

Lecture 1: CS/ECE 3810 Introduction

- Today's topics:
 - logistics
 - why computer organization is important
 - modern trends

Logistics

- See class web-page
<http://www.cs.utah.edu/classes/cs3810>
- Sign up for the cs3810 mailing list !!
- TAs: Devyani Ghosh, Josh McKee
Email: teach-cs3810
Office hours: TBA
- Textbook: Computer Organization – HW/SW Interface,
Patterson and Hennessy, 3rd edition

Course Organization

- 30% midterm, 40% final, 30% assignments
- ~10 assignments – you may skip one; assignments due at the start of class (either paper or electronic)
- Co-operation policy: you may discuss – you may not see someone else's written matter when writing your solution
- Print slides just before class

Why Computer Organization

- Yes, I know, required class...

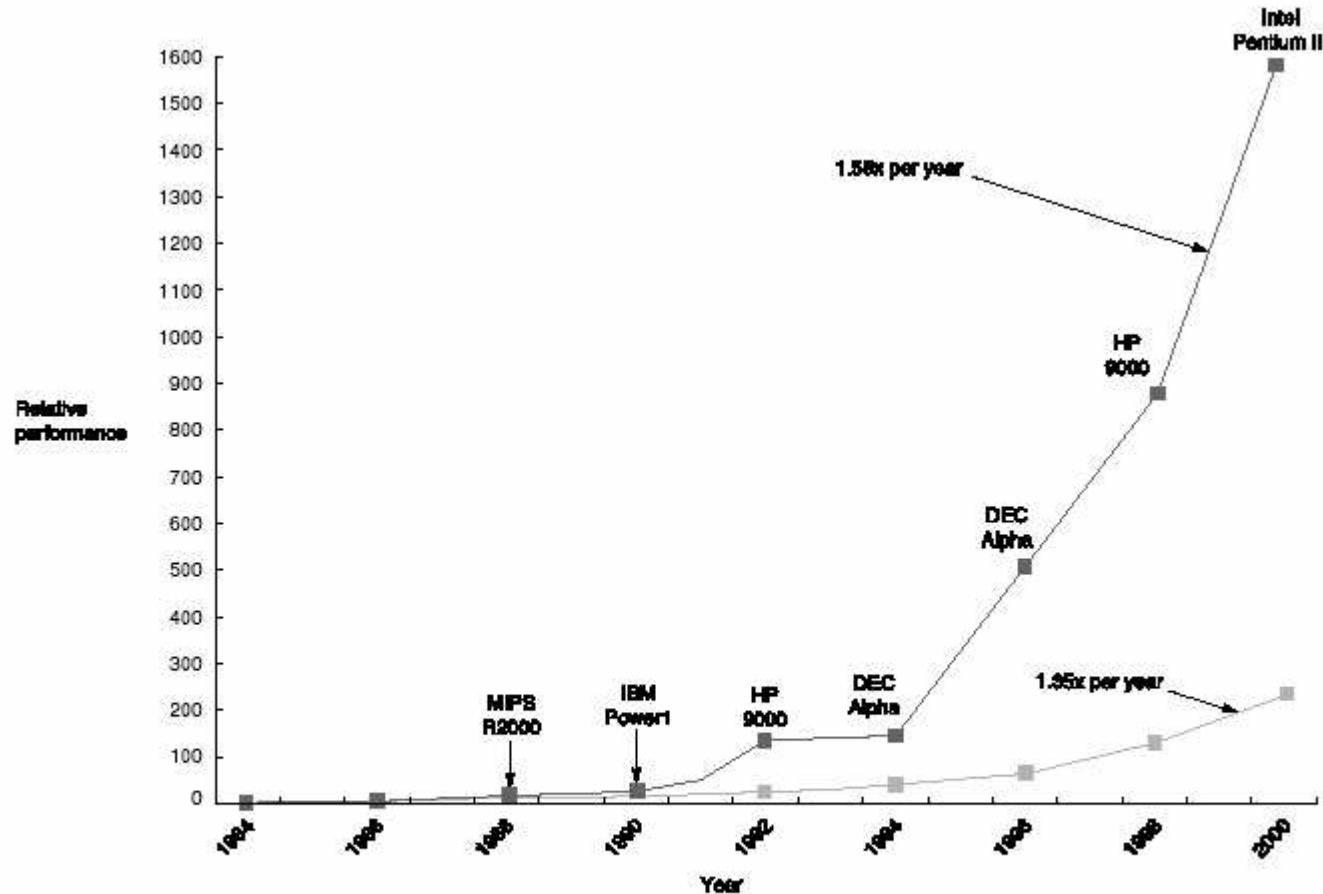
Why Computer Organization

- Embarrassing if you are a BS in CS/CE and can't make sense of the following terms: DRAM, pipelining, cache hierarchies, I/O, virtual memory
- Embarrassing if you are a BS in CS/CE and can't decide which processor to buy: 3 GHz P4 or 2.5 GHz Athlon (helps us reason about performance/power)
- Obvious first step for chip designers, compiler/OS writers
- Will knowledge of the hardware help me write better programs?

