

Courtesy of Prof. John D. Owens, ECE dept, UC-Davis

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?





1350 mph

132

178,200

3 hours Which has higher performance?

BAD/Sud

Concorde

- Time to do the task (Execution Time) execution time, response time, latency
- Tasks per day, hour, week, sec, ns ... (Performance) throughput, bandwidth
- Response time and throughput often are in opposition

Example Time of Concorde vs. Boeing 747? • Concorde is 1350 mph / 610 mph = 2.2 times faster = 6.5 hours / 3 hours Throughput of Concorde vs. Boeing 747 ? • Concorde is 178,200 pmph / 286,700 pmph = 0.62 "times faster" • Boeing is 286,700 pmph / 178,200 pmph = 1.60 "times faster" Boeing is 1.6 times ("60%") faster in terms of throughput Concord is 2.2 times ("120%") faster in terms of flying time We will focus primarily on execution time for a single job

· But sysadmins may use throughput as their primary metric!

Latency vs. Throughput Latency (Response Time) • 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?

Definitions

Performance is in units of things-per-time

- Miles per hour, bits per second, widgets per day ...
- Bigger is better
- If we are primarily concerned with response time:
 - Performance(x) = 1 / ExecutionTime(x)

"X is n times faster than Y" means

- n = Performance(X) / Performance(Y) = Speedup
- If X is 1.yz times faster than Y, we can informally say that X is yz% faster than Y. Speedup is better.

Execution Time

Elapsed Time

- counts everything (disk and memory accesses, I/O, etc.)
- a useful number, but often not good for comparison purposes

0.50 user

9.86 sys

CPU time

doesn't count I/O or time spent running other programs

- can be broken up into system time and user time
- % /usr/bin/time du -s
 81329656 .
- 104.44 real

Our focus: user CPU time

 time spent executing the lines of code that are "in" our program



How to Improve Performance

the clock rate.

decrease):

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

being equal) you can either (increase/

So, to improve performance (everything else

the # of required cycles for a program, or

the clock cycle time or, said another way,

Clock Speed Is Not The Whole Story

	SPECint95	SPECfp95
195 MHz MIPS R10000	11.0	17.0
400 MHz Alpha 21164	12.3	17.2
300 MHz UltraSPARC	12.1	15.5
300 MHz Pentium II	11.6	8.8
300 MHz PowerPC G3	14.8	11.4
135 MHz POWER2	6.2	17.6

[http://www.pattosoft.com.au/Articles/ModernMicroprocessors/]

Clock Rate

Comparing processor performance for same or different architectures using clocks rate is invalid. Ignores instruction count and CPI, e.g.:

• 2.4GHz AMD Athlon processor is faster than 3.4GHz Intel Pentium 4 executing floatingpoint code (P4 has higher CPI)





Example instruction latencies Imagine Stream Processor: On ALU: Other functional units: Integer adds: 2 cycles Integer multiply: 4 •FP adds: 4 cycles Integer divide: 22 •Logic ops (and, or, xor): 1 Integer remainder: 23 •Equality: 1 •FP multiply: 4 • < or >: 2 •FP divide: 17 • Shifts: 1 •FP sqrt: 16

•Float->int: 3

Int->float: 4
Select (a?b:c): 1

CPI

How many clock cycles, on average, does it take for every instruction executed?
We call this CPI ("Cycles Per Instruction").
Its inverse (1/CPI) is IPC ("Instructions Per Cycle").
CISC machines: this number is

• high(er)

RISC machines: this number is

low(er)







<u>Millions of Instructions Per Second</u> (MIPS)

MIPS = Clock Rate / CPI x 10⁶

Ignores program instruction count. Hence, valid only for comparing processors running same object code.

Can vary inversely with performance!

• E.g., optimizing compiler eliminates instructions with relatively low cycle count. Causes CPI to increase and MIPS to decrease, but performance increases.

Sometime called <u>Meaningless</u> <u>Indicator-of-</u> <u>Performance</u> <u>Statistic</u>

megaFLOPS

Floating point performance metric: <u>Million Floating</u> <u>Point Operations Per Second</u>

• for given program rate of fadd, fmult, etc.

Not valid for different architectures because available floating point operations may differ

- E.g., some include div, sine, sqrt; other synthesize these operations with other simpler floating point operations
- E.g., some include a single multiply/accumulate

Peak megaFLOPS: rate guaranteed not to exceed, supercomputers may achieve only a few percent of peak. Very misleading.

