

# Computer Science 146

## Computer Architecture

---

Fall 2019

Harvard University

Instructor: Prof. David Brooks

dbrooks@eecs.harvard.edu

Lecture 2: Performance Metrics, ISA Intro

---

Computer Science 146  
David Brooks

## Lecture Outline

---

- Performance Metrics
- Averaging
- Amdahl's Law
- Benchmarks
- The CPU Performance Equation
  - Optimal Pipelining Case Study
- Modern Processor Analysis
- Begin ISAs...

---

Computer Science 146  
David Brooks

# Performance Metrics

---

- Execution Time is often what we target
- Throughput (tasks/sec) vs. latency (sec/task)
- How do we decide the tasks? Benchmarks
  - What the customer cares about, real applications
  - Representative programs (SPEC, SYSMARK, etc)
  - Kernels: Code fragments from real programs (Linpack)
  - Toy Programs: Sieve, Quicksort
  - Synthetic Programs: Just a representative instruction mix (Whetstone, Dhrystone)

Better  


# Measuring Performance

---

- Total Execution Time:

$$\frac{1}{n} \sum_{i=1}^n \text{Time}_i$$

- This is *arithmetic* mean
  - This should be used when measuring performance in execution *times* (CPI)

# Measuring Performance

---

- Weighted Execution Time:

$$\sum_{i=1}^n \text{Weight}_i \times \text{Time}_i$$

- What if P1 and P2 are not run equally?

# Measuring Performance

---

- Normalized Execution Time
- Normalize to *reference* machine
- Can only use geometric mean (arithmetic mean can vary depending on the reference machine)

$$\sqrt[n]{\prod_{i=1}^n \text{ExecutionTimeRatio}_i}$$

- Problem: Ratio not Execution Time is the result

## Harmonic Mean: Motivation

---

- 30 mph for the first 10 miles
- 90 mph for the next 10 miles
- Average speed?  $(30+90)/2 = 60\text{mph}$
  
- WRONG! Average speed = total distance / total time
- $20/(10/30+10/90) = 45\text{mph}$

## Harmonic Mean

---

- Each program has  $O$  operations
- $n$  programs executed  $nO$  operations in  $\sum T_i$
- Execution rate is then

$$nO / \sum T_i = n / \sum (T_i / O) = n / \sum 1/P_i$$

where  $1/P_i$  is the rate of execution of program  $i$

- Harmonic mean should be used when measuring performance in execution *rates* (IPC)

## Amdahl's Law (Law of Diminishing Returns)

---

- Very Intuitive – Make the Common case fast

$$\text{Speedup} = \frac{\text{Execution Time for task without enhancement}}{\text{Execution Time for task using enhancement}}$$

$$\text{Execution time}_{\text{new}} = \text{Execution time}_{\text{old}} \times$$

$$\left( (1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}} \right)$$

---

Computer Science 146  
David Brooks

## Amdahl's Law Corollary

---

$$\text{Speedup}_{\text{Overall}} = \frac{1}{\left( (1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}} \right)}$$

$$\text{As Speedup}_{\text{Enhanced}} \gg 0, \text{Speedup}_{\text{Overall}} = \frac{1}{(1 - \text{Fraction}_{\text{enhanced}})}$$

---

Computer Science 146  
David Brooks

# Amdahl's Law Example

---

$$\text{Fraction}_{\text{enhanced}}=95\%, \text{Speedup}_{\text{Enhanced}}=1.1x$$

$$\text{Speedup}_{\text{Overall}}=1/((1-.95)+(.95/1.1))=1.094$$

$$\text{Fraction}_{\text{enhanced}}=5\%, \text{Speedup}_{\text{Enhanced}}=10x$$

$$\text{Speedup}_{\text{Overall}}=1/((1-.05)+(.05/10))=1.047$$

Make the common case fast!

$$\text{Fraction}_{\text{enhanced}}=5\%, \text{Speedup}_{\text{Enhanced}}=\text{Infinity}$$

$$\text{Speedup}_{\text{Overall}}=1/(1-.05)=1.052$$

# MIPS

---

- $\text{MIPS} = \text{instruction count}/(\text{execution time} \times 10^6)$   
 $= \text{clock rate}/(\text{CPI} \times 10^6)$

