

Lecture 41
Performance II



Lecturer PSOE Dan Garcia

www.cs.berkeley.edu/~ddgarcia

UWB...Ultra Wide Band! ⇒

The FCC moved one step closer to approving a standard for this technology which uses spread spectrum pulses to send its information. Imagine no data wires to ANY of your devices!



www.nytimes.com/2005/05/04/technology/techspecial/04markoff.html
CS61C L41 Performance II (1) Garcia © UCB

What is Time?

- Straightforward definition of time:
 - Total time to complete a task, including disk accesses, memory accesses, I/O activities, operating system overhead, ...
 - “real time”, “response time” or “elapsed time”
- Alternative: just time processor (CPU) is working only on your program (since multiple processes running at same time)
 - “CPU execution time” or “CPU time”
 - Often divided into system CPU time (in OS) and user CPU time (in user program)



CS61C L41 Performance II (2) Garcia © UCB

How to Measure Time?

- User Time ⇒ seconds
- CPU Time: Computers constructed using a clock that runs at a constant rate and determines when events take place in the hardware
 - These discrete time intervals called clock cycles (or informally clocks or cycles)
 - Length of clock period: clock cycle time (e.g., 2 nanoseconds or 2 ns) and clock rate (e.g., 500 megahertz, or 500 MHz), which is the inverse of the clock period; use these!



CS61C L41 Performance II (4) Garcia © UCB

Measuring Time using Clock Cycles (1/2)

- CPU execution time for a program
 - = Clock Cycles for a program
x Clock Cycle Time
- or
 - = $\frac{\text{Clock Cycles for a program}}{\text{Clock Rate}}$



CS61C L41 Performance II (5) Garcia © UCB

Measuring Time using Clock Cycles (2/2)

- One way to define clock cycles:
Clock Cycles for program
= Instructions for a program
(called “Instruction Count”)
x Average Clock cycles Per Instruction
(abbreviated “CPI”)
- CPI one way to compare two machines with same instruction set, since Instruction Count would be the same



CS61C L41 Performance II (6) Garcia © UCB

Performance Calculation (1/2)

- CPU execution time for program
 - = Clock Cycles for program
x Clock Cycle Time
- Substituting for clock cycles:
CPU execution time for program
= (Instruction Count x CPI)
x Clock Cycle Time
= Instruction Count x CPI x Clock Cycle Time



CS61C L41 Performance II (7) Garcia © UCB

Performance Calculation (2/2)

$$\text{CPU time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Cycle}}$$

$$\text{CPU time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Cycle}}$$

$$\text{CPU time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Cycle}}$$

$$\text{CPU time} = \frac{\text{Seconds}}{\text{Program}}$$

- Product of all 3 terms: if missing a term, can't predict time, the real measure of performance



How Calculate the 3 Components?

- **Clock Cycle Time:** in specification of computer (Clock Rate in advertisements)

• Instruction Count:

- Count instructions in loop of small program
- Use simulator to count instructions
- Hardware counter in spec. register
 - (Pentium II,III,4)

• CPI:

- Calculate: $\frac{\text{Execution Time}}{\text{Clock cycle time}} \times \text{Instruction Count}$

- Hardware counter in special register (PII,III,4)



Calculating CPI Another Way

- First calculate CPI for each individual instruction (add, sub, and, etc.)
- Next calculate frequency of each individual instruction
- Finally multiply these two for each instruction and add them up to get final CPI (the weighted sum)



Example (RISC processor)

Op	Freq _i	CPI _i	Prod	(% Time)
ALU	50%	1	.5	(23%)
Load	20%	5	1.0	(45%)
Store	10%	3	.3	(14%)
Branch	20%	2	.4	(18%)

Instruction Mix $\frac{2.2}{2.2}$ (Where time spent)

- What if Branch instructions twice as fast?



What Programs Measure for Comparison?

- Ideally run typical programs with typical input before purchase, or before even build machine
 - Called a "workload"; For example:
 - Engineer uses compiler, spreadsheet
 - Author uses word processor, drawing program, compression software
- In some situations its hard to do
 - Don't have access to machine to "benchmark" before purchase
 - Don't know workload in future
- Next: benchmarks & PC-Mac showdown!



