Lecture 1: CS/ECE 3810 Introduction

- Today's topics:
  - Why computer organization is important
  - Logistics
  - Modern trends

HIT RECORD

Why Computer Organization

Image credits: uber, extremetech, anandtech

Why Computer Organization

2018

Image credits: gizmodo

Why Computer Organization

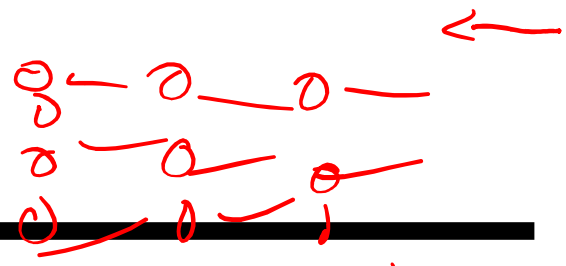

- Efficiency is important! You learn about big-O notation in Algorithms, you learn about the constants in this class
- The AI revolution owes a lot to fast hardware; a complete CS/CE education includes a deep understanding of hardware

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: 4.4 GHz Intel Core i9 or
  4.7 GHz AMD Ryzen 9 (reason about performance/power)

- Obvious first step for chip designers, compiler/OS writers

- Will knowledge of the hardware help you write better and more secure programs?

Must a Programmer Care About Hardware?

- Must know how to reason about program performance and energy and security
- 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?

Example

200x speedup for matrix vector multiplication

- Data level parallelism: 3.8x
- Loop unrolling and out-of-order execution: 2.3x
- Cache blocking: 2.5x
- Thread level parallelism: 14x

Further, can use accelerators to get an additional 100x.

Key Topics

- Moore's Law, power wall
- Use of abstractions
- Assembly language
- Computer arithmetic
- Pipelining
- Using predictions
- Memory hierarchies
- Accelerators
- Reliability and Security

Logistics

- See class web-page for syllabus/resources
  https://www.cs.utah.edu/~rajeev/cs3810
- TAs and office hours: on the class webpage
- Most communication on Canvas; email me directly to set up meetings, or meet me in office hours right before class
- Textbook: Computer Organization – HW/SW Interface, Patterson and Hennessy, 5th or 6th edition

Course Organization

- 30% midterm, 40% final, 30% assignments
- ~10 assignments – you may skip two; automatic 1.5 day extension until Wed/Fri late night; upload on Canvas
- Co-operation policy: you may discuss – you may not see someone else's written matter when writing your solution
- Exams are open-notes (1 page)
- Print slides just before class
- Screencast YouTube videos

HW1, posted Thu 1/11
due Thu 1/18
by Fri 1/19 11:59pm

Grading Policy

- Grading by rank (32% receive A or A-, 40% receive B+/B/B-)
- About 13% receive D+ or lower, scores below 70
- No tolerance for cheating (see class webpage)
- Rank in exams matters more than rank in homeworks

