

Multicore Processor Performance Analysis - A Survey

Dali Ismail, dalieismail@gmail.com (A paper written under the guidance of [Prof. Raj Jain](#))

[Download](#)



Abstract

Multicore Central Processing Units (CPU) are becoming the standard for the current era of processors through the significant level of performance that CPUs offer. This includes multiple multicore architectures, different levels of performance, and with the variety of architectures, it becomes necessary to compare multicore architectures to make sure that the performance aligns itself with the expected specifications. This paper surveys the proper ways of selecting multicore CPU performance techniques, examples of metrics used in multicore CPUs performance studies, factors that affect the performance of multicore CPU and the benchmarks that target multicore CPU aspects of performance.

Keyword: Multicore CPU, Performance analysis, Performance evaluation, Benchmarks, Evaluation techniques, Server virtualization, Performance measurement.

Table of Contents:

- [1. Introduction](#)
 - [1.1 The Need For Multicore CPU](#)
 - [1.2 The Need For Performance Analysis](#)
- [2. Evaluating Multicore CPU Performance](#)
 - [2.1 Evaluation Techniques](#)
 - [2.2 Evaluation Metrics](#)
 - [2.3 Factors Affect The Performance](#)
- [3. Multicore CPU Benchmarking](#)
 - [3.1. General Multicore CPU Benchmarks](#)
 - [3.2. Multicore CPU Power Benchmarks](#)
- [4. Example of Multicore CPU Performance Analysis](#)
- [5. Summary and Conclusion](#)
- [List of Acronyms](#)
- [References](#)

1. Introduction

Moore's™ law of in-chip performance doubling has become the standard for progress in the computer and semiconductor industry. Memory capacity and CPU speeds are two of the many digital electronic devices expected by Moore's™ law to experience increased development approximately every 18 months [[Moore's Law](#)]. Multicore CPUs have evolved in this process and have become an

essential part of our daily life. They are implemented in variety of devices, and their performance varies considerably by application. This situation draws the attention of researchers to evaluate CPU performance rigorously to obtain higher performance at lower cost. The following section discusses the need for multicore processors and the need for performance analysis.

1.1 The Need For Multicore CPU

The high performance speed achieved by multi-processors (multiple CPUs on different chips attached to the same motherboard), produce undesirably high power consumption, and as a result, alternative research trends encouraged the production of multicore CPUs in order to reduce power consumption, while simultaneously increasing the processing speed. The architecture of multicore CPUs provided the hungry applications and devices, speed and performance with lower power consumption.

1.2 The Need For Performance

Performance analysis is a criterion that defines the performance of a system, and is required at every stage of the computer system life-cycle, to ensure high performance at a given cost. The demand for performance analysis was derived by radical changes in a number of elements including, a. The present day computer user who is more demanding than computer users 20 years ago. b. The popularity of computer technology, which is no longer a secret, resulted in an inundation in the computer market of different computer manufactures, each differing in performance. Such changes require performance analysis that meet user's demands and help select the best alternative which provides higher performance at given cost implementing trade-offs between what each technique provides and the required criteria in mind.

In this paper we will represent the different ways used to evaluate multicore CPUs performance, and the goal of this paper is to help understand the proper method for selecting a suitable evaluation techniques, metrics, and measure of multicore CPUs. In section 2, we will see the techniques used to evaluate multicore CPUs and the considerations that should be considered when selecting the techniques, and we will mention some of the metrics used in the analysis of multicore CPUs and the factors commonly affecting multicore CPU performance. In section 3, we will explore the benchmarks used in measuring multicore CPUs performance with an example on performance analysis of multicore CPUs provided in section 4.

2. Evaluating multicore CPU performance

This section explores the different methods used in evaluating multicore CPUs performance, with the metrics, analysis, and factors, varying performance based on requirements.

2.1 Evaluation Techniques

The first step in performance evaluation is to select the proper evaluation technique. The main techniques are: analytical modeling, simulation and measurement. In evaluating multicore CPUs performance, the techniques used are depending on different considerations. However, we cannot trust the result of one technique unless we validate that result with other techniques. For example, we cannot trust the analytical modeling technique without validating the result with simulation or measurement. That is, we require at the use of two techniques to get an accurate result. The considerations for selecting the appropriate technique listed in table 1 [[Jain91](#)].

