18-742:
Computer Architecture & Systems

Introduction

Prof. Phillip Gibbons

Spring 2025, Lecture 1

Waitlist Status

- **As of Jan 12, 2025 at 8 pm:**
  36 registered, 3 on waitlist
  - Depending on drops, may be room for a few more students
  - Please meet me after class

- **Admittance priority:**
  - ECE PhD, CSD PhD, other SCS PhD
  - ECE Masters, ECE Undergrads
  - CS Masters, CS Undergrads
  - other Masters, other Undergrads

- **Priority among Masters students (and Undergrads) based on relevant courses taken (e.g., 213/513/613, 18-447, 15-418/618) and grades obtained**

Today's Topics

- **Course Overview**
  - No slides, just a walk through of the key points on the course webpages

- **Discussion of a Few Classic Papers**

“Cramming More Components onto Integrated Circuits”

Gordon Moore 1965

- Founder & CEO of Intel
- National Academy of Engineering
- Presidential Medal of Freedom
- National Inventors Hall of Fame
- AAAS Fellow

Wisdom from 1965

“Integrated circuits will lead to such wonders as home computers, automatic controls for automobiles, and personal portable communications equipment.”

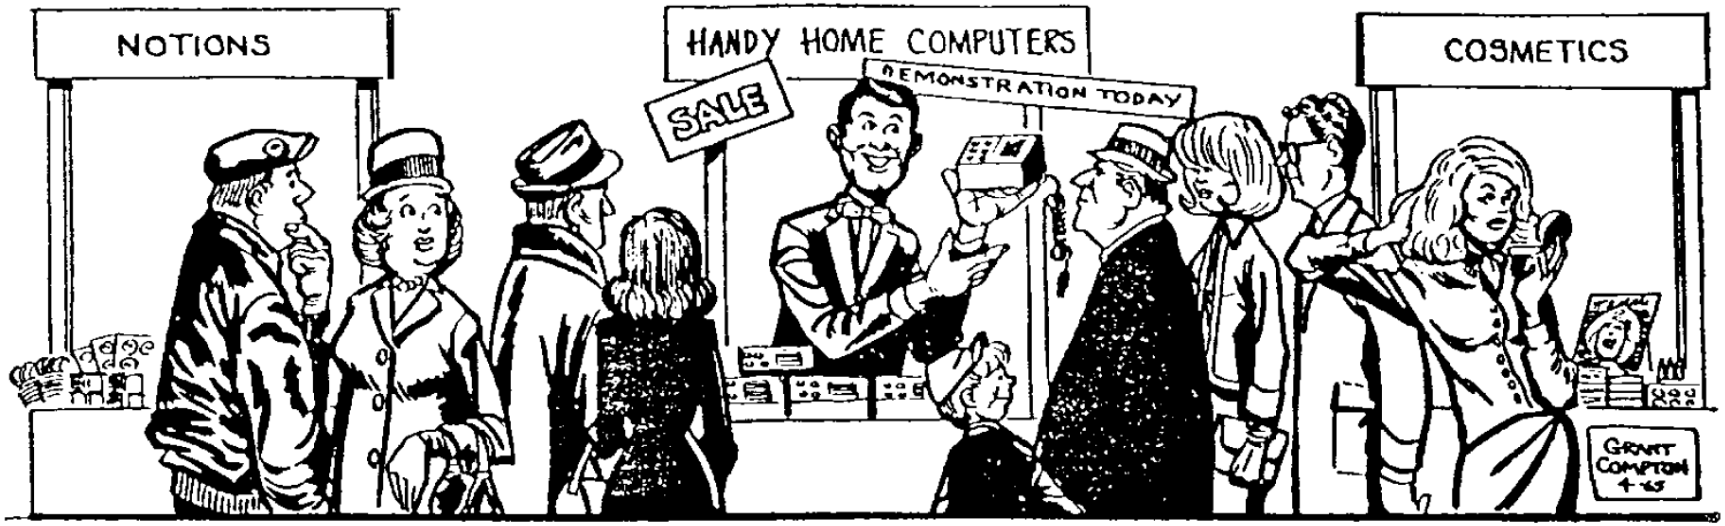

“Silicon is likely to remain the basic material.”

