

Modeling Virtual Machine Performance: Challenges and Approaches

Omesh Tickoo
Intel Corporation
omesh.tickoo@intel.com

Ravi Iyer
Intel Corporation
ravishankar.iyer@intel.com

Ramesh Illikkal
Intel Corporation
ramesh.g.illikkal@intel.com

Don Newell
Intel Corporation
donald.newell@intel.com

ABSTRACT

Data centers are increasingly employing virtualization and consolidation as a means to support a large number of disparate applications running simultaneously on server platforms. However, server platforms are still being designed and evaluated based on performance modeling of a single highly parallel application or a set of homogenous workloads running simultaneously. Since most future datacenters are expected to employ server virtualization, this paper takes a look at the challenges of modeling virtual machine (VM) performance on a datacenter server. Based on vConsolidate (a server virtualization benchmark) and latest multi-core servers, we show that the VM modeling challenge requires addressing three key problems: (a) modeling the contention of visible resources (cores, memory capacity, I/O devices, etc), (b) modeling the contention of invisible resources (shared microarchitecture resources, shared cache, shared memory bandwidth, etc) and (c) modeling overheads of virtual machine monitor (or hypervisor) implementation. We take a first step to addressing this problem by describing a VM performance modeling approach and performing a detailed case study based on the vConsolidate benchmark. We conclude by outlining outstanding problems for future work.

Keywords

Virtualization, Consolidation, CMP, servers, performance analysis, measurement, modeling.

1. INTRODUCTION

Traditionally, server platforms have been designed and evaluated based on individual parallel applications or benchmarks (TPC-C, TPC-E, TPC-W, SPECjbb, SPECjappserver, etc) as the focus. Developing the performance analysis capability for such commercial server workloads was by itself a significant challenge since the benchmarks had complex behaviors, required multiple client/server systems and were difficult to run on full-system simulators. However, industry and academic researchers coped with this problem by developing scaled-down execution-driven simulations [31] for these workloads, by trace-driven simulations [2, 11, 12, 24] where possible and also by developing analytical models driven by measurement/simulation experiments. Now, the recent emergence of virtualization introduces another level of complexity to the problem of server modeling and performance analysis. As datacenters rapidly adopt virtualization as a means to consolidate multiple applications on a server, it becomes critical that the performance behavior of virtual machines is well understood. In addition, it also becomes important that we develop the ability to estimate virtual machine performance on a datacenter

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

server. In this paper, we discuss the challenges to virtual machine performance estimation and introduce potential approaches for appropriate metrics and modeling techniques for this purpose.

The first challenge of virtualization was addressing the lack of a virtualization benchmark that can be used for consistent and repeatable server performance analysis. Recently, there have been specific industry benchmarks that have been developed (VMmark [29], vConsolidate [4]) to address this issue. In addition, there is also a SPEC committee [31] that is working on defining its first virtualization benchmark expected to release sometime in the future. In this paper, we focus our case studies on vConsolidate to show the challenges and the potential approach of server virtualization modeling.

The key considerations when modeling the performance of virtual machines (VMs) can be summarized as follows:

(a) VM performance is not only dependent on its own characteristics, but also dependent on the interference caused by the other virtual machines running on the same platform with it. We need a method to capture the effect of these interactions.

(b) The above interference can affect the use of (i) shared resources (e.g. core, memory capacity) that are *visible* to the operating system or virtual machine monitor directly or through performance counters and (ii) shared resources (cache space, memory bandwidth, etc) that are *invisible* to the operating system since they are transparent resources managed by the hardware. The modeling approach needs to be aware of both visible and invisible resource interference.

(c) the specifics of virtualization technology (both hardware virtualization and software virtualization) and the scheduling disciplines adopted by the virtual machine monitor could be quite different on any given platform. The modeling approach needs to take into account the virtualization technology as well as the scheduling heuristics required.

