COMPUTER ORGANIZATION AND DESIGN
The Hardware/Software Interface
5th Edition

Course Introduction

Dr. Curtis Nelson

Cptr380

Key Topics for This Quarter

- IC Design ↔ Computer Architecture
  - Moore's Law.
  - Power wall.
- High Level Language ↔ Assembly language
  - Instruction set design.
- Computer arithmetic.
- Pipelining.
- Branch prediction.
- Memory hierarchies.
- IO issues.

Today's Topics

- Why computer organization/architecture is important.
- Hardware ↔ Software relationship.
- Modern trends.

What Should You Get From This Class?

- A working vocabulary that includes
  - DRAM, pipelining, cache hierarchies, I/O busses, virtual memory, floating-point processing, etc.
- A knowledge base that helps you make intelligent choices i.e.
  - Deciding which processor to buy: 4 GHz Core I7 or 3.5 GHz Athlon (performance vs. power vs portability).
- A stepping stone toward the fields of IC design, compiler design, operating systems design.

Must a Programmer Care About Hardware?

- Must be knowledgeable about instruction set design.
- Must know how to make tradeoffs concerning speed, throughput, performance, and energy.
- Must know how to manage memory - if you understand how/where data is placed, you can help ensure that relevant data is nearby.
- Thread management - if you understand how threads are executed and interact with shared resources, you can write better multi-threaded programs.
  - Why should you care about multi-threaded programs?

Industry Example

- On a recent benchmark, a 200x speedup for matrix vector multiplication came as a result of:
  - Data level parallelism: 3.8x
  - Loop unrolling and out-of-order execution: 2.3x
  - Cache management: 2.5x
  - Thread level parallelism: 14x

Expected Student Proficiencies

- A high-level programming class or significant experience coding in C, C++, Python, Java, etc.
- Possess an overview of computer organization with some assembly level programming
  - High-level ⇒ Assembly ⇒ Machine Code
- Digital hardware
  - Combinational logic design (AND, OR, NAND, decoders, mux's, counters, shift registers, etc.)
  - Sequential design (latches, flip-flops, state machines, etc.)
- How many of you have programmed in VHDL or Verilog?

Microprocessor Performance Trends

- 50% improvement every year! What contributes to this improvement?

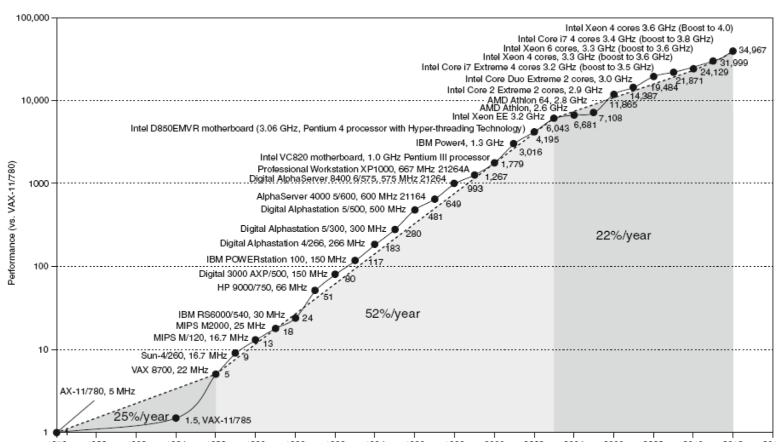

Transistor Size Over Time

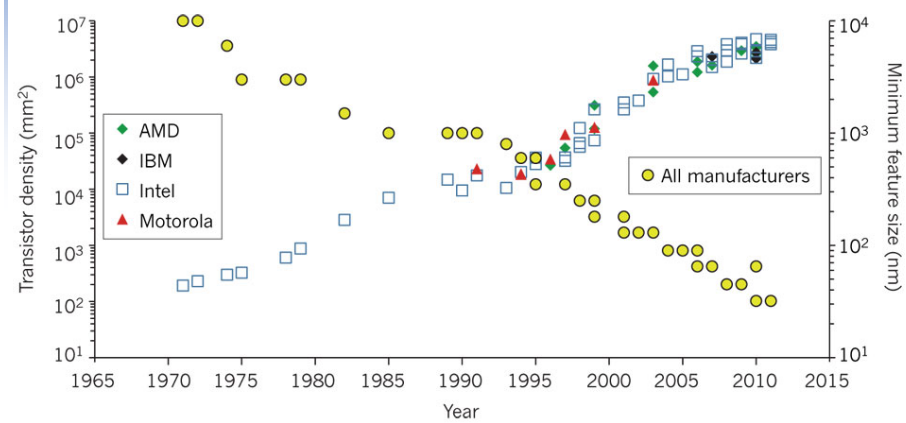

Important Trends

- We are running out of cost-effective ways to improve single thread performance.
- Power wall makes it harder to add complex features.
- Power wall makes it harder to increase frequency.