- Problems
  - ISAs are not equivalent, e.g. RISC vs. CISC
    - 1 CISC instruction may equal many RISC!
  - Programs use different instruction mixes
  - May be ok when comparing same benchmarks, same ISA, same compiler, same OS

# MFLOPS

---

- Same as MIPS, just FP ops
- Not useful either
  - FP-intensive apps needed
  - Traditionally, FP ops were slow, INT can be ignored
  - BUT, now memory ops can be the slowest!
- “Peak MFLOPS” is a common marketing fallacy
  - Basically, it just says #FP-pipes X Clock Rate

---

Computer Science 146  
David Brooks

# GHz

---

- Is this a metric? Maybe as good as the others...
- One number, no benchmarks, what can be better?
- Many designs *are* frequency driven

Processor	Clock Rate	SPEC FP2000
IBM POWER3	450 MHz	434
Intel PIII	1.4 GHz	456
Intel Pentium 4	2.4 GHz	833
Itanium-2	1.0 GHz	1356

---

Computer Science 146  
David Brooks

# Benchmark Suites

---

- SPEC CPU2000 (int and float) (Desktop, Server)
- EEMBC (“embassy”), SPECjvm (Embedded)
- TPC-C, TPC-H, SPECjbb, ECperf (Server)

---

Computer Science 146  
David Brooks

# SPEC CPU2000: Integer Benchmarks

---

164.gzip	C	Compression
175.vpr	C	FPGA Circuit Placement and Routing
176.gcc	C	C Programming Language Compiler
181.mcf	C	Combinatorial Optimization
186.crafty	C	Game Playing: Chess
197.parser	C	Word Processing
252.eon	C++	Computer Visualization
253.perlbnk	C	PERL Programming Language
254.gap	C	Group Theory, Interpreter
255.vortex	C	Object-oriented Database
256.bzip2	C	Compression
300.twolf	C	Place and Route Simulator

---

Computer Science 146  
David Brooks



## SPEC CPU2000: Floating Point Benchmarks

168.wupwise	Fortran 77	Physics / Quantum Chromodynamics
171.swim	Fortran 77	Shallow Water Modeling
172.mgrid	Fortran 77	Multi-grid Solver: 3D Potential Field
173.applu	Fortran 77	Parabolic / Elliptic Partial Differential Equations
177.mesa	C	3-D Graphics Library
178.galgel	Fortran 90	Computational Fluid Dynamics
179.art	C	Image Recognition / Neural Networks
183.quake	C	Seismic Wave Propagation Simulation
187.facerec	Fortran 90	Image Processing: Face Recognition
188.amp	C	Computational Chemistry
189.lucas	Fortran 90	Number Theory / Primality Testing
191.fma3d	Fortran 90	Finite-element Crash Simulation
200.sixtrack	Fortran 77	High Energy Nuclear Physics Accelerator Design
301.apsi	Fortran 77	Meteorology: Pollutant Distribution

Computer Science 146  
David Brooks

## Server Benchmarks

- TPC-C (Online-Transaction Processing, OLTP)

System	# / Processor	tpm	\$/tpm
Fujitsu PrimePower	128, 563MHz SPARC64	455K	\$28.58
HP SuperDome	64, 875MHz PA8700	423K	\$15.64
IBM p690	32, 1300MHz POWER4	403K	\$17.80


- TPC-H (Ad-hoc, decision support)

Computer Science 146  
David Brooks

# CPU Performance Equation

- Execution Time = **seconds/program**

$$\frac{\text{instructions}}{\text{program}} \times \frac{\text{cycles}}{\text{instruction}} \times \frac{\text{seconds}}{\text{cycle}}$$



Program Architecture (ISA) Compiler	Compiler (Scheduling) Organization (uArch) Microarchitects	Technology Physical Design Circuit Designers
---	--	--

Computer Science 146  
David Brooks


# Common Architecture Trick

- Instructions/Program (Path-length) is constant
  - Same benchmark, same compiler
  - Ok usually, but for some ideas compiler may change
- Seconds/Cycle (Cycle-time) is constant
  - “My tweak won’t impact cycle-time”
  - Often a bad assumption
  - Current designs are ~15-20FO4 Inverter Delays per cycle
- Just focus on Cycles/Instruction (CPI or IPC)
- Most academic architecture studies do just this!