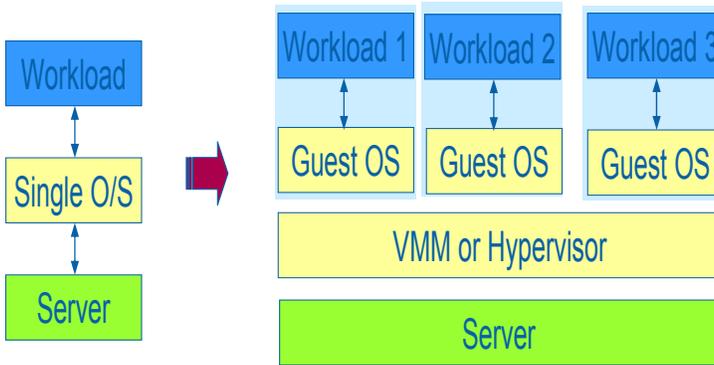
In this paper, we will expand primarily on the resource interference effects and describe how a modeling approach can potentially take into account both visible and invisible resource contention effects on performance. Specifically, the contributions of this paper are as follows:

(i) Using a virtualization/consolidation benchmark, we will show how VM performance depends heavily on visible & invisible resource interference caused by other VMs.

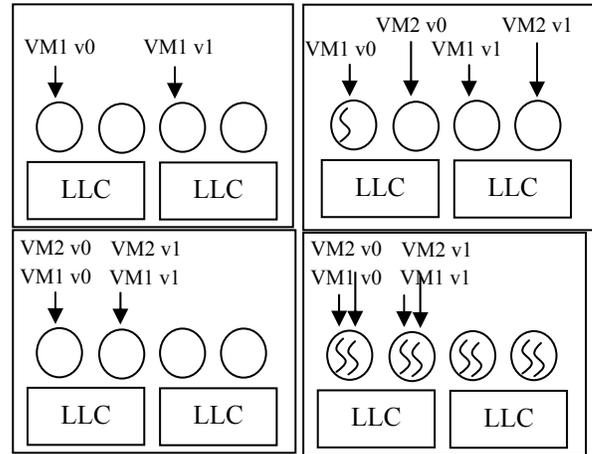
(ii) We will describe offline and online monitoring techniques needed to accurately characterize the behavior of multiple VMs sharing resources.

(iii) We will describe a potential modeling approach that employs either online or offline monitoring mechanisms to estimate the performance of the virtual machine.

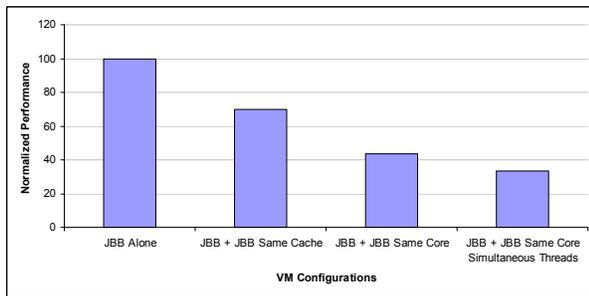
(iv) We will outline a detailed list of next steps required to achieve a complete model for VM performance estimation.



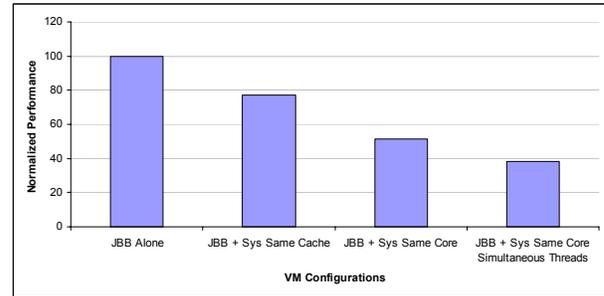
(a) From Dedicated Workloads to Consolidated Workloads



(b) VM Performance Scenarios



(c) SpecJBB Performance (1 copy vs 2 copies)



(d) SpecJBB Performance (Alone vs with Sysbench)

