Migration, Assignment, and Scheduling of Jobs in Virtualized Environment

Seung-Hwan Lim
Jae-Seok Huh
Youngjae Kim
Chita R. Das

The Pennsylvania State University
Oak Ridge National Laboratory

OAK RIDGE National Laboratory

PENNSTATE

A Virtualized Environment

Basis of Cloud Computing

Decouples Operating System Instances from Hardware

Enables migration of OS instances (Virtual Machines)

Challenges in Virtualized Environment

Obstacle	Opportunity
1 Availability/Business Continuity	Use Multiple Cloud Providers
2 Data Lock-In	Standardize APIs, Compatible SW to enable Surge or Hybrid Cloud Computing
3 Data Confidentiality and Auditability	Deploy Encryption, VLANS, Firewalls
4 Data Transfer Bottlenecks	FedExing Disks; Higher BW Switches
5 Performance Unpredictability	**Improved VM support; Flash memory; Gang Schedule VMs**
6 Scalable Storage	Invent Scalable Store
7 Bugs in Large Distributed Systems	Invent Debugger that relies on Distributed VMs
8 Scaling Quickly	Invent Auto-Scaler that relies on ML; Snapshots for Conservation
9 Reputation Fate Sharing	Offer reputation-guarding services like those for email
10 Software Licensing	Pay-for-use licenses

From Armbrust et al., A view of cloud computing. *Commun. ACM,* April 2010

Unpredictable Performance May Cause

Propagation Effects

Performance variance among jobs may create cascaded effects in all the related jobs

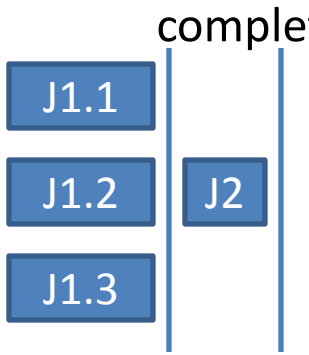

Desired

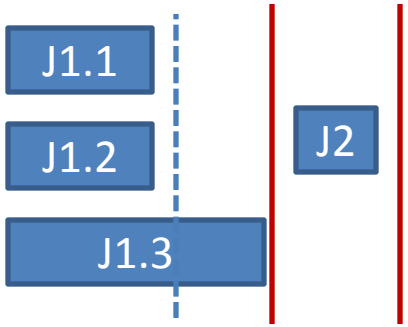

With performance variance

Repels performance critical applications.

VM assignment/scheduling schemes consider performance

Job Scheduling

An optimization problem to find assignment of jobs to the given set of machines so as *to* minimize the *total completion time* of jobs