Criterion	Analytical Modeling	Simulation	Measurements
Stage	Any	Any	Post Prototype
Time	Small	Medium	Varies
Tools	Analyst	Computer Language	Instrumentation
Accuracy	Low	Moderate	Varies
Trade-Off Evaluation	Easy	Moderate	Difficult
Cost	Low	Medium	High
Saleability	Low	Medium	High

Table 1: Evaluation Techniques Criteria

Table 1 shows the consideration in order of importance and the result may be wrong or misleading in all cases. For example, analytical modeling can be done at any stage of the system life-cycle. Although it takes smaller amounts of time than simulation and measurement because they vary with time, analytical modeling needs no tools for analysis; unfortunately it may give less accurate results. However, it is easy to do trade-off evaluation, as it is “cost less” in terms of capital compared to the other techniques. Unfortunately the saleability for products with just analytical modeling performance result is low. After knowing the evaluation techniques and the criterion for each technique the selection should be based upon the criterion and trade-offs that can be made between these criterions to get the required result.

2.2 Evaluation Metrics

Performance metrics are the measurements of the system performance or activity [[Metrics](#)], and the metrics selection depends upon the services provided by the system because metrics quantify the required output of the system. Metrics can be classified in three main classes:

- HB (Higher is Better): include metrics like throughput, higher values for this metric are preferred.
- LB (Lower is Better): include metrics like execution time, where lower values for this metrics preferred.
- NB (Nominal is Best): include metrics like power, with the median values being optimal in this case.

In this subsection we will list metrics that are used in evaluating multicore CPUs performance for applications.

- Throughput: Average rate of successful processes [[Sharma09](#)].
- Response time: The time that the user finishes the request and the system starts a response.
- Execution time: The time needed to complete program execution [[Monchiero06](#)].
- Energy: The power needed to run a program [[Monchiero06](#)].
- Memory bandwidth: the rate of data sustained from the CPU core to the RAM (Random Access Memory) [[Kayi07](#), [Memory](#)].

- Memory latency: the time delay between the memory controller signaled the memory module to access data from the RAM and the time the data become available for the memory module, also known as CAS (Column Address Strobe) latency [[Monchiero06](#), [CAS](#)].
- Percent of memory contention: the percentage among the cores trying to access the RAM at the same time [[Monchiero06](#)].

This list is just an example of metrics used in multicore CPUs performance evaluation. However, metrics are related to three criteria (1) time (2) rate and (3) resources, which are the criteria that can be measured and used to determine the system performance. Performance metrics varies upon the services provided by the system, in multicore CPU systems metrics can be chosen based on the purpose of the performance analysis and which type of performance requirements are required for the program to run efficiently by taking advantage of multicore CPU architectures.

2.3 Factors Affects The Performance

Factors are the performance parameters that we want to study to see their effects on the system. Factors also depend upon the required performance needed to utilize the CPU and get the expected outcome from it. In this subsection we are going to show examples of factors that affect multicore CPU performance.

- Memory: Memory architecture used and memory speed can affect the performance of the multicore CPU [[Sharma09](#)].
- Scalability: Affects performance based on the rate of increase of the workloads (i.e. tasks) [[Sharma09](#), [Carpenter07](#)].
- I/O bandwidth: Can affect the performance by utilizing the CPU cores which leads to more resources consumption without performance increasing [[Sharma09](#)].
- Inter-core communication: The interaction between cores in multicore CPU's can be implemented by various mechanisms, affecting overall CPU performance due to shared workloads between cores [[Sharma09](#)].
- Operating system (OS): OS is the manager for the CPU and it assigned tasks to cores based on a scheduling mechanism, affecting the multicore CPU performance [[Sharma09](#), [Pase05](#)].
- CPU clock speed: Clock speed has an impact on processor performance with slow clock speed reducing throughput [[Pase05](#)].
- Numbers of cores: Affect the CPU performance as multicore architecture workload is divided between the cores [[Pase05](#)].
- Cache coherent: Multicore architectures uses different caching mechanisms as the cache is shared among the cores, causing cache coherent to affect CPU performance [[Kavi07](#), [Kumar05](#), [Chang06](#), [Zheng04](#), [Yeh83](#)].

These are some examples of the factors affecting multicore CPU performance, and for the analysis of each factor under study we will define ways to optimize the performance by analyzing the effects of the factors and interpret the result to get the optimal expected performance.