Moore's Law

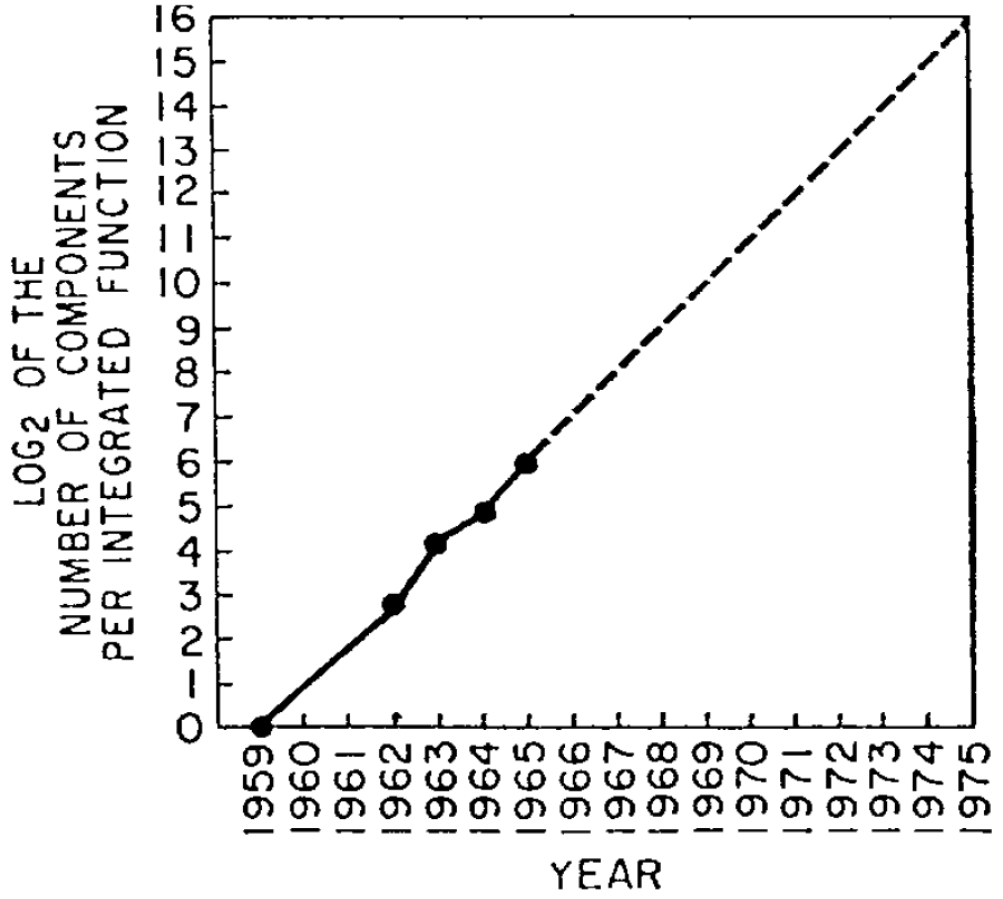

Number of components [transistors] per integrated function [integrated circuit] will double every year (for at least ten years).

Moore revised in 1975 to doubling every two years.