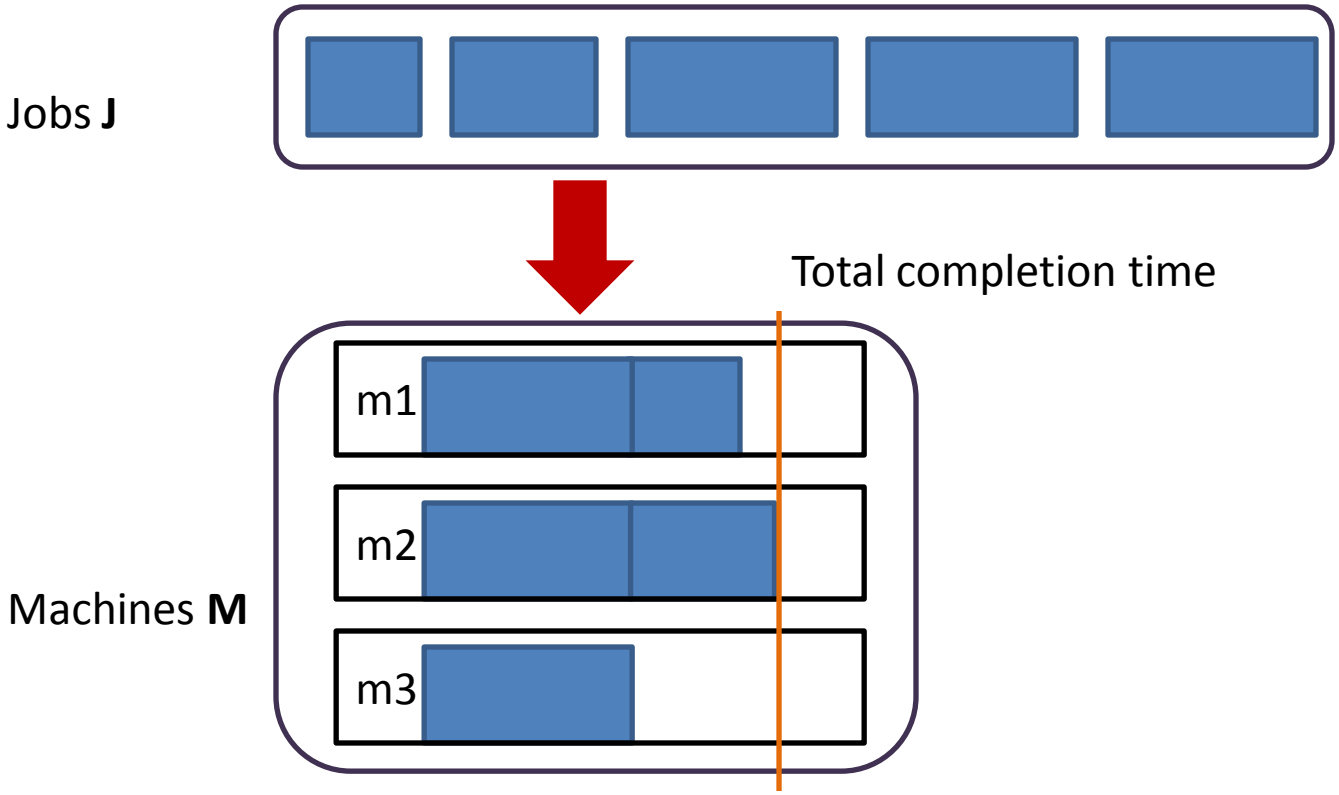

During (Re-)Assigning VMs

A set of VMs migrates

Migration policy determines the amount of time to reassign VMs and hence impacts performance

