UNIVERSITY of VIRGINIA

Considering Resource Demand Misalignments To Reduce Resource Over-Provisioning in Cloud Datacenters

Liuhua Chen

Dept. of Electrical and Computer Eng.

Clemson University, USA

Haiying Shen

Dept. of Computer Science

University of Virginia, USA

Cloud Computing

- Cloud computing: large groups of remote servers networked to allow centralized data storage and online access to computer services or resources

Cloud Providers

Cloud Customers

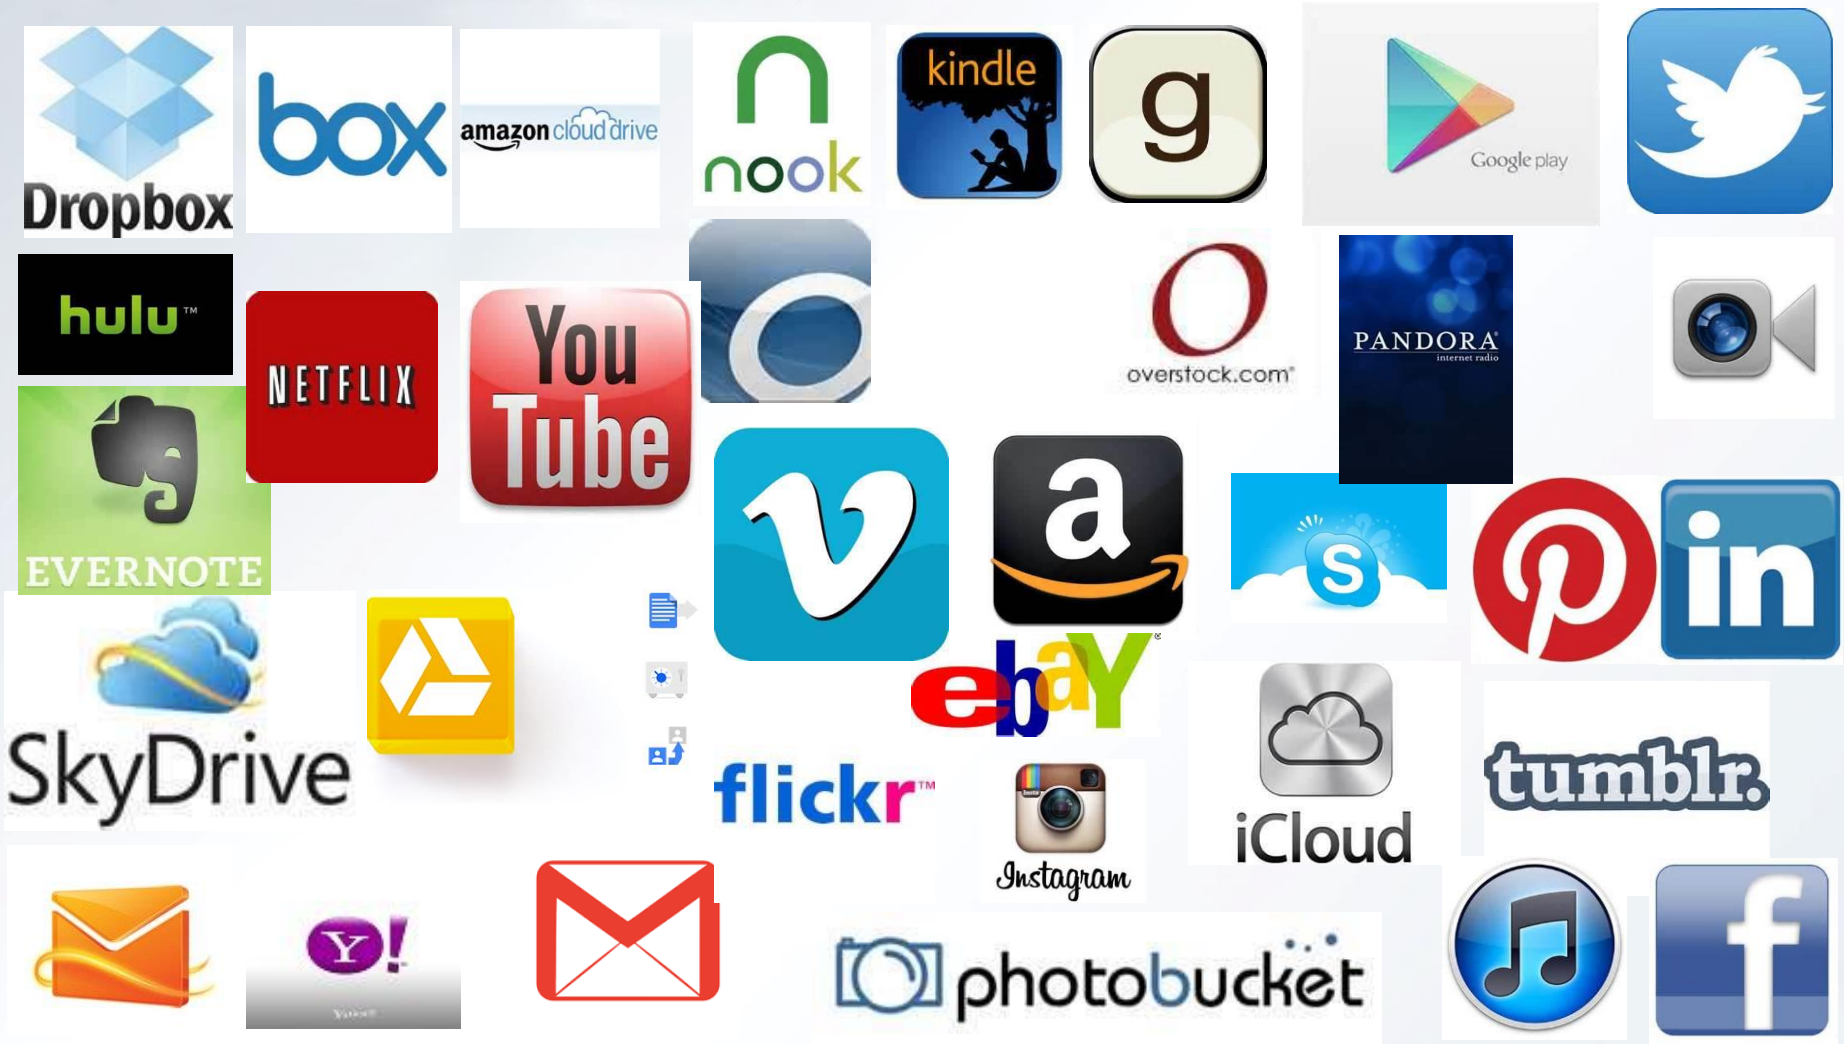

Research Problem and Goal