Economic Argument Based on Yield

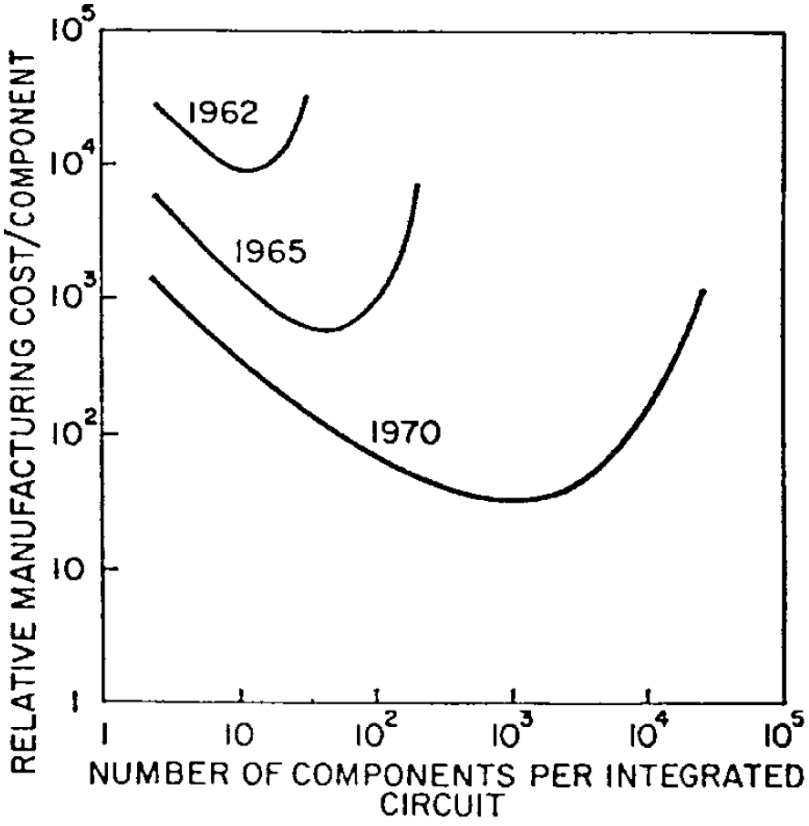

"No fundamental obstacle to achieving yields of 100%. Only the engineering effort is needed."