Computer Science 146  
David Brooks

# CPU Performance Equation: Case Study

- Fundamental Architecture Decision: How do we choose the number of pipeline stages?



$$\frac{\text{instructions}}{\text{program}} \times \frac{\text{cycles}}{\text{instruction}} \times \frac{\text{seconds}}{\text{cycle}}$$

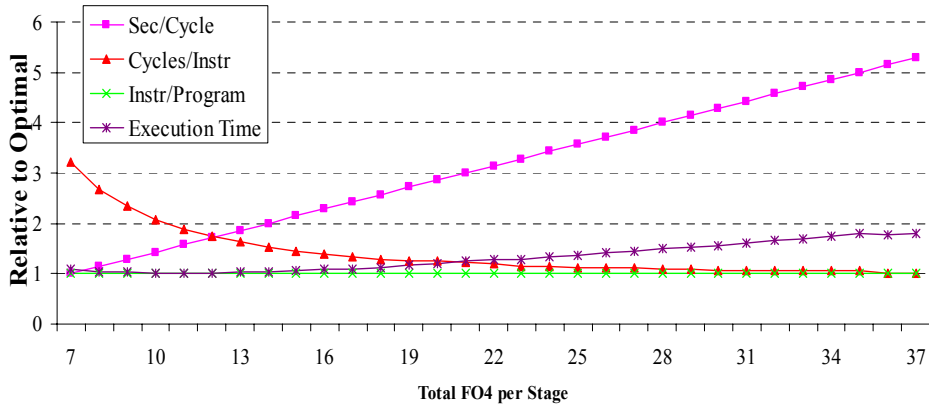
Computer Science 146  
David Brooks

# Measuring Instruction Counts and CPI

- Hardware counters on processors
- Instrumented execution (ATOM)
- Instruction-level interpretation (instruction counts)
  
- Execution-driven simulation
  - Detailed simulation of execution core and memory hierarchy