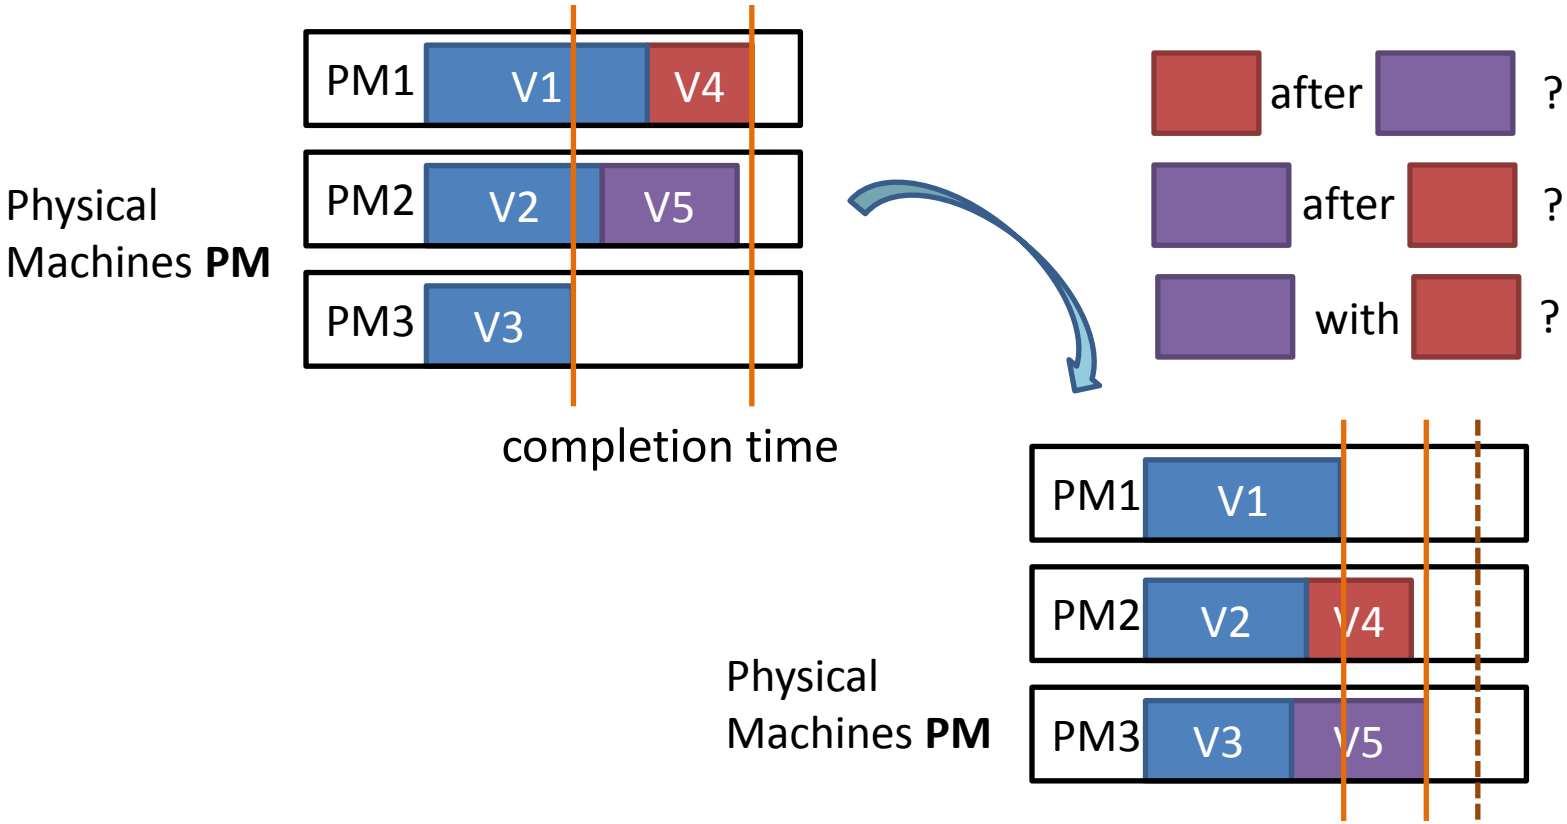

VM Migration

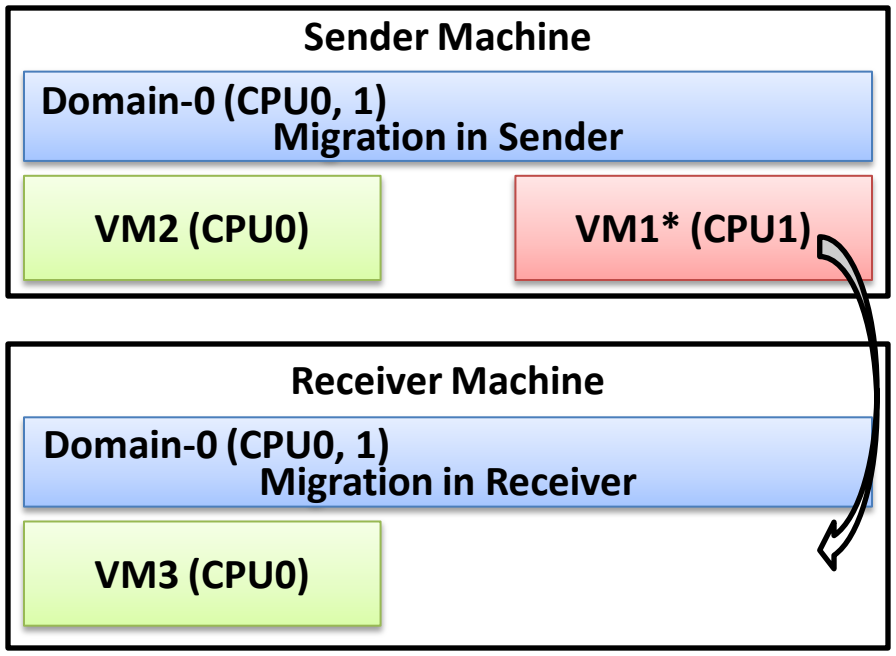

VM1* migrates

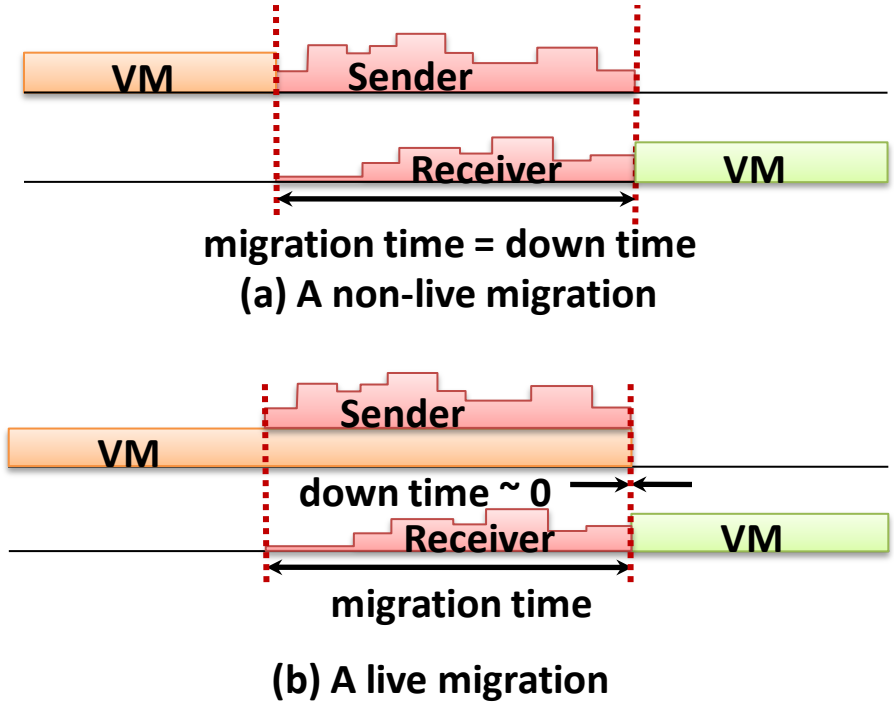

Types of VM migrations

Cost of VM Migration

When a set of VM migrates, how do we minimize T while bounding β?

