


CS252  
Graduate Computer Architecture  
Lecture 2

Review of Instruction Sets, Pipelines, and Caches

Prof. David Culler  
Electrical Engineering and Computer Sciences  
University of California, Berkeley

<http://www.eecs.berkeley.edu/~culler/course/cs252-c08>



1/28/08

CS252-663 Lec2

Review, #1

- Technology is changing rapidly:
 

	Capacity	Speed
Logic	2x in 3 years	2x in 3 years
DRAM	4x in 3 years	2x in 10 years
Disk	4x in 3 years	2x in 10 years
Processor	(n.a.)	2x in 1.5 years
- What was true five years ago is not necessarily true now.
- Execution time is the REAL measure of computer performance!
  - Not clock rate, not CPI
- "X is n times faster than Y" means:

$$\frac{\text{ExTime}(y)}{\text{ExTime}(X)} = \frac{\text{Performance}(X)}{\text{Performance}(Y)}$$

1/28/08

CS252-663 Lec2

2

Amdahl's Law

$$\text{ExTime}_{\text{new}} = \text{ExTime}_{\text{old}} \times \left[ (1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}} \right]$$

$$\text{Speedup}_{\text{overall}} = \frac{\text{ExTime}_{\text{old}}}{\text{ExTime}_{\text{new}}} = \frac{1}{(1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}}}$$

Best you could ever hope to do:

$$\text{Speedup}_{\text{maximum}} = \frac{1}{(1 - \text{Fraction}_{\text{enhanced}})}$$



1/28/08

CS252-663 Lec2

3

Today:  
Quick review of everything you should have learned

1/28/08

CS252-663 Lec2

4

Computer Performance



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

	Inst Count	CPI	Clock Rate
Program	X		
Compiler	X	(X)	
Inst. Set.	X	X	
Organization		X	X
Technology			X

1/28/08

CS252-663 Lec2

5

Cycles Per Instruction (Throughput)

"Average Cycles per Instruction"

$$\text{CPI} = (\text{CPU Time} \times \text{Clock Rate}) / \text{Instruction Count} = \text{Cycles} / \text{Instruction Count}$$

$$\text{CPU time} = \text{Cycle Time} \times \sum_{j=1}^n \text{CPI}_j \times I_j$$

$$\text{CPI} = \sum_{j=1}^n \text{CPI}_j \times F_j \quad \text{where } F_j = \frac{I_j}{\text{Instruction Count}}$$

"Instruction Frequency"

1/28/08

CS252-663 Lec2

6

## Example: Calculating CPI bottom up

Run benchmark and collect workload characterization (simulate, machine counters, or sampling)

Op	Freq	Cycles	CPI()	(% Time)
ALU	50%	1	.5	(33%)
Load	20%	2	.4	(27%)
Store	10%	2	.2	(13%)
Branch	20%	2	.4	(27%)
			1.5	

Typical Mix of instruction types in program

Design guideline: Make the common case fast

MIPS 1% rule: only consider adding an instruction if it is shown to add 1% performance improvement on reasonable benchmarks.

1/28/00

CS252-685 Lec3

7

## Example: Branch Stall Impact

- Assume CPI = 1.0 ignoring branches (ideal)
- Assume solution was stalling for 3 cycles
- If 30% branch, Stall 3 cycles on 30%

Op	Freq	Cycles	CPI()	(% Time)
Other	70%	1	.7	(37%)
Branch	30%	4	1.2	(63%)

⇒ new CPI = 1.9

- New machine is  $1/1.9 = 0.52$  times faster (i.e. slow!)

1/28/00

CS252-685 Lec3

8

## SPEC: System Performance Evaluation Cooperative

- First Round 1988
  - 10 programs yielding a single number ("SPECmarks")
- Second Round 1989
  - SPECint92 (8 integer programs) and SPECfp92 (14 floating point programs)
    - Compiler flags unlimited. March 93 of DEC 4000 Model 610:
 

```
apix: unix.c/dad- (spar, has_bcopy, "bcopy(a,b,c) = bcopy(b,a,c) "
```

```
base5: /add- (add, doom-rac) /sg-a/cr-4/cr-200
```

```
base7: /ccacc/sg-a/cr-4/cr2-200/1c-51.a
```
- Third Round 1995
  - new set of programs: SPECint95 (8 integer programs) and SPECfp95 (10 floating point)
  - "benchmarks useful for 3 years"
  - Simple flag setting for all programs: SPECint\_base95, SPECfp\_base95
- Fourth Round 2000: 28 apps
  - analysis and simulation programs
  - Compression: bzip2, gzip,
  - Integrated circuit layout, ray tracing, lots of others

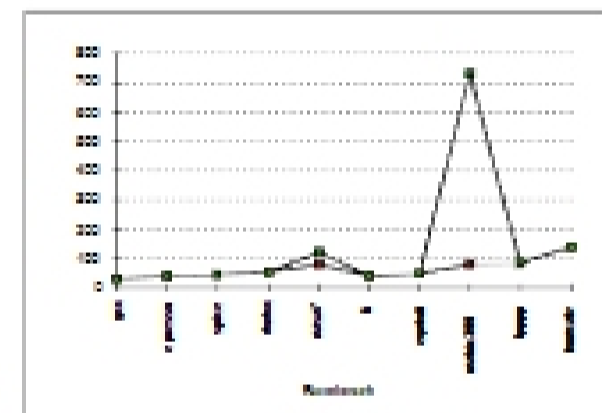
1/28/00

CS252-685 Lec3

9

## SPEC First Round

- One program: 99% of time in single line of code
- New front-end compiler could improve dramatically



1/28/00


10

## Integrated Circuits Costs

$$IC \text{ cost} = \frac{\text{Die cost} + \text{Testing cost} + \text{Packaging cost}}{\text{Final test yield}}$$

$$\text{Die cost} = \frac{\text{Wafer cost}}{\text{Dies per Wafer} \times \text{Die yield}}$$

$$\text{Dies per wafer} = \frac{\pi (\text{Wafer\_diam}/2)^2}{\text{Die\_Area}} = \frac{\pi \times \text{Wafer\_diam}}{\sqrt{2} \cdot \text{Die\_Area}} = \text{Test\_die}$$



$$\text{Die Yield} = \text{Wafer\_yield} \times \left\{ 1 + \left( \frac{\text{Defect\_Density} \times \text{die\_area}}{\pi} \right) \right\}^{-1}$$

Die Cost goes roughly with die area<sup>4</sup>

1/28/00

CS252-685 Lec3

11

## A "Typical" RISC

- 32-bit fixed format instruction (3 formats)
- 32 32-bit GPR (R0 contains zero, DP take pair)
- 3-address, reg-reg arithmetic instruction
- Single address mode for load/store:
  - base + displacement
  - no indirection
- Simple branch conditions
- Delayed branch

see: SPARC, MIPS, HP PA-Risc, DEC Alpha, IBM PowerPC, CDC 6600, CDC 7600, Cray-1, Cray-2, Cray-3

1/28/00

CS252-685 Lec3

12