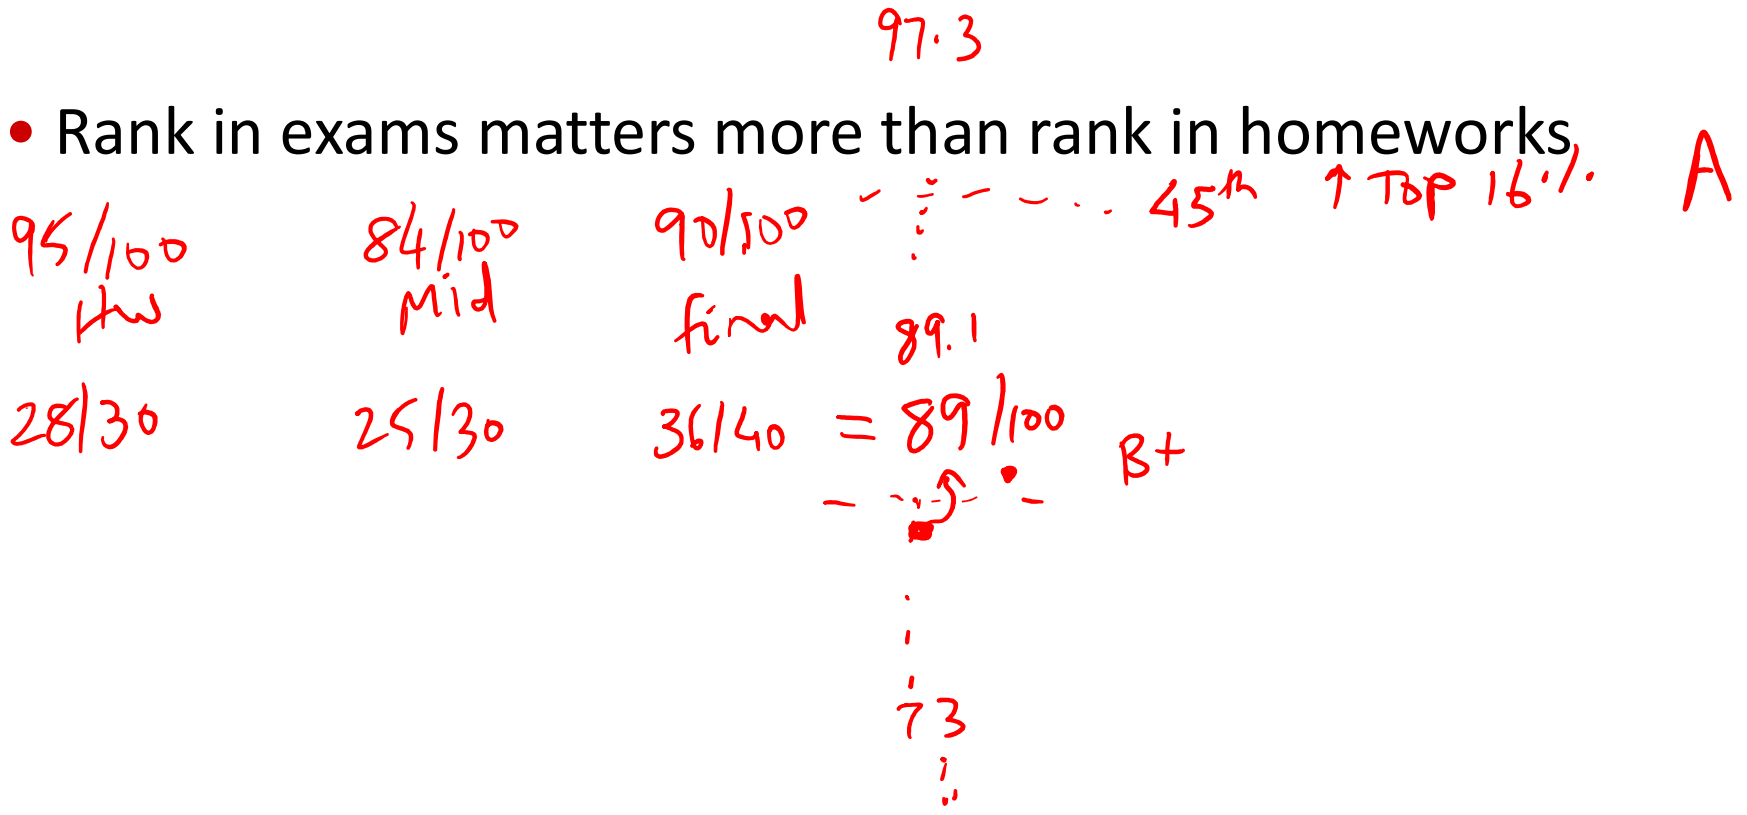

Microprocessor Performance

