WWW.MWFTR.COM

EECE 417 Computer Systems Architecture

Department of Electrical and Computer Engineering Howard University

Charles Kim

Spring 2007

Computer Organization and Design (3rd Ed)

-The Hardware/Software Interface

by

David A. Patterson John L. Hennessy

Chapter 4 - Part A

Assessing and Understanding Performance

- Measure, Report, and Summarize
- What determines the performance of a comp
 - Is hardware performance the key?
 - What is the role of software in performance?
 - one tough job
- Make intelligent choices
- See through the marketing hype
- Questions:





•Which of these airplanes has the best performance?

<u>Airplane</u>	Passengers	Range (mi)	Speed (mph)	
D : 727 100	101	(20)	500	
Boeing 737-100	101	630	598	
Boeing 747	470	4150	610	
BAC/Sud Concor	rde 132	4000	1350	
Douglas DC-8-50) 146	8720	544	

- •In Speed?
- •In Cruising Range?
- •In Passenger Capacity?

Computer Performance: TIME, TIME, TIME

• Response Time (latency, "execution time")

The time between the start and completion of a task

- How long does it take for my job to run?
- How long does it take to execute a job?
- How long must I wait for the database query?
- Throughput

The total amount of work done in a given time.

- How many jobs can the machine run at once?
- What is the average execution rate?
- How much work is getting done?
- If we upgrade a machine with a new faster processor what do we increase?
 - Response Time, Throughput, or both?
- If we add a new machine to the lab what do we increase?
 - Response Time, Throughput, or both?

- Elapsed Time (Wall-clock time, response time)
 - counts everything (disk and memory accesses, I/O, etc.)
 - a useful number, but often not good for comparison purposes
- CPU (execution) time
 - doesn't count I/O or time spent running other programs
 - can be broken up into system CPU time, and user CPU time (difficult, though)
 - User CPU time
 - time spent executing the lines of code that are "in" our program
 - System CPU time
 - time spent in the operating system performing tasks on be half of the program

Execution Time & Performance

- For some program running on machine X.
 Performance X = Execution Time X
- "X is n times faster than Y"

- Example:
 - machine A runs a program in 20 seconds
 - machine B runs the same program in 25 seconds
 - "A is 1.5 times faster than B"

Performance A = 1.51 Performance B

CPU Performance - clock

- Clock Cycle ("ticks"):
 - discrete time interval
 - indicate when to start activities
- Clock Cycle Time
 - time between ticks
 - Time of a cycle
- Clock Rate (frequency)
 - Inverse of clock cycle

– cycles per second (1 Hz. = 1 cycle/sec)

A 4 Ghz. clock has a
$$\frac{1}{4 GH_z} = \frac{1}{4 X 10^{\circ}} = 25 X 10^{\circ} = 250 X 10^{\circ}$$

clock cycle time $\frac{1}{4 GH_z} = \frac{1}{4 X 10^{\circ}} = 250 X 10^{\circ} = 250 X 10^{\circ}$

tick tick

• Simple formula for CPU performance:

• Alternative formula:

CPU execution time for a program = CPU clock cycles for a program Clock Rate

- CPU performance comes from:
 - reduced length of clock cycle
 - reduced number of clock cycles required for a program

Different numbers of cycles for different instructions



- Multiplication takes more time than addition
- Floating point operations take longer than integer ones
- Accessing memory takes more time than accessing registers

Performance Improvement - Example

 Our favorite program runs in 10 seconds on computer A, which has a 4 GHz clock. We are trying to help a computer designer build a new machine B, that will run this program in 6 seconds. The designer can use new (or perhaps more expensive) technology to substantially increase the clock rate, but has informed us that this increase will affect the rest of the CPU design, causing machine B to require 1.2 times as many clock cycles as machine A for the same program. What clock rate should we tell the designer to target?"