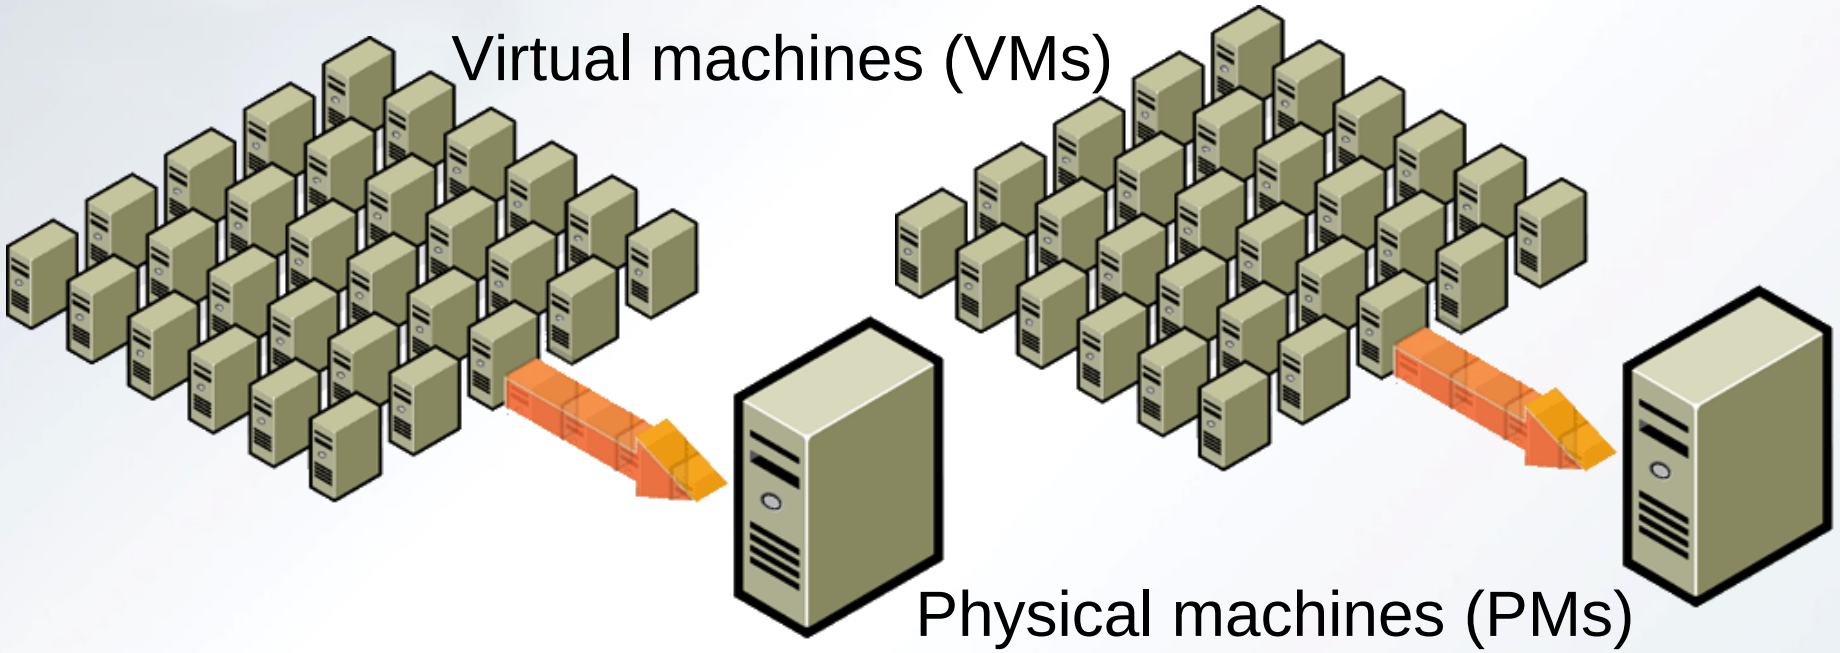

Motivation

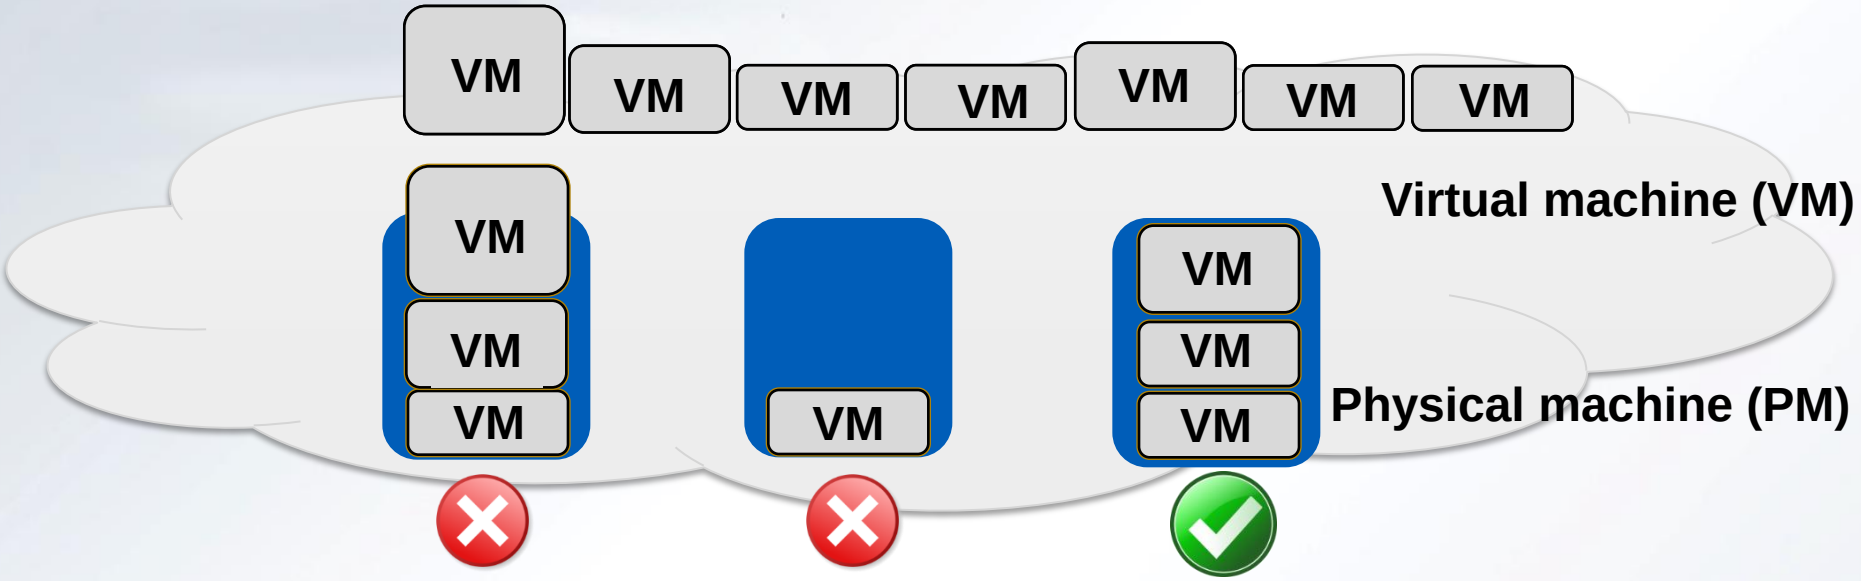

- Over-loaded PMs→low QoS→SLO violation→penalty
- Under-loaded PMs→ resource waste→high system cost
- Problem: reduce over-loaded and under-loaded PMs
- Goal: high QoS, high resource efficiency, high profit

Initial Complementary VM Consolidation

- Previous work (CompVM):

L. Chen and H. Shen, Consolidating Complementary VMs with Spatial/Temporal-awareness in Cloud Datacenters, *Proc. of the 33rd Annual IEEE International Conference on Computer Communications (INFOCOM'14),* Toronto, Canada, 2014

