

CS 755 – System and Network Architectures and Implementation

Module 0 - Introduction

Martin Karsten

mkarsten@uwaterloo.ca

Distributed System

- Leslie Lamport:

“A distributed system is one in which the failure of a computer you didn't even know existed can render your own computer unusable.”

Distributed System - Description

- A collection of computing components
 - with non-negligible communication latency
 - possibly non-deterministic
 - with potential partial failures
 - communication
 - components
- ... that appears as coherent system
 - *distribution transparency*

Scaling Dimensions

- time:
cycle, memory, disk, network, human, daily,...
- space:
core, chip, board, chassis, building, campus,...
- administrative:
uniform vs. cooperation vs. competition

Tentative Schedule

- today: organization, overview, background
- 5 lectures covering networking basics
- 5 classes with lectures/paper discussion
- final exam

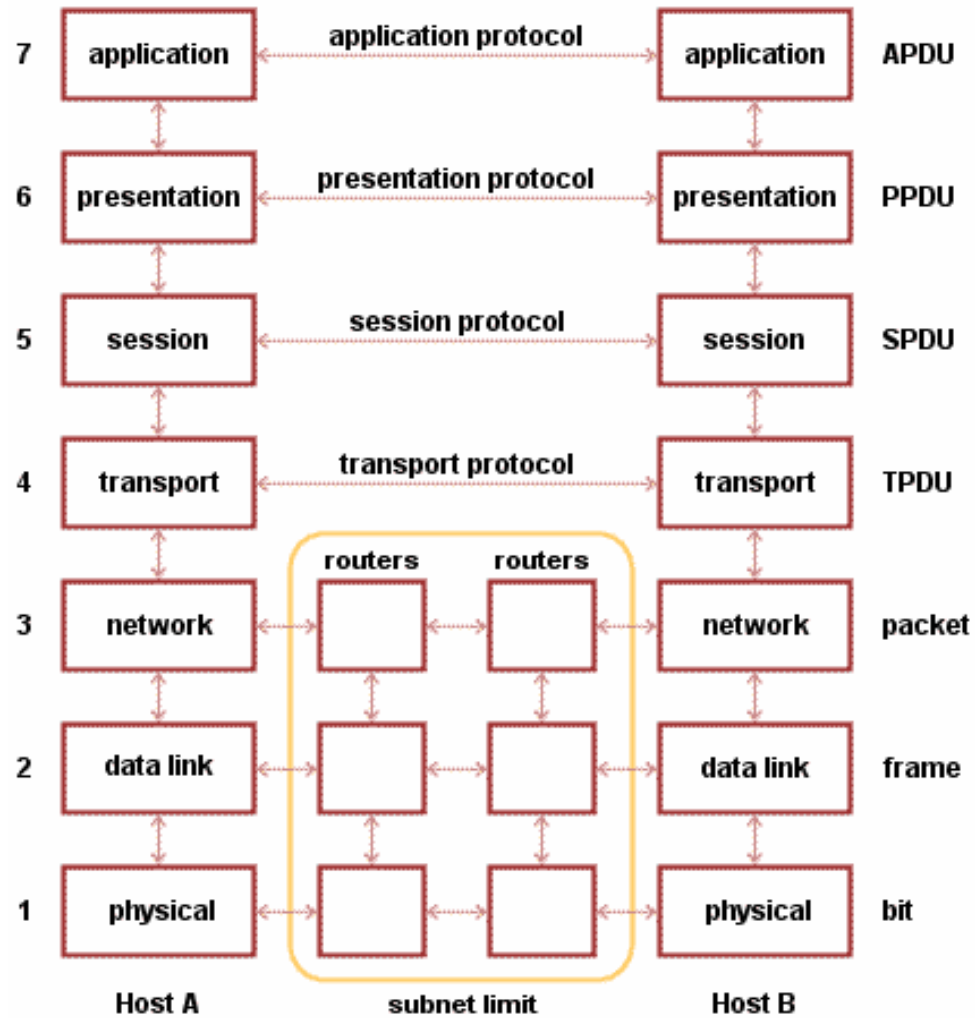
Basics

- Sep 17: Channels
 - communication between neighbours
- Sep 24: Network
 - communication in a network (graph)
- Oct 1: Transport
 - end-to-end challenges