Figure 1: Performance Effects of Virtualization and Consolidation on CMP Servers

2. PLATFORM OVERVIEW

Our measurement-based evaluation was conducted on an dual-socket Intel server platform that has two processor sockets (Intel Xeon series) and is populated with 16GB of memory. Within each processor socket, there are 4 cores running at 3 GHz. Each pair of cores share a 4MB cache, adding up to a total of 8MB on each socket. Figure 3 illustrates this measurement platform configuration. On this platform, we run vConsolidate on top of Xen 3.3. Commonly available tools with Xen such as sar, xentop, xm info, xentrace have been used to get details such as cpu usage from each VM. To get architectural metrics such as CPI (cycles per instruction) and MPI (misses per instruction) we used a tool that reads CPU performance counters. To measure the per-thread contention effects we used a similar set-up as above with the latest Nehalem based processors offering multithreading capability.

In this paper, we study the implications of server consolidation by employing a recent server consolidation benchmark called vConsolidate [4]. The vConsolidate (vCon) benchmark consists of four key virtual machines (VMs): (a) a compute intensive workload/application, (b) a database workload, (c) a web server workload, and (d) a mail server workload. To emulate a real world environment, an idle virtual machine was also added to the mix. The compute intensive VM runs a modified version of SPECjbb2005 [17]. Typically SPECjbb2005 is a cpu intensive workload that consumes as much cpu as it possibly can. However, in vConsolidate environment, SPECjbb has been modified to consume roughly 75% of the cpu or so, by inserting in random sleeps every few milliseconds. The database virtual machine runs Sysbench [19], which is an OLTP workload executing transactions

either a MySQL or a MSSQL database. The web server VM runs Webbench [22] which uses either the Apache webserver or the Microsoft IIS webserver. The mail server virtual machine runs Microsoft Exchange workload that executes transactions on Outlook with 500 users logged in simultaneously.

3. VM PERFORMANCE EFFECTS

One of the key challenges when modeling the performance of virtual machines is that the application no longer has platform resources dedicated solely for its consumption. Figure 1(a) illustrates the transition from a dedicated execution (where one application used to run on a server previously) to a shared execution (where multiple guest OSes or VMs now run on the same server). Figure 1(b) illustrates a typical server platform with multiple cores and the following four scenarios for the execution of a virtual machine (say VM1) with two virtual cpus (v0 and v1):

- VM1 runs in dedicated mode with v0 running on c0 and v1 running c2. Since c0 and c2 have last-level caches of their own, these two virtual cpus do not share cache resources.
- VM1 runs in shared mode with VM2 running on sibling cores (virtual cpus of these VMs share cache). This causes cache contention between the VMs and affects performance.
- VM1 runs in shared mode with VM2 scheduled to run on the same core (time sliced scheduling). This causes core as well as cache contention between the two VMs and affects performance.
- VM1 runs in shared mode with VM2 scheduled to run on the same core but on simultaneous threads made possible by hyperthreading [30]. This causes core contention (because of thread contention within the core for shared pipeline resources) as

well as cache contention between the two VMs and affects performance.

To measure the performance in various scenarios, we ran experiments using virtual machines within the vConsolidate benchmark (specifically SPECjbb and Sysbench). We chose SPECjbb as VM1 as well as VM2 to collect the data shown in Figure 1(c). We also chose SPECjbb as VM1 and Sysbench as VM2 for the data collected in Figure 1(d). The first three configurations (illustrated in Figure 1b) were collected on a Core 2 Duo – based server platform with two cores on each die sharing a last-level cache. The performance effect of the last configuration (effect of sharing simultaneous threads) was measured on the latest Intel Nehalem-based server platform since the Core 2 Duo processor does not support threads, whereas future platforms are expected to.

As shown in the Figure 1(c) and Figure 1(d), the performance of the virtual machine is significantly affected by the sharing of cache and core resources. The performance can drop by as much as 20 to 30% due to cache contention alone. The performance further drops by another 30% when sharing cores. The performance can deteriorate further by ~25% when SMT threads are being simultaneously used on the same core. Overall, the performance of a virtual machine can vary from its performance during dedicated execution to its performance when running with other virtual machines either on a sibling core, a sibling thread, the same core or even the same thread. It should be noted that some of these resources are visible (e.g. core) and some are invisible (cache space, pipeline resources shared by simultaneous threads). Overall, the combined performance effect can be as high as 3X reduction due to resource contention. In addition, the performance depends on which virtual machine monitor you use.