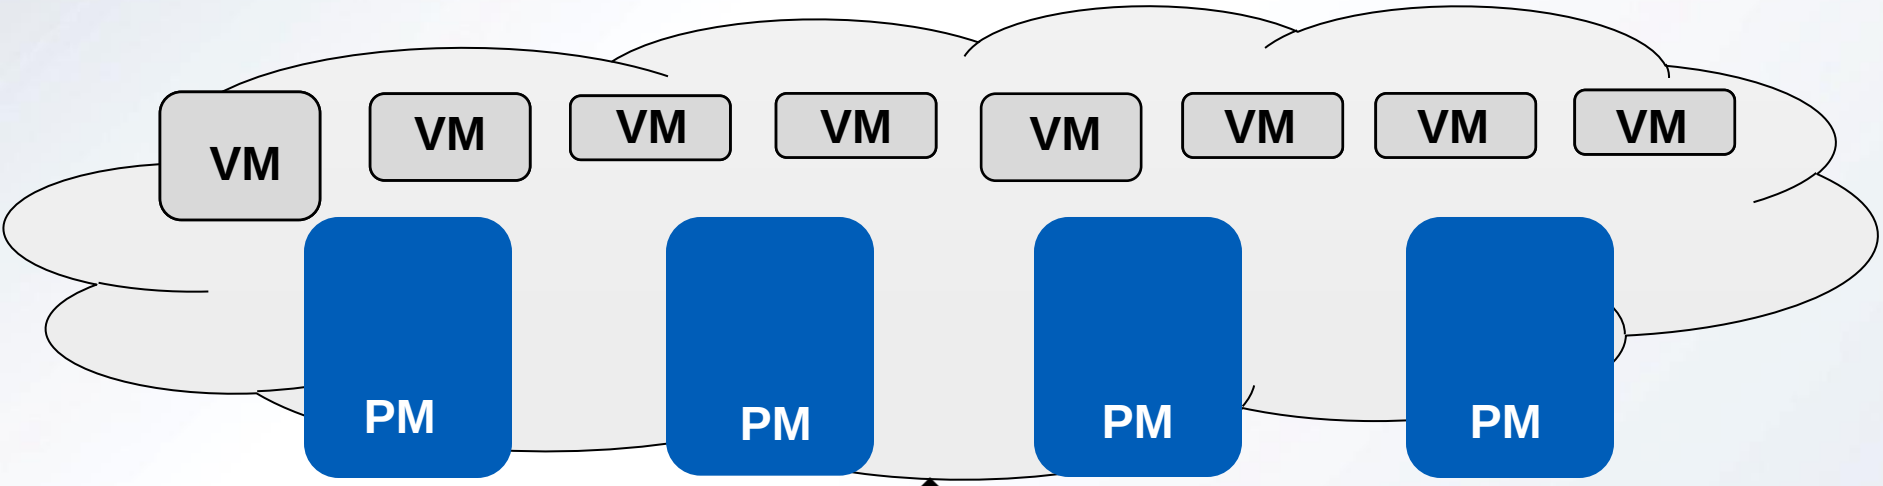

- How to achieve  ?

Initial Complementary VM Consolidation – Motivation

- CompVM: load balance in the long term
  consolidate complementary VMs

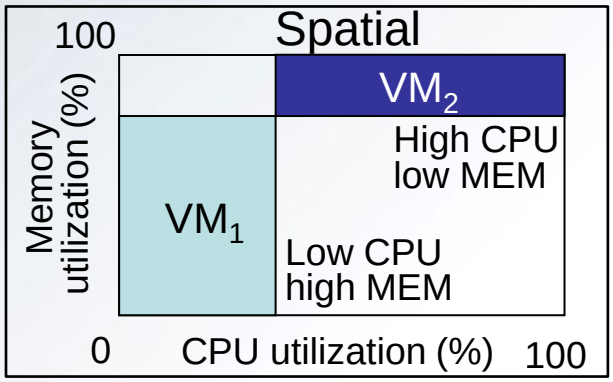

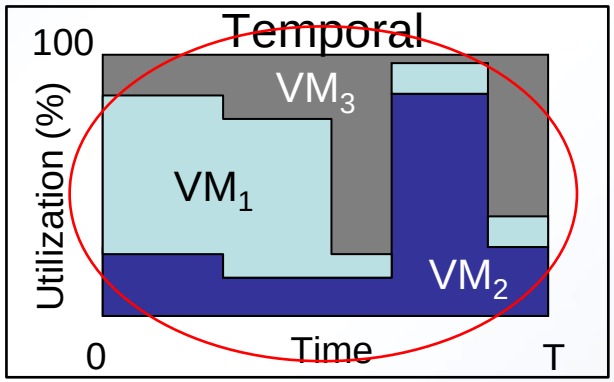

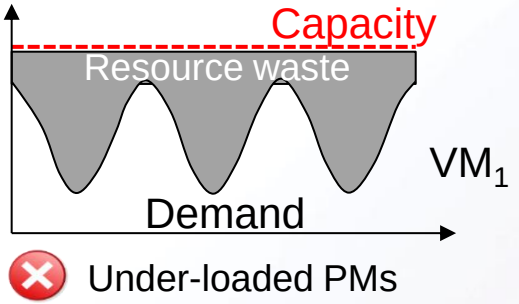

Under-loaded PMs

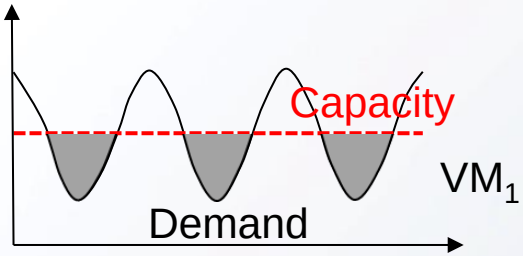

Over-loaded PMs

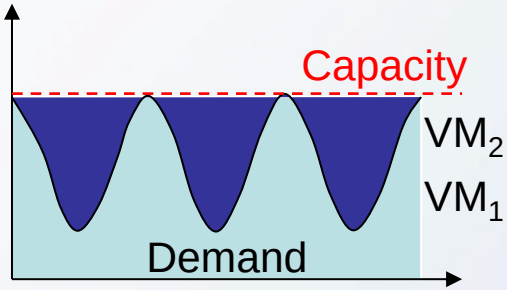

No under/over-loaded PMs

Patterns?