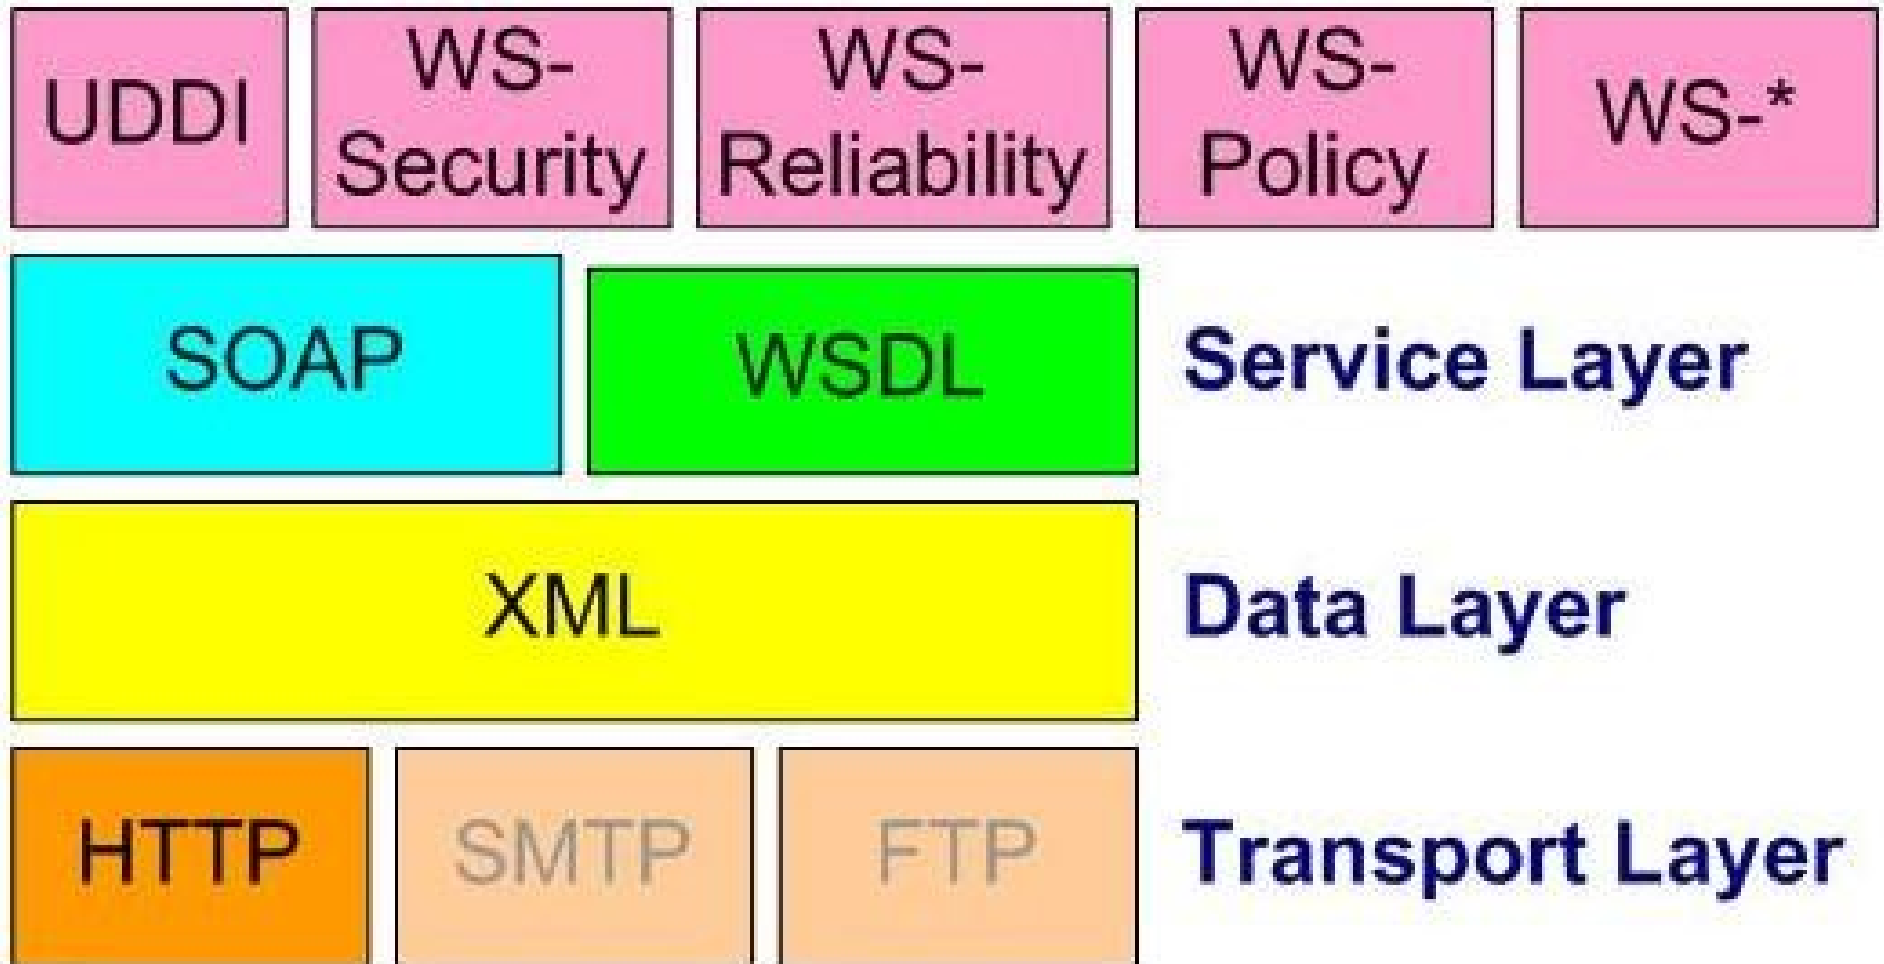
Basics

- Oct 8: Remote Services
 - service beyond basic communication
- Oct 15: Naming and Mobility
 - labelling and finding communication entities

Network Model



Web Services Model



Advanced

- Oct 22: Ordering and Consistency
 - difficult: timing in distributed system
- Oct 29: Fault Tolerance
 - even worse: components might fail

