Implications of Cache Asymmetry on Server Consolidation Performance

Presenter: Omesh Tickoo

Padma Apparao, Ravi Iyer, Don Newell

*Hardware Architecture Lab

Intel Corporation

Outline

- Server Consolidation
- Asymmetric Caches
- Performance Implications
- Measurement-Based Analysis
- Conclusions / Future Work

Server Consolidation

- **Motivation**
  - Virtualization and consolidation are a growing trend in datacenters
  - Majority of servers expected to run consolidated workloads within few years

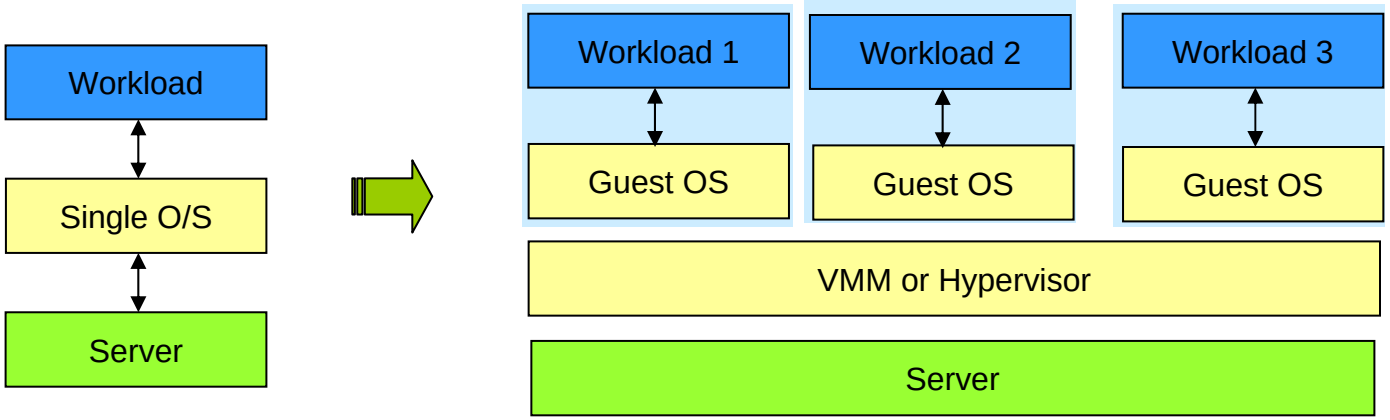

- **Problem**
  - Performance analysis of consolidation scenarios is challenging
    - Different virtualization overheads depending on VMM & platform virtualization support
    - Resource contention (core, cache, memory, etc) between VMs affects performance
- **Focus**
  - Server consolidation performance as a function of cache contention & asymmetry

Why study asymmetry?

- CMP platforms today have symmetric caches
  - But space in cache is asymmetrically allocated depending on demand from virtual machines

    => Virtual Asymmetry

- Future CMP platforms may have asymmetric caches
  - Asymmetry to reduce cache space domination of die area
  - Asymmetry due to process variability / faults

