


CMSC 611: Advanced Computer Architecture

Benchmarking

Performance Benchmarks

- Many widely-used benchmarks are small programs that have significant locality of instruction and data reference
- Universal benchmarks can be misleading since hardware and compiler vendors do optimize their design for these programs
- The best types of benchmarks are real applications since they reflect the end-user interest
- Architectures might perform well for some applications and poorly for others
- Compilation can boost performance by taking advantage of architecture-specific features
- Application-specific compiler optimization are becoming more popular

Effect of Compilation



App. and arch. specific optimization can dramatically impact performance

The SPEC Benchmarks

- SPEC stands for System Performance Evaluation Cooperative suite of benchmarks
 - Created by a set of companies to improve the measurement and reporting of CPU performance
- SPEC2006 is the latest suite that consists of 12 integer (C or C++) and 17 floating-point (Fortran, C, and C++) programs
 - Customized SPEC suites assess performance of graphics and transaction systems.
- Since SPEC requires running applications on real hardware, the memory system has a significant effect on performance

Performance Reports

Hardware	
Model number	Powerstation 550
CPU	41.67-MHz POWER 4164
FPU (floating point)	Integrated
Number of CPU	1
Cache size per CPU	64K data/8k instruction
Memory	64 MB
Disk subsystem	2 400-MB SCSI
Network interface	N/A
Software	
OS type and revision	AIX Ver. 3.1.5
Compiler revision	AIX XL C/6000 Ver. 1.1.5 AIX XL Fortran Ver. 2.2
Other software	None
File system type	AIX
Firmware level	N/A
System	
Tuning parameters	None
Background load	None
System state	Multi-user (single-user login)



Guiding principle is *reproducibility* (report environment & experiments setup)

The SPEC Benchmarks

$$\text{SPEC ratio} = \frac{\text{Execution time on SUN SPARCstation 10/40}}{\text{Execution time on the measure machine}}$$

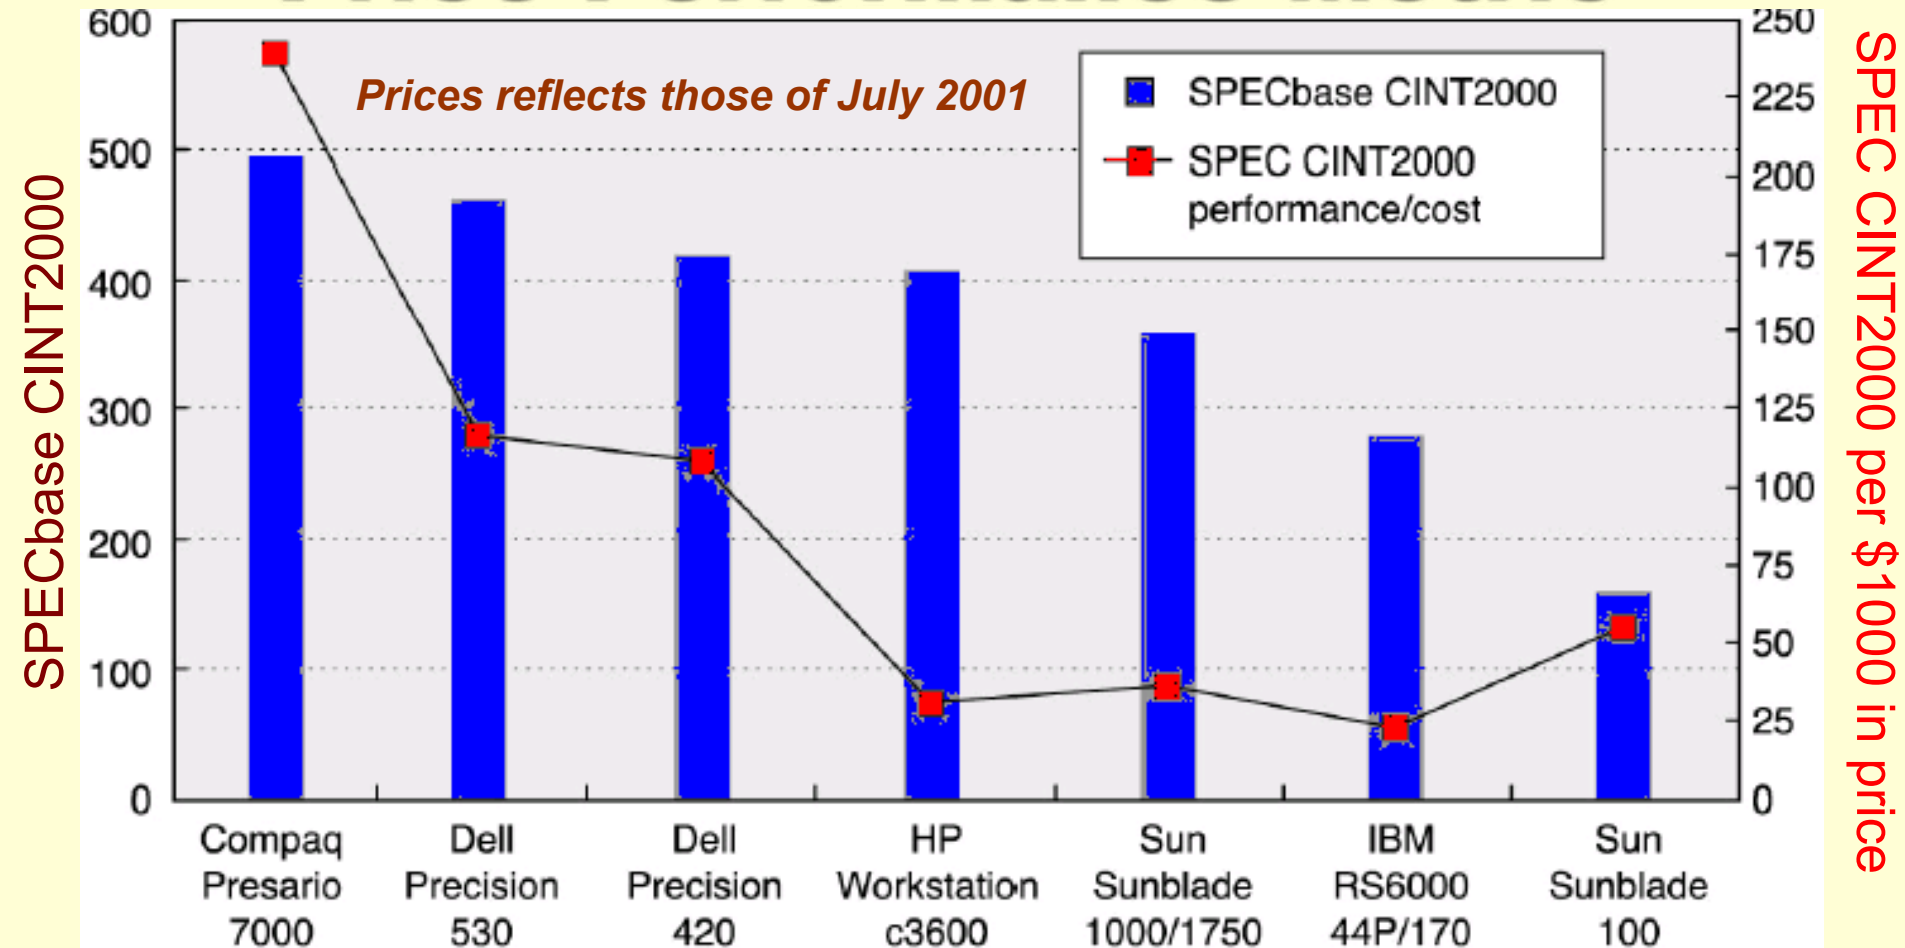
- Bigger numeric values of the SPEC ratio indicate faster machine
- 1997 “historical” reference machine