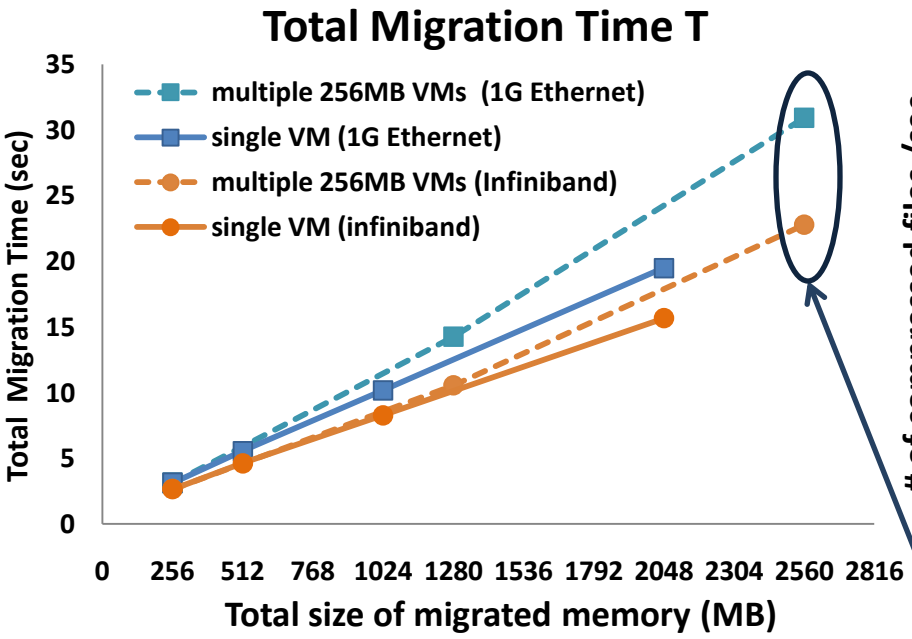

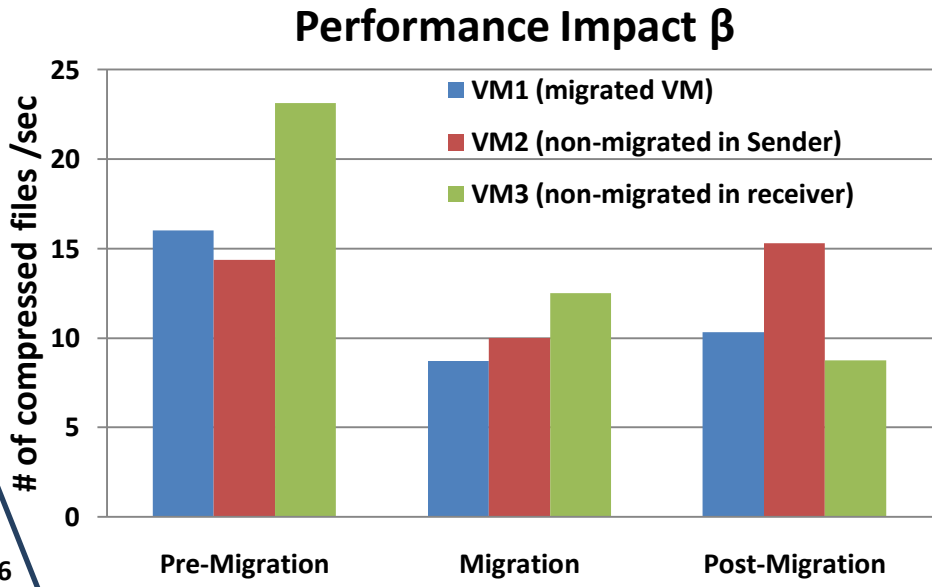

Workload : Compressing 256KB files.

Migration policy decides total migration time

Migration impacts performance

Slightly faster than ten sequential migrations (30.9sec < 10x3.1 sec, 22.8 sec< 10x2.6sec) , but with greater performance impact

Minimizing time T while bounding performance variation β

Desired is an accurate estimator of T and β
when **multiple jobs** contend for **multiple resources** *

*A generic model for *n* job and *m* resources has been developed.