Moore's Law in Action

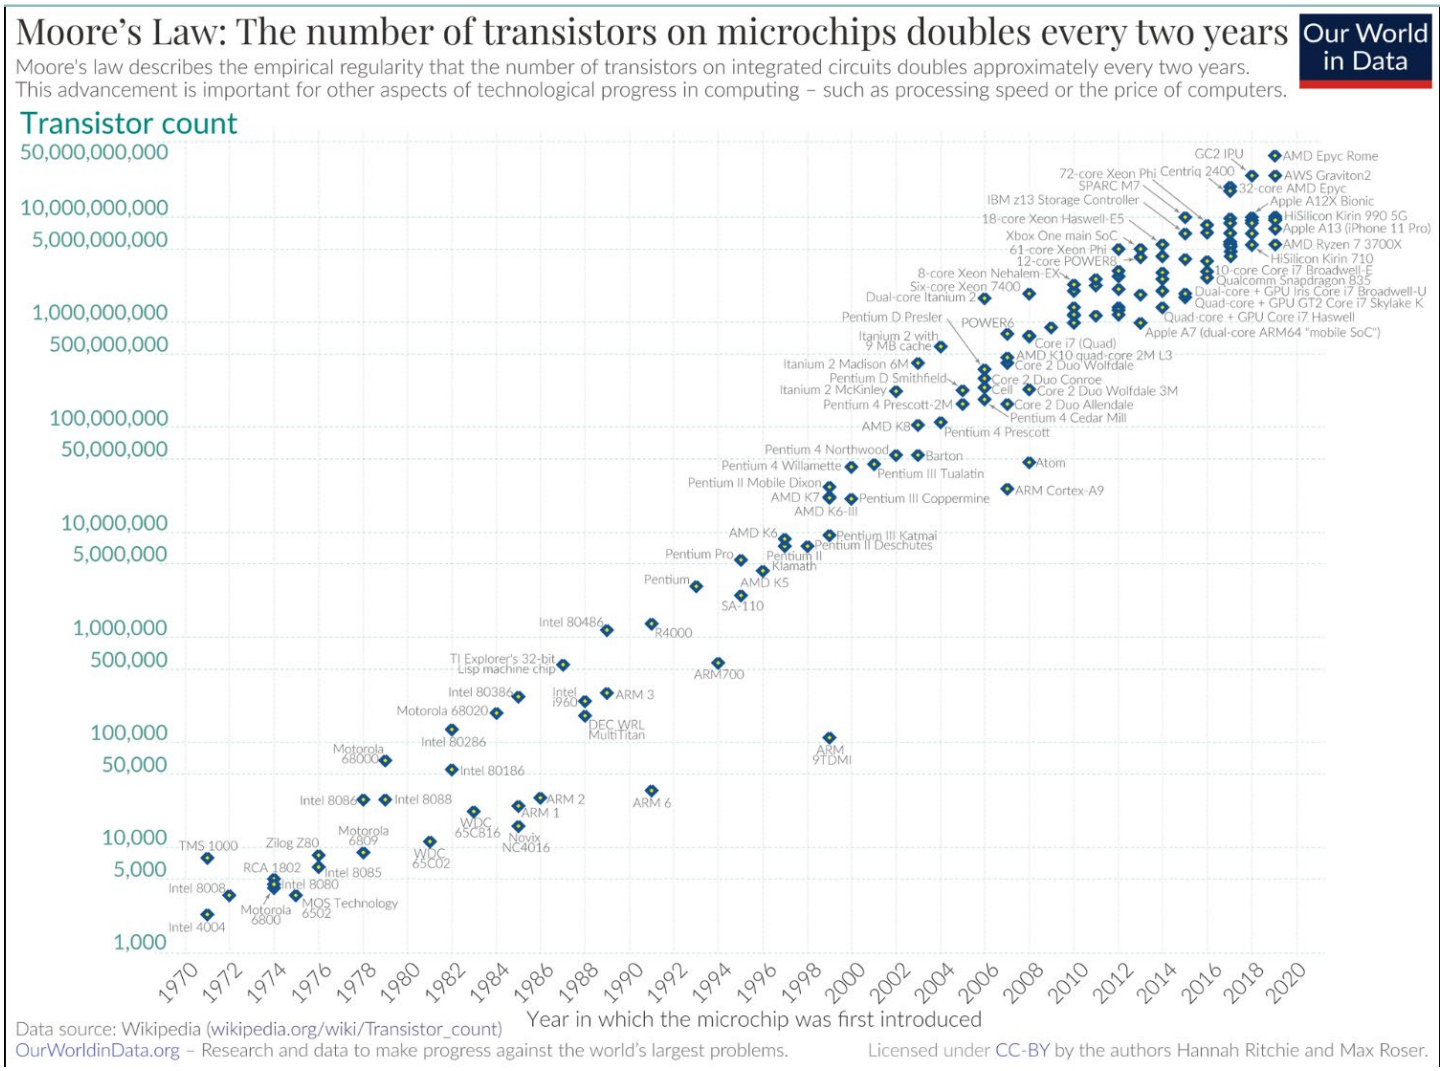

"Rather than becoming something that chronicled the progress of the industry, it became something that drove it."

“Validity of the Single Processor Approach to Achieving Large Scale Computing Capabilities”

Gene Amdahl 1967

- IBM Mainframe pioneer
- Founded Amdahl Corporation
- National Academy of Engineering
- Amdahl’s Law

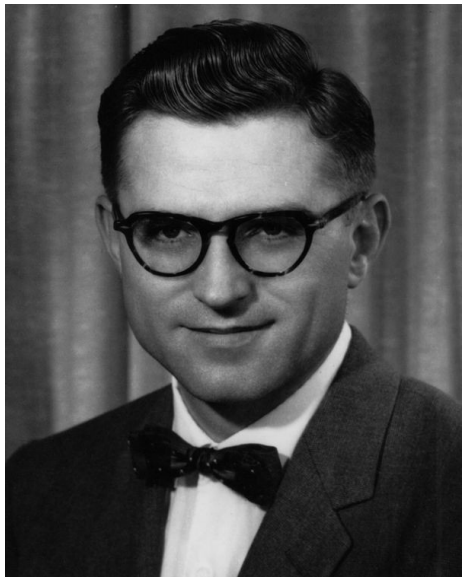

1922-2015

Amdahl's Law (Travel Analogy)