Initial Complementary VM Consolidation – VM Utilization Pattern

- Measurement:
  - MapReduce jobs: TeraSort, TestDFSIO read/write
  - Google cluster trace, PLANETLAB trace
- Periodic resource utilization patterns exist in many VMs running
  - the same short-term job
  - a long-term job

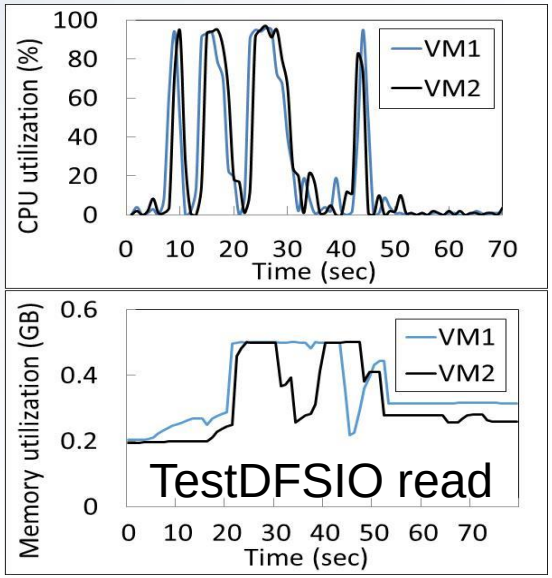

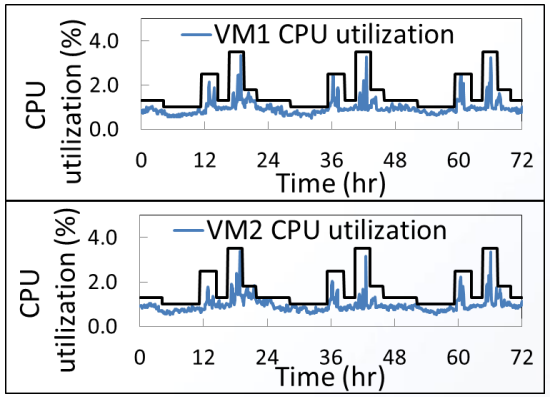

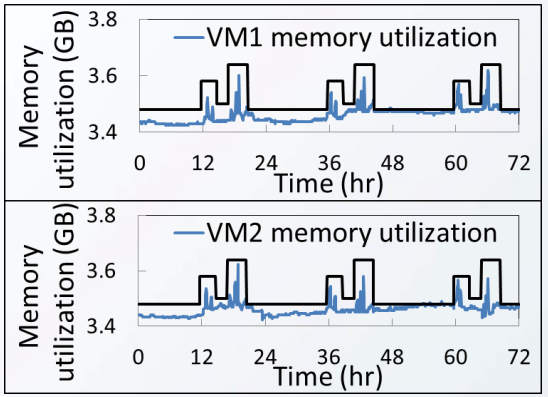