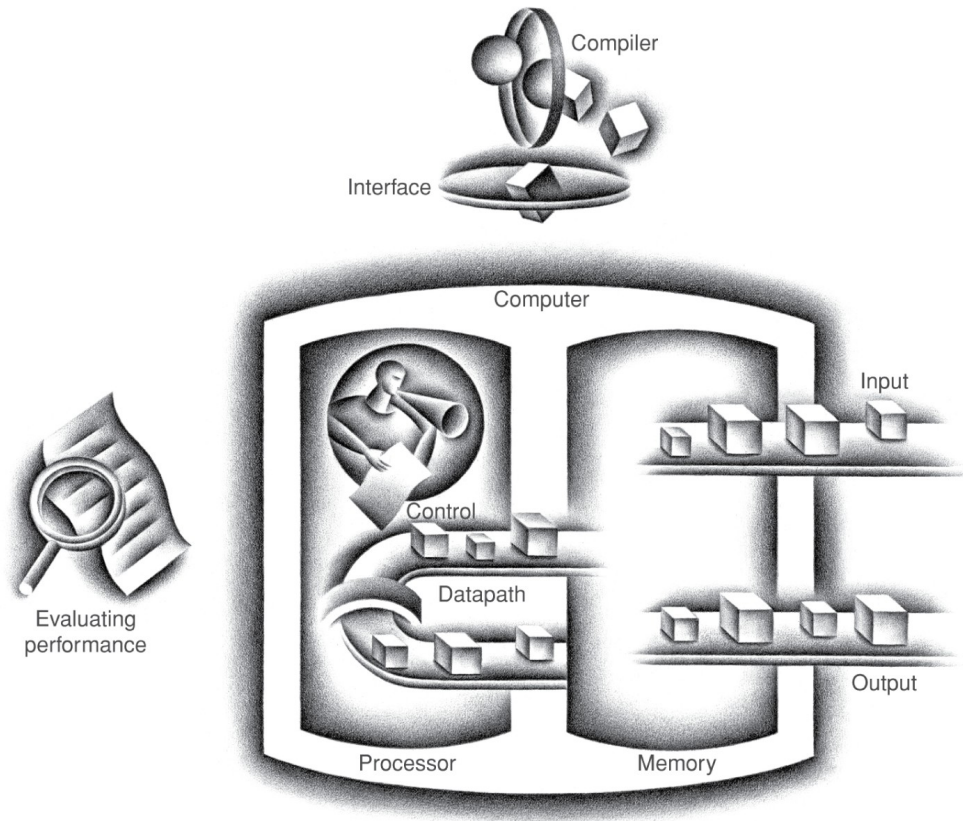
Advanced

- Nov 5: Storage & Replication
 - essential distributed service
- Nov 12: Cloud Services
 - novel paradigm?
- Nov 19: Big Data
 - important application area

Important Dates

- before Oct 1:
 - send me priority list of papers for presentation (student must be registered in Quest)
 - FCFS selection
- Oct 17, 5pm: 1st critical review due
- Nov 21, 5pm: 2nd critical review due
- Nov 26: last class & final exam

Model of a Computer

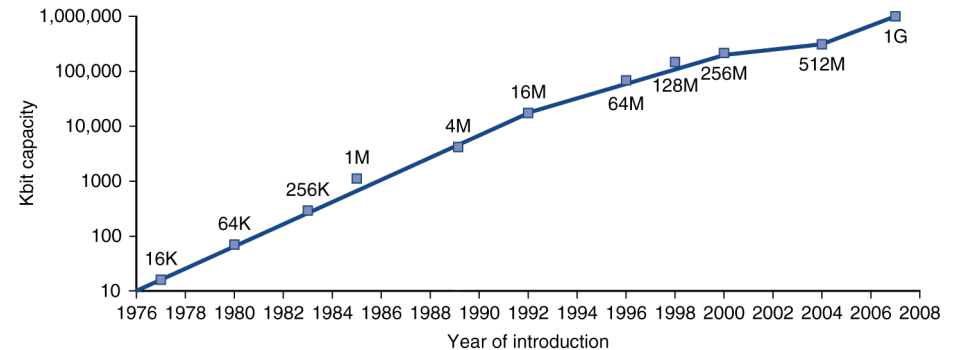


- von Neumann model
- CPU
 - control & data path
- I/O
 - user, storage, network
- memory

program & data
stored in memory

Technology Trends

- electronics technology continues to evolve
 - increased capacity and performance
 - reduced cost



DRAM capacity

Year	Technology	Relative performance/cost
1951	Vacuum tube	1
1965	Transistor	35
1975	Integrated Circuit	900
1995	Very large scale IC (VLSI)	2,400,000
2005	Ultra large scale IC	6,200,000,000

Performance: Latency vs. Throughput

- Tim Horton's
 - time to coffee vs. customers/hour
 - low latency => high throughput
 - but not vice versa
 - faster coffee makers vs. more (and more space)
- latency (response time)
 - completion time of specific task
- throughput
 - total work done over time period

Performance

- reduce latency?
 - faster processor
 - better algorithm (software)
 - more processors (needs parallelization)
 - generally increases throughput
- increase throughput?
 - more processors
 - rearrange system components (scheduling):
often increases latency

Efficiency Matters

- network-centric computing, Internet
 - large data centers
- hardware cheap, but
 - power consumption → heat
 - heat → cooling → more power consumption
 - money and environment costs
- often:
software performance (throughput) ~ efficiency