    => Physical Asymmetry

Cache Asymmetry

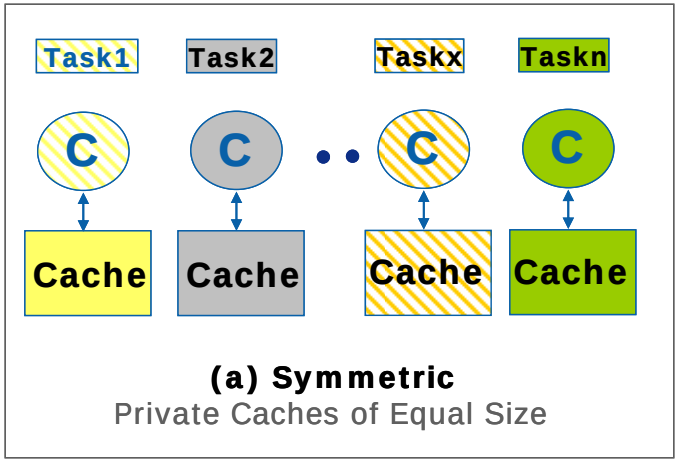

(a) Symmetric
Private Caches of Equal Size

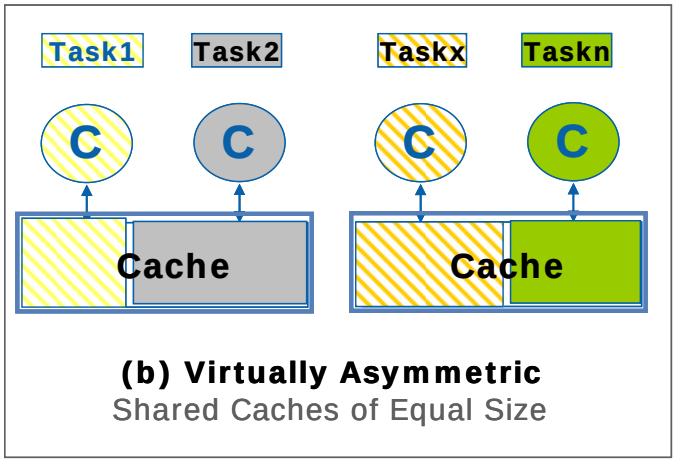

(b) Virtually Asymmetric
Shared Caches of Equal Size

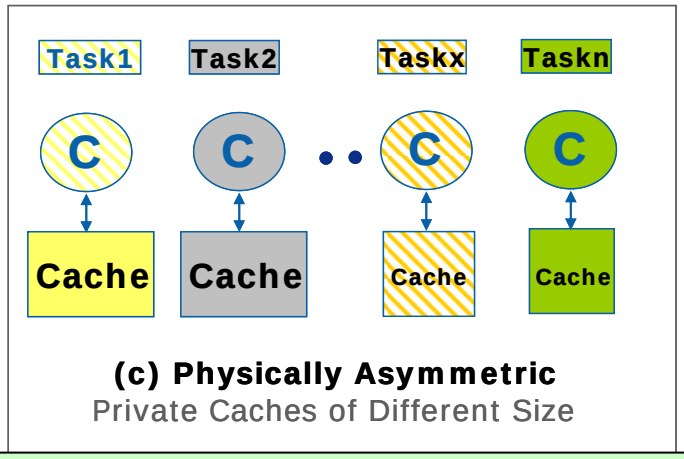

(c) Physically Asymmetric
Private Caches of Different Size

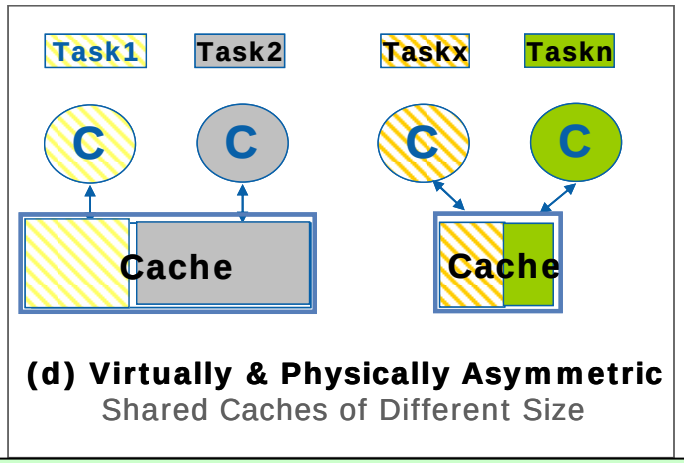

(d) Virtually & Physically Asymmetric
Shared Caches of Different Size

What are the implications on server consolidation performance?

