



EECS 252 Graduate Computer Architecture

Lec 1 - Introduction

David Patterson
Electrical Engineering and Computer Sciences
University of California, Berkeley

<http://www.eecs.berkeley.edu/~pattsrn>
<http://www-inst.eecs.berkeley.edu/~cs252>

Outline

- Computer Science at a Crossroads
- Computer Architecture v. Instruction Set Arch.
- What Computer Architecture brings to table

4/16/2008

CS252-s06, Lec 01-intro

3



Crossroads: Conventional Wisdom in Comp. Arch

- Old Conventional Wisdom: Power is free, Transistors expensive
 - New Conventional Wisdom: **"Power wall"** Power expensive, Xtors free (Can put more on chip than can afford to turn on)
 - Old CW: Sufficiently increasing Instruction Level Parallelism via compilers, innovation (Out-of-order, speculation, VLIW, ...)
 - New CW: **"ILP wall"** law of diminishing returns on more HW for ILP
 - Old CW: Multiplies are slow, Memory access is fast
 - New CW: **"Memory wall"** Memory slow, multiplies fast (200 clock cycles to DRAM memory, 4 clocks for multiply)
 - Old CW: Uniprocessor performance 2X / 1.5 yrs
 - New CW: Power Wall + ILP Wall + Memory Wall = **Brick Wall**
 - Uniprocessor performance now 2X / 5(?) yrs
- ⇒ Sea change in chip design: multiple "cores"
(2X processors per chip / ~ 2 years)
» More simpler processors are more power efficient


4/16/2008

CS252-s06, Lec 01-intro

4



Crossroads: Uniprocessor Performance



4/16/2008


CS252-s06, Lec 01-intro

5



Sea Change in Chip Design

- Intel 4004 (1971): 4-bit processor, 2312 transistors, 0.4 MHz, 10 micron PMOS, 11 mm² chip
- RISC II (1983): 32-bit, 5 stage pipeline, 40,760 transistors, 3 MHz, 3 micron NMOS, 60 mm² chip
- 125 mm² chip, 0.065 micron CMOS = 2312 RISC II+FPU+Icache+Dcache
 - RISC II shrinks to ~ 0.02 mm² at 65 nm
 - Caches via DRAM or 1 transistor SRAM (www.t-ram.com) ?
 - Proximity Communication via capacitive coupling at > 1 TB/s ? (Ivan Sutherland @ Sun / Berkeley)



• Processor is the new transistor?

4/16/2008

CS252-s06, Lec 01-intro

6

Déjà vu all over again?

- Multiprocessors imminent in 1970s, '80s, '90s, ...
- "... today's processors ... are nearing an impasse as technologies approach the speed of light.."
 - David Mitchell, *The Transputer: The Time Is Now* (1989)
- Transputer was premature
 - ⇒ Custom multiprocessors strove to lead uniprocessors
 - ⇒ Procrastination rewarded: 2X seq. perf. / 1.5 years
- "We are dedicating all of our future product development to multicore designs. ... This is a sea change in computing"
 - Paul Otellini, President, Intel (2004)
- Difference is all microprocessor companies switch to multiprocessors (AMD, Intel, IBM, Sun; all new Apples 2 CPUs)
 - ⇒ Procrastination penalized: 2X sequential perf. / 5 yrs
 - ⇒ Biggest programming challenge: 1 to 2 CPUs

4/16/2008

CS252-s06, Lec 01-intro

7



Problems with Sea Change

- Algorithms, Programming Languages, Compilers, Operating Systems, Architectures, Libraries, ... not ready to supply Thread Level Parallelism or Data Level Parallelism for 1000 CPUs / chip,
- Architectures not ready for 1000 CPUs / chip
 - Unlike Instruction Level Parallelism, cannot be solved by just by computer architects and compiler writers alone, but also cannot be solved *without* participation of computer architects
- This edition of CS 252 (and 4th Edition of textbook *Computer Architecture: A Quantitative Approach*) explores shift from Instruction Level Parallelism to Thread Level Parallelism / Data Level Parallelism