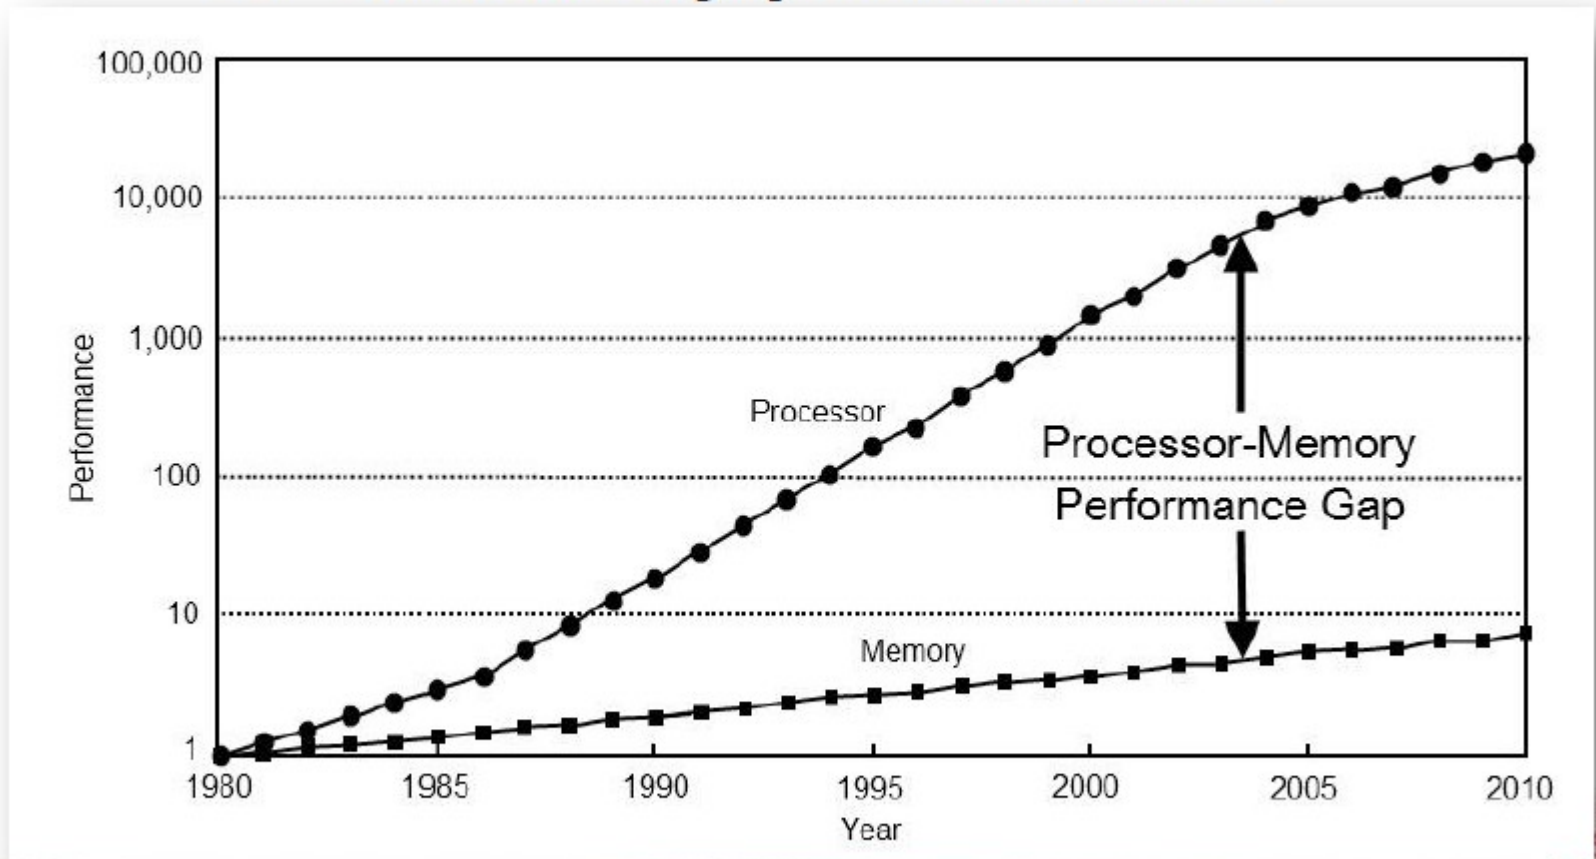
Moore's Law

- transistor density doubles every two years
 - every year 1959-1975
- in the past
 - transistor density translated into processing power
 - almost double speed every 2 years...
 - reduce latency, increase throughput
- recently: memory wall
- more recently: power wall