Consolidation Benchmark

- vConsolidate

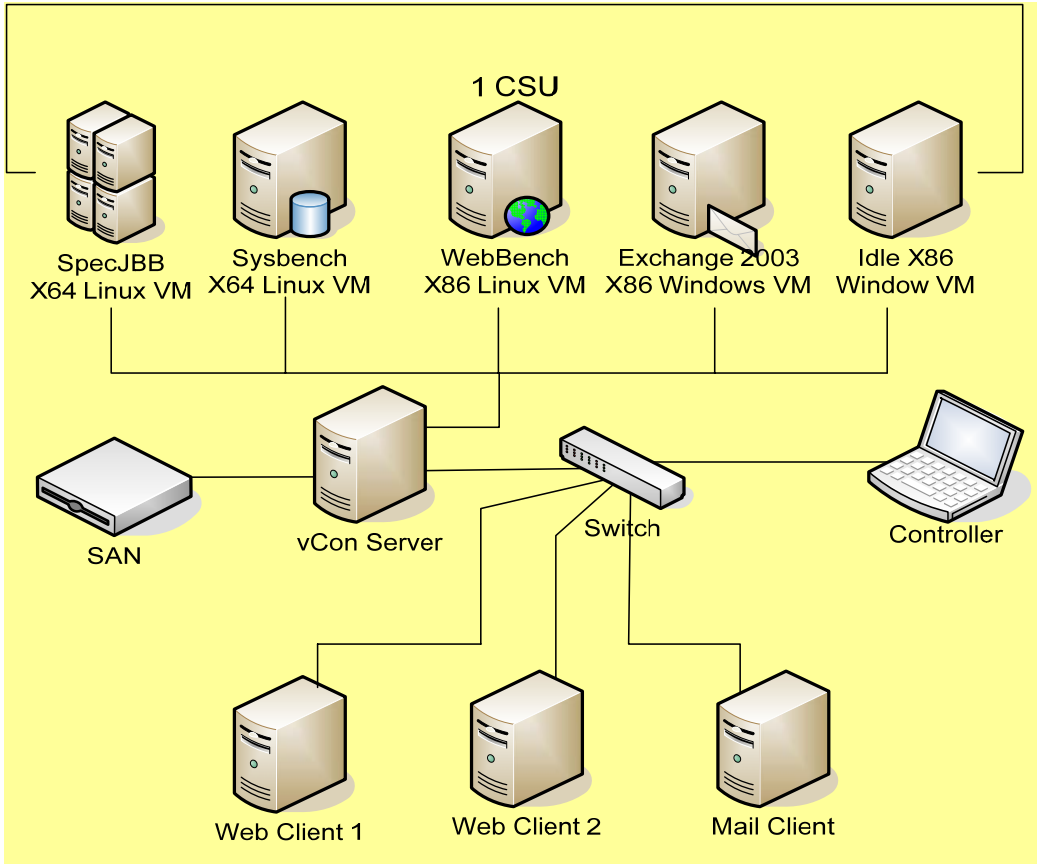

5 VMs
-- SPECjbb VM
-- Sysbench VM
-- Webbench VM
-- MailServer VM
-- Idle VM

VM/Workload	Vcpus Configuration	Memory Configuration in MB
Java/SPECjbb (bops/sec)	2	2056
Database/Sysbench (Tx/sec)	2	1544
Web/Webench (Tx/sec)	2	1544
Mail/Exchange (hits/sec)	1	1544
Idle	1	418

Platform Configuration

Hardware

- Intel Xen 5400 series
  - Quadcore per socket
  - 6MB+6MB $ per socket
  - Used 4MB, 3MB, 2MB cache configs also to create physical asymmetry