Google cluster trace

Initial Complementary VM Consolidation – Utilization Pattern Detection

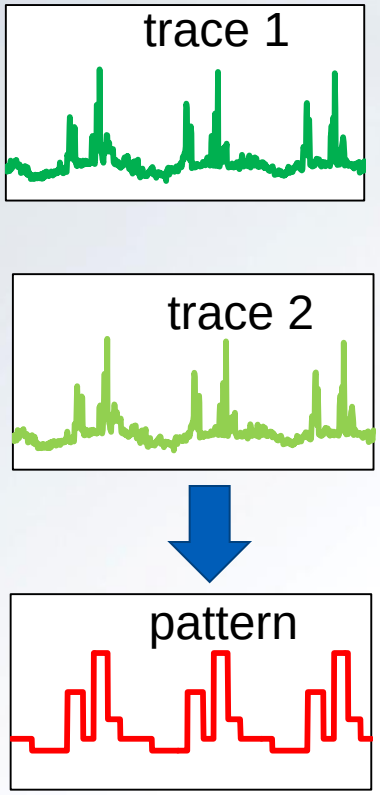

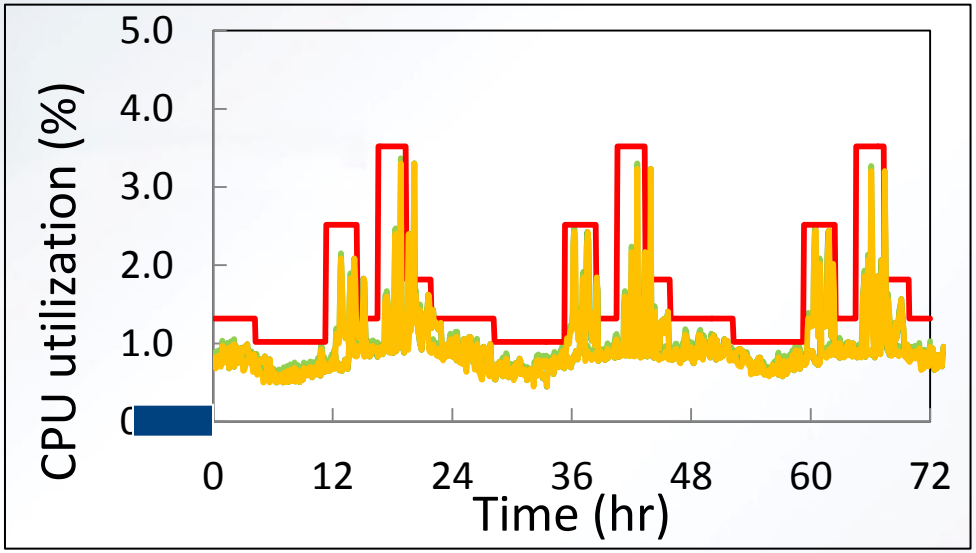

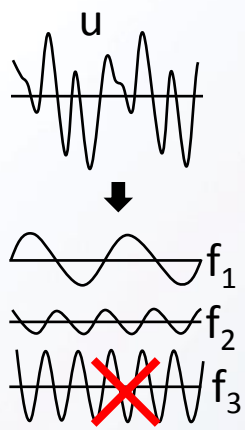

Initial Complementary VM Consolidation – VM Allocation Method

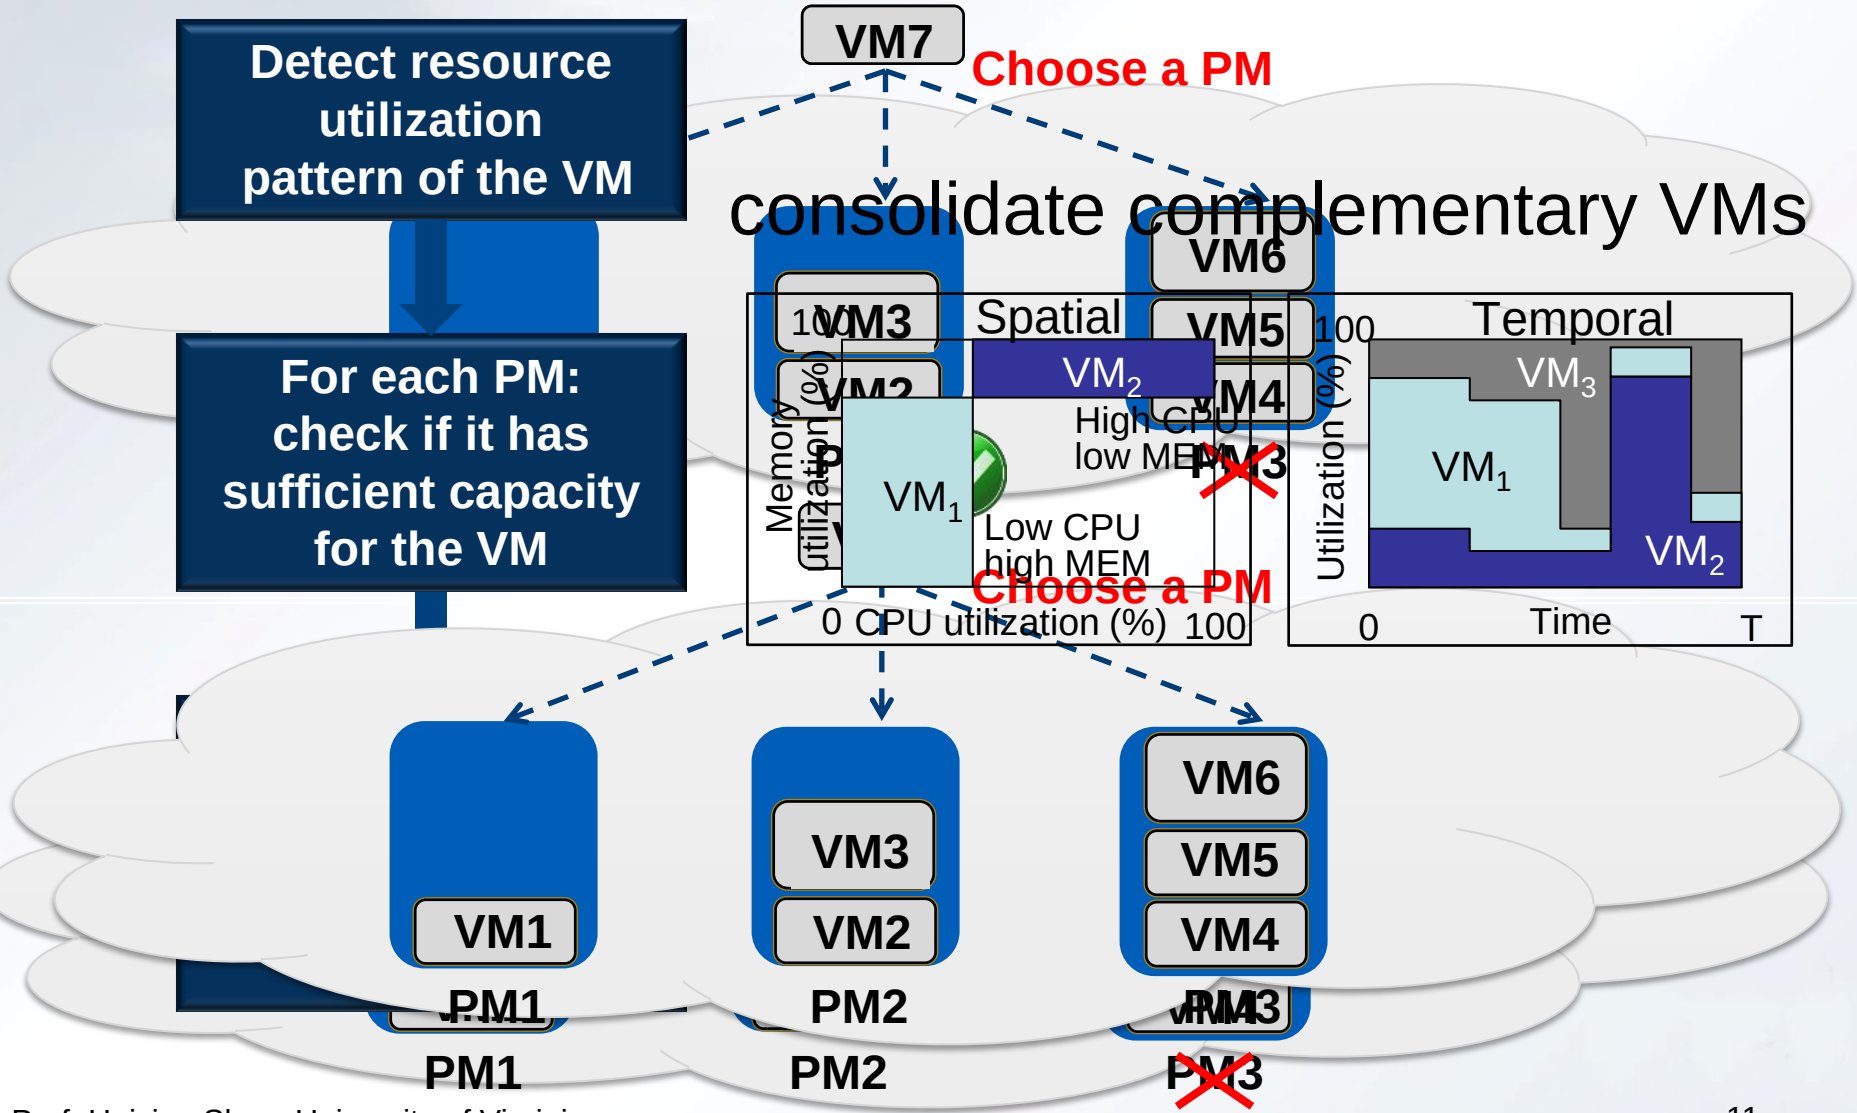

Reducing Prediction Error

Subject of this paper:

How to reduce resource utilization prediction error

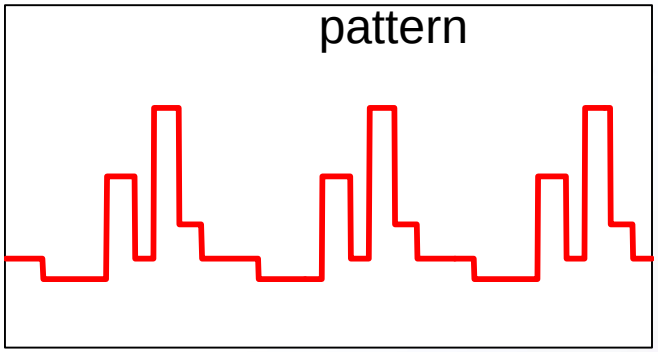

VM Consolidation – Reduce Provisioned Resource

- Pulse deviation yields a pattern with a pulse width larger than the actual pulse width
- Resource over-provisioning
- Not revealed or studied before

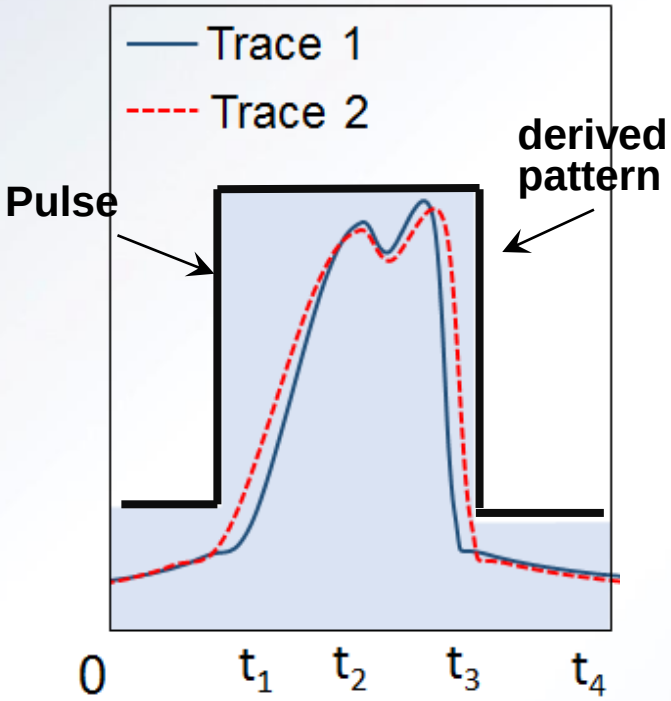

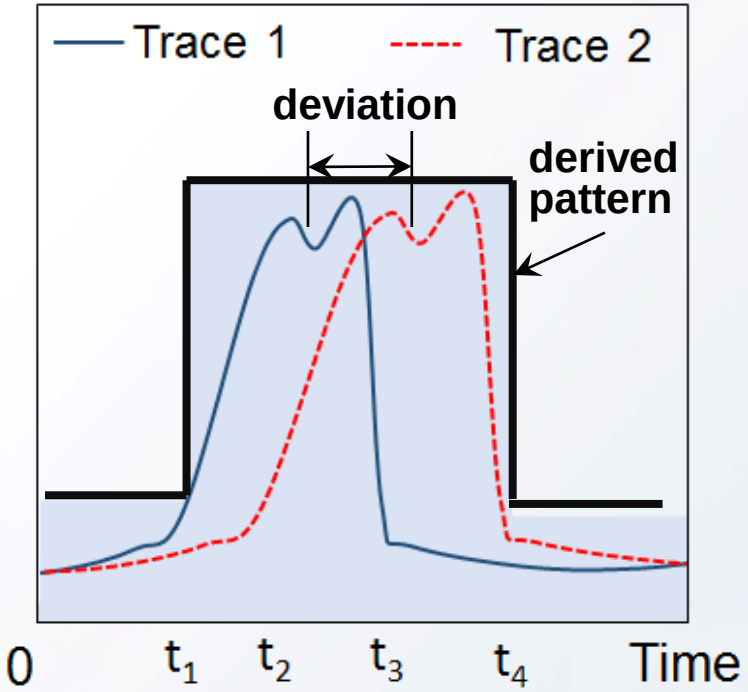

VM Consolidation – Trace Study

- Pulse deviations are common

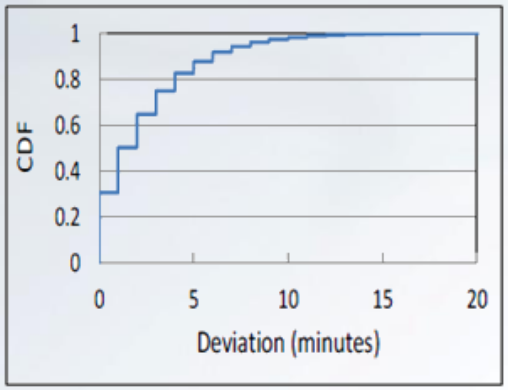

(a) Google CPU trace.

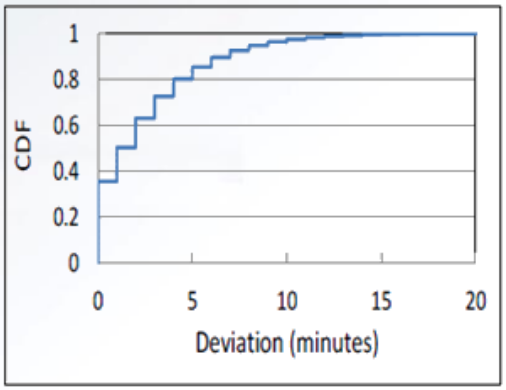

(b) Google memory trace.