The Performance Equation

Time = Clock Speed * CPI * Instruction Count • = seconds/cycle * cycles/instr * instrs/program • => seconds/program

"The only reliable measure of computer performance is time."

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.





Aspects of CPU Performance					
CPU time = Sec	onds = Instructio	ons x Cycle	s x Seconds		
Pro	gram Program	n instru	ction Cycle		
	Instr count	CPI	Clock rate		
Program	-				
Compiler					
Instruction Set					
Organization					
Technology					
		I			

CPU time = Se Pr	conds = Instructio	= Instructions x Cycles x Seconds Program Instruction Cycle		
***************************************	Instr count	CPI	Clock rate	
Program	X			
Compiler	x	x		
nstruction Set	X	x	x	
rganization		x	x	
echnology			x	



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

Pitfall:

expecting the improvement of one aspect of a computer to increase performance by an amount proportional to the size of improvement.



Undergrad Productivity

Average ECE student spends:

- 4 hours sleeping
- 2 hours eating
- 18 hours studying (yeah ... right!)

Magic pill gives you all sleeping, eating in 1 minute!

How much more productive can you get?

Undergrad Productivity

1 Speedup (with E) = ------(1-F) + F/S F = accelerated fraction = 0.25 (6 hrs/24 hrs) S = speedup = 6 hrs / 1 minute = 360

Overall speedup: • 1 / [(1-0.25) + (0.25/360)] • ~= 1 / (1-0.25) • ~= 1.33

33% more productive!

Benchmarks

Widely used programs for assessing execution time Results must be reproducible => input is specified along with program

Benchmark types:

- Application programs: Most accurate example, SPEC benchmarks suite.
- Kernels: Key inner loops extracted from applications (Program typically spend 90% of time in 10% of code).
 Examples, Linpack & Livermore loops for supercomputers
- Toy benchmarks: Small enough to be entered by hand (e.g., 100 lines).
- Synthetic Benchmarks: Non-real programs based on instruction mix statistics

Problems with Small Benchmarks

- Does not fully exercise memory hierarchy (e.g., entire program data set may fit in cache)
- Encourages "benchmark specific optimizations" (cheating?) in compiler or processor hardware that do not benefit real applications => do not accurately predict application performance.

Proprietary Benchmarks

Results cannot be reproduced, leads to unproductive debate

"Benchmark wars erupt again"

By Frank Williams, Chicago Tribune

Tuesday, October 6, 1998

The benchmark wars in the world of computer hyping have erupted again. The latest major offensive is by PC Magazine. In reality this is but a belated counterattack to the Apple offensive of last November. Wintel platform proponents claim to have "proof" that their machines are faster than Macintoshes. Apple supporters claim to have "proof" that the processor inside every new Mac is twice as fast as those in Wintels...



Reporting Performance

Result should be reproducible implies must specify:

- processor model
- memory configuration (including cache, main & disk)
- compiler version
- OS version

Best is to report execution time for all applications

• Allows user of data to focus on data most relevant

Basis of Evaluation

Pros		Cons
 representative 	Actual Target Workload	 very specific non-portable difficult to run, or measure hard to identify
 portable widely used improvements useful n reality 	Full Application Benchmarks	 less representative
• easy to run, early in design cycle	Small "kernel" benchmarks	• easy to "fool"
 identify peak capability and potential pottlenecks 	Microbenchmarks	 "peak" may be a long way from application performance



SPEC	'95
Benchmark	Description
go	Artificial intelligence: plays the game of Go
m88ksim	Motorola 88k chip simulator; runs test program
qcc	The Gnu C compiler generating SPARC code
compress	Compresses and decompresses file in memory
li	Lisp interpreter
ijpeg	Graphic compression and decompression
perl	Manipulates strings and prime numbers in the special-purpose programming language Perl
vortex	A database program
tomcatv	A mesh generation program
swim	Shallow water model with 513 x 513 grid
su2cor	quantum physics; Monte Carlo simulation
hydro2d	Astrophysics; Hydrodynamic Naiver Stokes equations
marid	Multiorid solver in 3-D potential field
applu	Parabolic/elliptic partial differential equations
trub3d	Simulates isotropic, homogeneous turbulence in a cube
apsi	Solves problems regarding temperature, wind velocity, and distribution of pollutant
foooo	Quantum chemistry
wave5	Plasma physics; electromagnetic particle simulation