Must a Programmer Care About Hardware?

- Memory management: if we understand how/where data is placed, we can help ensure that relevant data is nearby
 - Thread management: if we understand how threads interact, we can write smarter multi-threaded programs
- Why do we care about multi-threaded programs?

Microprocessor Performance



50% improvement every year!!

What contributes to this improvement?

© 2003 Elsevier Science (USA). All rights reserved.

Modern Trends

- Historical contributions to performance:
 - Better processes (faster devices) ~20%
 - Better circuits/pipelines ~15%
 - Better organization/architecture ~15%

In the future, bullet-2 will help little and bullet-3 will not help much for a single core!

	Pentium	P-Pro	P-II	P-III	P-4	Itanium	Montecito
Year	1993	95	97	99	2000	2002	2005
Transistors	3.1M	5.5M	7.5M	9.5M	42M	300M	1720M
Clock Speed	60M	200M	300M	500M	1500M	800M	1800M

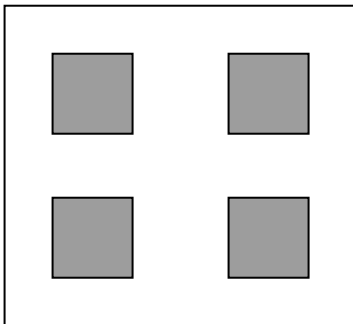
Moore's Law in action

At this point, adding transistors
to a core yields little benefit

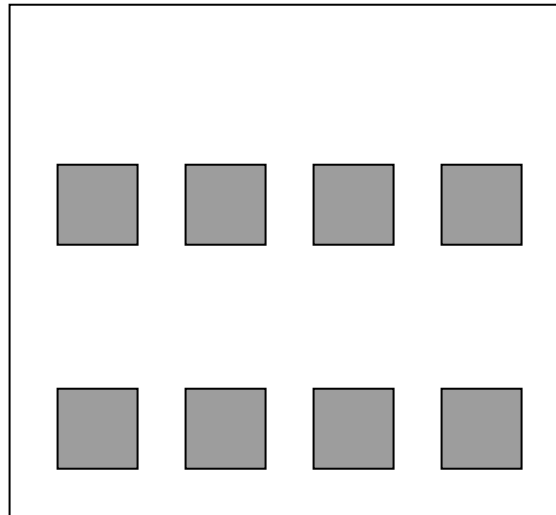
What Does This Mean to a Programmer?

- In the past, a new chip directly meant 50% higher performance for a program
- Today, one can expect only a 20% improvement, unless... the program can be broken up into multiple threads
- Expect `#threads` to emerge as a major metric for software quality

