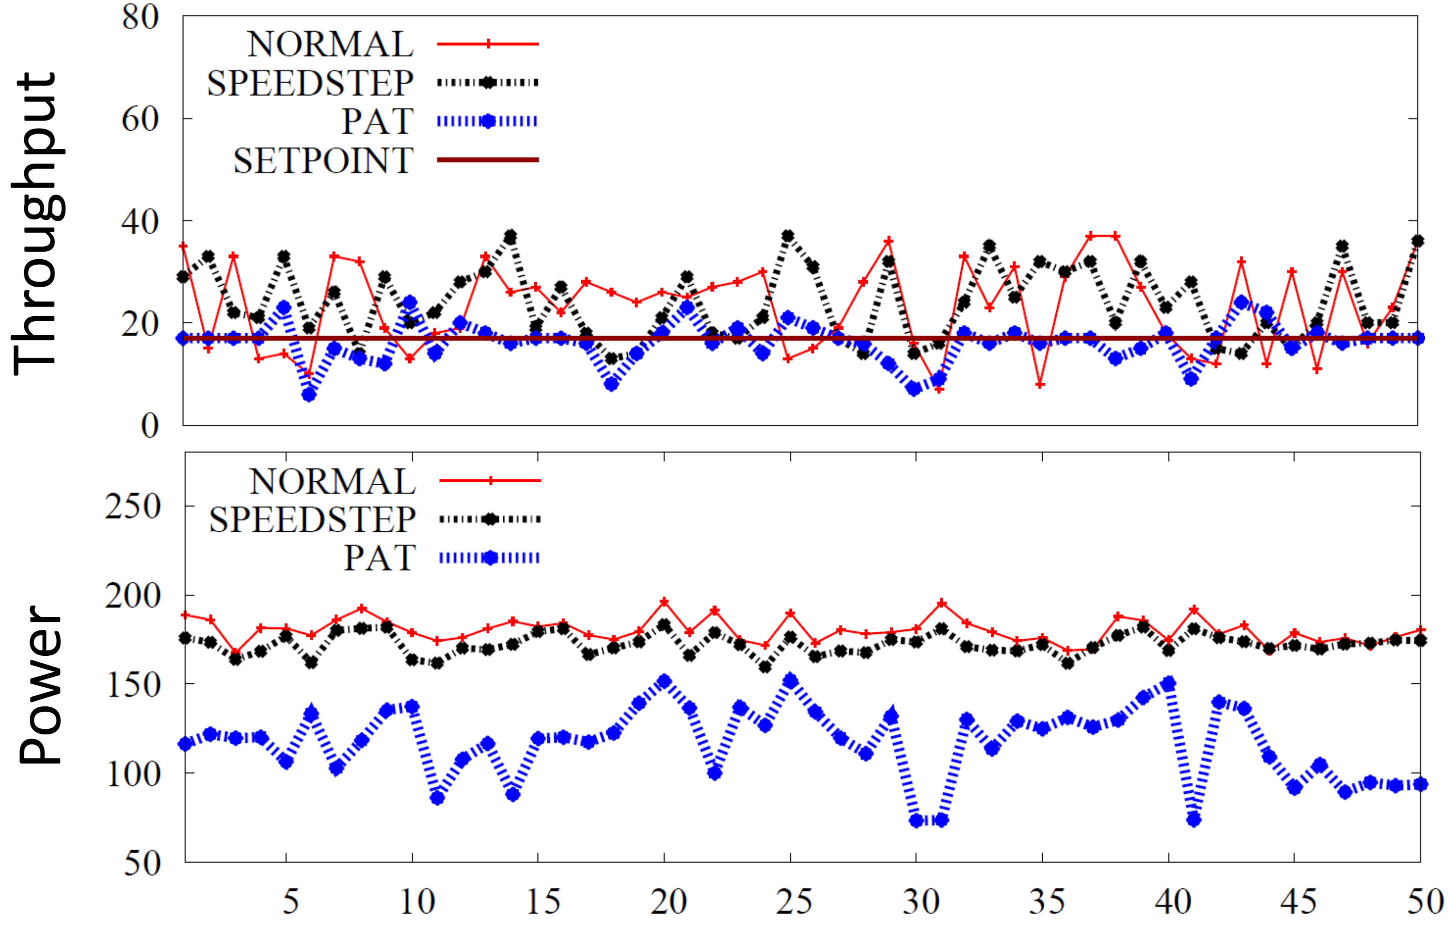

Backup: Errors From Modeling

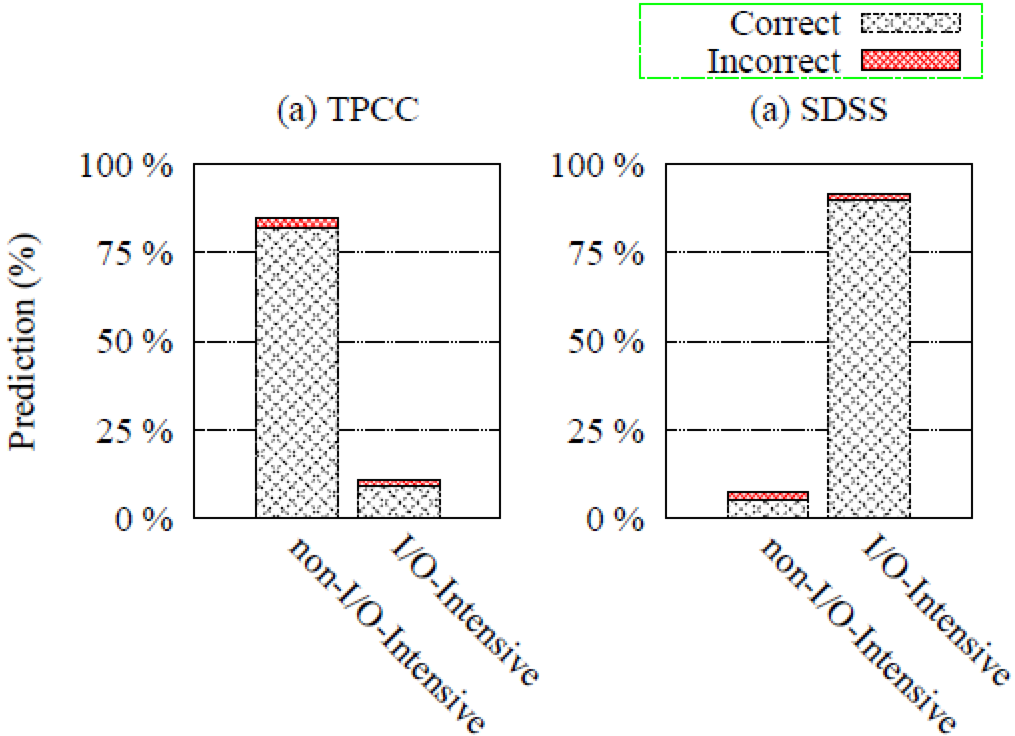

- FWC could produce almost 10% error in the system, The designed maximum overshoot is 40% in PAT.