S	PEC CPU	2000 (FP)	
	Benchmark	Description	
	wupwise	Physics/Quantum Chromodynamics (F77)	
	swim	Shallow water modeling (F77)	
	mgrid	Multi-grid solver: 3D potential field (F77)	
	applu	Parabolic/elliptic PDEs (F77)	
	mesa	3D graphics library (C)	
	galgel	Computational fluid dynamics (F90)	
	art	Image processing / neural networks (F90)	
	equake	Seismic wave propagation simulation (C)	
	facerec	Image processing: face recognition (F90)	
	ammp	Computational chemistry (C)	
	lucas	Number theory / primality test (F90)	
	fma3D	Finite-element crash simulation (F90)	
	sixtrack	High-energy nuclear physics accelerator design (F77)	
	apsi	Meteorology: pollutant distribution (F77)	



Choosing SPEC Benchmarks				
Good benchmarks • Have many users • Exercise significant hardware resources • Solve interesting technical problems • Generate published results • Add variety to benchmark	Bad benchmarks Can't be ported in a reasonable time Are not compute bound (instead I/O bound) Have unchanged workload from previous SPEC suite Are code fragments instead of complete appe			
suite	Are redundant			

[http://spec.unipv.it/cpu2000/papers /COMPUTER_200007-abstract.JLH.html]

- ads es
- s
- Do different work on different platforms

Calculating Overall SPEC Values						
	Time on A	Time on B	A, norm to A	B, norm to A	A, norm to B	B, norm to B
Prog 1	1	10	1	10	0.1	1
Prog 2	1000	100	1	0.1	10	1
Arith mean	500,5	55	1	5.05	5.05	1

Calcu	Calculating Overall SPEC Values						
	Time on A	Time on B	A, norm to A	B, norm to A	A, norm to B	B, norm to B	
Prog 1	1	10	1	10	0.1	1	
Prog 2	1000	100	1	0.1	10	1	
Arith mean	500.5	55	1	5.05	5.05	1	
Geom mean	31.6	31.6	1	1	1	1	

Calculating Overall SPEC Values

- Normalize execution times to reference machine (SPEC '95: Sun SparcStation 10/40)
- Average normalized execution times
- But what do we mean by "average"?

Calculating The Mean

Arithmetic mean (n items): $\sum_{i=1}^{\infty} \Sigma$ (Execution Time Ratio);] / n

Geometric mean (n items): $[\prod_{i=1 \text{ to } n} (\text{Execution Time Ratio})_i]^{1/n}$

Averaging normalized execution times requires the geometric mean. Using it, mean of ratios and ratio of means gives the same result.

The geometric mean does *not* predict relative execution time.

Discussion Section

CPI Example

Suppose we have two implementations of the same instruction set architecture (ISA).

For some program,

- Machine A has a clock cycle time of 10 ns and a CPI of 2.0
- Machine B has a clock cycle time of 20 ns and a CPI of 1.2
- What machine is faster for this program, and by how much?

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

of Instructions Example

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
- 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.
- $\ensuremath{^\circ}$ 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 example

Two different compilers are being tested for a 100 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.

MIPS example

Which sequence will be faster according to MIPS?

Which sequence will be faster according to execution time?

	ille (Reg / r	leg)		
Ор	Freq	Cycles	CPI(i)	% Time
ALU	50%	1	.5	23%
Load	20%	5	1.0	45%
Store	10%	3	.3	14%
Branch	20%	2	.4	18%
1	ypical Mix		2.2	

Example (RISC processor)

How much faster would the machine be if a better data cache reduced the average load time to 2 cycles?

Example (RISC processor)

How does this compare with using branch prediction to shave a cycle off the branch time?

Example (RISC processor)

What if two ALU instructions could be executed at once?

(different than deleting half the ALU ops!)

Example (Amdahl's Law 1)

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?

Example (Amdahl's Law 2)

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 floatingpoint instructions?

Example (Amdahl's Law 2)

We are looking for a benchmark to show off the new floating-point unit described in Part 2, 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?

Example (Compiler Optimization)

You want to understand the performance of a specific program on your 3.3 GHz machine. You collect the following statistics for the instruction mix and breakdown:

Instruction Class	Frequency (%)	Cycles
Arithmetic/logical	50	1
Load	20	2
Store	10	2
Jump	10	1
Branch	10	3

Part A

Calculate the CPI and MIPS for this program.

Part B

Your compiler team reports they can eliminate 20% of ALU instructions (i.e. 10% of all instructions). What is the speedup?

Part C

With the compiler improvements, what is the new CPI and MIPS?