4. VM PERFORMANCE MODELING APPROACHES

In the previous subsection, we showed how virtual machine performance can be significantly affected by resource contention with other virtual machine running along with it. In order to model the VM performance, it is important to accurately characterize all shared resource contention effects in the platform. In this section, we will describe how this can be achieved for two modeling modes:

(a) Offline modeling: This assumes that the workload performance can be measured on several platform configurations alone as well as pairwise with other virtual machines. It also assumes that the workload (each VM at a minimum) can be traced to run through core and cache simulators to collect behavioral information. Here, the intent is to fully calibrate a model in order to predict the performance of a virtual machine on a future server architecture or configuration.

(b) Online modeling: This assumes that no offline analysis is available and the characterization of the VM performance effects needs to be done online (on production servers with perhaps some brief initial runs on test configurations within the datacenter). Here, the intent is to develop a model that will predict the performance of a VM on a different platform (if it was migrated to that platform within the datacenter).

In addition, it is important to consider the following key visible and invisible resources that will be contended for:

(a) Visible Resources: Platform resources that are currently exposed (or visible) to the OS/VMM include core, memory

capacity and potential I/O devices. In this paper we focus primarily on core usage since that is a dominant factor out of these three.

(b) Invisible Resources: Platform resources that are currently transparent (or invisible) to the OS/VMM include shared cache space and core pipeline resources. In other words, the cache space occupied by one VM or another cannot be monitored today by an OS/VMM. Similarly, the pipeline resources used by one thread (running one VM) versus another thread (running another VM) cannot be observed today. The effect of cache and core contention (i.e. misses per instruction (MPI) and instructions per cycle (IPC) can however be obtained through performance monitoring counters).

The proposed approach for the VM modeling effort is the resource estimation of a virtual platform architecture that each virtual machine ends up with when running on a physical server platform. A virtual platform architecture (VPA) is defined as the set of resources used by (or solely made available to) the virtual machine of interest. The VPA resources include key resources such as number of cores, the core frequency and utilization, the cache space at each level and the memory frequency and bandwidth. In this paper, we focus on a VPA proof of concept that consists of one visible resource (the core) and one invisible resource (shared last-level cache space). In other words, the estimated VPA will produce (a) the number of cores and core utilization that a VM is given, (b) the cache space that a VM is given. Once the VPA core usage and the cache usage are available, the problem of estimating the performance of the virtual machine is mapped back to the simple problem of estimating the performance of an application running on a dedicated platform (with VPA as the dedicated platform).

A. VPA Core Utilization Estimation

Since the core is a visible resource, the VMM can easily track the VPA core utilization on the platform for every VMM. The cost of tracking the core utilization is negligible and provides an accurate accounting. However, in order to predict the core utilization on a different platform configuration, a simplistic approach to start with is as follows:

(i) Monitor a VM's core utilization when running alone (VM_x -Alone-Util)

- In an offline mode, this requires running the VM by itself on a platform and monitoring the core utilization using performance counters.
- In an online mode, this requires monitoring the VM core utilization and keeping track of the maximum core utilization (especially when no other VM is running).

(ii) Predict a VM's core utilization on a different consolidated platform by scaling down the above VM_x core utilization by the ratio of number of available physical CPUs over the total utilization of all VMs on the target platform (VM_x -Cons-Util)

$$VM_x\text{-Cons-Util} = \min \left[VM_x\text{-Alone-Util}, \frac{VM_x\text{-Alone-Util} * \text{PhysicalCPUs}}{VM_{all}\text{-Alone-Util}} \right]$$

This simple model allows for reasonable estimation of the VPA core utilization of a VM. In the future, we plan to extend this simple model by taking VMM scheduling heuristics into account.