3. Multicore CPU Benchmarking

Multicore CPUs are designed for a variety of applications, (virtualization, Games, and Embedded systems), with this kind of diversity, measuring the performance for multicore systems became a necessity to ensure that the performance delivered as required by the system. Different tools are used to measure multicore CPU performance. Tools like profiling, which is used to monitor and observe

system performance behavior rather than measuring. By measuring elapse time for the processes in multicore architectures having multiple threads, results in high-level information of processing speeds which increases performance as result of parallelization [[Prinslow11](#)]. Utilizing benchmark tools often results in better measurements that become more relevant and accurate to system profiling. In this section we will introduce benchmarking approaches used to measure multicore CPU performance.

3.1 General multicore CPU Benchmarks

Multicore CPU benchmarks must target the aspect of concurrency from the parallelism prospective, which can be represented by the throughput in data and computational workloads [[Levy09](#)]. This subsection will introduce benchmarks that been used to measure multicore CPU performance. In this subsection we are going to list two general benchmarks used to measure multicore CPU performance.

- SPEC CPU2006: Measures multicore processors by running multiple copies of the same program (which is determined by the test implementer) concurrently and determine the throughput of the test [[Domeika09](#)].
- EEMBC MultiBench software: A set of multi-context benchmarks used in embedded systems, uses MITH (Multi-Instance Test Harness) which creates workloads from different items in systems and can be portable (runs on different embedded system) [[Domeika09](#)].

3.2 Multicore CPU power benchmarks

More recently, following the evolution of mobile devices and computers, power efficiency has become an important aspect that has required new approaches to measure multicore CPU performance. Benchmark companies developed new benchmarks depending on power to measure the performance. In this subsection we are going to list two power benchmarks with can be considered as the industry standard power benchmarks.

- EEMBC EnergyBench: Uses average power used (in watts) and energy (in Joules per iteration) as a performance metric to measure the performance. Initial it runs a multicore EEMBC benchmark and provides power measurement simultaneously, measuring performance and power by monitoring the power rails on system board at execution time [[Domeika09](#)].
- BDTI Benchmark suite: Used to estimate or measure processor power efficiency. In estimating power, it uses consistent assumptions and conditions to give accurate estimations that are helpful in making decisions. In measuring performance it uses vendor benchmarks for the specific targets of measurement [[Domeika09](#)].

Benchmark results for multicore CPU performance depend on the test run by the benchmark to measure the performance of the multicore CPU for specific applications, and by defining the reasoning for measurements we can relate different multicore CPU performances to each other based on the benchmark used to utilize and measure the multicore CPU performance.

4. Exmample of Multicore CPU Performance Analysis

This section will introduce an example of performance analysis processes for multicore CPUs that will assist in selecting the proper CPU for a machine specification.

Server virtualization of Multicore CPU: Intel IT (Information Technology) team evaluated server performances based on three Intel multicore CPU servers (A Four-socket server based on Quad-Core

Intel Xeon CPU X7350 with 16 cores, a dual-socket server based on Quad-Core Intel Xeon CPU X5355 with 8 cores and a dual-socket server based on Intel Dual-Core Xeon CPU 5160 with four cores) [Carpenter07]. In comparing the performance of the multicore CPUs, the Intel IT team targeted the speed of the CPUs and the power efficiency. This became a major concern as mentioned earlier in section one in the paper [Carpenter07]. Due to CPU clock speed, runtime used to measure the performance on each CPU. The data was normalized. Furthermore, the normalized workload consists of VMs (Virtual Machines) and a copy of a synthetic CPU intensive DB application in each VM. W-M/Job (Watt-minute per job) metrics was utilized to measure CPU power efficiency with an increasing workload to test the scalability factors of the CPUs [Carpenter07]. The results from [Carpenter07] show that the three servers different levels of scalability in terms of power consumption. As the VMs number increased the run time remains constant until the workload equals the number of cores. After the number of VMs exceeds the number of cores, the run time begins to increase, figure 1 shows the result of the servers based on Intel multicore CPU run times.