VMM

  - Xen 3.1

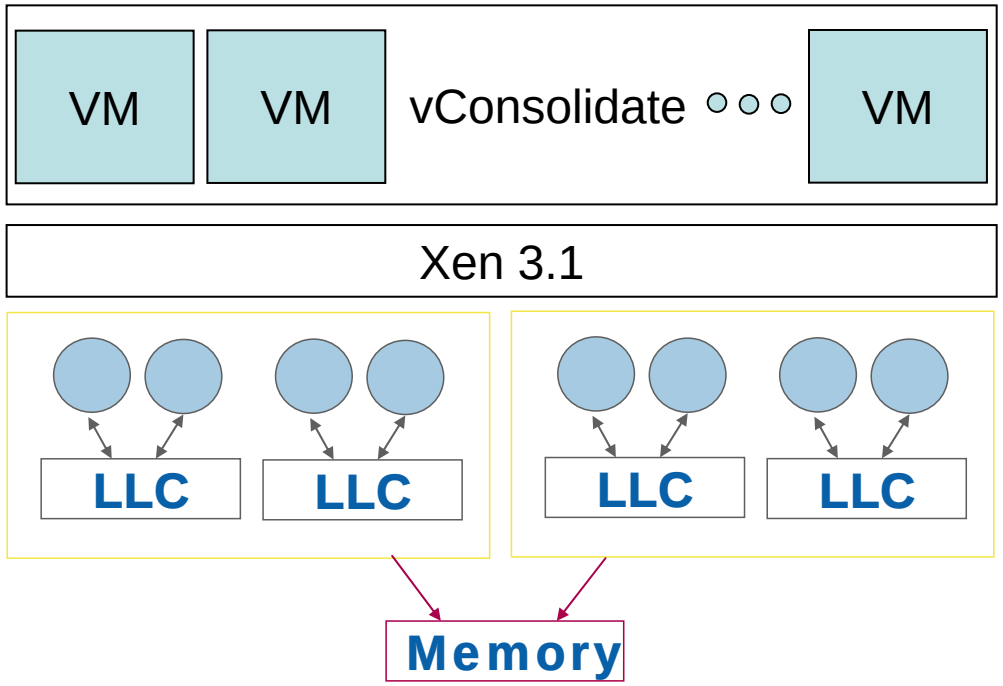

Analyzing Implications

- Four Key Configurations
  - 1 Virtual Machine
    - On physically symmetric cache
    - On physically asymmetric cache
  - Multi-Virtual Machine
    - On physically symmetric cache
      - But virtually asymmetric
    - On physically asymmetric cache
      - But virtually asymmetric also

1VM / Symmetric Caches

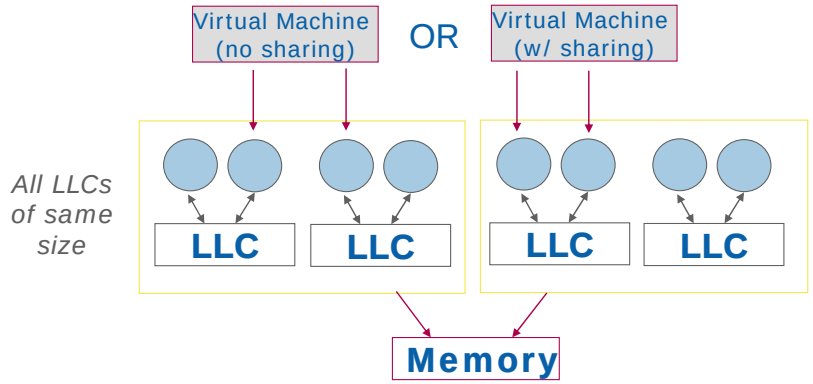

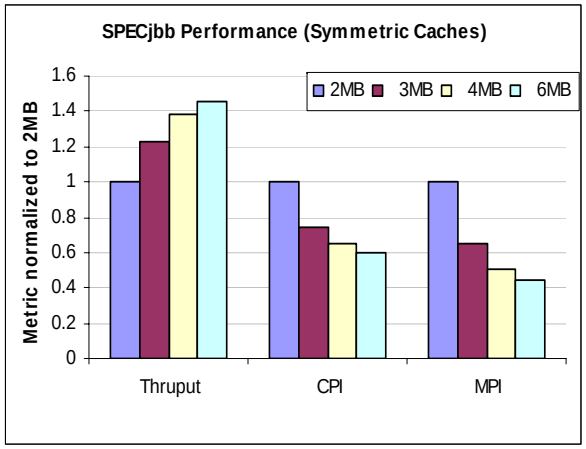

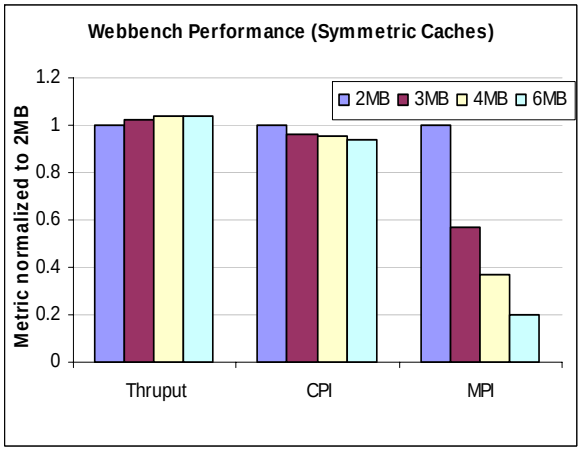

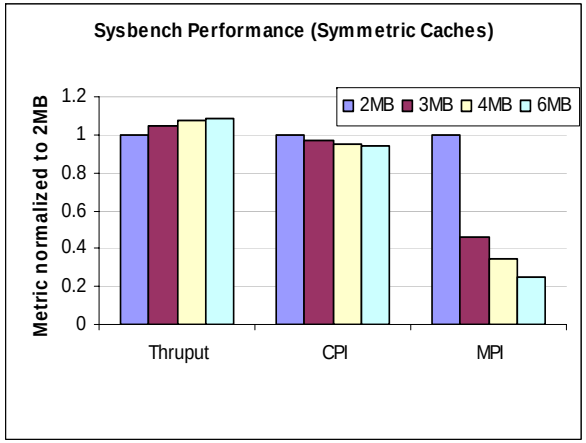

SPECjbb2005 most sensitive to cache – 50% perf improvement from 2MB to 6MB
Sysbench and Webbench show less than 10% improvement