4-way multi-core



8-way multi-core

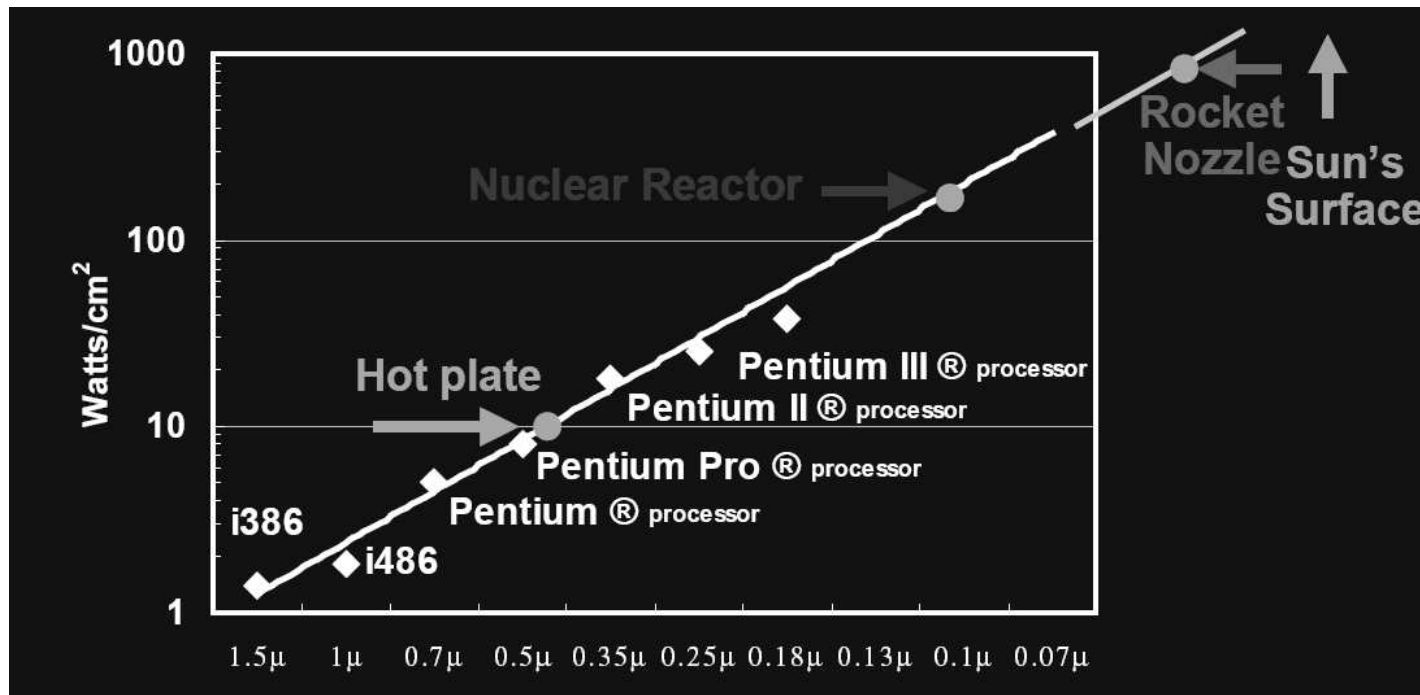


Challenges for the Hardware Designers

Challenges for the Hardware Designers

Major concerns:

- The performance problem (especially scientific workloads)
- The power dissipation problem (especially embedded processors)
- The temperature problem
- The reliability problem



The HW/SW Interface

Application software

Systems software
(OS, compiler)

Hardware

$a[i] = b[i] + c;$

↓ Compiler

```
lw    $15, 0($2)
add   $16, $15, $14
add   $17, $15, $13
lw    $18, 0($12)
lw    $19, 0($17)
add   $20, $18, $19
sw    $20, 0($16)
```

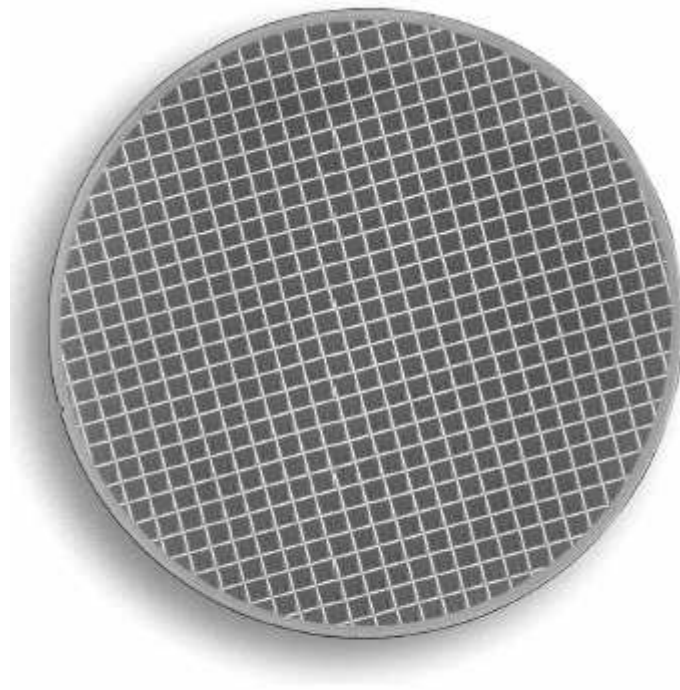
↓ Assembler

```
000000101100000
110100000100010
...
```

Computer Components

- Input/output devices
- Secondary storage: non-volatile, slower, cheaper
- Primary storage: volatile, faster, costlier
- CPU/processor

Wafers and Dies



© 2003 Elsevier Science (USA). All rights reserved.

Manufacturing Process

- Silicon wafers undergo many processing steps so that different parts of the wafer behave as insulators, conductors, and transistors (switches)
- Multiple metal layers on the silicon enable connections between transistors
- The wafer is chopped into many dies – the size of the die determines yield and cost

Processor Technology Trends

- Shrinking of transistor sizes: 250nm (1997) → 130nm (2002) → 70nm (2008) → 35nm (2014)
- Transistor density increases by 35% per year and die size increases by 10-20% per year... functionality improvements!
- Transistor speed improves linearly with size (complex equation involving voltages, resistances, capacitances)
- Wire delays do not scale down at the same rate as transistor delays

Memory and I/O Technology Trends

- DRAM density increases by 40-60% per year, latency has reduced by 33% in 10 years (the memory wall!), bandwidth improves twice as fast as latency decreases
- Disk density improves by 100% every year, latency improvement similar to DRAM
- Networks: primary focus on bandwidth; 10Mb → 100Mb in 10 years; 100Mb → 1Gb in 5 years

Power Consumption Trends

- Dyn power \propto activity x capacitance x voltage² x frequency
- Capacitance per transistor and voltage are decreasing, but number of transistors and frequency are increasing at a faster rate
- Leakage power is also rising and will soon match dynamic power
- Power consumption is already around 100W in some high-performance processors today

Next Class

- Topics: MIPS instruction set architecture (Chapter 2)
- Visit the class web-page
<http://www.cs.utah.edu/classes/cs3810>
- Sign up for the mailing list
- Pick up CADE Lab passwords

Title

- Bullet