		Speed-Up
• Flying jet non-stop from PIT -> LHR:	7.5 Hours	1
• Or, old fashioned SST way:		
– Fly jet from PIT -> JFK: 1.5 Hours		
– Fly SST from JFK -> LHR: 3.5 Hours	**5 Hours**	**1.5x**
• Or, Using FTL travel:		
– Fly jet from PIT -> JFK: 1.5 Hours		
– Fly FTL from JFK -> LHR: .01 Hours	**1.51 Hours**	**~5x**

- **Best possible speed-up is 5X, even with FTL because have to get to New York.**

*** FTL = Faster than light**

Amdahl's Law

- **Overall problem**
  - T Total sequential time required
  - f Fraction of total that can be sped up ($0 \leq f \leq 1$)
  - S Speedup factor

- **Resulting Performance**
  - $T_S = f\,(T/S) + (1\text{-}f)T$
    - Portion which can be sped up runs S times faster
    - Portion which cannot be sped up stays the same
  - Minimal possible time given maximum possible speedup
    - When $S = \infty$, $T_\infty = (1\text{-}f)T$
  - Speedup = T / T_S

$$\text{Speedup}_{\text{enhanced}}(f, S) = \frac{1}{(1-f) + \frac{f}{S}}$$

Amdahl's Law (Travel Analogy)

		Speed-Up
• **Flying jet non-stop from PIT -> LHR:**	**7.5 Hours**	**1**
• **Using FTL travel:**		
– Fly jet from PIT -> JFK: 1.5 Hours		
– Fly FTL from JFK -> LHR: .01 Hours	**1.51 Hours**	**~5x**

- **Best possible speed-up is 5X, even with FTL because have to get to New York.**