Multi-VM / Virtually Asymmetry

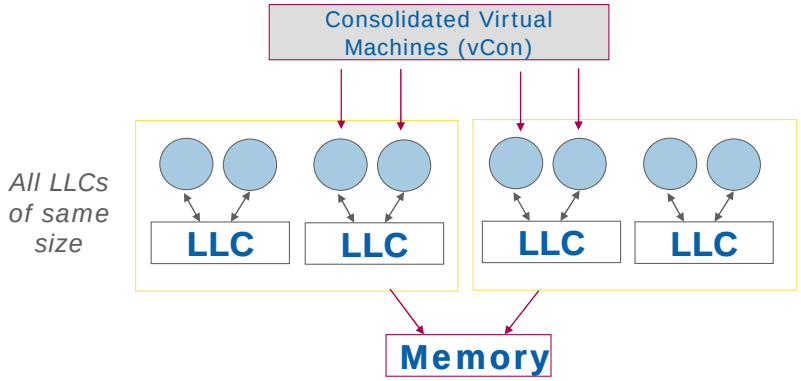

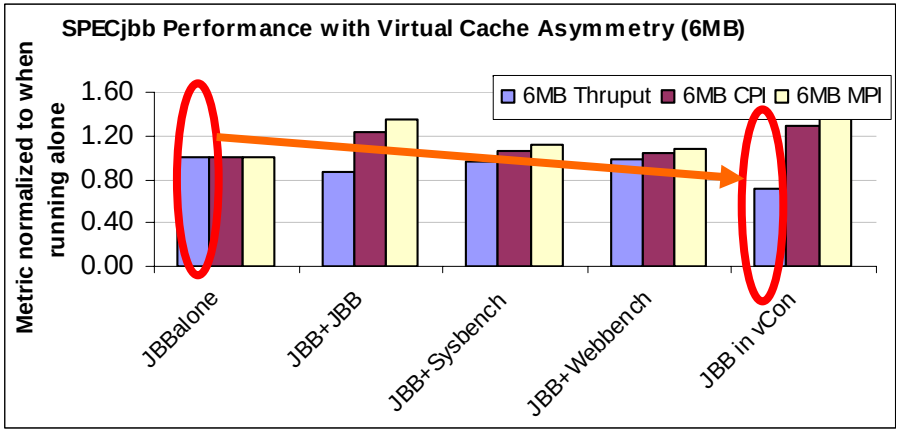

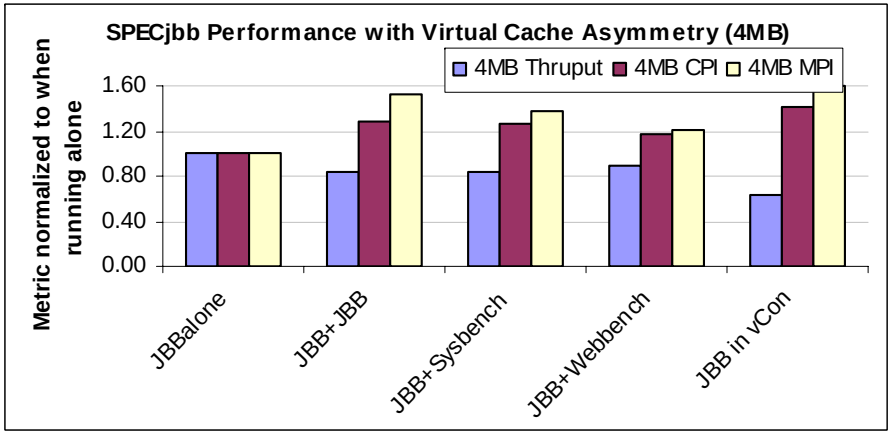

Consolidation causes causes ~30% loss in performance
Cache Interference => 20%
Core Inteference => 9%