CPI – Clock Cycles per Instruction

- Execution time & number of instructions
 - They are related
 - Execution time depends of the number of instructions
 - Each instruction requires different number of clock cycles

(Executione Time) = (Number of Instructions) × (Average Time Per Instruction)

- CPI (Clock cycles per Instruction)
 - "Average number of clock cycles each instruction takes to execute"

(CPU clock cycle)= (Number of Instr.) x (CPI)

 Suppose we have two implementations For some program,

Machine A has a clock cycle time of 250 ps and a CPI of 2.0 Machine B has a clock cycle time of 500 ps and a CPI of 1.2

What machine is faster for this program, and by how much?

CPI Example

- CPI is not the same
 - Memory system dependent
 - Instruction type dependent
 - Application dependent



CPI Example for Code Segments

• A compiler designer is trying to decide between two code sequences for a particular machine. Based on the hardware implementation, there are three different classes of instructions: Class A, Class B, and Class C, and they require one, two, and three CPIs (respectively).

The first code sequence has 5 instructions: 2 of A, 1 of B, and 2 of C The second sequence has 6 instructions: 4 of A, 1 of B, and 1 of C.

Which sequence will be faster? How much? What is the CPI for each sequence?

Another Example surrounding CPI – p.253

 A given application in Java runs 15 seconds on a desktop processor. A new Java complier is released that requires only 0.6 as many instructions as the old compiler. Unfortunately it increases the CPI by 1.1. How fast can we expect the application to run using this new compiler?

Benchmarks

- Ideal Condition of Performance Evaluation
 - Same workload
 - different computer systems
 - Compare the execution time

Alternative Approach

- Evaluation of computer systems using a set of benchmarks
- "Benchmark"
 - program specifically chosen to measure performance.
 - forms a workload for the prediction of performance.
 - The best benchmarks are real applications.

Small benchmarks

- nice for architects and designers
- easy to standardize
- can be abused

Benchmark

- Different types of benchmarks for different classes and applications
 - Desktop CPU performance (SPEC CPU)
 - Servers CPU-oriented benchmark or Web-oriented benchmark (SPEC Web)
 - Embedded-Computing (EEMBC)
- SPEC (System Performance Evaluation Corporation)
 - companies have agreed on a set of real program and inputs
 - valuable indicator of performance (and compiler technology)
- EEMBC (EDN Embedded Microprocessor Benchmark Consortium)
 - Variable Types of benchmark





Benchmark Games

An embarrassed Intel Corp. acknowledged Friday that a bug in a • software program known as a compiler had led the company to overstate the speed of its microprocessor chips on an industry benchmark by 10 percent. However, industry analysts said the coding error...was a sad commentary on a common industry practice of "cheating" on standardized performance tests... The error was pointed out to Intel two days ago by a competitor, Motorola ...came in a test known as SPECint92...Intel acknowledged that it had "optimized" its compiler to improve its test scores. The company had also said that it did not like the practice but felt to compelled to make the optimizations because its competitors were doing the same thing...At the heart of Intel's problem is the practice of "tuning" compiler programs to recognize certain computing problems in the test and then substituting special handwritten pieces of code...

Saturday, January 6, 1996 New York Times

• Compiler "enhancements" and performance



SPEC



Benchmarks

- 🗷 CPU
- Graphics/Applications
- HPC/OMP
- 🖉 Java Client/Server
- 🗷 Mail Servers
- 🕷 Network File System
- 🕷 Power and Performance
- 🗷 Virtualization
- 🕷 Web Servers

CPU

CPU2006

[benchmark info] [published results] [support] [order benchmark]

Designed to provide performance measurements that can be used to compare compute-intensive workloads on different computer systems, SPEC CPU200⁵ contains two benchmark suites: CINT2006 for measuring and comparing compute-intensive integer performance, and CFP2006 for measuring and comparing compute-intensive floating point performance.

CPU2000

[benchmark info] [published results] [support] [order benchmark]

Designed to provide performance measurements that can be used to compare compute-intensive workloads on different computer systems, SPEC CPU2000 contains two benchmark suites: CINT2000 for measuring and comparing compute-intensive integer performance, and CFP2000 for measuring and comparing compute-intensive floating point performance. The current version is CPU2000 V1.3.