Benchmarks

- Obviously, apparent speed of processor depends on code used to test it
- Need industry standards so that different processors can be fairly compared
- Companies exist that create these benchmarks: "typical" code used to evaluate systems
- Need to be changed every 2 or 3 years since designers could (and do!) target for these standard benchmarks



Example Standardized Benchmarks (1/2)

- Standard Performance Evaluation Corporation (SPEC) SPEC CPU2000
 - CINT2000 12 integer (gzip, gcc, crafty, perl, ...)
 - CFP2000 14 floating-point (swim, mesa, art, ...)
- All relative to base machine
Sun 300MHz 256Mb-RAM Ultra5_10, which gets score of 100
- www.spec.org/osg/cpu2000/
- They measure
 - System speed (SPECint2000)
 - System throughput (SPECint_rate2000)



CS61C L41 Performance II (14)

Garcia © UCB

Example Standardized Benchmarks (2/2)

- SPEC
 - Benchmarks distributed in source code
 - Members of consortium select workload
 - 30+ companies, 40+ universities
 - Compiler, machine designers target benchmarks, so try to change every 3 years
 - The last benchmark released was SPEC 2000
 - They are still finalizing SPEC 2005

CINT2000	CFP2000	CFP2000	CFP2000
gzip C Compression	swim C	Fortran77	Physcis / Quantum Chromodynamics
perl C Perl	gcc C Compiler	Fortran77	Shellin Water Modeling
gcc C C	gcc C Compiler	Fortran77	Multi-grid Solver: 3D Potential Field
gcc C C	gcc C Compiler	Fortran77	Parabolic / Elliptic Partial Differential Equations
crafty C Game Play/Al Chess	mesa C	Fortran90	3-D Graphics Library
parsec C Word Processing	gcc C	Fortran90	Computational Fluid Dynamics
swim C++ Computer Visualization	art C	Fortran90	Image Recognition / Neural Networks
perl C Perl	gcc C	Fortran90	Relativistic Neutron Propagation Simulation
gcc C C	gcc C Compiler	Fortran90	Image Processing / Face Recognition
gcc C C	gcc C Compiler	Fortran90	Computational Chemistry
gcc C C	gcc C Compiler	Fortran90	Huber Theory / Fractality Testing
gcc C C	gcc C Compiler	Fortran90	Water-solvent Crack Simulation
gcc C C	gcc C Compiler	Fortran77	High Energy Nuclear Physics Accelerator Design
gcc C C	gcc C Compiler	Fortran77	Seismology: Pollutant Distribution



CS61C L41 Performance II (15)

Garcia © UCB

Example PC Workload Benchmark

- PCs: Ziff-Davis Benchmark Suite
 - “Business Winstone is a system-level, application-based benchmark that measures a PC’s overall performance when running today’s top-selling Windows-based 32-bit applications... **it doesn’t mimic what these packages do; it runs real applications through a series of scripted activities** and uses the time a PC takes to complete those activities to produce its performance scores.
 - Also tests for CDs, Content-creation, Audio, 3D graphics, battery life

<http://www.etestinglabs.com/benchmarks/>



CS61C L41 Performance II (16)

Garcia © UCB

Performance Evaluation

- Good products created when have:
 - Good benchmarks
 - Good ways to summarize performance
- Given sales is a function of performance relative to competition, should invest in improving product as reported by performance summary?
- If benchmarks/summary inadequate, then choose between improving product for real programs vs. **improving product to get more sales; Sales almost always wins!**



CS61C L41 Performance II (17)

Garcia © UCB

Performance Evaluation: The Demo

If we’re talking about performance, let’s discuss the ways shady salespeople have fooled consumers (so that you don’t get taken!)

5. Never let the user touch it
4. Only run the demo through a script
3. Run it on a stock machine in which “no expense was spared”
2. Preprocess all available data
1. Play a movie



CS61C L41 Performance II (18)

Garcia © UCB

PC / PC / Mac Showdown!!! (1/4)