4/16/2008

CS252-s06, Lec 01-intro

8



ISA vs. Computer Architecture

- Old definition of computer architecture = instruction set design
 - Other aspects of computer design called implementation
 - Insinuates implementation is uninteresting or less challenging
- Our view is computer architecture >> ISA
- Architect's job much more than instruction set design; technical hurdles today *more* challenging than those in instruction set design
- Since instruction set design not where action is, some conclude computer architecture (using old definition) is not where action is
 - We disagree on conclusion
 - Agree that ISA not where action is (ISA in CA:AQA 4/e appendix)

4/16/2008

CS252-s06, Lec 01-intro

13

Comp. Arch. is an Integrated Approach

- What really matters is the functioning of the complete system
 - hardware, runtime system, compiler, operating system, and application
 - In networking, this is called the “End to End argument”
- Computer architecture is not just about transistors, individual instructions, or particular implementations
 - E.g., Original RISC projects replaced complex instructions with a compiler + simple instructions

4/16/2008

CS252-s06, Lec 01-intro

14

New Project opportunity this semester

- FPGAs as New Research Platform
- As ~ 25 CPUs can fit in Field Programmable Gate Array (FPGA), 1000-CPU system from ~ 40 FPGAs?
 - 64-bit simple “soft core” RISC at 100MHz in 2004 (Virtex-II)
 - FPGA generations every 1.5 yrs; 2X CPUs, 2X clock rate
- HW research community does logic design (“gate shareware”) to create out-of-the-box, Massively Parallel Processor runs standard binaries of OS, apps
 - Gateware: Processors, Caches, Coherency, Ethernet Interfaces, Switches, Routers, ... (IBM, Sun have donated processors)
 - E.g., 1000 processor, IBM Power binary-compatible, cache-coherent supercomputer @ 200 MHz; fast enough for research

4/16/2008

CS252-s06, Lec 01-intro

26

RAMP

- Since goal is to ramp up research in multiprocessing, called **R**esearch **A**ccelerator for **M**ultiple **P**rocessors
 - To learn more, read “RAMP: Research Accelerator for Multiple Processors - A Community Vision for a Shared Experimental Parallel HW/SW Platform,” Technical Report UCB//CSD-05-1412, Sept 2005
 - Web page ramp.eecs.berkeley.edu

4/16/2008

CS252-s06, Lec 01-intro

27

Why RAMP Good for Research?

	SMP	Cluster	Simulate	RAMP
Cost (1000 CPUs)	F (\$40M)	C (\$2M)	A+ (\$0M)	A (\$0.1M)
Cost of ownership	A	D	A	A
Scalability	C	A	A	A
Power/Space (kilowatts, racks)	D (120 kw, 12 racks)	D (120 kw, 12 racks)	A+ (.1 kw, 0.1 racks)	A (1.5 kw, 0.3 racks)
Community	D	A	A	A
Observability	D	C	A+	A+
Reproducibility	B	D	A+	A+
Flexibility	D	C	A+	A+
Credibility	A+	A+	F	A
Perform. (clock)	A (2 GHz)	A (3 GHz)	F (0 GHz)	C (0.2 GHz)
GPA	C	B+	B	A ₂₈

4/16/2008

CS252-s06, Lec 01-intro

28




RAMP 1 Hardware

- Completed Dec. 2004 (14x17 inch 22-layer PCB)

- Module:

- FPGAs, memory, 10GigE conn.
- Compact Flash
- Administration/maintenance ports:
 - » 10/100 Enet
 - » HDMI/DVI
 - » USB
- ~4K/module w/o FPGAs or DRAM



□ Called "BEE2" for Berkeley Emulation Engine 2


4/16/2008

CS252-s06, Lec 01-intro

29



Multiple Module RAMP 1 Systems



- 8 compute modules (plus power supplies) in 8U rack mount chassis
 - 500-1000 emulated processors
- Many topologies possible
- 2U single module tray for developers
- Disk storage: disk emulator + Network Attached Storage

• FPGA
• MGT link
• LVCMOS link

4/16/2008

CS252-s06, Lec 01-intro

30



Vision: Multiprocessing Watering Hole



Parallel file system Dataflow language/computer Data center in a box
Thread scheduling Security enhancements Internet in a box
Multiprocessor switch design Router design Compile to FPGA
Fault insertion to check dependability Parallel languages