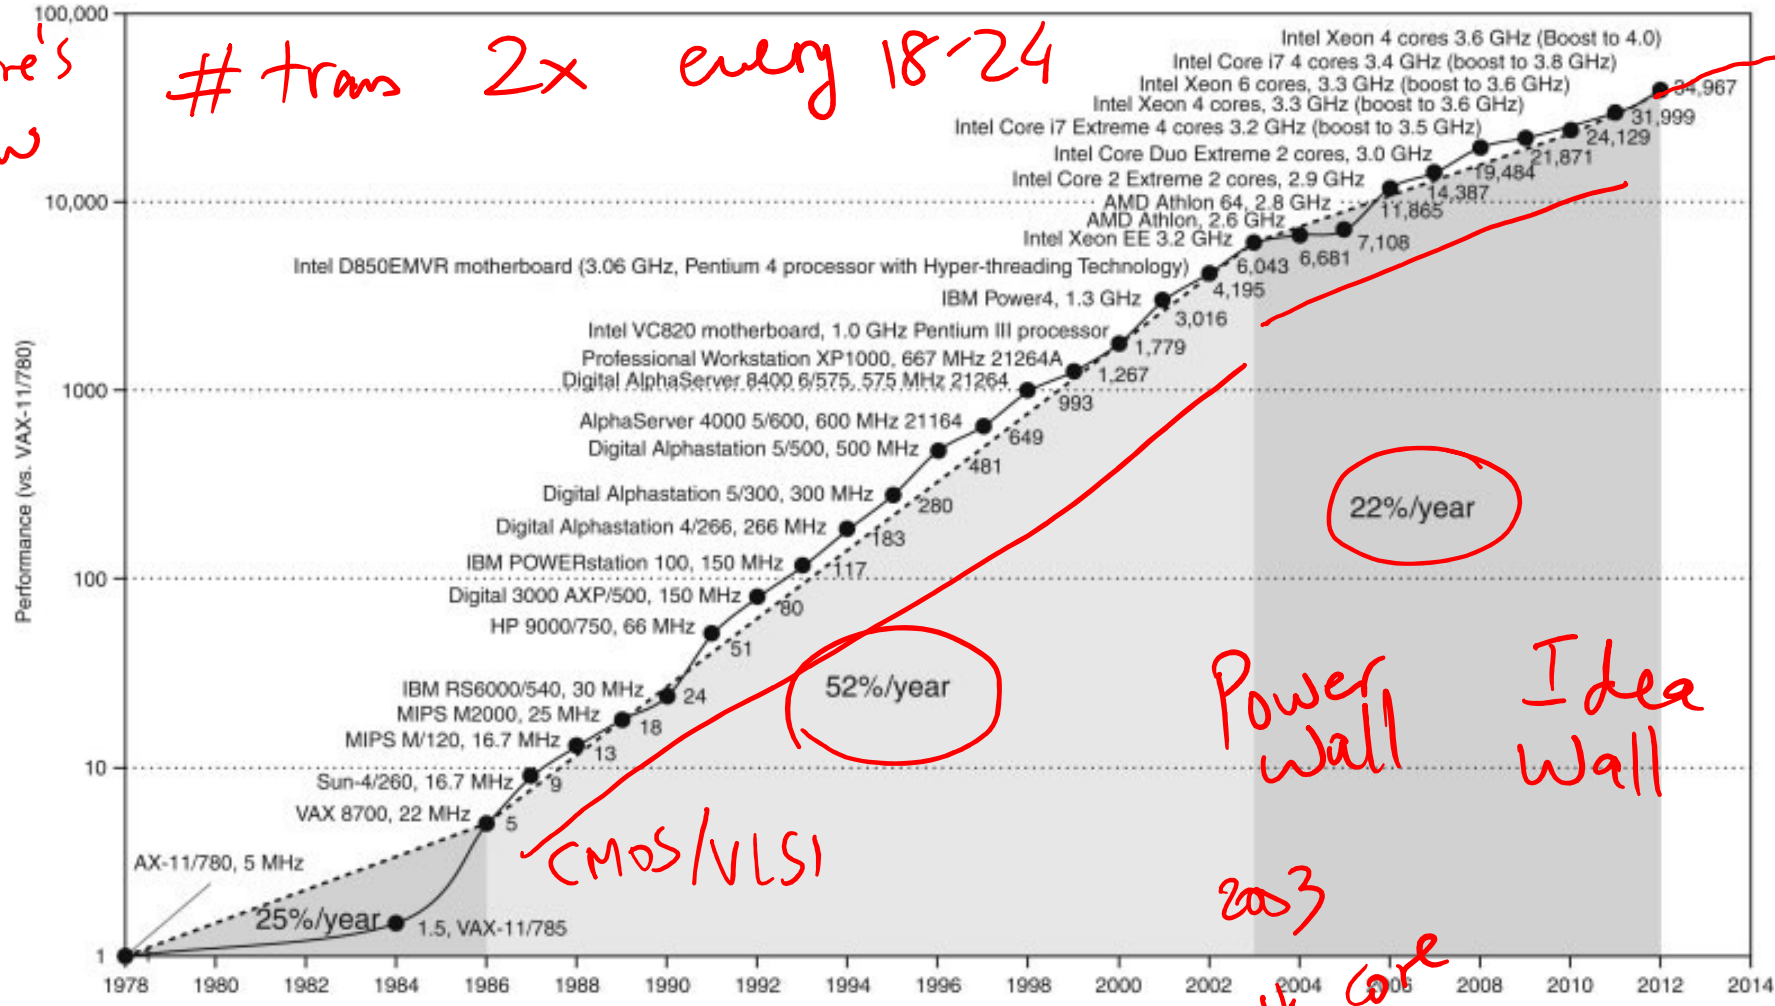

Source: H&P Textbook

50% improvement every year!!

What contributes to this improvement?

Why the lower improvement?