- CPU95
 [Retired]
- CPU92
 [Retired]

SPEC CPU2006 - Benchmarks

Floating Point Benchmarks

Integer Benchmarks

400.perlbench	С	PERL Programming Language
<u>401.bzip2</u>	С	Compression
<u>403.gcc</u>	С	C Compiler
<u>429.mcf</u>	С	Combinatorial Optimization
445.gobmk	С	Artificial Intelligence: go
456.hmmer	С	Search Gene Sequence
458.sjeng	С	Artificial Intelligence: chess
462.libquantum	С	Physics: Quantum Computing
464.h264ref	С	Video Compression
471.omnetpp	C++	Discrete Event Simulation
<u>473.astar</u>	C++	Path-finding Algorithms
483.xalancbmk	C++	XIML Processing

410.bwaves	Fortran	Fluid Dynamics		
416.gamess	Fortran	Quantum Chemistry		
433.milc	С	Physics: Quantum Chromodynamics		
434.zeusmp	Fortran	Physics / CFD		
435.gromacs	C/Fortran	Biochemistry/Molecular Dynamics		
436.cactusADM	C/Fortran	Physics / General Relativity		
437.1eslie3d	Fortran	Fluid Dynamics		
444.namd	C++	Biology / Molecular Dynamics		
<u>447.deal∏</u>	C++	Finite Element Analysis		
450.soplex	C++	Linear Programming, Optimization		
453.povray	C++	Image Ray-tracing		
454.calculix	C/Fortran	Structural Mechanics		
459.GemsFDTD Fortran		Computational Electromagnetics		
465.tonto	Fortran	Quantum Chemistry		
<u>470.1bm</u>	С	Fluid Dynamics		
<u>481.wrf</u>	C/Fortran	Weather Prediction		
482.sphinx3	С	Speech recognition		

Metrics of Test

13. What metrics can be measured?

After the benchmarks are run on the system under test (SUT), a ratio for each of them is calculated using the run time on the SUT and a SPEC-determined <u>reference time</u>. From these ratios, the following metrics are calculated:

CINT2006 (for integer compute intensive performance comparisons):

- SPECint2006: The geometric mean of twelve normalized ratios one for each integer benchmark when the benchmarks are compiled with peak tuning.
- SPECint_base2006. The geometric mean of twelve normalized ratios when the benchmarks are compiled with base tuning.
- SPECint_rate 2006: The geometric mean of twelve normalized throughput ratios when the benchmarks are compiled with peak tuning.
- SPECint_rate_base2006: The geometric mean of twelve normalized throughput ratios when the benchmarks are compiled with base tuning.

Q14. What is the difference between a "base" metric and a "peak" metric?

In order to provide comparisons across different computer hardware, SPEC provides the benchmarks as source code. Thus, in order to run the benchmarks, they must be compiled. There is agreement that the benchmarks should be compiled the way users compile programs. But how do users compile programs?

- Some users might experiment with many different compilers and compiler flags to achieve the best performance, and may be willing to develop multi-step make processes and "training" workloads.
- Other users might prefer the relative simplicity of using a single set of switches and a single-step make process.

In addition to the above, a wide range of other types of usage models could also be imagined, ranging in a continuum from -Odebug at the low end, to inserting directives and/or re-writing the source code at the high end. Which points on this continuum should SPEC CPU2006 allow?

SPEC recognizes that any point chosen from that continuum might seem arbitrary to those whose interests lie at a different point. Nevertheless, choices must be made.

For CPU2006, SPEC has chosen to allow two types of compilation:

- The base metrics (e.g. SPECint_base2006) are required for all reported results and have stricter guidelines for compilation. For example, the same flags must be used in the same order for all benchmarks of a given language. This is the point closer to those who might prefer a relatively simple build process.
- The possk metrics (e.g. SPECint2006) are optional and have less strict requirements. For example, different compiler options may be used on each benchmark, and feedback-directed optimization is allowed. This point is closer to those who may be willing to invest more time and effort in development of build procedures.