- RAMP attracts many communities to shared artifact
 - ⇒ Cross-disciplinary interactions
 - ⇒ Accelerate innovation in multiprocessing
- RAMP as next Standard Research Platform?
(e.g., VAX/BSD Unix in 1980s, x86/Linux in 1990s)

4/16/2008

CS252-s06, Lec 01-intro

31



What Computer Architecture brings to Table

- Other fields often borrow ideas from architecture
- Quantitative Principles of Design
 1. Take Advantage of Parallelism
 2. Principle of Locality
 3. Focus on the Common Case
 4. Amdahl's Law
 5. The Processor Performance Equation
- Careful, quantitative comparisons
 - Define, quantify, and summarize relative performance
 - Define and quantify relative cost
 - Define and quantify dependability
 - Define and quantify power
- Culture of anticipating and exploiting advances in technology
- Culture of well-defined interfaces that are carefully implemented and thoroughly checked

4/16/2008

CS252-s06, Lec 01-intro

37

1) Taking Advantage of Parallelism

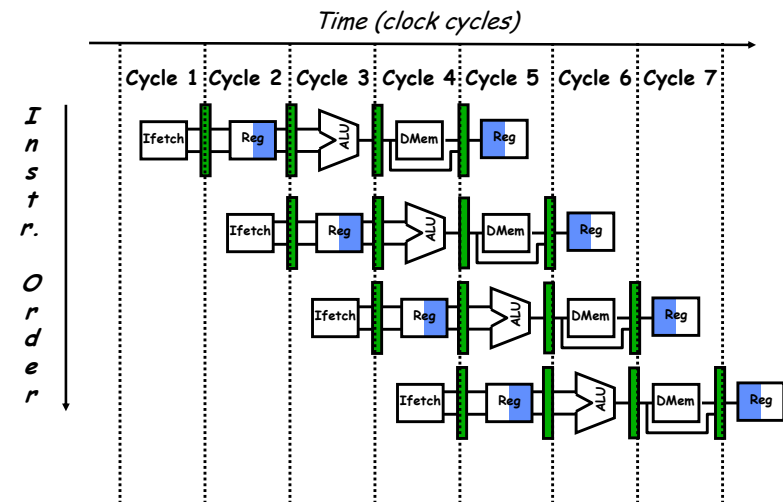
- Increasing throughput of server computer via multiple processors or multiple disks
- Detailed HW design
 - Carry lookahead adders uses parallelism to speed up computing sums from linear to logarithmic in number of bits per operand
 - Multiple memory banks searched in parallel in set-associative caches
- **Pipelining**: overlap instruction execution to reduce the total time to complete an instruction sequence.
 - Not every instruction depends on immediate predecessor \Rightarrow executing instructions completely/partially in parallel possible
 - Classic 5-stage pipeline:
 - 1) Instruction Fetch (Ifetch),
 - 2) Register Read (Reg),
 - 3) Execute (ALU),
 - 4) Data Memory Access (Dmem),
 - 5) Register Write (Reg)

4/16/2008

CS252-s06, Lec 01-intro

38

Pipelined Instruction Execution



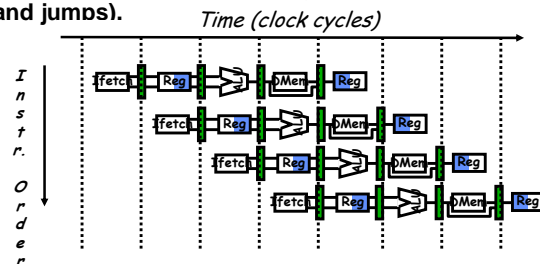
4/16/2008

CS252-s06, Lec 01-intro

39

Limits to pipelining

- **Hazards** prevent next instruction from executing during its designated clock cycle
 - **Structural hazards**: attempt to use the same hardware to do two different things at once
 - **Data hazards**: Instruction depends on result of prior instruction still in the pipeline
 - **Control hazards**: Caused by delay between the fetching of instructions and decisions about changes in control flow (branches and jumps).




