

## Lecture 15

- Performance
- Midterm review !!

1

## Performance



- Now we'll discuss issues related to performance:
  - Latency/Response Time/Execution Time vs. Throughput
  - How do you make a reasonable performance comparison?
  - The 3 components of CPU performance
  - The 2 laws of performance

2

## Why know about performance

- Purchasing Perspective:**
  - Given a collection of machines, which has the
    - Best Performance?
    - Lowest Price?
    - Best Performance/Price?
- Design Perspective:**
  - Faced with design options, which has the
    - Best Performance Improvement?
    - Lowest Cost?
    - Best Performance/Cost?
- Both require**
  - Basis for comparison
  - Metric for evaluation

3

## Many possible definitions of performance

- Every computer vendor will select one that makes them look good. How do you make sense of conflicting claims?



**Q:** Why do end users need a new performance metric?  
**A:** End users who rely only on MHz as an indicator for performance do not have a complete picture of PC processor performance and may pay the price of missed expectations.

*4x*  
*9/10*

4

## Two notions of performance

Plane	DC to Paris	Speed	Passengers	Throughput (pmp)
747	6.5 hours	610 mph	470	286,700
Concorde	3 hours	1350 mph	132	178,200

- Which has higher performance?
  - Depends on the **metric**
    - Time to do the task (Execution Time, Latency, Response Time)
    - Tasks per unit time (Throughput, Bandwidth)
  - Response time and throughput are often in opposition

*parallelism*

5

## Some Definitions

- Performance is in units of things/unit time
  - E.g., Hamburgers/hour
  - Bigger is better
- If we are primarily concerned with response time
  - Performance(x) =  $\frac{1}{\text{execution\_time}(x)}$
- Relative performance: "X is N times faster than Y"

$$N = \frac{\text{Performance}(X)}{\text{Performance}(Y)} = \frac{\text{execution\_time}(Y)}{\text{execution\_time}(X)} = \frac{T_{\text{old}}}{T_{\text{new}}} = 5$$

*New is 5 times*

6

*SPECint* *SPECfp* *SPECint* *SPECfp* *SPECint* *SPECfp*

### Basis of Comparison

- When comparing systems, need to fix the workload
  - Which workload?

Workload	Pros	Cons
Actual Target Workload	Representative	Very specific ✓ Non-portable ✓ Difficult to run/measure ✓
Full Application Benchmarks	Portable Widely used Realistic	Less representative
Small "Kernel" or "Synthetic" Benchmarks	Easy to run Useful early in design	Easy to "fool"
Microbenchmarks	Identify peak capability and potential bottlenecks	Real application performance may be much below peak

10

7

### Benchmarking

- Some common benchmarks include:
  - Adobe Photoshop for image processing
  - BAPCo Sysmark for office applications
  - Unreal Tournament 2003 for 3D games
  - SPEC2000 for CPU performance
- The best way to see how a system performs for a variety of programs is to just show the execution times of all of the programs.
- Here are execution times for several different Photoshop 5.5 tasks, from <http://www.tech-report.com>

8

### Summarizing performance

- Summarizing performance with a single number can be misleading—just like summarizing four years of school with a single GPA!
- If you must have a single number, you could **sum** the execution times. This example graph displays the total execution time of the individual tests from the previous page.
- A similar option is to find the **average** of all the execution times. For example, the 800MHz Pentium III (in yellow) needed 227.3 seconds to run 21 programs, so its average execution time is  $227.3/21 = 10.82$  seconds.
- A **weighted** sum or average is also possible, and lets you emphasize some benchmarks more than others.

9

### The components of execution time

- Execution time can be divided into two parts.
  - User time** is spent running the application program itself.
  - System time** is when the application calls operating system code.
- The distinction between user and system time is not always clear, especially under different operating systems.
- The Unix **time** command shows both.

```

salary.125 > time distill 05-examples.ps
Distilling 05-examples.ps (449,119 bytes)
10.8 seconds (0:11)
449,119 bytes PS => 94,999 bytes PDF (21%)
10.61u 0.98s 0:15.15 76.5%
  
```

10

### Three Components of CPU Performance

$$\text{CPU time}_{x,p} = \text{Instructions executed}_{x,p} \cdot \text{CPI}_{x,p} \cdot \text{Clock cycle time}_x$$

Cycles Per Instruction

*CPI = 1 single-cycle*

*CPIs ≥ 1 pipelined*

11

### Instructions Executed

- Instructions executed:
  - We are not interested in the **static instruction count**, or how many lines of code are in a program.
  - Instead we care about the **dynamic instruction count**, or how many instructions are actually executed when the program runs.
- There are three lines of code below, but the number of instructions executed would be 2001.

```

li $a0, 1000
Ostrich: sub $a0, $a0, 1
         bne $a0, $0, Ostrich
  
```

1000 times

12

### CPI (Review)

- The average number of clock cycles per instruction, or **CPI**, is a function of the machine and program.
  - The CPI depends on the actual instructions appearing in the program— a floating-point intensive application might have a higher CPI than an integer-based program.
  - It also depends on the CPU implementation. For example, a Pentium can execute the same instructions as an older 80486, but faster.
- Initially we assumed each instruction took one cycle, so we had CPI = 1.
  - The CPI can be >1 due to memory stalls and slow instructions.
  - The CPI can be <1 on machines that execute more than 1 instruction per cycle (superscalar).