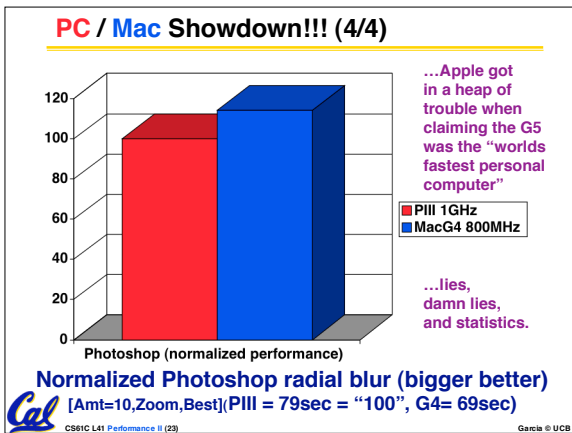
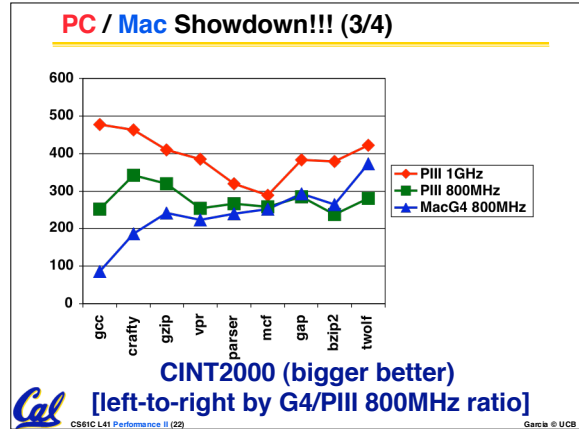
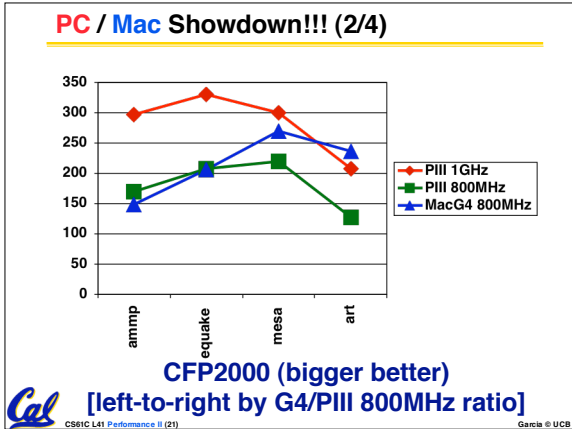
- PC
 - 1 GHz Pentium III
 - 256 Mb RAM
 - 512KB L2 Cache
 - No L3
 - 133 MHz Bus
 - 20 GB Disk
 - 16MB VRAM
- PC 800MHz PIII
- Mac
 - 800 MHz PowerbookG4
 - 1 Gb RAM
 - 2 512Mb SODIMMs
 - 32KB L1Inst, L1Data
 - 256KB L2 Cache
 - 1Mb L3 Cache
 - 133 MHz Bus
 - 40 GB Disk
 - 32MB VRAM

Let’s take a look at SPEC2000 and a simulation of a real-world application.



CS61C L41 Performance II (20)

Garcia © UCB



- ### Administrivia
- HKN evaluations on Monday
 - Final survey in lab this week
 - Last semester's final + solutions online
 - Final exam review
 - Sunday, 2005-05-08 @ 2pm in 10 Evans
 - Final exam
 - Tuesday, 2005-05-14 @ 12:30-3:30pm in 220 Hearst Gym
 - Same rules as Midterm, except you get 2 double-sided handwritten review sheets (1 from your midterm, 1 new one) + green sheet **[Don't bring backpacks]**
- CS61C L41 Performance II (24) Garcia © UCB

Peer Instruction

A. **Performance is a stinking business**; easily corruptible and (in general) you won't hear honest reports from a company if they have a vested interest in the results.

B. The **Sieve of Eratosthenes**, **Puzzle** and **Quicksort** were early effective benchmarks.

C. A program runs in 100 sec. on a machine, mult accounts for 80 sec. of that. If we want to make the program run 6 times faster, we need to up the speed of mults by **AT LEAST 6**.

	ABC
1:	FFF
2:	FFT
3:	FTF
4:	FTT
5:	FFF
6:	FTF
7:	FTT
8:	TTT

CS61C L41 Performance II (25) Garcia © UCB

"And in conclusion..."

$$\text{CPU time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Cycle}}$$

- Latency v. Throughput
- Performance doesn't depend on any single factor: need Instruction Count, Clocks Per Instruction (CPI) and Clock Rate to get valid estimations
- User Time: time user waits for program to execute; depends heavily on how OS switches between tasks
- CPU Time: time spent executing a single program: depends solely on design of processor (datapath, pipelining effectiveness, caches, etc.)
- Benchmarks
 - Attempt to predict perf, Updated every few years
 - Measure everything from simulation of desktop graphics programs to battery life
- Megahertz Myth
 - MHz ≠ performance, it's just one factor

CS61C L41 Performance II (27) Garcia © UCB