Backup Characterization Study – Energy and

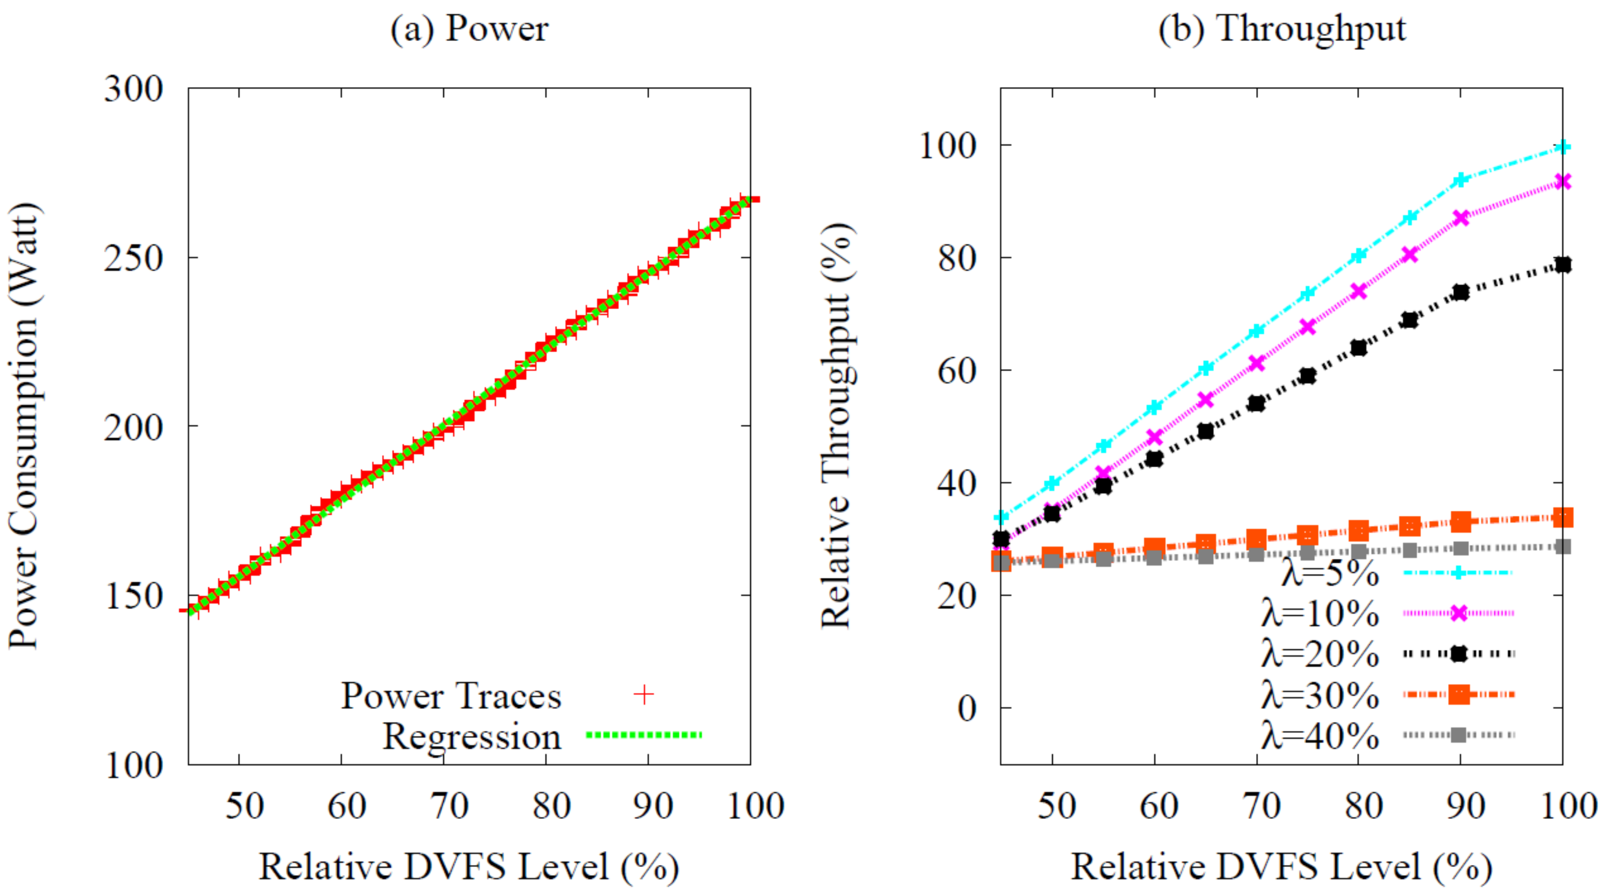

- DBMS throughput and CPU frequency can be approximated as a linear model

Backup: Rule-based Workload Classifier

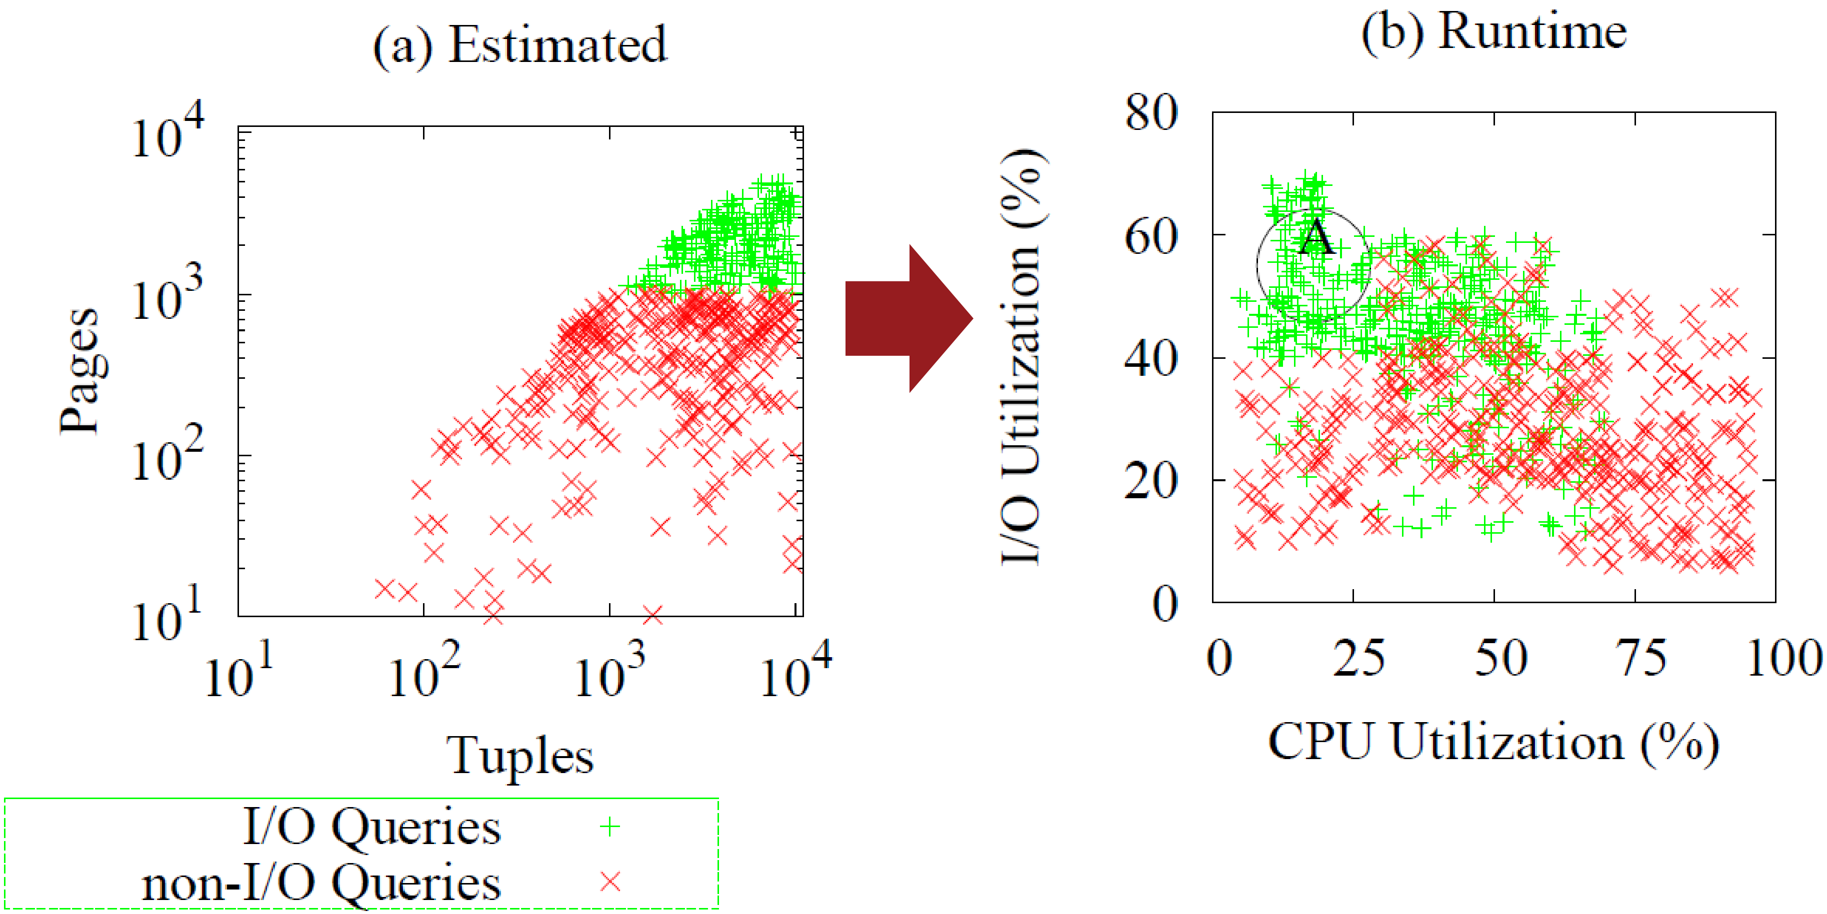

Outline

- Problem Overview & Motivation
- Characterization Study: DBMS and Energy
- Overall Framework
- Fuzzy Workload Classifier
- The Controller Design
- Evaluation
- Conclusion

Fuzzy Workload Classifier

- One fuzzy rule base
  - If $[d^i_{CPU}, d^i_{I/O}] \in$ cluster X_j,
    then $[u^i_{CPU}, u^i_{I/O}] = M_j [d^i_{CPU}, d^i_{I/O}] + N_j$

- The procedure of workload classification
  - Evaluation: $[d^i_{CPU}, d^i_{I/O}]$ -> $[u^i_{CPU}, u^i_{I/O}]$ per X_j
  - Implication calculation: $\frac{\Sigma(p_j)-p_j}{\Sigma(p_j)}$ -> t^i_j of X_j
  - Aggregation result: $[\sum t^i_j u^i_{CPU}, \sum t^i_j u^i_{I/O}]^T$