Performance of Shared Systems with Multiple Resources

Total completion time may ***not*** be linear to individual completion times

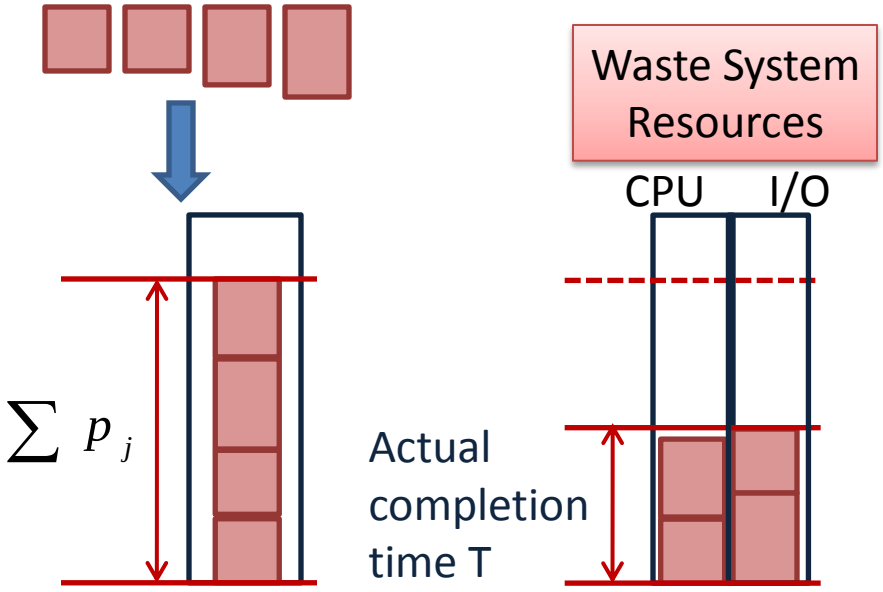

Bin packing or scheduling algorithms use ***linear*** relation

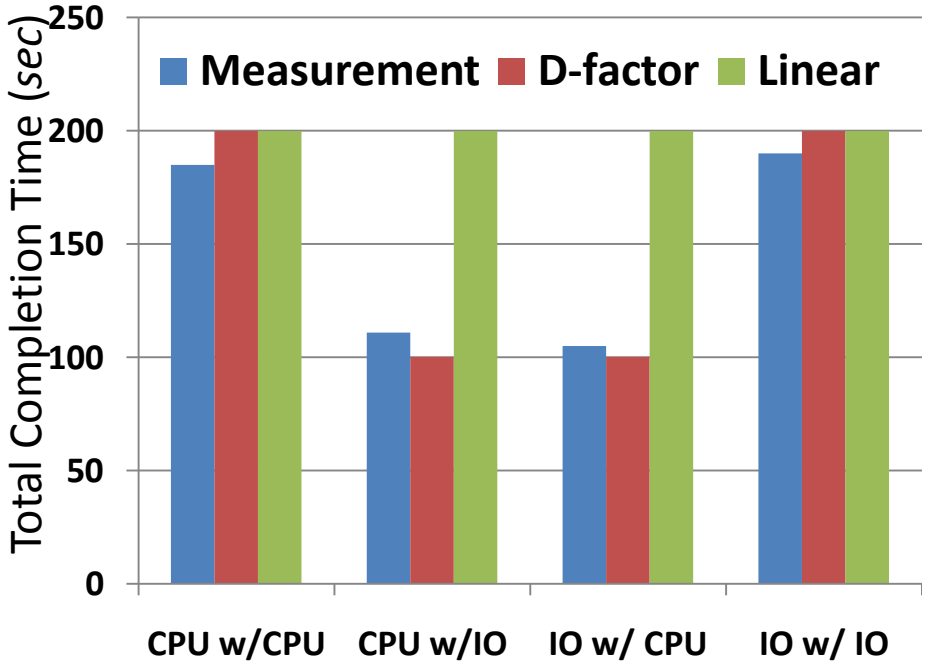

Both CPU job and I/O job take 100 *sec* without presence of other workloads

Estimating Completion Time *T*

Assume *two* 2-resource-busy jobs
with their loading vectors, access *probability* of each resource, $\boldsymbol{p_i} = (p_i, 1\text{-}p_i)$*.
Then, expanded completion time T of each job is given by

Without other jobs

$$T_1 = \tau_1(1 + p_1 p_2 + (1 - p_1)(1 - p_2)) \le \tau_1 + \tau_2$$
$$T_2 = \tau_2(1 + p_1 p_2 + (1 - p_1)(1 - p_2)) \le \tau_1 + \tau_2$$