13

### Execution time, again

$$\text{CPU time}_{x,p} = \text{Instructions executed}_p \cdot \text{CPI}_{x,p} \cdot \text{Clock cycle time}_x$$

- The easiest way to remember this is match up the units:

$$\frac{\text{Seconds}}{\text{Program}} = \frac{\text{Instructions}}{\text{Program}} \cdot \frac{\text{Clock cycles}}{\text{Instructions}} \cdot \frac{\text{Seconds}}{\text{Clock cycle}}$$

- Make things faster by making any component smaller!!

	Program	Compiler	ISA	Organization	Technology
Instruction Executed	✓	✓	✓	✗	✗
CPI	✓	✓	✓	✓	✗
Clock Cycle Time	✗	✗	✗	✓	✓

- Often easy to reduce one component by increasing another

14

### Example: Comparing across ISAs

- Intel's Itanium (IA-64) ISA is designed to facilitate executing multiple instructions per cycle. If an Itanium processor achieves an average CPI of .3 (3 instructions per cycle), how much faster is it than a Pentium4 (which uses the x86 ISA) with an average CPI of 1? (assume same freq)

- Itanium is three times faster
- Itanium is one third as fast
- Not enough information

Same program

different ISAs

15

### Improving CPI

- Many processor design techniques we'll see improve CPI
  - Often they only improve CPI for certain types of instructions

$$\text{CPI} = \sum_{i=1}^n \text{CPI}_i \cdot F_i \quad \text{where } F_i = \frac{I_i}{\text{Instruction Count}}$$

- $F_i$  = Fraction of instructions of type  $i$

- First Law of Performance:

**Make the common case fast**

16

### Example: CPI improvements

- Base Machine:

Op Type	Freq (fi)	Cycles	CPIi
ALU	50%	3	1.5
Load	20%	5	1
Store	10%	3	0.3
Branch	20%	2	0.4

$0.25$   
 $0.6$   
 $0.2$   
 $2.8$

- How much faster would the machine be if:
  - we added a cache to reduce average load time to 3 cycles?
  - we added a branch predictor to reduce branch time by 1 cycle?
  - we could do two ALU operations in parallel?

17

### Amdahl's Law

- Amdahl's Law states that optimizations are limited in their effectiveness.

$$\text{Execution time after improvement} = \frac{\text{Time affected by improvement}}{\text{Amount of improvement}} + \text{Time unaffected by improvement}$$

- For example, doubling the speed of floating-point operations sounds like a great idea. But if only 10% of the program execution time  $T$  involves floating-point code, then the overall performance improves by just 5%.

$$\text{Execution time after improvement} = \frac{0.10 T}{2} + 0.90 T = 0.95 T$$

- What is the maximum speedup from improving floating point?

- Second Law of Performance:

**Make the fast case common**

$$T_{\text{new}} = \frac{0.10 T}{2} + 0.90 T = 0.95 T$$

## Summary

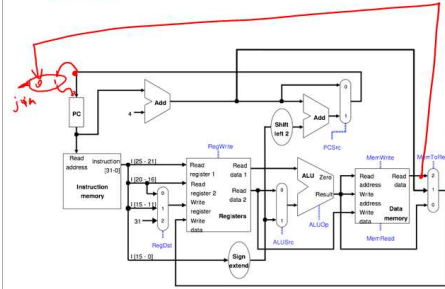
- **Performance** is one of the most important criteria in judging systems.
- There are two main measurements of performance.
  - **Execution time** is what we'll focus on.
  - **Throughput** is important for servers and operating systems.
- Our main performance equation explains how performance depends on several factors related to both hardware and software.
 
$$\text{CPU time}_{x,p} = \text{Instructions executed}_p \cdot \text{CPI}_{x,p} \cdot \text{Clock cycle time}_x$$
- It can be hard to measure these factors in real life, but this is a useful guide for comparing systems and designs.
- **Amdahl's Law** tell us how much improvement we can expect from specific enhancements.
- The best **benchmarks** are real programs, which are more likely to reflect common instruction mixes.

19

Not Gih

## Midterm Review -- datapath

Let's add a Jump Indirect instruction to the single-cycle processor developed in class. This I-format instruction `jalr imm16(rn)` will cause the processor to jump to the address stored in the word at memory location `imm16 + R[rn]` (the same address computed by `lw` and `sw`).



20

## hazards

```
add s7 <- s3, s4
or s8 <- s2, s4
lw s6 <- s4
addi s7 <- s3, 0x5
sub s8 <- s8, s4
```



(a) Draw arrows on the instructions above indicating all the data dependencies.

(b) Reorder the instructions into a new schedule that will execute without any stalls on a 5-stage pipelined processor with forwarding. For reference, you may refer to the pipeline diagram below.

```
add
or
lw
sub
addi
```

(c) Reorder the instructions into a new schedule that will execute without any stalls on a 5-stage pipelined processor without forwarding. For reference, you may refer to the pipeline diagram below.



21

## Calling conventions?

22