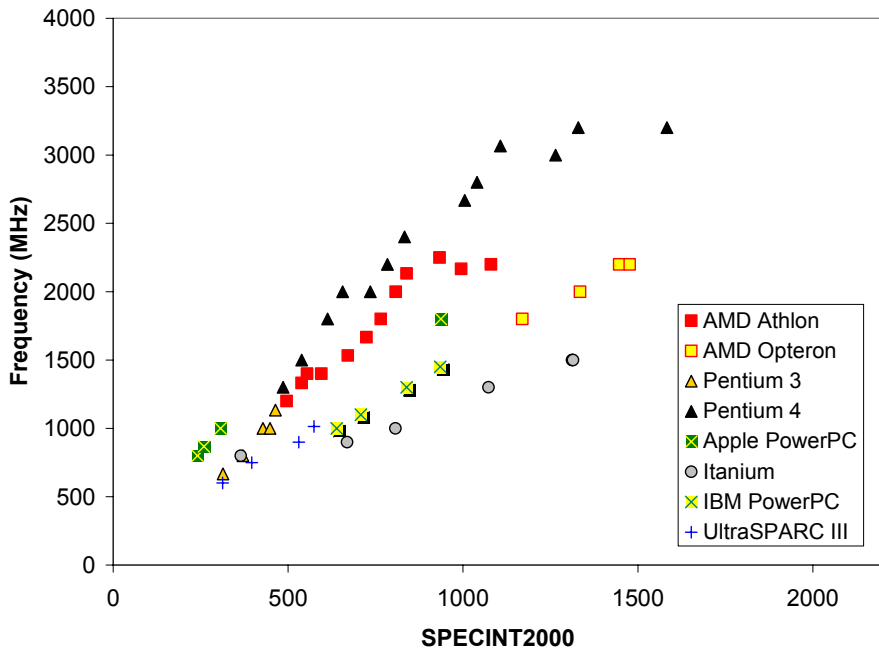
Computer Science 146  
David Brooks

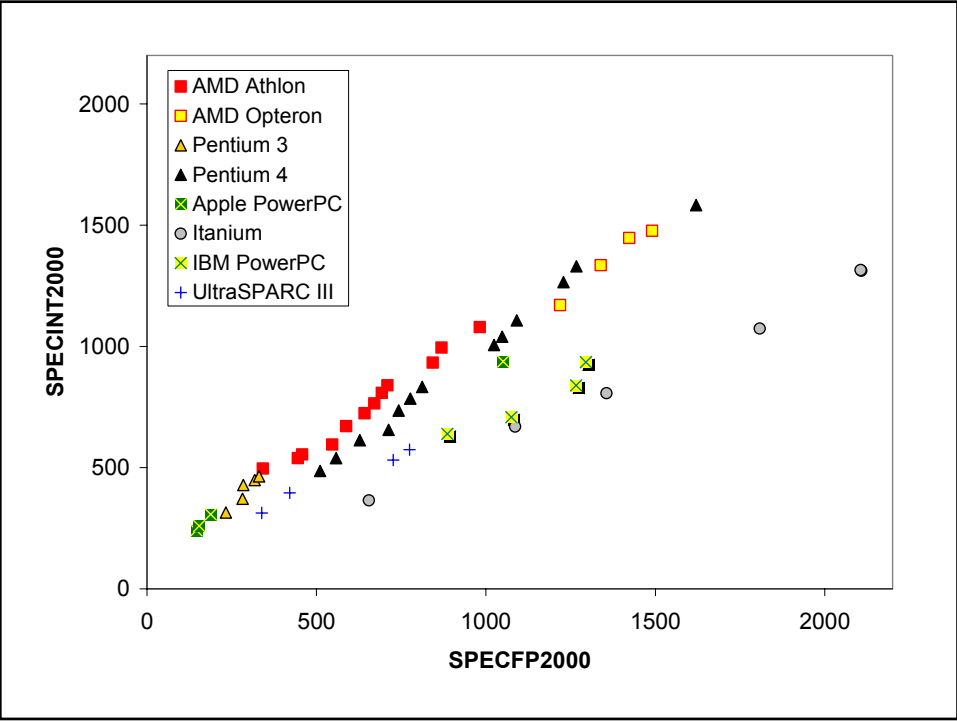
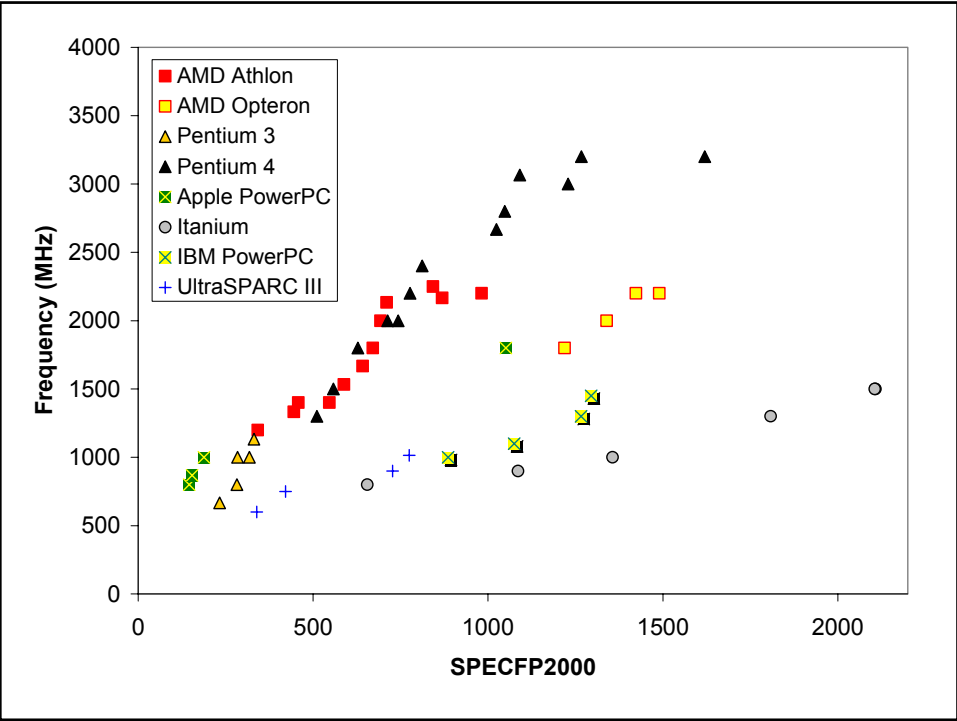
# Pipelining Case Study

