

Performance

- Measure, Report, and Summarize
- Make intelligent choices
- See through the marketing hype
- Key to understanding underlying organizational motivation

Why is some hardware better than others for different programs?

*What factors of system performance are hardware related?
(e.g., Do we need a new machine, or a new operating system?)*

How does the machine's instruction set affect performance?

Which of these airplanes has the best performance?

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

•How much faster is the Concorde compared to the 747?

•How much bigger is the 747 than the Douglas DC-8?

Computer Performance: TIME, TIME, TIME

- **Response Time** (latency)
 - 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**
 - 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 processor what do we increase?*
- *If we add a new machine to the lab what do we increase?*

Execution Time

- **Elapsed Time**
 - counts everything (*disk and memory accesses, I/O, etc.*)
 - a useful number, but often not good for comparison purposes
- **CPU time**
 - doesn't count I/O or time spent running other programs
 - can be broken up into system time, and user time
- **Our focus: user CPU time**
 - time spent executing the lines of code that are "in" our program

Book's Definition of Performance

- For some program running on machine X,

$$\text{Performance}_X = 1 / \text{Execution time}_X$$

- "X is n times faster than Y"

$$\text{Performance}_X / \text{Performance}_Y = n$$


- Problem:
 - machine A runs a program in 20 seconds
 - machine B runs the same program in 25 seconds

Clock Cycles

- Instead of reporting execution time in seconds, we often use cycles

$$\frac{\text{seconds}}{\text{program}} = \frac{\text{cycles}}{\text{program}} \times \frac{\text{seconds}}{\text{cycle}}$$

- Clock "ticks" indicate when to start activities (one abstraction):



- cycle time = time between ticks = seconds per cycle
- clock rate (frequency) = cycles per second (1 Hz. = 1 cycle/sec)

A 200 Mhz. clock has a $\frac{1}{200 \times 10^6} \times 10^9 = 5$ nanoseconds cycle time

How to Improve Performance


$$\frac{\text{seconds}}{\text{program}} = \frac{\text{cycles}}{\text{program}} \times \frac{\text{seconds}}{\text{cycle}}$$

So, to improve performance (everything else being equal) you can either

- _____ the # of required cycles for a program, or
- _____ the clock cycle time or, said another way,
- _____ the clock rate.

How many cycles are required for a program?

- Could assume that # of cycles = # of instructions




This assumption is incorrect,

different instructions take different amounts of time on different machines.

Why? hint: remember that these are machine instructions, not lines of C code

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
- *Important point: changing the cycle time often changes the number of cycles required for various instructions (more later)*

Example

(See pages 60-61.)

- Our favorite program runs in 10 seconds on computer A, which has a 400 Mhz 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 clockrate, 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 clockcycles as machine A for the same program. What clockrate should we tell the designer to target?"

- Don't Panic, can easily work this out from basic principles

$$\begin{aligned} \text{CPU execution time} &= \text{CPU clock cycles for the program} \times \text{Clock cycle time} \\ \text{for a program} &= \frac{\text{CPU clock cycles for the program}}{\text{Clock rate}} \end{aligned}$$

Now that we understand cycles

- A given program will require
 - some number of instructions (machine instructions)
 - some number of cycles
 - some number of seconds
- We have a vocabulary that relates these quantities:
 - cycle time (seconds per cycle)
 - clock rate (cycles per second)
 - CPI (cycles per instruction)
 - a floating point intensive application might have a higher CPI*
 - MIPS (millions of instructions per second)
 - this would be higher for a program using simple instructions*

Performance

- Performance is determined by execution time
- Do any of the other variables equal performance?
 - # of cycles to execute program?
 - # of instructions in program?
 - # of cycles per second?
 - average # of cycles per instruction?
 - average # of instructions per second?
- Common pitfall: thinking one of the variables is indicative of performance when it really isn't.

CPI Example

(See pages 62-63.)

- Suppose we have two implementations of the same instruction set architecture (ISA). For some program,

Machine A has a clock cycle time of 1 ns. and a CPI of 2.0

Machine B has a clock cycle time of 2 ns. and a CPI of 1.2

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

$$\text{CPU clock cycles for a program} = \text{Instruction count for the program} \cdot \text{Average clock cycles per instruction}$$

$$\begin{aligned} \text{CPU time} &= \text{Instruction count} \cdot \text{CPI} \cdot \text{Clock cycle time} \\ &= \frac{\text{Instruction count} \cdot \text{CPI}}{\text{Clock rate}} \end{aligned}$$

- *If two machines have the same ISA, which of our quantities (e.g., clock rate, CPI, execution time, # of instructions, MIPS) will always be identical?*

of Instructions Example

(See pages 64-65.)

