


inst.eecs.berkeley.edu/~cs61c
UCB CS61C : Machine Structures

Lecture 36 – Performance
2010-04-23




Lecturer SOE
Dan Garcia

How fast is your computer?

CRAY XT5-HE IS FASTEST SUPERCOMPUTER!


Every 6 months (Nov/June), the fastest supercomputers in the world face off. The fastest computer is now the Jaguar, a Linux Cray XT5-HE Opteron Six Core 2.6 GHz, with 224,256 cores, achieves 2.3 PFlops. They use LINPACK floating point benchmark (A x = B).



www.top500.org/lists/2009/11
www.nccs.gov/computing-resources/jaguar/

Why Performance? Faster is better!

- **Purchasing Perspective:** given a collection of machines (or upgrade options), which has the
 - best performance ?
 - least cost ?
 - best performance / cost ?
- **Computer Designer Perspective:** faced with design options, which has the
 - best performance improvement ?
 - least cost ?
 - best performance / cost ?
- **All require basis for comparison and metric for evaluation!**
 - Solid metrics lead to solid progress!




CS61C L16 Performance (2) Garcia, Spring 2010 © UCB

Two Notions of “Performance”

| Plane | DC to Paris | Top Speed | Passengers | Throughput (pmph) |
|------------------|-------------|-----------|------------|-------------------|
| Boeing 747 | 6.5 hours | 610 mph | 470 | 286,700 |
| BAD/Sud Concorde | 3 hours | 1350 mph | 132 | 178,200 |


- Which has higher performance?
 - Interested in time to deliver 100 passengers?
 - Interested in delivering as many passengers per day as possible?
- In a computer, time for one task called Response Time or Execution Time
- In a computer, tasks per unit time called Throughput or Bandwidth



CS61C L16 Performance (3) Garcia, Spring 2010 © UCB

Definitions


- Performance is in units of things per sec
 - bigger is better
- If mostly concerned with response time
 - $performance(x) = \frac{1}{execution_time(x)}$
- “F(aster) is n times faster than S(low)” means:

$$n = \frac{performance(F)}{performance(S)} = \frac{execution_time(S)}{execution_time(F)}$$


CS61C L16 Performance (4) Garcia, Spring 2010 © UCB

Example of Response Time v. Throughput


- **Time of Concorde vs. Boeing 747?**
 - Concorde is 6.5 hours / 3 hours = 2.2 times faster
 - Concorde is 2.2 times (“120%”) faster in terms of flying time (response time)
- **Throughput of Boeing vs. Concorde?**
 - Boeing 747: 286,700 pmph / 178,200 pmph = 1.6 times faster
 - Boeing is 1.6 times (“60%”) faster in terms of throughput
- We will focus primarily on response time.



CS61C L16 Performance (5) Garcia, Spring 2010 © UCB

Words, Words, Words...

- Will (try to) stick to “n times faster”; its less confusing than “m % faster”
- As faster means both decreased execution time and increased performance, to reduce confusion we will (and you should) use “improve execution time” or “improve performance”



CS61C L16 Performance (6) Garcia, Spring 2010 © 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)

Cal

CS63C L16 Performance (7)

Gerds, Spring 2010 © UCS

How to Measure Time?

- **Real Time** ⇒ Actual time elapsed
- **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., 1/2 nanoseconds or 1/2 ns) and clock rate (e.g., 2 gigahertz, or 2 GHz), which is the inverse of the clock period; use these!

Cal

CS63C L16 Performance (8)

Gerds, Spring 2010 © UCS

Measuring Time using Clock Cycles (1/2)

- **CPU execution time for a program**
 - Units of [seconds / program] or [s/p]
- = Clock Cycles for a program x Clock Period**
 - Units of [s/p] = [cycles / p] x [s / cycle] = [c/p] x [s/c]
- Or
 - = $\frac{\text{Clock Cycles for a program [c / p]}}{\text{Clock Rate [c / s]}}$**

Clock



Cal

CS63C L16 Performance (9)

Gerds, Spring 2010 © UCS

Measuring Time using Clock Cycles (2/2)

- One way to define clock cycles:
 - Clock Cycles for program [c/p]**
 - = Instructions for a program [i/p]** (called "Instruction Count")
 - x Average Clock cycles Per Instruction [c/i]** (abbreviated "CPI")
- CPI one way to compare two machines with same instruction set, since Instruction Count would be the same