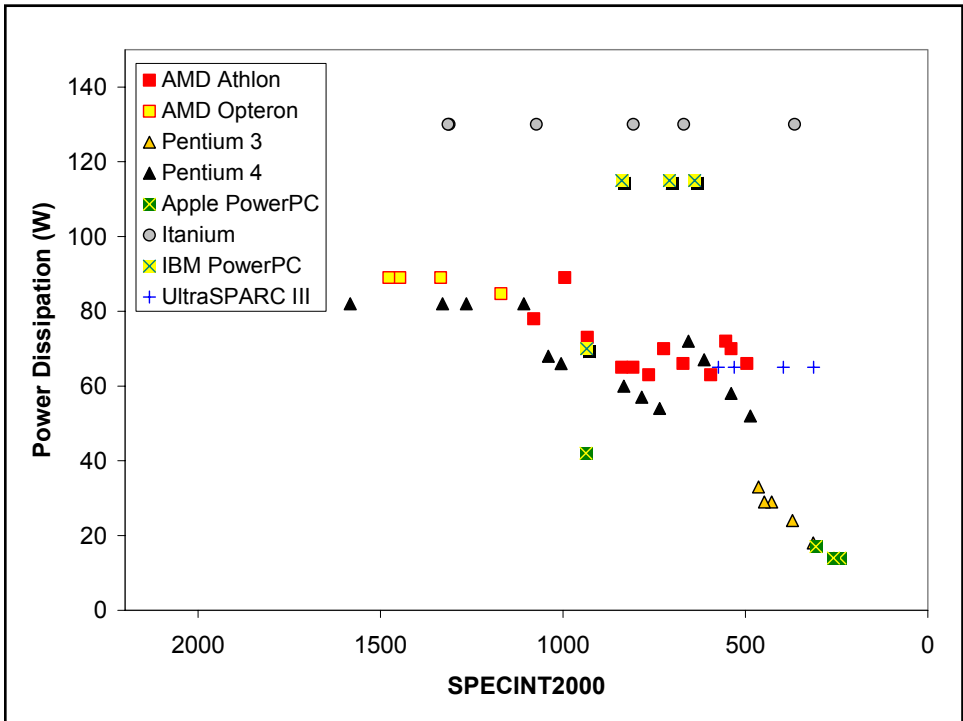
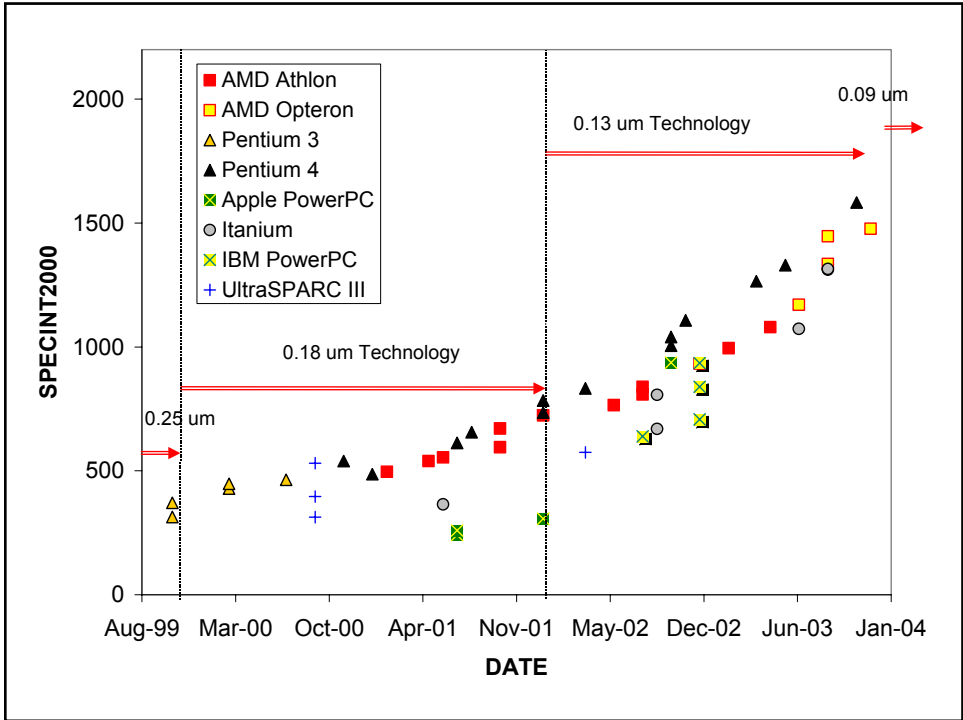