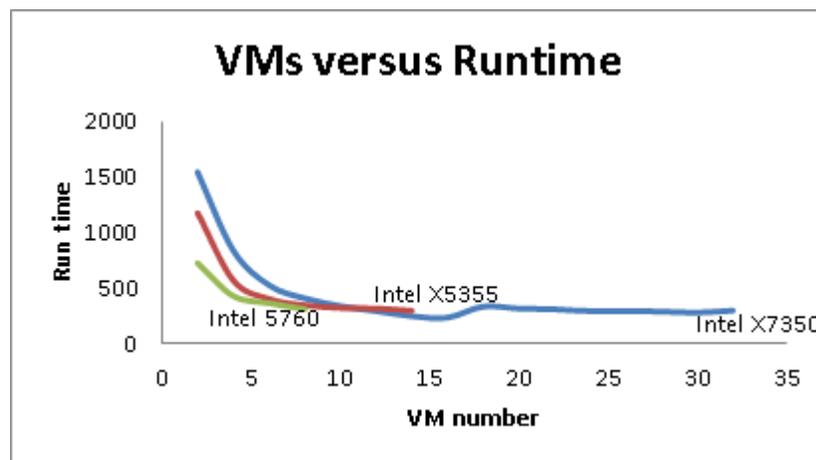


Figure 2: Run Time

We can observe that the Intel X7350 CPU run time is almost constant, until the number of VM reach 16, then it starts to increase. Alternatively, in the case of the Intel X5355, the run time starts to increase around 6 VMs, and around 4 VMs were running the Intel 5760. Another approach that is useful in performance evaluation of the CPUs is to measure the power consumption of the different CPUs based upon increasing workload to test the scalability [Carpenter07]. The result of the test shows that Quad-core Intel Xeon CPU X7350 based servers consumed more power than its alternatives due to the larger number of cores. Consuming at an average of 495 W (Watts) on 2 VMs running, 478 W for Quad-core Intel Xeon CPU X5355, and average of 330 W for Dual-core Intel Xeon CPU 5160. As the number of VMs increased, the servers became more power efficient which can be observed from figure 2. The Quad-core Intel Xeon CPU X7350 based servers with the maximum workload showed the power consumption per job decrease from the start and that is due to scalability of the CPU. Figure 2 shows the increasing of power consumption due to the number of VMs.

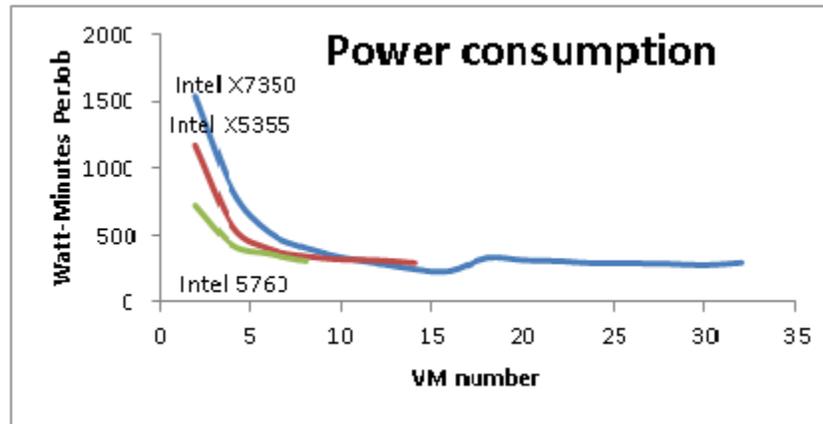


Figure 2: Power Consumption

This section introduced examples of how performance analysis works. By comparing the multicore CPUs, we can use the result to help us make proper decisions in terms of selecting the appropriate CPU for a required performance level.

5. Summary and Conclusion

As a result of Moore's law [Moore's Law], CPU performance is increasing rapidly. The number of cores on the chip increases at each release of a new generation of a CPU. With multicore CPU becoming not only faster but more power efficient depends on the required demand. To maintain relevant profiling of these systems, we need to evaluate the CPU depending on the workload we expect to process. We have addressed the techniques used to evaluate multicore CPU performance, metrics, factors, benchmark tools used to measure multicore CPU performance, and we provided an example of performance evaluation for multicore CPUs. Different approaches used in multicore CPU performance analysis depending on the purpose of the study. We are expecting to see new approaches in multicore CPU performance analysis as multicore CPU production increases to new levels.

List of Acronyms

- CPU - Central Processing Unit
- RAM - Random Access Memory
- HB - Higher is Better
- LB - Lower is Better
- NB - Nominal is Best
- CAS - Column Address Strobe
- OS - Operating System
- IT - Information Technology

References

1. [Jain91] Jain R., "The Art of Computer System Performance Analysis", Wiley-Interscience, New York, NY, April 1991, ISBN: 0471503361.