4/16/2008

CS252-s06, Lec 01-intro

40

2) The Principle of Locality

- The Principle of Locality:
 - Program access a relatively small portion of the address space at any instant of time.
- Two Different Types of Locality:
 - **Temporal Locality** (Locality in Time): If an item is referenced, it will tend to be referenced again soon (e.g., loops, reuse)
 - **Spatial Locality** (Locality in Space): If an item is referenced, items whose addresses are close by tend to be referenced soon (e.g., straight-line code, array access)
- Last 30 years, HW relied on locality for memory perf.



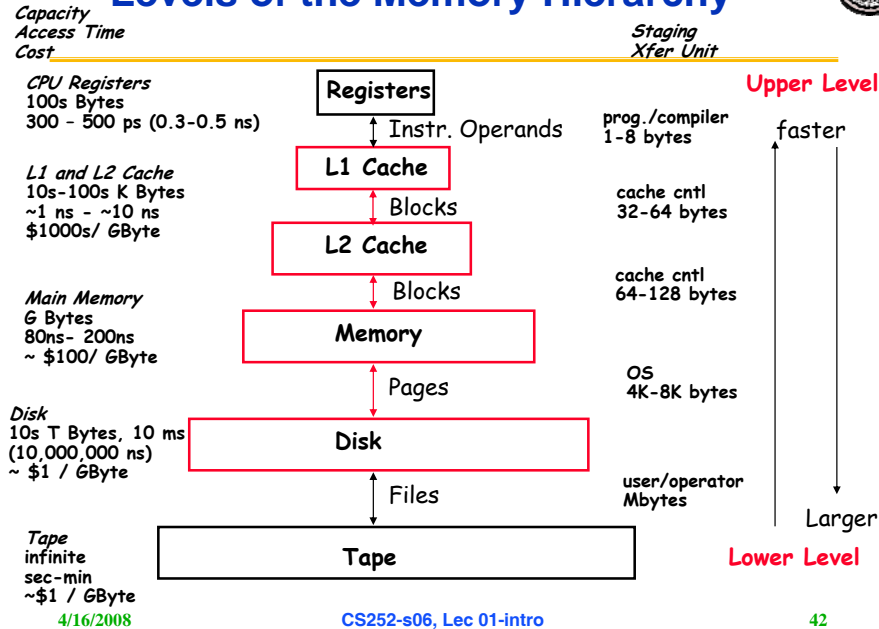
4/16/2008

CS252-s06, Lec 01-intro

41



Levels of the Memory Hierarchy



3) Focus on the Common Case

- Common sense guides computer design
 - Since its engineering, common sense is valuable
- In making a design trade-off, favor the frequent case over the infrequent case
 - E.g., Instruction fetch and decode unit used more frequently than multiplier, so optimize it 1st
 - E.g., If database server has 50 disks / processor, storage dependability dominates system dependability, so optimize it 1st
- Frequent case is often simpler and can be done faster than the infrequent case
 - E.g., overflow is rare when adding 2 numbers, so improve performance by optimizing more common case of no overflow
 - May slow down overflow, but overall performance improved by optimizing for the normal case
- What is frequent case and how much performance improved by making case faster => **Amdahl's Law**

4/16/2008

CS252-s06, Lec 01-intro

43




4) 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:

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



4/16/2008

CS252-s06, Lec 01-intro

44

Amdahl's Law example

- New CPU 10X faster
- I/O bound server, so 60% time waiting for I/O

$$\begin{aligned} \text{Speedup}_{\text{overall}} &= \frac{1}{(1 - \text{Fraction}_{\text{enhanced}}) + \frac{\text{Fraction}_{\text{enhanced}}}{\text{Speedup}_{\text{enhanced}}}} \\ &= \frac{1}{(1 - 0.4) + \frac{0.4}{10}} = \frac{1}{0.64} = 1.56 \end{aligned}$$

- Apparently, its human nature to be attracted by 10X faster, vs. keeping in perspective its just 1.6X faster

4/16/2008

CS252-s06, Lec 01-intro

45

5) Processor performance equation



CPI
 inst count Cycle time

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


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

4/16/2008

CS252-s06, Lec 01-intro

46

What's a Clock Cycle?



- Old days: 10 levels of gates
- Today: determined by numerous time-of-flight issues + gate delays
 - clock propagation, wire lengths, drivers

4/16/2008

CS252-s06, Lec 01-intro

47