- Caveat: Fixed architecture sizings (only latencies vary)

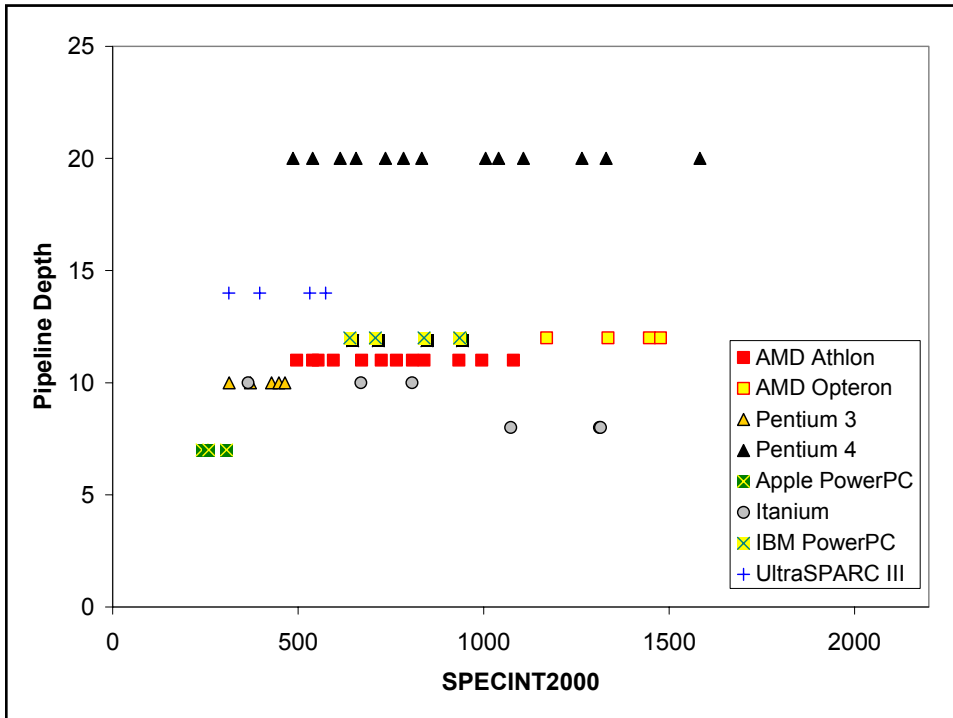
Computer Science 146  
David Brooks











## For next time

- Some CISC vs. RISC commentary
- Multimedia ISAs
- Paper available on webpage
  - <http://www.eecs.harvard.edu/cs146/>
  - **R. Lee, "Subword Parallelism with MAX-2", In IEEE Micro, August 1996.**
- Read through Ch1 and Ch2...
- Compiler/ISA overview, may start pipelining review