2. [Moore's Law] "Moore's Law", Wikipedia the free encyclopedia, http://en.wikipedia.org/wiki/Moore's_law.
3. [Metrics] "Performance Metrics", Wikipedia the free encyclopedia, http://en.wikipedia.org/wiki/Performance_metric.
4. [Sharma09] Sharma A. , Kumble M. , Moktali P. R. , Siri H. , "Performance analysis of Multicore Systems", Intel, March 2009, <http://software.intel.com/en-us/articles/performance-analysis-of-multicore-systems-4>
5. [Monchiero06] Monchiero M. , Canal R. , Gonzalez A. , "Design Space Exploration for Multicore Architectures: A Power/Performance/Thermal View", Proceedings of the 20th annual international conference on Supercomputers, Pages 177 - 186, ACM New York, NY, 2006, ISBN: 1-59593-282-8.
6. [Kayi07] Kayi A. , Yao W. , EL-Ghazawi T. , Newby G. , "Experimental Evaluation of Emerging Multi-core Architectures", Parallel and Distributed Processing Symposium, IEEE International, March 2007, <http://cecs.uci.edu/~papers/ipdps07/pdfs/PMEO-PDS-21-paper-1.pdf>.
7. [Memory] "Memory Bandwidth", Wikipedia the free encyclopedia, http://en.wikipedia.org/wiki/Memory_bandwidth
8. [CAS] "CAS Latency", Wikipedia the free encyclopedia, http://en.wikipedia.org/wiki/CAS_latency
9. [Carpenter07] Carpenter R. E. , "Comparing Multi-Core Processors for Server Virtualization", Intel Corporation, August 2007, http://www.multicoreinfo.com/research/papers/whitepapers/multicore_virtualization.pdf
10. [Pase05] Pase D. M. , Eckl M. A. , "A Comparison of Single-Core and Dual-Core Opteron Processor Performance for HPC", Technical report, IBM Developer Works, IBM Corporation, 2005, ftp://ftp.support.lotus.com/eserver/benchmarks/wp_Dual_Core_072505.pdf
11. [Kumar05] Kumar R. , Zyuban V. , Tullsen D. M. , "Interconnections in Multi-core Architectures: Understanding Mechanism, Overheads and Scaling", ISCA '05 Proceeding 32th annual international symposium on Computer Architecture, Pages 408 - 419, IEEE Computer Society, Washington, DC, 2005, ISBN: 0-7695-2270-X.
12. [Chang06] Chang J. , Sohi G. S. , "Cooperative Caching for Chip Multiprocessors", ISCA '06 Proceeding 33th annual international symposium on Computer Architecture, Pages 264 - 276, IEEE Computer Society, Washington, DC, 2006, ISBN: 0-7695-2608-X.
13. [Zheng04] Zheng Y. , Davis B. T. , Jordan M. , "Performance Evaluation of Exclusive Cache Hierarchies", ISPASS '04 Proceedings of the 2004 IEEE International Symposium on Performance Analysis of Systems and Software, Pages 89 - 96, IEEE Computer Society, Washington, DC, 2004, ISBN: 0-7803-8385-0.
14. [Yeh83] Yeh P. C. C. , Patel J. H. , Davison E. S. , "Shared Cache for Multiple-Stream Computer Systems", IEEE Transactions on Computers, Vol. 32, Issue 1, pages 38 - 47, IEEE Computer Society, Washington, DC, 1983.
15. [Prinslow11] Prinslow G. , "Overview of Performance Measurement and Analytical Modeling Techniques for Multi-core Processors", 2011, <http://www.cse.wustl.edu/~jain/cse567-11/ftp/multicore/index.html>
16. [Levy09] Levy M. , "Multicore Processor Performance Analysis: Revealing the Truth", October 2009, <http://eecatalog.com/multicore/2008/10/09/multicore-processor-performance-analysisrevealing-the-truth/>
17. [Domeika09] Domeika M. , "Evaluating the Performance of Multi-Core Processors", September 2009, <http://www.embedded.com/design/embedded/4008794/Evaluating-the-performance-of-multi-core-processors-Part-1>

Last Modified: April 24, 2013

This and other papers on latest advances in Performance Analysis and Modeling are available on line at <http://www.cse.wustl.edu/~jain/cse567-13/index.html>

[Back to Raj Jain's Home Page](#)