Table 1: SPECjbb VPA Core Utilization

specJBB	4MB	2MB	1MB
Estimated	115%	124%	126%
Measured	122%	124%	121%
Error (%)	-5.58	0.39	3.91

For the vConsolidate measurement, we applied the estimation model to platforms configured with different cache sizes (1MB, 2MB, 4MB) as described earlier. We also instrumented the Xen VMM to measure the actual utilization of the SPECjbb virtual machine and compare it to the estimated utilization. Table 1 shows the utilization (in %) as well as the accuracy (computed as relative error). For example, the estimated utilization is 1.15 cores (115% utilization) for a 4MB configuration, whereas the measured utilization is 1.22 cores (122%) for the same configuration. With such a simplistic model, we find that the error is around 6% or less. The error is, in part, attributed to software scheduling that is not considered in the model above.

B. VPA Cache Space Estimation

Since the shared last-level cache is an invisible resource and cannot be directly monitored by OS/VMMs, we propose the following approach to estimating VPA cache space (per VM) as follows:

(a) **Profile the execution of a VM by capturing the percentage of time it runs with other VMs on sibling cores (sharing the same cache).**

- One approach to this is to instrument the VMM at schedule/deschedule points to keep track of where a VM was running. However, this is useful only in hindsight for understanding the performance effects of other VMs in a platform configuration already established.
- In order to predict the performance on a future platform configuration (with different VMs), it is possible to statistically approximate the percentage of time that a VM runs with other VMs.

$$P(\text{VM}_x + \text{VM}_y) = \frac{\text{VM}_y \text{Util}}{(\text{VM}_{\text{allCons}} \text{Util} - \text{VM}_x \text{Cons-Util})}$$

Table 2: Execution Probability of VMs running cores sharing a L2 Cache with SPECjbb

	specJBB	sysbench	webbench	MMB	Domain-0
4MB	18%	38%	32%	3%	9%
2MB	37%	19%	31%	3%	10%
1MB	36%	17%	32%	4%	12%

Table 2 presents the fraction of SPECjbb VM’s execution time that it spends running with another virtual machine contending on a shared cache. For example, a SPECjbb virtual cpu runs with a SysBench virtual cpu sharing the same shared cache for an estimated 38% of its overall execution in the platform configuration with 4MB of shared cache. On the other hand, it spends an estimated 18% of its overall execution time running with the other SPECjbb virtual cpu in the same configuration. It should be noted that we have validated these fractions (and found them to be within 2% error) by profiling the execution of vConsolidate on Xen.

(b) **Estimate the pairwise VPA cache space for a VM when it runs with another sibling VM (assuming two CPUs sharing a cache)**

- In the offline mode, it is possible to capture a trace for each VM and then run every pair of traces (for different VMs) through a shared cache model to create a table of VPA cache space usage per VM.
- In the online mode, it is not possible today to monitor the cache space usage directly. However, proposals like [24] that provide accurate cache space monitoring in future architectures might address this problem in the future. Instead,

the VMM can monitor the effect of cache interference by keeping track of the misses per instruction (collected through performance counters) as a function of the other VMs running on sibling cores. Even here, it is required that this estimation be done on a test configuration where all pairs of VMs are run simultaneously and measured.

Figure 2 shows the effect on normalized MPIs (cache Misses-per-Instruction) of SPECjbb and sysbench from pairwise workload runs. The normalization is done with respect to the MPI values when the workload was running alone. As shown in Figure 2, for a 4MB cache configuration, one SPECjbb virtual CPU’s MPI increases by as much as 62% when running with another SPECjbb virtual CPU, whereas it increases by 58% and 34% when running with Sysbench and Webbench virtual CPUs respectively.