Google Cluster trace 100 jobs, 29920 tasks

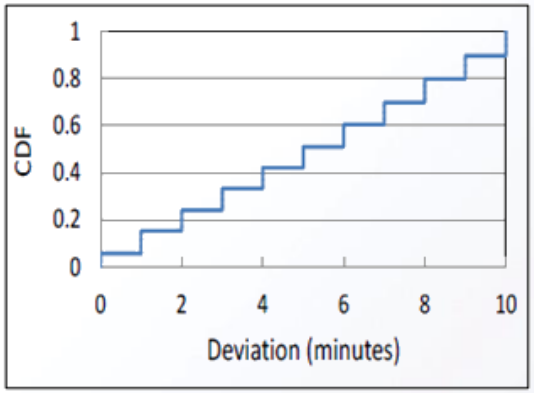

PlanetLab trace

1000 jobs, 4695 tasks

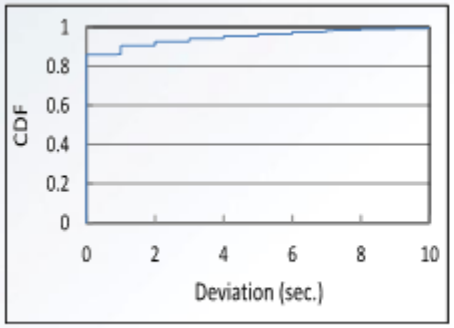

(a) CPU utilization

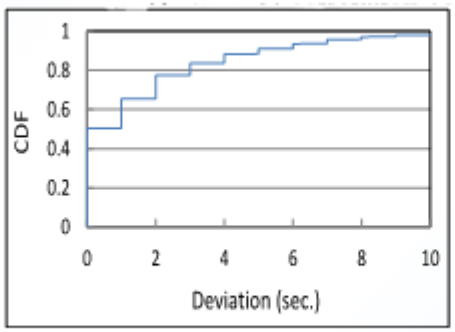

(b) RX bandwidth utilization

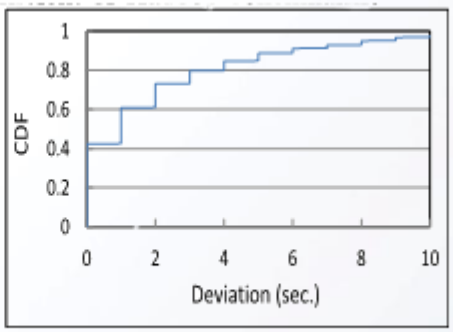

(c) I/O utilization

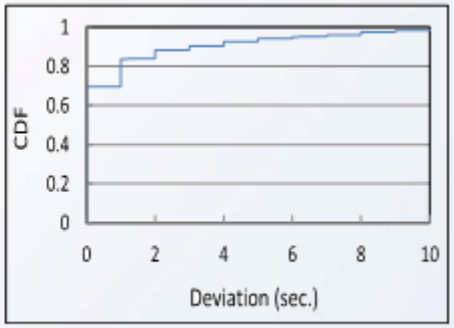

(d) Memory utilization

MapReduce benchmarks on a HPC cluster (Wordcount, Grep, Terasort, TestDFSIO and PiEstimator)

Average task execution time is around 100 minutes/seconds

VM Consolidation – Trace Study

- Resource efficiency: demand/capacity
- Even using CompVM, the resource efficiency still needs to improve

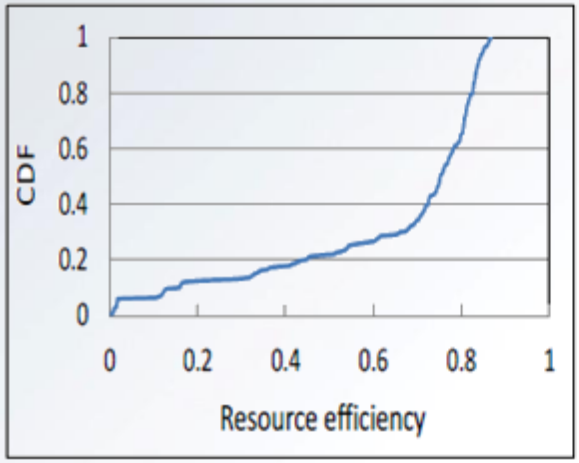

(a) CPU efficiency.

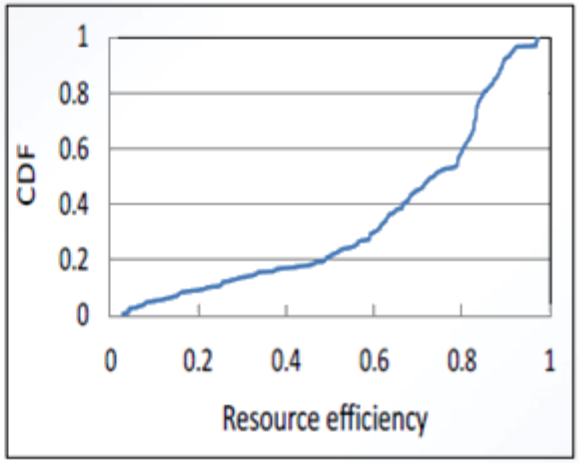

(b) Memory efficiency.

Google Cluster trace

100 jobs, 1550 tasks

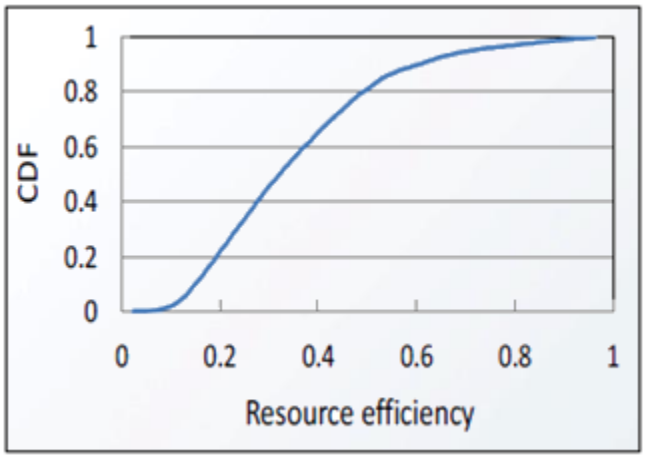

PlanetLab trace

1000 jobs, 4695 tasks

VM Consolidation – Pattern Refinement Methodology

- Pattern refinement methods
  - Lowering cap: lower each value in the pattern by c_{high}-c_{low}
  - Reducing pulse width
  - Optimal base provisioning

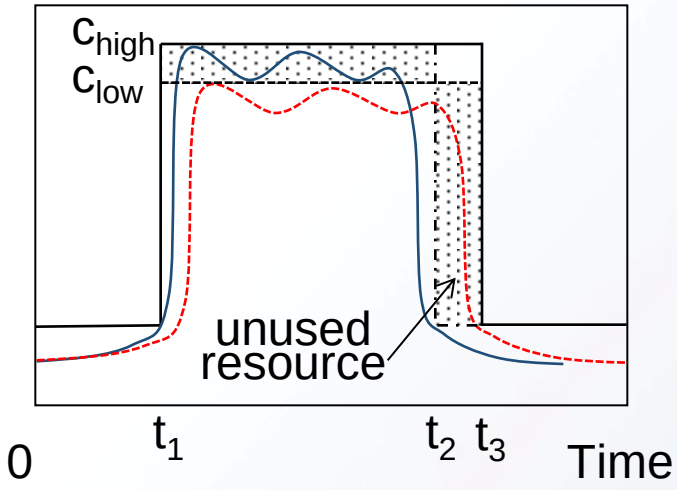

VM Consolidation – Pattern Refinement Methodology

- Pattern refinement methods
  - Lowering cap
  - Reducing pulse width: postpone the pulse from t_1 to t_3
  - Optimal base provisioning