1VM / Physical Asymmetry

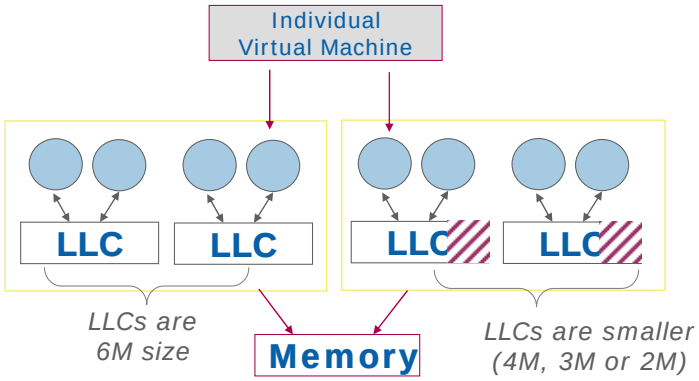

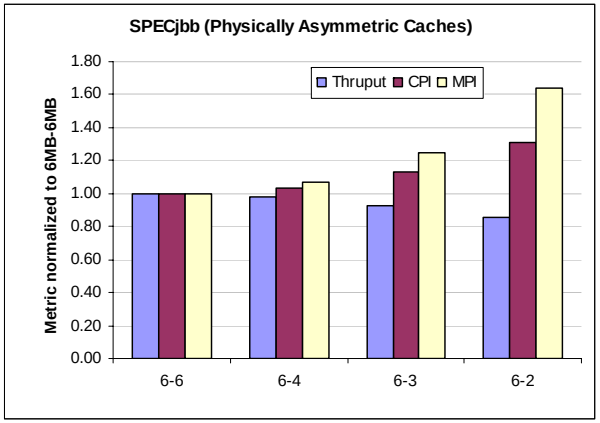

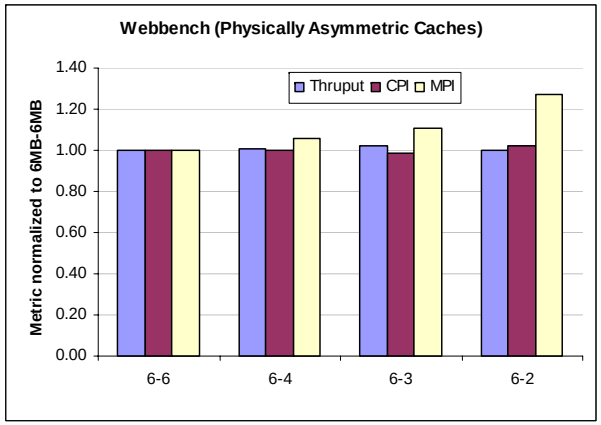

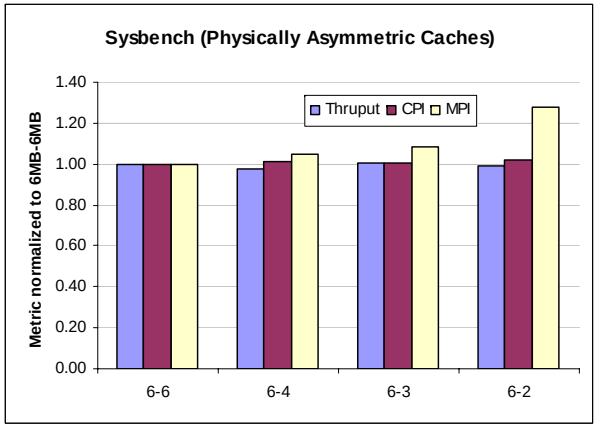

SPECjbb2005 is affected the most
Sysbench and Webbench are not affected much

Multi-VM / Virtual+Physical Asymmetry

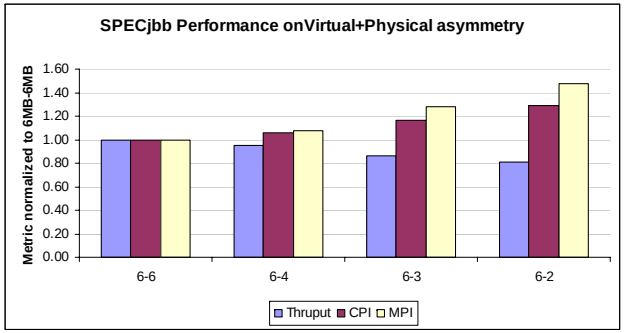

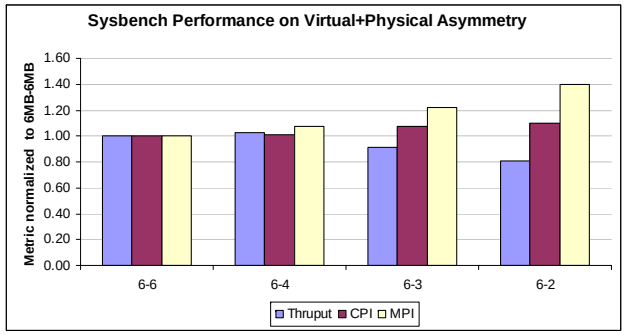

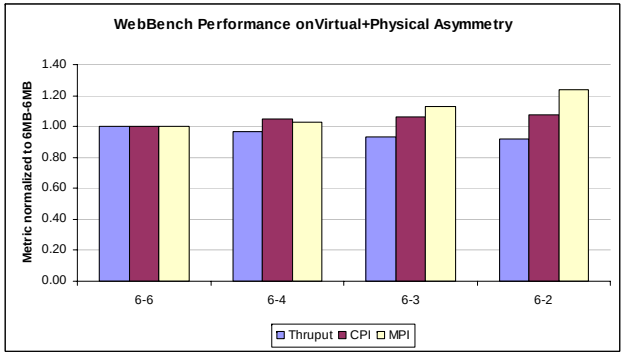

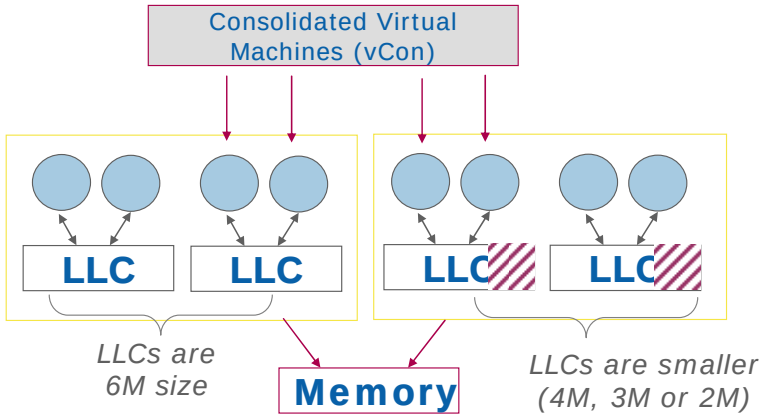

SPECjbb is affected the most (as expected)

Sysbench and Webbench are not affected much

Opportunity to move Sysbench and Webbench to smaller cache cores

=> can improve performance of SPECjbb?

Inferences

- Asymmetry-Aware Scheduling
  - Virtual Asymmetry
    - Monitor usage and interference
    - Modify VMM scheduler to take this into account
  - Physical Asymmetry
    - Monitor usage in large and small cores
    - Modify VMM scheduler to affinitize
      - Cache-sensitive VMs to large-cache-cores
      - Cache-insensitive VMs to small-cache-cores

Affinitization Experiment:
Affinitize one vcpu to large core
Leave the other vcpu floating

JBB	vcpu0 (affinitized to 6MB)	vcpu1 (floating)	% benefit
CPI	1.51	1.80	19%
MPI	0.0051	0.0070	39%
Sysbench	vcpu0 (affinitized to 6MB)	vcpu1 (floating)	% benefit
CPI	2.51	2.96	18%
MPI	0.0016	0.0020	25%
Webbench	vcpu0 (6MB cache)	vcpu1 (floating)	% benefit
CPI	2.59	2.88	11%
MPI	0.0023	0.0026	11%

Allows for detection of sensitivity for Improved scheduling

Summary

- Presented cache asymmetry
  - Symmetric
  - Virtual Asymmetry
  - Physical Asymmetry
  - Virtual + Physical Asymmetry
- Studied the implications of cache asymmetry on a consolidation workload
  - Using vConsolidate & asymmetric CMP platform
- Showed cache contention overheads and overall cache sensitivity
- Discussed the potential for asymmetry-aware scheduling