Adaptive Virtual Machine Management in the Cloud: A Performance-Counter-Driven Approach

Chen Liu and Gildo Torres • Department of Electrical and Computer Engineering • Clarkson University • {cliu, torresg}@clarkson.edu

Problem:

- Cloud computing and virtualization technologies have become pervasive
- Co-running virtual machines (VM) introduce contention on shared resources
- **OS scheduler is unaware of potential contention on processor resources among concurrent VMs**

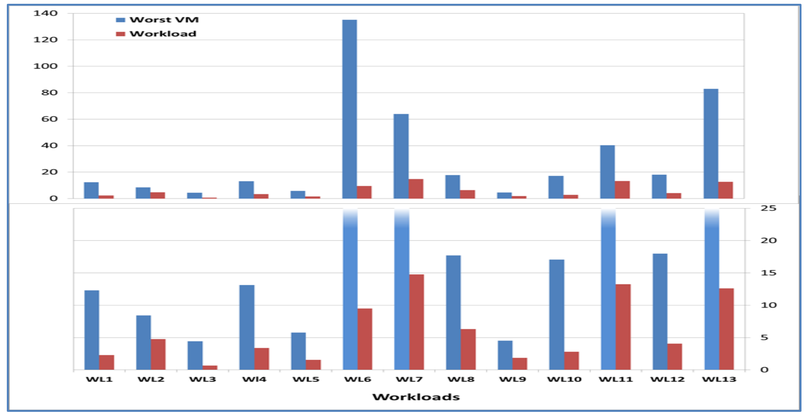

Key Idea:

We propose to adaptively assist the OS scheduler to monitor and manage the co-running VMs:

- Use **Hardware Performance Counters** (HPC) to collect runtime information of VMs
- Construct **statistical model** to represent the phases of execution of each VM
- Assist **Linux scheduler** by scheduling/migrating VMs based on dynamic evaluation of resource demands

Benefits:

- **Dynamic model** to capture VM's runtime behavior
- Manage VM-to-Core mapping in an **adaptive and autonomous** fashion

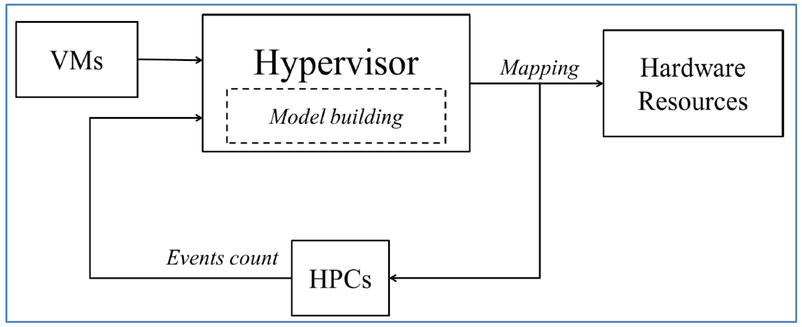

Technical Approach:

- **Platform:** Kernel–based Virtual Machine (KVM)
- **Hardware Events:** Monitor Last-Level Cache (LLC) misses, L1 misses, and committed instructions
- **Runtime Model:** Regression model based on memory access behavior to capture the phase change of VMs

Proposed Scheme:

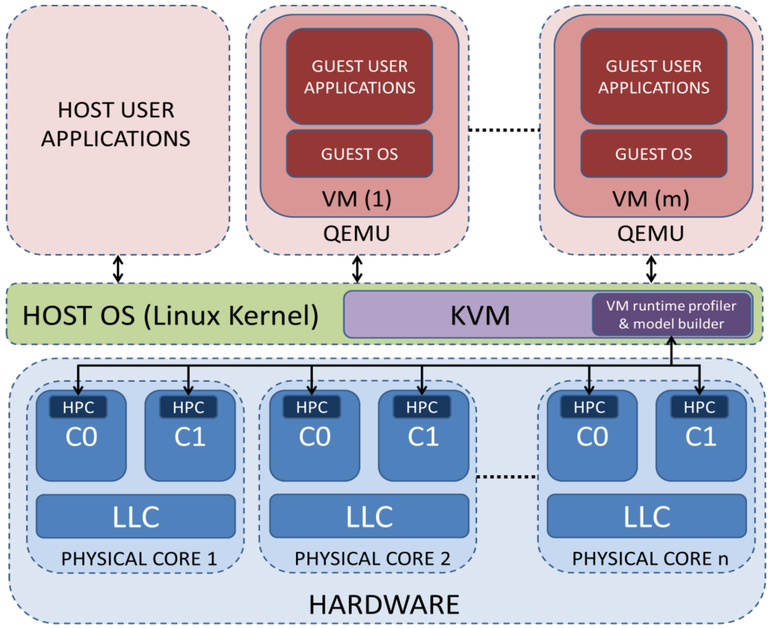

The profiler keeps track of each VM by monitoring their hardware counters

General Implementation:

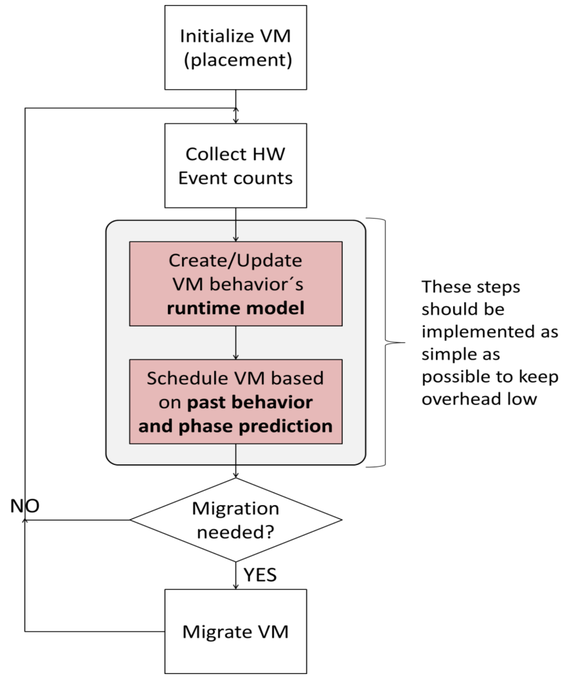

Expected Results:

- Mitigating resource contention among VMs
- Improvement of overall system throughput

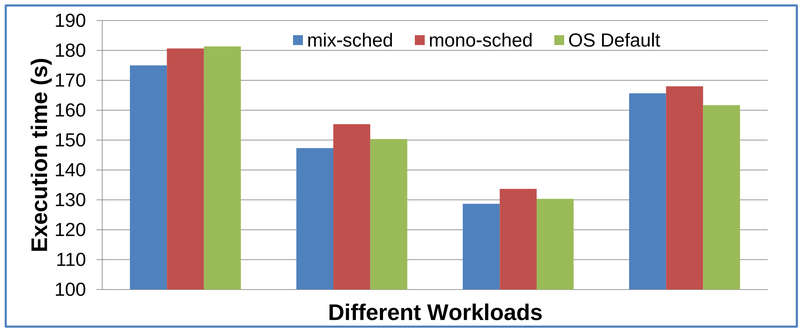

Future Work:

- Perform dynamic phase detection and prediction
- Machine learning-based online performance model