Cal

CS63C L16 Performance (10)

Gerds, Spring 2010 © UCS

Performance Calculation (1/2)

- CPU execution time for program [s/p]
 - = Clock Cycles for program [c/p]**
 - x Clock Cycle Time [s/c]**
- Substituting for clock cycles:
 - CPU execution time for program [s/p]
 - = (Instruction Count [i/p] x CPI [c/i])**
 - x Clock Cycle Time [s/c]**
- = Instruction Count x CPI x Clock Cycle Time**

Cal

CS63C L16 Performance (11)

Gerds, Spring 2010 © UCS

Performance Calculation (2/2)

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

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

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

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

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

Cal

CS63C L16 Performance (12)

Gerds, Spring 2010 © UCS

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: $\frac{\text{Execution Time}}{\text{Instruction Count}}$ / Clock cycle time
 - Hardware counter in special register (PII,III,4)



CS50C L16 Performance (18)

©ards, Spring 2010 © UCS

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)**



CS50C L16 Performance (18)

©ards, Spring 2010 © UCS

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 | | | 2.2 | (Where time spent) |

- What if Branch instructions twice as fast?



CS50C L16 Performance (18)

©ards, Spring 2010 © UCS

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



CS50C L16 Performance (18)

©ards, Spring 2010 © UCS

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 ~5 years since designers could (and do!) target for these standard benchmarks**



CS50C L16 Performance (17)

©ards, Spring 2010 © UCS

Example Standardized Benchmarks (1/2)

- **Standard Performance Evaluation Corporation (SPEC) SPEC CPU2006**
 - CINT2006 12 integer (perl, bzip, gcc, go, ...)
 - CFP2006 17 floating-point (povray, bwaves, ...)
 - All relative to base machine (which gets 100) Sun Ultra Enterprise 2 w/296 MHz UltraSPARC II
 - They measure
 - System speed (SPECint2006)
 - System throughput (SPECint_rate2006)
 - www.spec.org/osg/cpu2006/



CS50C L16 Performance (18)

©ards, Spring 2010 © UCS

Example Standardized Benchmarks (2/2)

- SPEC

- Benchmarks distributed in source code
- Members of consortium select workload
 - 30+ companies, 40+ universities, research labs
- Compiler, machine designers target benchmarks, so try to change every 5 years
- SPEC CPU2006:

| | | | |
|------------------|---|------------------|---------|
| CSPEC2006 | | CSPEC2006 | |
| perlbench | C | perl | Fortran |
| compbench | C | comp | Fortran |
| hpc1 | C | hpc1 | Fortran |
| hpc2 | C | hpc2 | Fortran |
| hpc3 | C | hpc3 | Fortran |
| hpc4 | C | hpc4 | Fortran |
| hpc5 | C | hpc5 | Fortran |
| hpc6 | C | hpc6 | Fortran |
| hpc7 | C | hpc7 | Fortran |
| hpc8 | C | hpc8 | Fortran |
| hpc9 | C | hpc9 | Fortran |
| hpc10 | C | hpc10 | Fortran |
| hpc11 | C | hpc11 | Fortran |
| hpc12 | C | hpc12 | Fortran |
| hpc13 | C | hpc13 | Fortran |
| hpc14 | C | hpc14 | Fortran |
| hpc15 | C | hpc15 | Fortran |
| hpc16 | C | hpc16 | Fortran |
| hpc17 | C | hpc17 | Fortran |
| hpc18 | C | hpc18 | Fortran |
| hpc19 | C | hpc19 | Fortran |
| hpc20 | C | hpc20 | Fortran |
| hpc21 | C | hpc21 | Fortran |
| hpc22 | C | hpc22 | Fortran |
| hpc23 | C | hpc23 | Fortran |
| hpc24 | C | hpc24 | Fortran |
| hpc25 | C | hpc25 | Fortran |
| hpc26 | C | hpc26 | Fortran |
| hpc27 | C | hpc27 | Fortran |
| hpc28 | C | hpc28 | Fortran |
| hpc29 | C | hpc29 | Fortran |
| hpc30 | C | hpc30 | Fortran |
| hpc31 | C | hpc31 | Fortran |
| hpc32 | C | hpc32 | Fortran |
| hpc33 | C | hpc33 | Fortran |
| hpc34 | C | hpc34 | Fortran |
| hpc35 | C | hpc35 | Fortran |
| hpc36 | C | hpc36 | Fortran |
| hpc37 | C | hpc37 | Fortran |
| hpc38 | C | hpc38 | Fortran |
| hpc39 | C | hpc39 | Fortran |
| hpc40 | C | hpc40 | Fortran |
| hpc41 | C | hpc41 | Fortran |
| hpc42 | C | hpc42 | Fortran |
| hpc43 | C | hpc43 | Fortran |
| hpc44 | C | hpc44 | Fortran |
| hpc45 | C | hpc45 | Fortran |
| hpc46 | C | hpc46 | Fortran |
| hpc47 | C | hpc47 | Fortran |
| hpc48 | C | hpc48 | Fortran |
| hpc49 | C | hpc49 | Fortran |
| hpc50 | C | hpc50 | Fortran |
| hpc51 | C | hpc51 | Fortran |
| hpc52 | C | hpc52 | Fortran |
| hpc53 | C | hpc53 | Fortran |
| hpc54 | C | hpc54 | Fortran |
| hpc55 | C | hpc55 | Fortran |
| hpc56 | C | hpc56 | Fortran |
| hpc57 | C | hpc57 | Fortran |
| hpc58 | C | hpc58 | Fortran |
| hpc59 | C | hpc59 | Fortran |
| hpc60 | C | hpc60 | Fortran |
| hpc61 | C | hpc61 | Fortran |
| hpc62 | C | hpc62 | Fortran |
| hpc63 | C | hpc63 | Fortran |
| hpc64 | C | hpc64 | Fortran |
| hpc65 | C | hpc65 | Fortran |
| hpc66 | C | hpc66 | Fortran |
| hpc67 | C | hpc67 | Fortran |
| hpc68 | C | hpc68 | Fortran |
| hpc69 | C | hpc69 | Fortran |
| hpc70 | C | hpc70 | Fortran |
| hpc71 | C | hpc71 | Fortran |
| hpc72 | C | hpc72 | Fortran |
| hpc73 | C | hpc73 | Fortran |
| hpc74 | C | hpc74 | Fortran |
| hpc75 | C | hpc75 | Fortran |
| hpc76 | C | hpc76 | Fortran |
| hpc77 | C | hpc77 | Fortran |
| hpc78 | C | hpc78 | Fortran |
| hpc79 | C | hpc79 | Fortran |
| hpc80 | C | hpc80 | Fortran |
| hpc81 | C | hpc81 | Fortran |
| hpc82 | C | hpc82 | Fortran |
| hpc83 | C | hpc83 | Fortran |
| hpc84 | C | hpc84 | Fortran |
| hpc85 | C | hpc85 | Fortran |
| hpc86 | C | hpc86 | Fortran |
| hpc87 | C | hpc87 | Fortran |
| hpc88 | C | hpc88 | Fortran |
| hpc89 | C | hpc89 | Fortran |
| hpc90 | C | hpc90 | Fortran |
| hpc91 | C | hpc91 | Fortran |
| hpc92 | C | hpc92 | Fortran |
| hpc93 | C | hpc93 | Fortran |
| hpc94 | C | hpc94 | Fortran |
| hpc95 | C | hpc95 | Fortran |
| hpc96 | C | hpc96 | Fortran |
| hpc97 | C | hpc97 | Fortran |
| hpc98 | C | hpc98 | Fortran |
| hpc99 | C | hpc99 | Fortran |
| hpc100 | C | hpc100 | Fortran |



CSPEC L36 Performance (1)

Gerds, Spring 2010 © UCS

Performance Evaluation: The Demo

If we're talking about performance, let's discuss the ways shady salespeople have fooled consumers (so 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



CSPEC L36 Performance (2)

Gerds, Spring 2010 © UCS

Megahertz Myth Marketing Movie

Peer Instruction

- 1) The Sieve of Eratosthenes and Quicksort were early effective benchmarks.
- 2) 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.

- | | |
|----|----|
| 1 | 2 |
| a) | FF |
| b) | FT |
| c) | TF |
| d) | TT |



CSPEC L36 Performance (3)

Gerds, Spring 2010 © UCS

"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



CSPEC L36 Performance (4)

Gerds, Spring 2010 © UCS