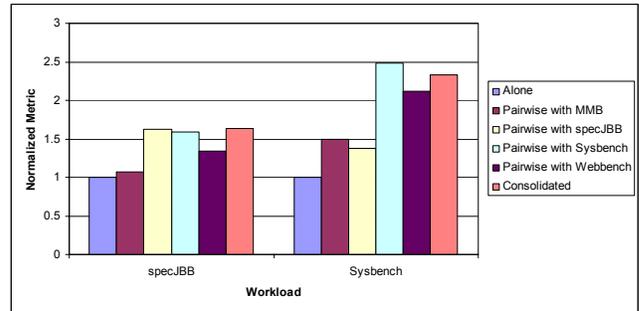


Figure 2: MPI Increase in Pairwise/Consolidated Execution

(c) **Estimate the consolidated VPA cache space that is a weighted average of the pairwise VPA cache space (with weights being the fraction of execution time).**

Table 3: Overall Cache Contention Effect for Consolidation

	Est. MPI	Measur. MPI
specJBB	0.0077	0.0085
sysbench	0.0041	0.0043

Table 3 shows that the estimated MPI (for SPECjbb and sysbench) comes reasonably close to the measured MPI during consolidation. The cache contention effect can also be translated to prediction of effective cache size by looking up the cache size based on offline MPI profiling experiments. For example, Figure 3 shows the MPI of SPECjbb running alone and in consolidated mode as a function of cache size. By looking up the miss rate curve (for 4MB), we can find the effective cache size (~3.5MB).

(c) **Use the estimated MPI to derive the CPI for the virtual machine**

The cache contention effects can now be translated into estimated CPI values (cycles per instruction) for each virtual machine in the consolidated environment using the performance data and based on the VMA core utilization. Due to space constraints we skip the exact procedure here. Table 4 shows that the estimated CPI values match the measured values with low error (~5%).

Table 4: CPI Estimation for SPECjbb

	SPECjbb Single VM	specJBB in vCon
Estimated CPI	1.549	2.231
Measured CPI	1.543	2.234

(d) **Estimate performance loss from cache contention (as a multiplier) as follows:**

$$\text{CacheContentionPerfMultiplier} = \frac{\text{VM}_x \text{-Alone-CPI}}{\text{VM}_x \text{-Cons-CPI}}$$

The factor calculated above represents the performance loss due to cache contention.

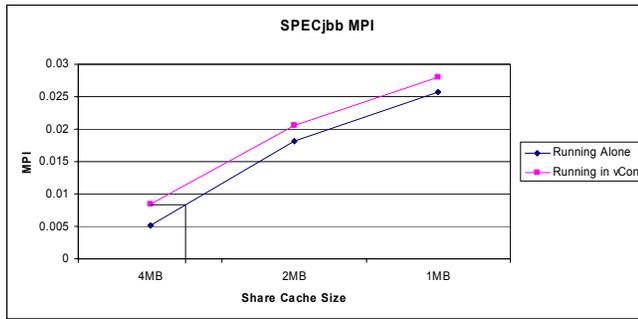


Figure 3: SPECjbb cache occupancy analysis

5. CONCLUSIONS AND FUTURE WORK

In this paper we presented the VM performance modeling challenges and highlighted that visible and invisible resource interference cause a significant performance degradation and need to be considered when modeling VM performance. We then proposed VPA estimation as an approach to modeling VM performance. Using a case study of vConsolidate (a server virtualization benchmark), we showed that VPA core and cache usage estimation is possible with offline models. With online models, only VPA core usage estimation is possible and cache usage estimation needs to be approximated as cache performance effects.

Future work in this area is as follows. We plan to extend the VPA modeling approach to include other resources (core pipeline resources, memory bandwidth, etc). We also plan to test out the modeling approach on a wider set of configurations and put together a analytical model that can be calibrated for this purpose. Last but not least, we would also like to study other virtualization benchmarks and VMMs to ensure that this approach proves valid for a wider range of workload and VMM implementations.