 \mathbb{R}

Test Results

All Published SPEC CPU20	06 Results - Mozilla Firefox							\ge
Ble Edit Yew Go Bookman	ks Tools Field							5
🖕 · 🌳 · 🔗 🔞 😚) 🔯 http://www.spec.org/cpu2006/results/cint2006.html			۷	G 60	GL.		
📄 Free Hotmail 📄 Customize Link	s 🗋 Windows Marketplace 📄 Windows Media 📄 Windows							
Hewlett-Packard Company	HP Integrity rx6600 (1.6GHz/24MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	2	1	14.5	15.7	
Hewlett-Packard Company	HP Integrity rx2620 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	2	1	13.4	14.5	
			Pre	cessor		Results		
Test Sponsor	System Name	Cores	Chips	Cores/ Chip	Threads/ Core	Base	Peak	
Hewlett-Packard Company	HP Integrity rz4640 (1.6GHz/24MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	A.	1	13.8	15.1	
Hewlett-Packard Company	HP Integrity rx3600 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	2	1	13.9	15.1	
Hewlett-Packard Company	HP Integrity rz6600 (1.6GHz/24MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	12.2	12.9	
Hewlett-Packard Company	HP Integrity rx2620 (1.4GHz/12MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	2	1	11.7	12.6	
Hewlett-Packard Company	HP Integrity rx2620 (1.4GHz/12MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	9.73	10.2	
Hewlett-Packard Company	HP Integrity rx3600 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	11.6	12.3	
Hewlett-Packard Company	HP Integrity rx2620 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	11.4	12.0	
Hewlett-Packard Company	HP Integrity rz4640 (1.6GHz/24MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	11.8	12.5	
Hewlett-Packard Company	HP Integrity rx3600 (1.4GHz/12MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	10.0	10.6	
Hewlett-Packard Company	HP Integrity rx2660 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	1	1	2	1	11.7	12.3	
Hewlett-Packard Company	HP Integrity rx2660 (1.6GHz/18MB Dual-Core Intel Itanium 2) HTML CSV Text PDF PS Config	2	1	2	1	13.9	15.2	
Heszlett-Packard Company	HP Integrity rx3600 (1.4GHz/12MB Dual-Core Intel Itanium 2)	2	1	2	1	12.2	13.2	4

🕅 1.

SPEC 2000

Does doubling the clock rate double the performance? Can a machine with a slower clock rate have better performance?



Benchmark and power mode

Experiment

 Phone a major computer retailer and tell them you are having trouble deciding between two different computers, specifically you are confused about the processors strengths and weaknesses

(e.g., Pentium 4 at 2Ghz vs. Celeron M at 1.4 Ghz)

- What kind of response are you likely to get?
- What kind of response could you give a friend with the same question?

Comparing and Summarizing Performance

(4) Executin Times of 2 programs prog 1: Perf / Perf = 10 on 2 different computers prog 2: Perfo/Perfa = 10 prg 1 | 10 prog 2 1000 prj1+2: $\operatorname{Perf}_{\mathcal{B}}/\operatorname{Perf}_{\mathcal{A}} = \frac{1001}{100} = 9.1$ L Arithmetre Mean (AM) offere. $\operatorname{Amishimetre}_{i=1}^{n} \operatorname{Time}_{i}$ 100 Total Total Time for 2 programs. (Q) Collectively, what's the relative performance of * hidden? assumption? * hidden * time A&B? Weighted Arithmetic Mean of execution time each run an equal number of times. -> Right assumption ?

measurement of execution time				
	A	B		
program 1	Z sec	4 sec		
prigrom 2	5 sec	2 sec		

- Which of the following statements is true?
 - a. A is faster than B for program 1
 - b. A I faster than B for program 2
 - c. A is faster than B for a workload with equal numbers of executions of program 1 and 2
 - d. A is faster than B a workload with twice as many executions of program 1 as of program 2.