From original completion time

Linear estimation

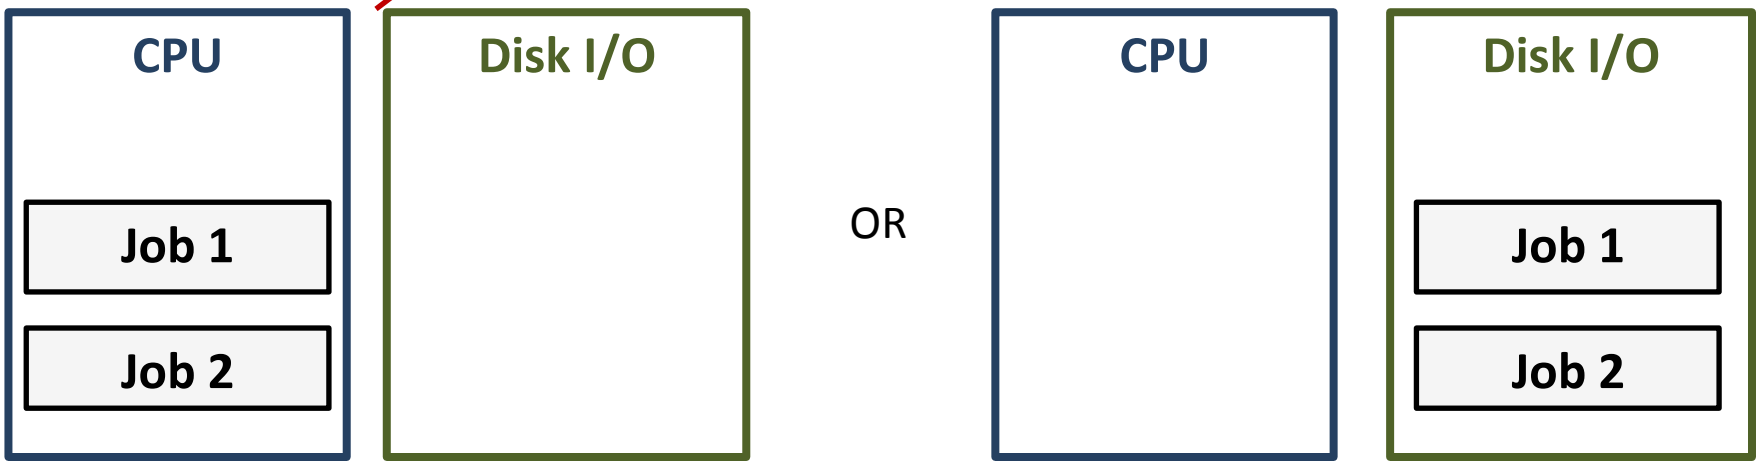

**Without* resource monitoring, loading vectors can be constructed (Algorithm 1)

Estimating Performance Variation β

Consider slow-down of jobs in the system as the performance variation

Given τ_j, define performance impact β by

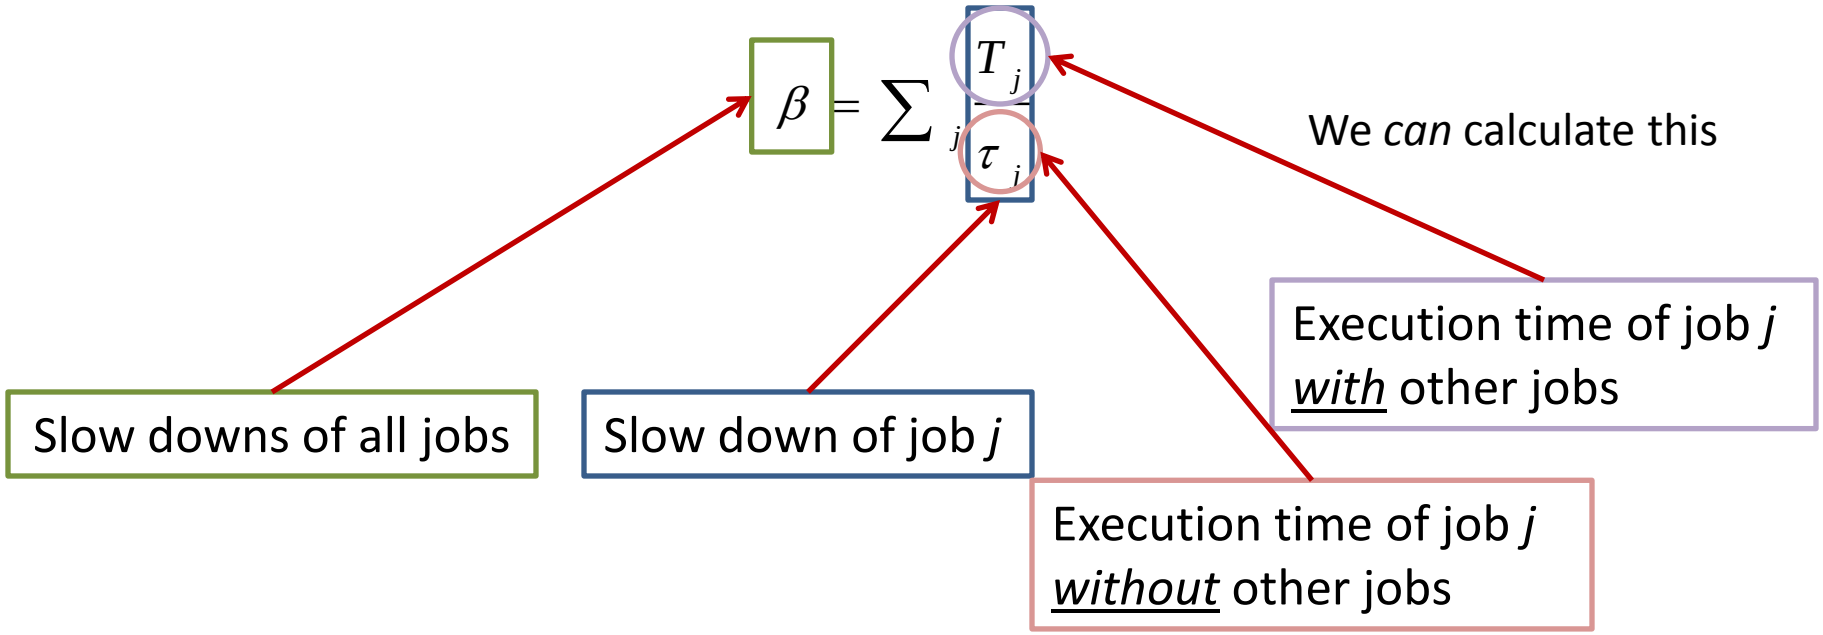

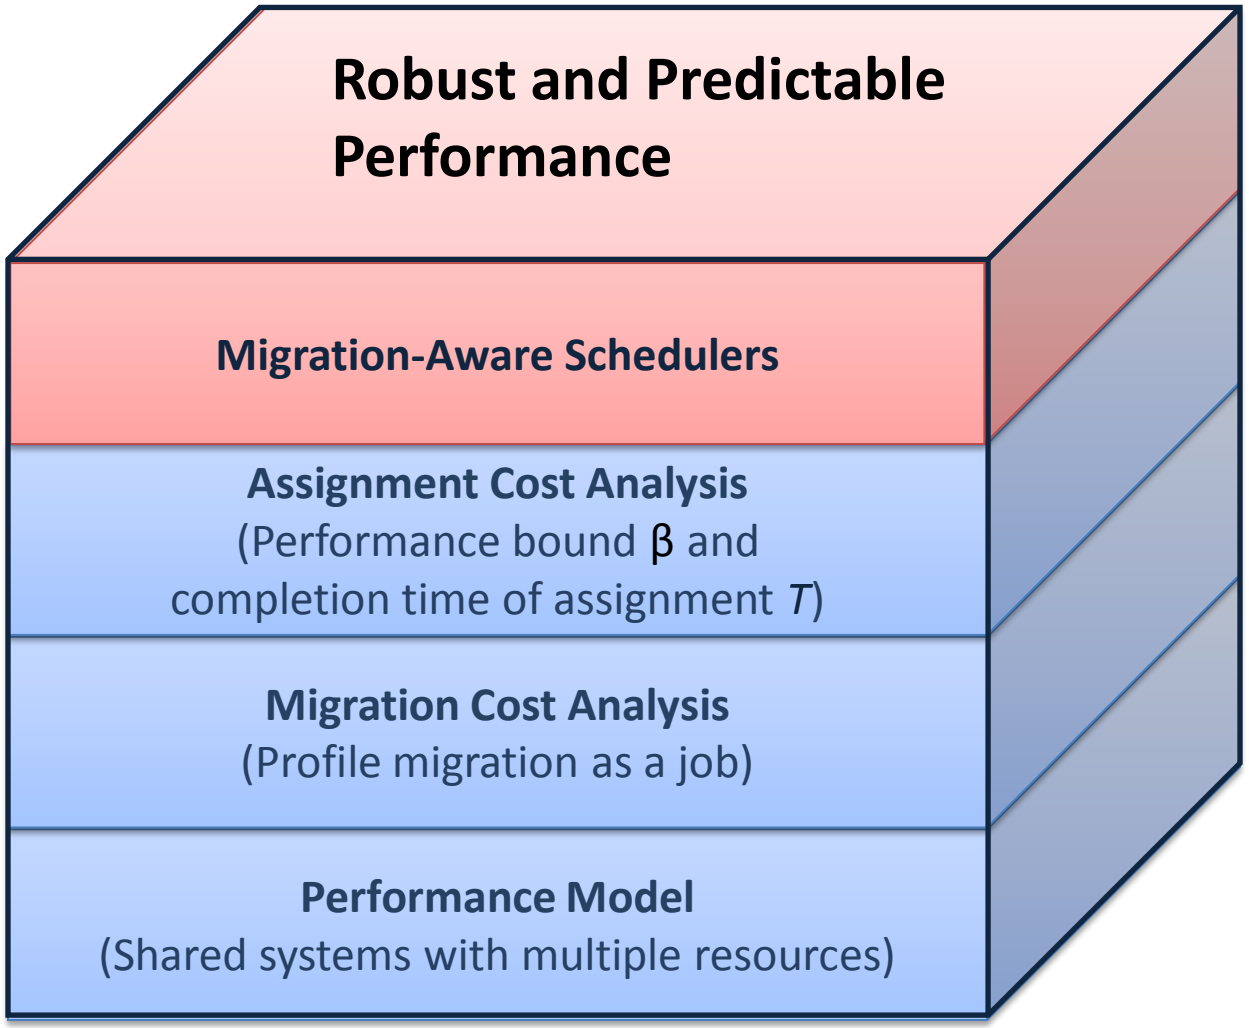
Robust and Predictable Performance
Migration-Aware Schedulers
Assignment Cost Analysis
(Performance bound β and completion time of assignment T)
Migration Cost Analysis
(Profile migration as a job)
Performance Model
(Shared systems with multiple resources)

Migration, Assignment, and Scheduling of Jobs in Virtualized Environment

Seung-Hwan Lim
(seulim@cse.psu.edu)

PENNSTATE

OAK RIDGE National Laboratory