REFERENCES

- [1] Amazon Elastic Compute Cloud (EC2), <http://www.amazon.com/ec2/>
- [2] R. Iyer, et al, "Datacenter-on-Chip Architectures," *Intel Technology Journal*, 2007.
- [3] P. Barham, et al. Xen and the Art of Virtualization. In *Proc. of the ACM Symposium on Operating Systems Principles (SOSP)*, Oct 2003.
- [4] J. P. Casazza, M. Greenfield, K. Shi, Redefining Server Performance Characterization for Virtualization Benchmarking, *Intel Technology Journal*, Volume 10, Issue 03
- [5] D. Chandra, F. Guo, et al. Predicting inter-thread cache contention on a chip multiprocessor architecture", *11th High Performance Computer Architecture (HPCA)*, Feb 2005.
- [6] Google App Engine, <http://appengine.google.com>
- [7] HP Utility Computing, <http://h71028.www7.hp.com/enterprise/cache/308072-0-0-0-121.html>
- [8] Intel Xeon 5400 Series, <ftp://download.intel.com/products/processor/xeon/dc54kprodbrief.pdf>
- [9] Intel Virtualization Technology Specification for the IA-32 Intel Architecture, April 2005 <http://www.intel.com/technology/platformtechnology/virtualization/>
- [10] Intel Corporation, "World's first quad-core processors for desktop and mainstream processors," <http://www.intel.com/quad-core/>
- [11] R. Iyer. CQoS: A Framework for Enabling QoS in Shared Caches of CMP Platforms. *18th Annual International Conference on Supercomputing (ICS'04)*, July 2004.
- [12] R. Iyer, L. Zhao, et al., "QoS Policies and Architecture for Cache/Memory in CMP Platforms", *the ACM SIGMETRICS Conference on Measurement and Modeling of Computer Systems (SIGMETRICS)*, June 2007
- [13] S. Kim, D. Chandra, and Y. Solihin. Fair Cache Sharing and Partitioning in a Chip Multiprocessor Architecture. In *Proc. of 13th Int'l Conf. on Parallel Arch. & Complication Techniques(PACT)*, Sept 2004.
- [14] Microsoft Live Mesh, <http://www.mesh.com>
- [15] P. Ranganathan and N. Jouppi. Enterprise IT Trends and Implications on Architecture Research. In *Proc. of the 11th International Symposium on High Performance Computer Architecture (HPCA)*, Feb 2005
- [16] M. Rosenblum, T. Garfinkel. Virtual Machine Monitors: Current Technology and Future Trends. *IEEE Trans on Computers*, 2005.
- [17] SPECjbb2005, <http://www.spec.org/jbb2005/>
- [18] Sun network.com (Sun Grid), <http://www.network.com>
- [19] Sysbench <http://sysbench.sourceforge.net/>
- [20] Twenty Experts Define Cloud Computing, http://cloudcomputing.syscon.com/read/612375_p.htm
- [21] R. Uhlig, et al., "Intel Virtualization Technology," *IEEE Computer*, 2005
- [22] Webbench <http://cs.uccs.edu/~cs526/webbench/webbench.htm>
- [23] Xen: The Xen Virtual Machine Monitor. <http://www.cl.cam.ac.uk/Research/SRG/netos/xen/architecture.html>
- [24] L. Zhao et.al CacheScouts: Fine-Grain Monitoring of Shared Caches in CMP Platforms, PACT 2007
- [25] L. Cherkasova and R. Gardner, "Measuring CPU Overhead for I/O Processing in the Xen Virtual Machine Monitor," *Proceedings of the USENIX Annual Technical Conference*, April 2005.
- [26] N. Enright Jerger, et al, "Evaluation of Server Consolidation Workloads for Multi-core Designs," *IISWC-2007*
- [27] C. A. Waldspurger *Memory Resource Management in VMware ESX Server*, Proceedings of the 5th Symposium on Operating Systems Design and Implementation, Boston, Massachusetts, USA, December 9-11, 2002.
- [28] Microsoft Corporation www.microsoft.com
- [29] VMware Vmark <http://www.vmware.com/products/vmmark/results.html>
- [30] Intel Hyperthreading Technology. <http://www.intel.com/technology/platform-technology/hyper-threading/index.htm>
- [31] R. A. Hankins, T. A. Diep, M. Annavaram, B. Hirano, H. Eri, H. Nueckel, and J. P. Shen. Scaling and characterizing database workloads: Bridging the gap between research and practice. In *Proceedings of the 36th International Symposium on Microarchitecture. IEEE Computer Society Press*.
- [31] SPEC, <http://www.spec.org>