$$\text{Speedup}_{\text{enhanced}}(f, S) = \frac{1}{(1-f) + \frac{f}{S}}$$

T=7.5, f=(7.5-1.5)/7.5=0.8

S= $\infty \Rightarrow$ T_∞ = 1/(1-f) = 1/(1-0.8) = 5

Amdahl's Law

$$\text{Speedup}_{\text{enhanced}}(f, S) = \frac{1}{(1-f) + \frac{f}{S}}$$

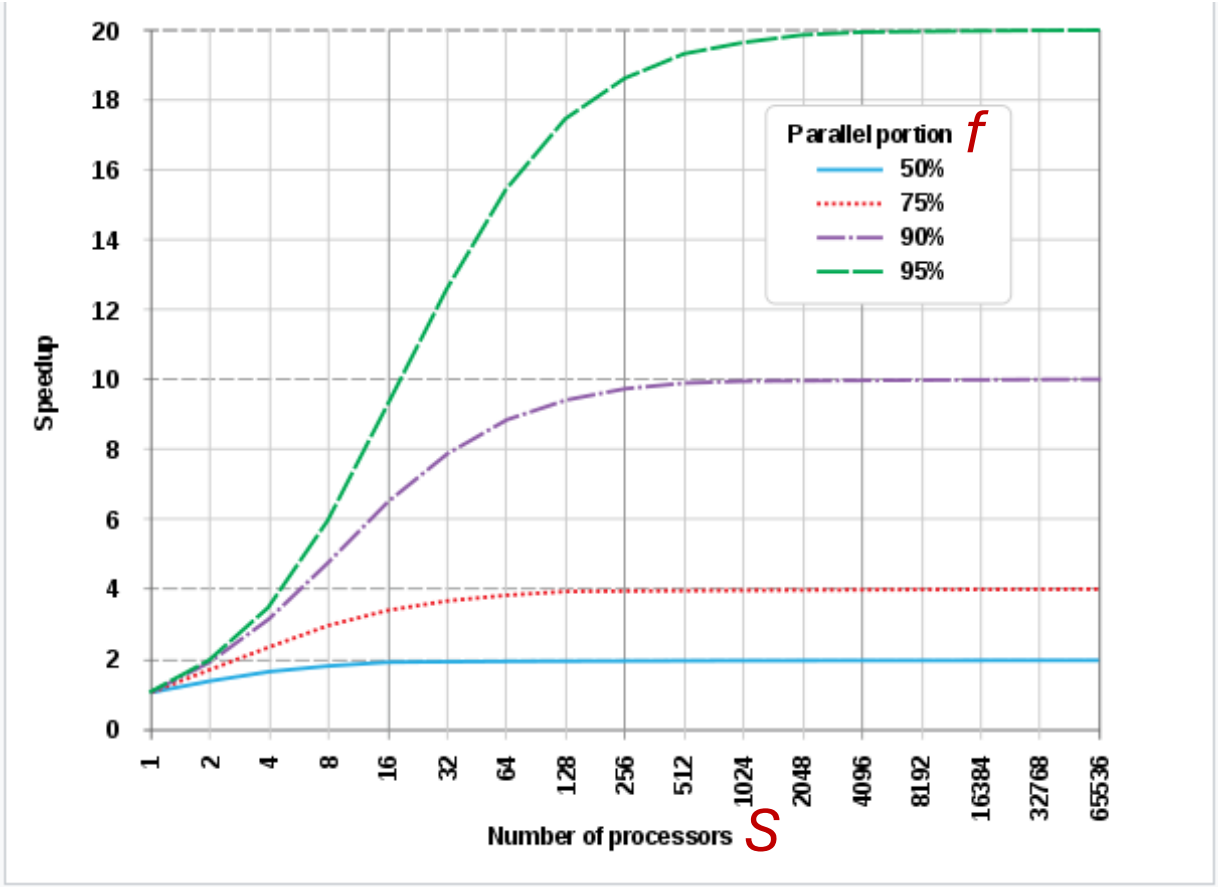

Lessons: Parallelize the common case (where f is large).
Also improve the serial case.

Weak Scaling: Grow problem size in proportion to increase in cores.

You and Your Research

Richard Hamming 1986

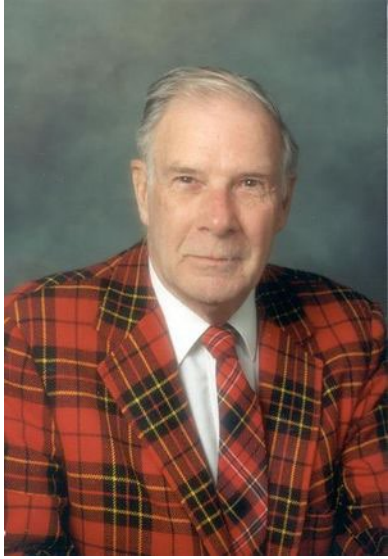

1915-1998

- Hamming distance
- Hamming codes (first error correcting codes)
- Turing Award winner 1968
- “The purpose of computing is insight not numbers”

Q: Why do so few scientists make significant contributions and so many are forgotten in the long run?

How to be a Great Scientist

- **"Luck favors the prepared mind" – Pasteur**
- **Key Characteristic: Courage**
  - When an opportunity opens up, get after it and pursue it
- **Most great scientists have tremendous drive**
  - must be intelligently applied
- **Knowledge and productivity are like compound interest**
- **What are the important problems in your field?**
  - and must have plan of attack

How to be a Great Scientist

- **Know yourself, your strengths and weaknesses**
- **Work with the door open**
  - Get interrupted, but also exposed to what's important
- **Never again solve an isolated problem except as characteristic of a class**
- **Should get into a new field every 7 years**
- **Need to sell your work, via good writing, formal talks, and informal talks**

To Read for Wednesday

"Dark Silicon and the End of Multicore Scaling"

Hadi Esmaeilzadeh, Emily Blem, Renee St. Amant, Karthikeyan Sankaralingam, Doug Berger 2011

Optional Further Reading:

"Scaling the Bandwidth Wall: Challenges in and Avenues for CMP Scaling"

Brian Rogers, Anil Krishna, Gordon Bell, Ken Vu, Xiaowei Jiang, Yan Solihin 2009

"Clearing the Clouds: A Study of Emerging Scale-out Workloads on Modern Hardware"

Michael Ferdman et al. 2012