SPEC95 for Pentium and Pentium Pro



- The performance measured may be different on other Pentium-based hardware with different memory system and using different compilers
 - At the same clock rate, the SPECint95 measure shows that Pentium Pro is 1.4-1.5 times faster while the SPECfp95 shows that it is 1.7-1.8 times faster
 - When the clock rate is increased by a certain factor, the processor performance increases by a lower factor

Price-Performance Metric



- Different results are obtained for other benchmarks, e.g. SPEC CFP2000
- With the exception of the Sunblade price-performance metrics were consistent with performance

Historic Perspective

- In early computers most instructions of a machine took the same execution time
 - The measure of performance for old machines was the time required performing an individual operation (e.g. addition)
- New computers have diverse set of instructions with different execution times
 - The relative frequency of instructions across many programs was calculated
 - The average instruction execution time was measured by multiplying the time of each instruction by its frequency
- The average instruction execution time was a small step to MIPS that grew in popularity

Using MIPS

- MIPS = Million of Instructions Per Second
 - one of the simplest metrics
 - valid only in a limited context

$$\text{MIPS (native MIPS)} = \frac{\text{Instruction count}}{\text{Execution time} \times 10^6}$$

- There are three problems with MIPS:
 - MIPS specifies the instruction execution rate but not the capabilities of the instructions
 - MIPS varies between programs on the same computer
 - MIPS can vary inversely with performance (see next example)

The use of MIPS is simple and intuitive, faster machines have bigger MIPS

Example

Consider the machine with the following three instruction classes and CPI:

Instruction class	CPI for this instruction class
A	1
B	2
C	3

Now suppose we measure the code for the same program from two different compilers and obtain the following data:

Code from	Instruction count in (billions) for each instruction class		
	A	B	C
Compiler 1	5	1	1
Compiler 2	10	1	1

Assume that the machine's clock rate is 500 MHz. Which code sequence will execute faster according to MIPS? According to execution time?

Answer:

Using the formula:
$$\text{CPU clock cycles} = \sum_{i=1}^n \text{CPI}_i \times C_i$$

Sequence 1: CPU clock cycles = $(5 \times 1 + 1 \times 2 + 1 \times 3) \times 10^9 = 10 \times 10^9$ cycles
Sequence 2: CPU clock cycles = $(10 \times 1 + 1 \times 2 + 1 \times 3) \times 10^9 = 15 \times 10^9$ cycles

Example (Cont.)

Using the formula: Execution time = $\frac{\text{CPU clock cycles}}{\text{Clock rate}}$

Sequence 1: Execution time = $(10 \times 10^9) / (500 \times 10^6) = 20$ seconds

Sequence 2: Execution time = $(15 \times 10^9) / (500 \times 10^6) = 30$ seconds

Therefore compiler 1 generates a faster program

Using the formula: MIPS = $\frac{\text{Instruction count}}{\text{Execution time} \times 10^6}$

Sequence 1: MIPS = $\frac{(5 + 1 + 1) \times 10^9}{20 \times 10^6} = 350$

Sequence 2: MIPS = $\frac{(10 + 1 + 1) \times 10^9}{30 \times 10^6} = 400$

Although compiler 2 has a higher MIPS rating, the code from generated by compiler 1 runs faster