- **Note:**
$$\text{CPU clock cycles for a program} = \sum_{i=1}^n (\text{CPI}_i \cdot C_i)$$

where: C_i = count of number of instructions of class i executed

CPI_i = average number of cycles per instruction for instruction class i

- 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 cycles (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?

MIPS as a Performance Metric

- **MIPS = million instructions per second**
- **Variants:**

$$\text{native MIPS} = \frac{\text{Instruction count}}{\text{Execution time} \cdot 10^6} \quad \text{"default"}$$

$$\text{relative MIPS} = \frac{\text{Time}_{\text{reference}}}{\text{Time}_{\text{unrated}}} \cdot \text{MIPS}_{\text{reference}}$$

peak MIPS = Same as *native MIPS*, but obtained by choosing an instruction mix that minimizes the CPI, even if the instruction mix is totally impractical.

MIPS example

(See pages 78-79.)

- **Two different compilers are being tested for a 500 MHz. machine with three different classes of instructions: Class A, Class B, and Class C, which require one, two, and three cycles (respectively). Both compilers are used to produce code for a large piece of software.**

The first compiler's code uses 5 million Class A instructions, 1 million Class B instructions, and 1 million Class C instructions.

The second compiler's code uses 10 million Class A instructions, 1 million Class B instructions, and 1 million Class C instructions.

- **Which sequence will be faster according to MIPS?**
- **Which sequence will be faster according to execution time?**


Benchmarks

- Performance best determined by running a real application
 - Use programs typical of expected workload
 - Or, typical of expected class of applications
e.g., compilers/editors, scientific applications, graphics, etc.
- Small benchmarks
 - nice for architects and designers
 - easy to standardize
 - can be abused
- **SPEC (System Performance Evaluation Cooperative)**
 - companies have agreed on a set of real program and inputs
 - can still be abused (Intel' s “other” bug)
 - valuable indicator of performance (and compiler technology)

SPEC '89

(See page 68.)

- Compiler “enhancements” and performance



SPEC '95

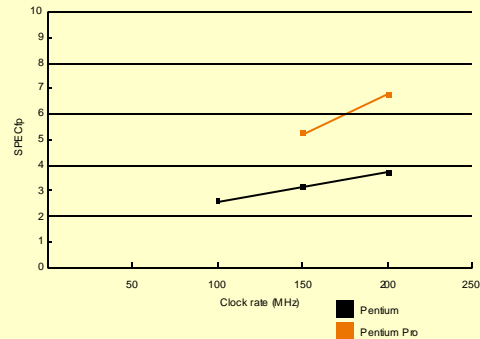
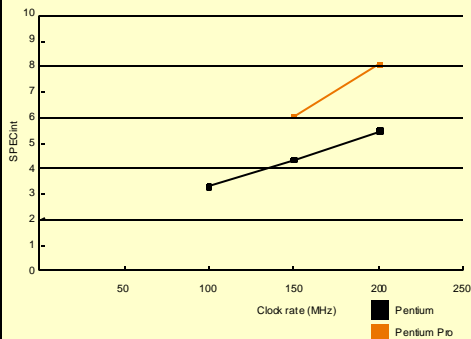
(See page 72.)

Benchmark	Description
go	Artificial intelligence: plays the game of Go
m88ksim	Motorola 88k chip simulator: runs test program
gcc	The Gnu C compiler generating SPARC code
compress	Compresses and decompresses file in memory
li	Lisp interpreter
ljpeg	Graphic compression and decompression
perl	Manipulates strings and prime numbers in the special purpose programming language Perl
vortex	A database program
lomcatv	A mesh generation program
swim	Shallow water model with 513 x 513 grid
su2cor	quantum physics: Monte Carlo simulation
hydro2d	Astrophysics: Hydrodynamic Navier Stokes equations
mgrid	Multigrid solver in 3-D potential field
apollu	Parabolic/elliptic partial differential equations
trub3d	Simulates isotropic homogeneous turbulence in a cube
apsi	Solves problems regarding temperature, wind velocity, and distribution of pollutant
lqppp	Quantum chemistry
wave5	Plasma physics: electromagnetic particle simulation

SPEC '95

(See pages 73-74.)

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



Amdahl's Law

(See page 75.)

Execution Time After Improvement =

Execution Time Unaffected + (Execution Time Affected / Amount of Improvement)

- **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?"

How about making it *5 times* faster?

- *Principle: Make the common case fast*

Example

(See pages 101-102.)

- **[Exercise 2.41]** 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?
- **[Exercise 2.42]** 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?

Remember

- Performance is specific to a particular program/s
 - Total execution time is a consistent summary of performance
- For a given instruction set 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
- Pitfall: expecting improvement in one aspect of a machine's performance to affect the total performance
- You should not always believe everything you read! Read carefully!