Microprocessor Performance

clock

3 GHz Freq

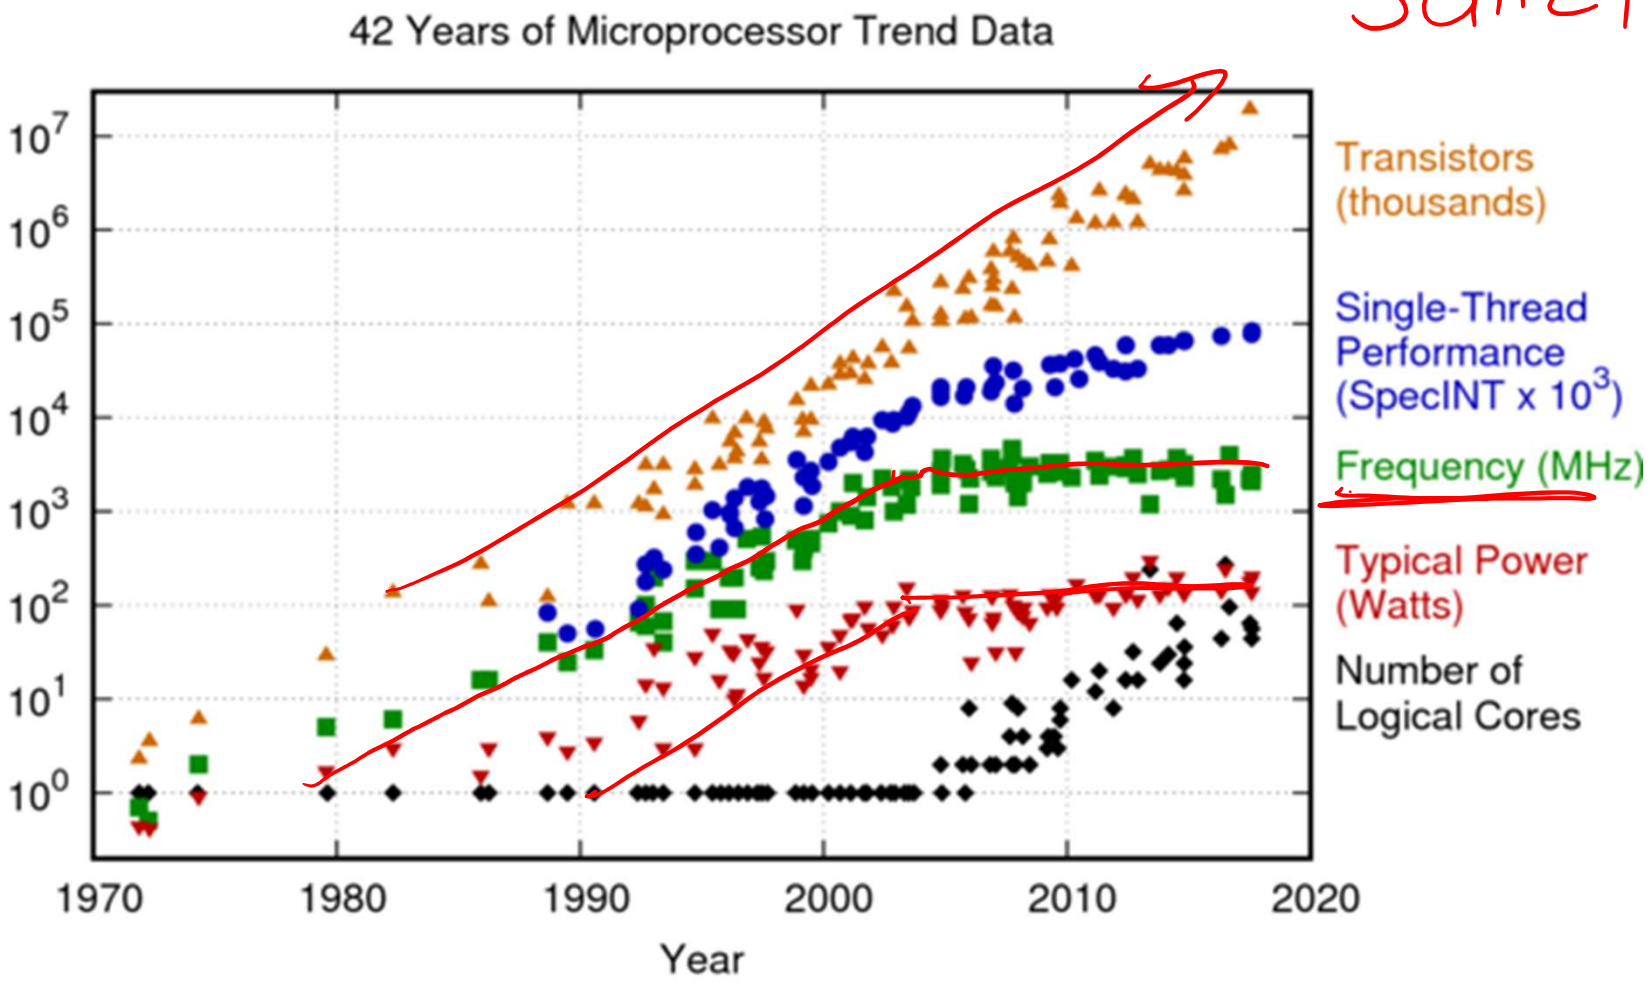

Original data up to the year 2010 collected and plotted by M. Horowitz, F. Labonte, O. Shacham, K. Olukotun, L. Hammond, and C. Batten
New plot and data collected for 2010-2017 by K. Rupp

Source: karlrupp.net

Power Consumption Trends

- Dyn power α 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 #trans and voltage)