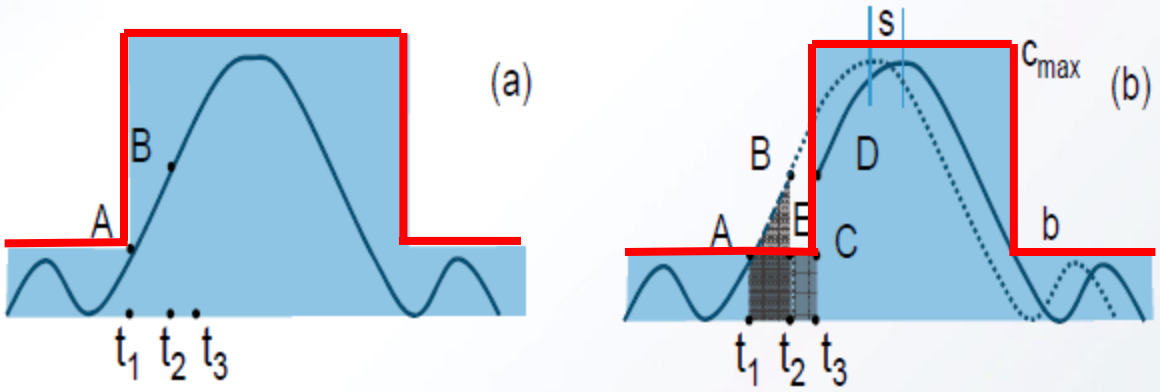

VM Consolidation – Pattern Refinement Methodology

- Pattern refinement methods
  - Lowering cap
  - Reducing pulse width
  - Optimal base provisioning

    refine pattern based on optimal b value that maximizes resource efficiency

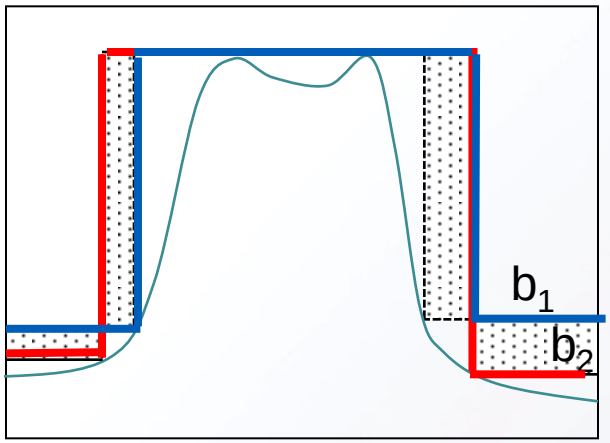

VM Consolidation – Pattern Refinement Methodology

- Pattern refinement methods
  - Lowering cap
  - Reducing pulse width
  - Optimal base provisioning
- Risk violating SLOs

VM Consolidation – Performance Evaluation

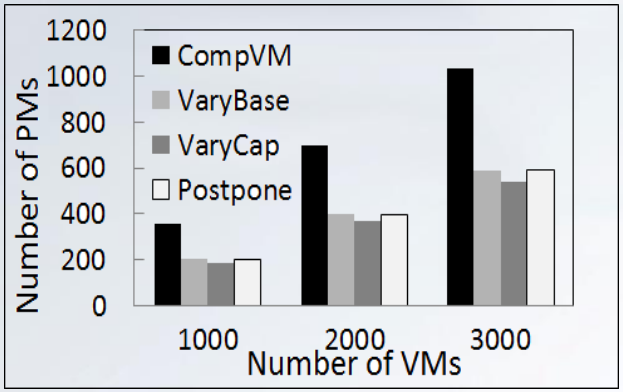

Google Cluster trace

Improve by 40%

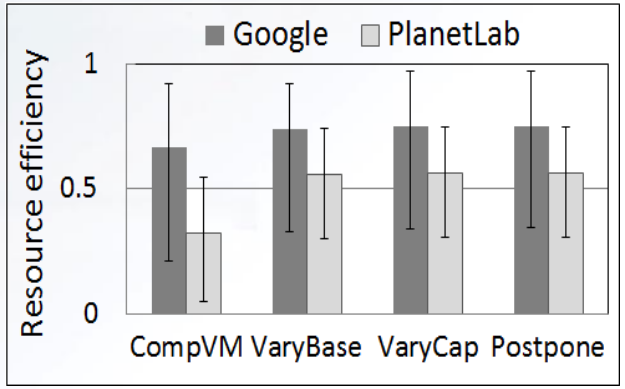

Improve by 10%, 70%

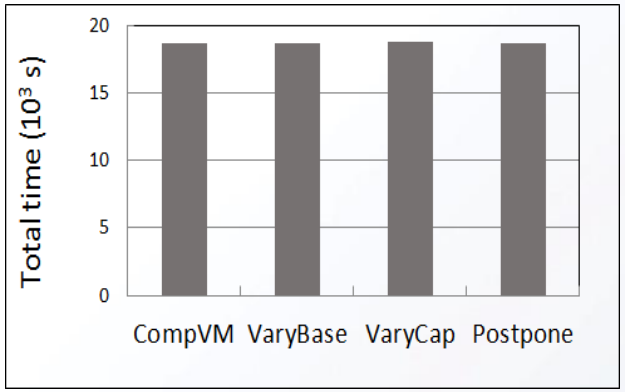

Cluster experiment

2000 VMs, CloudSim, Palmetto HPC cluster, Traces: NAS Parallel Benchmark, Google cluster PLANETLAB

Pattern refinement yields higher resource efficiency without compromising VM performance by handling pulse deviations!

Conclusion

- Trace study
  - Pulse deviations are common
  - Even using CompVM, the resource efficiency still needs to improve
- Pattern refinement methods
  - Lowering cap
  - Reducing pulse width
  - Optimal base provisioning
- Experiments
  - Higher resource efficiency without compromising VM performance

Future Work

- Consider other factors (e.g., SLOs) in VM consolidation

- Consider VM migration

Thank you!
Questions & Comments?

Haiying Shen

hs6ms@virginia.edu

Associate Professor

Department of Computer Science

University of Virginia