Important 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 #1 will eventually disappear.

	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

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

Challenges to Hardware Designers

- Find efficient ways to:
  - Boost single-thread performance.
  - Improve data sharing.
  - Manage the memory system.
  - Accelerate important parts of the computer.
  - Reduce system energy per instruction.

Levels of Program Code

- High-level language program (e.g. C++)

```
swap (int v[], int k)
(int temp;
        temp = v[k];
        v[k] = v[k+1];
        v[k+1] = temp;
)
```

- Assembly language program (for MIPS)

```
swap:  sll    $2, $5, 2
       add    $2, $4, $2
       lw     $15, 0($2)
       lw     $16, 4($2)
       sw     $16, 0($2)
       sw     $15, 4($2)
       jr     $31
```

- Machine (object, binary) code (for MIPS)

```
000000 00000 00101 0001000010000000
000000 00100 00010 0001000000100000
. . .
```

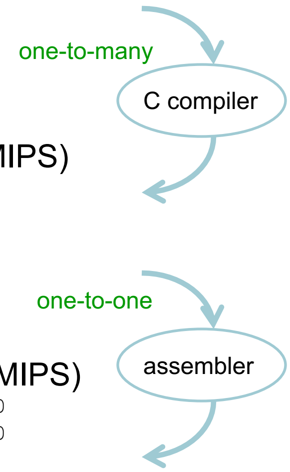

Intel Core i7 Wafer

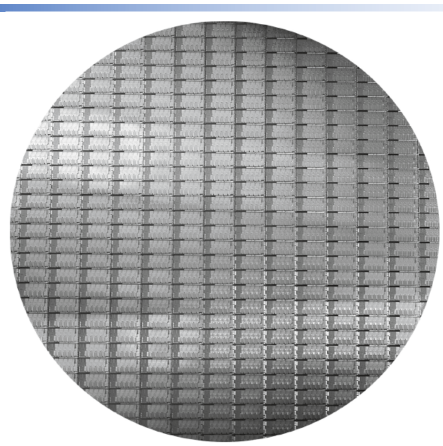

- 300mm wafer, 280 chips, 32nm technology.
- Each chip is 20.7 x 10.5 mm.

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) → 15nm (2018) → 9 nm (2020)
- Transistor density increases are slowing and commercially feasible die size increases by 10-20% per year… functionality improvements.
- Transistor speed improves linearly with size (complex equation involving voltages, resistances, capacitances).
- All shrinking (scaling) is not proportional, however.
  - Current densities rise.
  - Wires do not scale down the same way transistors do.

Memory and I/O Technology Trends

- DRAM density increases are slowing, access time has gone down by 33% in 10 years, bandwidth improves twice as fast as latency.
  - **Bandwidth** refers to the amount of information that something, like a connection to the internet, can handle in a given amount of time.
- 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

- Dynamic power = activity x capacitance x voltage^2 x frequency.
- Voltage and frequency are somewhat constant now, while capacitance per transistor is decreasing and number of transistors (activity) is increasing.
- Leakage power is also rising (function of the number of transistors and voltage).

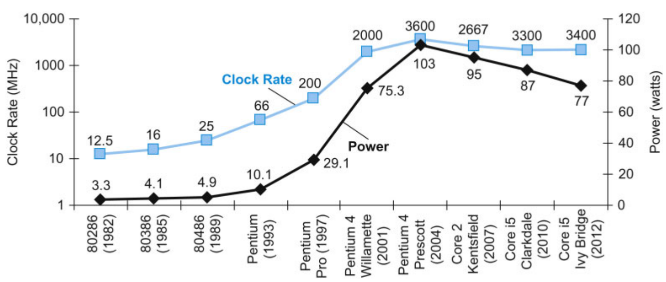

Marriage of Hardware and Software

- Overall computer performance is related to:
  - Architecture – better overall organization
    - What are some architectural improvements over the last 20 years?
  - Software – better languages, compilers, algorithms, etc.
    - What are some software improvements over the last 20 years?
  - Hardware – better processes, manufacturing, designs, etc.
    - What are some hardware improvements over the last 20 years?

Next Class

- Computer performance (Chapter 1)
- Visit the class web-page
  http://gab.wallawalla.edu/~curt.nelson/cptr380/index 2020.html
- Check out HW#1