EEMBC



Benchmark Scores & Software Licensing

Automotive

Consumer

Digital Entertainment

GrinderBench

Networking

Network Storage

Office Automation

Telecom

PowerEnergy



Embedded Processors Target

Digital Entertainment Applications

Digital Entertainment Represents Many Market Segments

In car entertainment

Set top boxes

Mobile phones

PDAs

Portable audio products

Still and video digital cameras



EEMBC - digital entertainment benchmark

Digital Entertainment Benchmark Details - Dynamic

MP3 Decode

- EEMBC has more focus on reprogrammable solutions for mobile phones, PDAs, etc.
- MPEG-4 Video encode and decode benchmarks
 - Focus on mobile market
 - Also used in set top boxes for lower bit-rate streams
- MPEG-2 Video encode and decode benchmarks
 - More high-end focused
 - Fixed and floating point versions of encode

Cryptography

- AES, DES, RSA, and Huffman
- Rapidly rising demand for random numbers
- Across the net, applications need cryptography

Digital Entertainment Benchmark Details - Static

- Digital photography manipulation
 - RGB to YIQ Conversion
 - RGB to HPG Conversion
 - RGB to CMYK
 - JPEG compression and decompression
- Multiple datasets represents a variety of workloads

Test Harness



Consolidated Score

Components of Consolidated Scores for Digital Entertainment



Example Performance

Overall Performance Comparison



- "Laws of Variable Proportions"
 - in a production system with fixed and variable inputs (say factory size and labor), beyond some point, each additional unit of variable input yields less and less additional output.
- Illustration Example:

"Suppose a program runs in 100 seconds on a machine, with multiply responsible for 80 seconds of this time. How much do we have to improve the speed of multiplication if we want the program to run 4 times faster?"

• Amdahl's Law (E) - "the performance enhancement possible with a given (F) improvement is limited by the amount that the improved feature (N) is used" E : Expected perf. enhancementN: Amount of improvementF: Affected portion byimprovement.

Amdahl's Law Example

• Back to the example: "Suppose a program runs in 100 seconds on a machine, with multiply responsible for 80 seconds of this time. How much do we have to improve the speed of multiplication if we want the program to run 4 times faster?" M?

Example

• Suppose we enhance a machine making all floating-point instructions run five times faster. If the execution time of some benchmark before the floating-point enhancement is 10 seconds, what will the speedup be if half of the 10 seconds is spent executing floating-point instructions?

• We are looking for a benchmark to show off the new floating-point unit described above, and want the overall benchmark to show a speedup of 3. One benchmark we are considering runs for 100 seconds with the old floating-point hardware. How much of the execution time would floating-point instructions have to account for in this program in order to yield our desired speedup on this benchmark?

Problem of using only subset to measure performance

- Problem of using only one of
 - Clock rate
 - Instruction count
 - CPI
- Problem of using 2 of three factors
- Alternative to time
 - MIPS (million instructions per second)

- Problems

- Instruction counts differ with different instruction sets
- MIPS varies between programs
- MIPS can vary inversely with performance.

Example of MIPS problem (p.268)

CPI Compter 1 5 Instruction (in Billions) Compilerz (1) Which code sequence will execute faster according to Mips? 10 (1) According to execution time? t = 10x10 $C_{1} = (5 \times 1 + 1 \times 2 + 1 \times 3) \times 10^{9} = 10 \times 10^{9}$ $C_{2} = (10 \times 1 + 1 \times 2 + 1 \times 3) \times 10^{9} = 15 \times 10^{9}$ $\frac{7\times10^{9}}{2.5\times10^{6}} = 2800, \quad \text{MIPS}_{2} = \frac{15\times10^{9}}{3.75\times10^{6}} = 3200.$ MIPS1 =

- Performance is specific to a particular program/s
 - Total execution time is a consistent summary of performance
- For a given architecture performance increases come from:
 - increases in clock rate (without adverse CPI affects)
 - improvements in processor organization that lower CPI
 - compiler enhancements that lower CPI and/or instruction count
 - Algorithm/Language choices that affect instruction count
- Pitfall: expecting improvement in one aspect of a machine's performance to affect the total performance