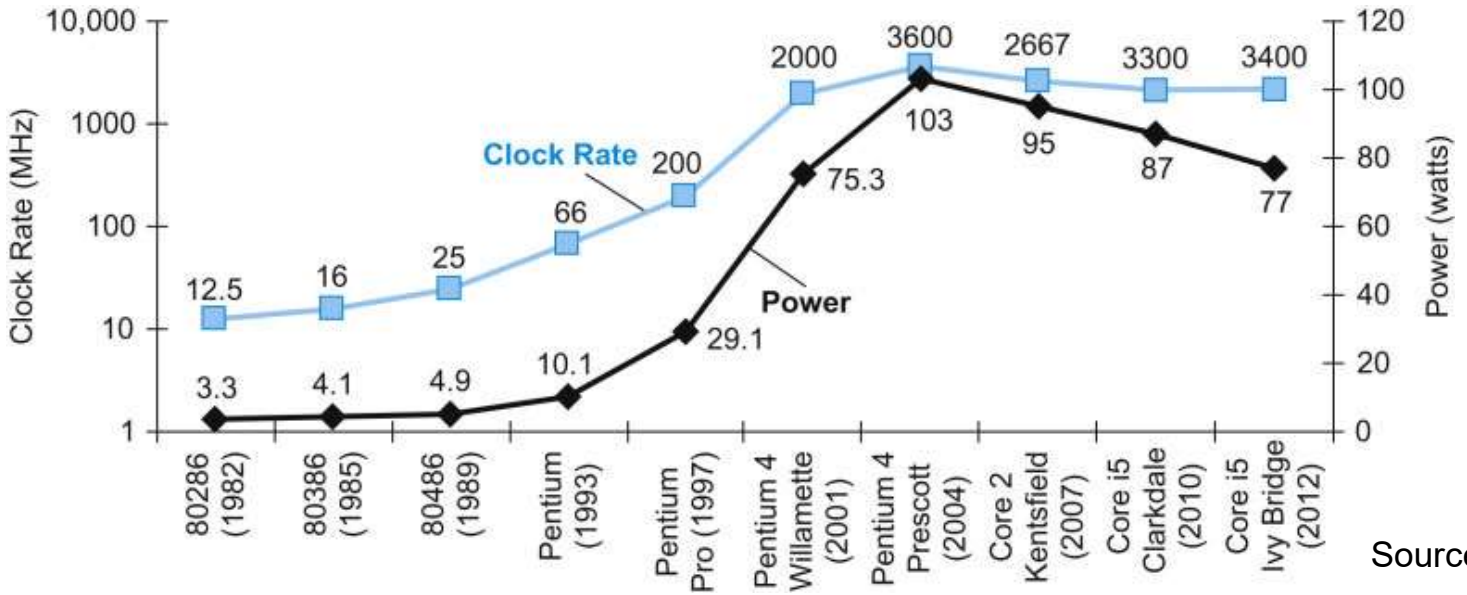

Source: H&P Textbook

Important Trends

- Running out of ideas to improve single thread performance
- Power wall makes it harder to add complex features
- Power wall makes it harder to increase frequency
- Technology scaling likely to end soon
- Additional performance provided by: more cores, occasional spikes in frequency, accelerators

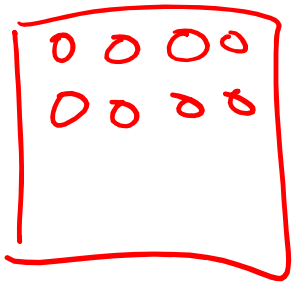

Turbo boost

dim silicon
all 8 active
2GHz

or dark silicon
4 out of 8 off
3GHz

Summary

- Three roadblocks: power, ideas, technology scaling
- Fixed power budget because of cooling constraints; implies that frequency can't be increased; discourages complex ideas
- End of voltage (Dennard) scaling in early 2010s; the end of Moore's Law also imminent
- Has led to dark silicon and dim silicon (occasional turbo)

Next Class

- Topics: Trends, Performance, MIPS instruction set architecture (Chapter 2)

- Visit the class web-page
  https://www.cs.utah.edu/~rajeev/cs3810