Memory Wall

Developm

CPU/Memory performance

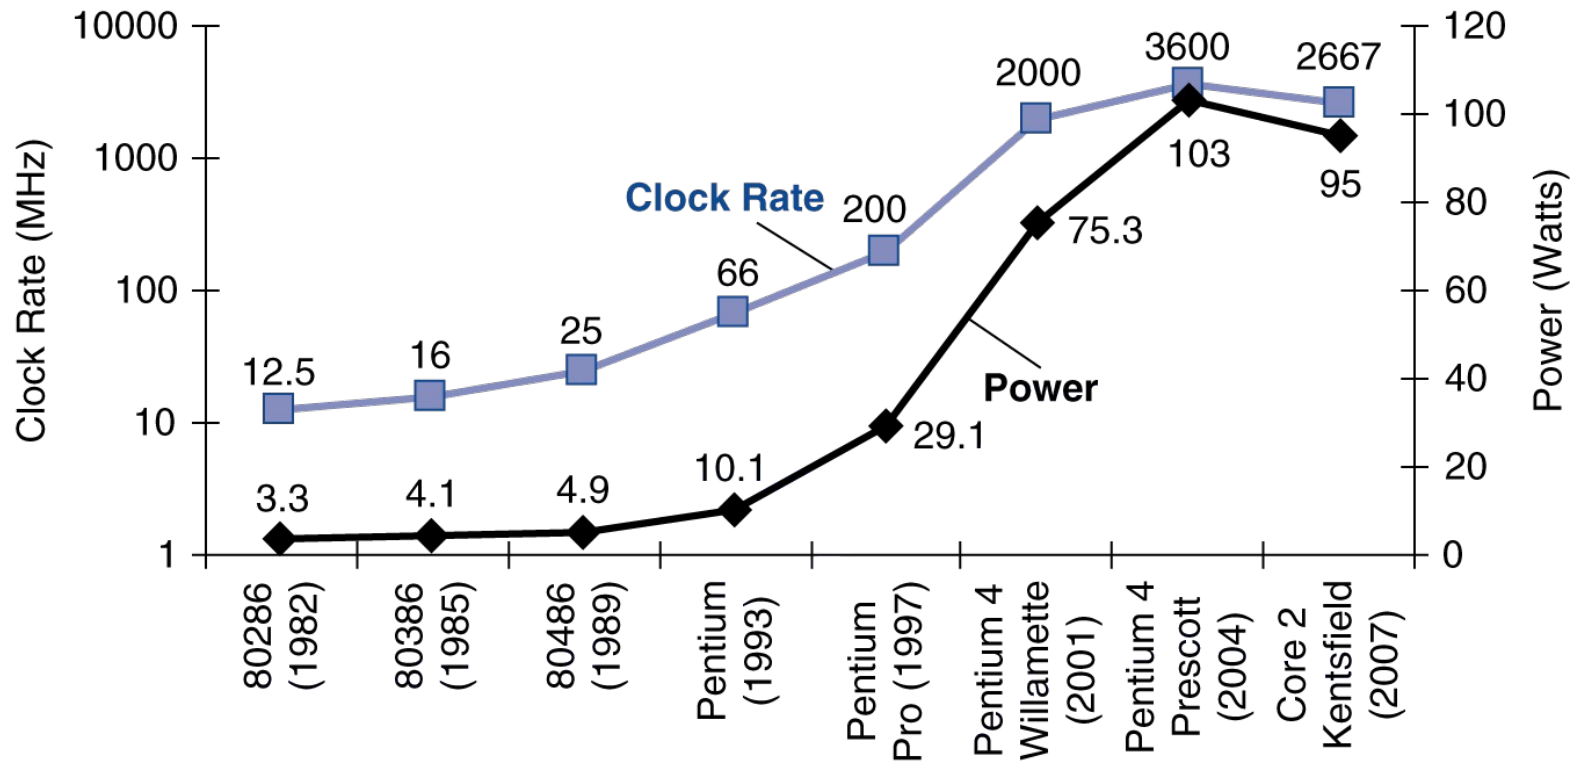


3

Slide 17

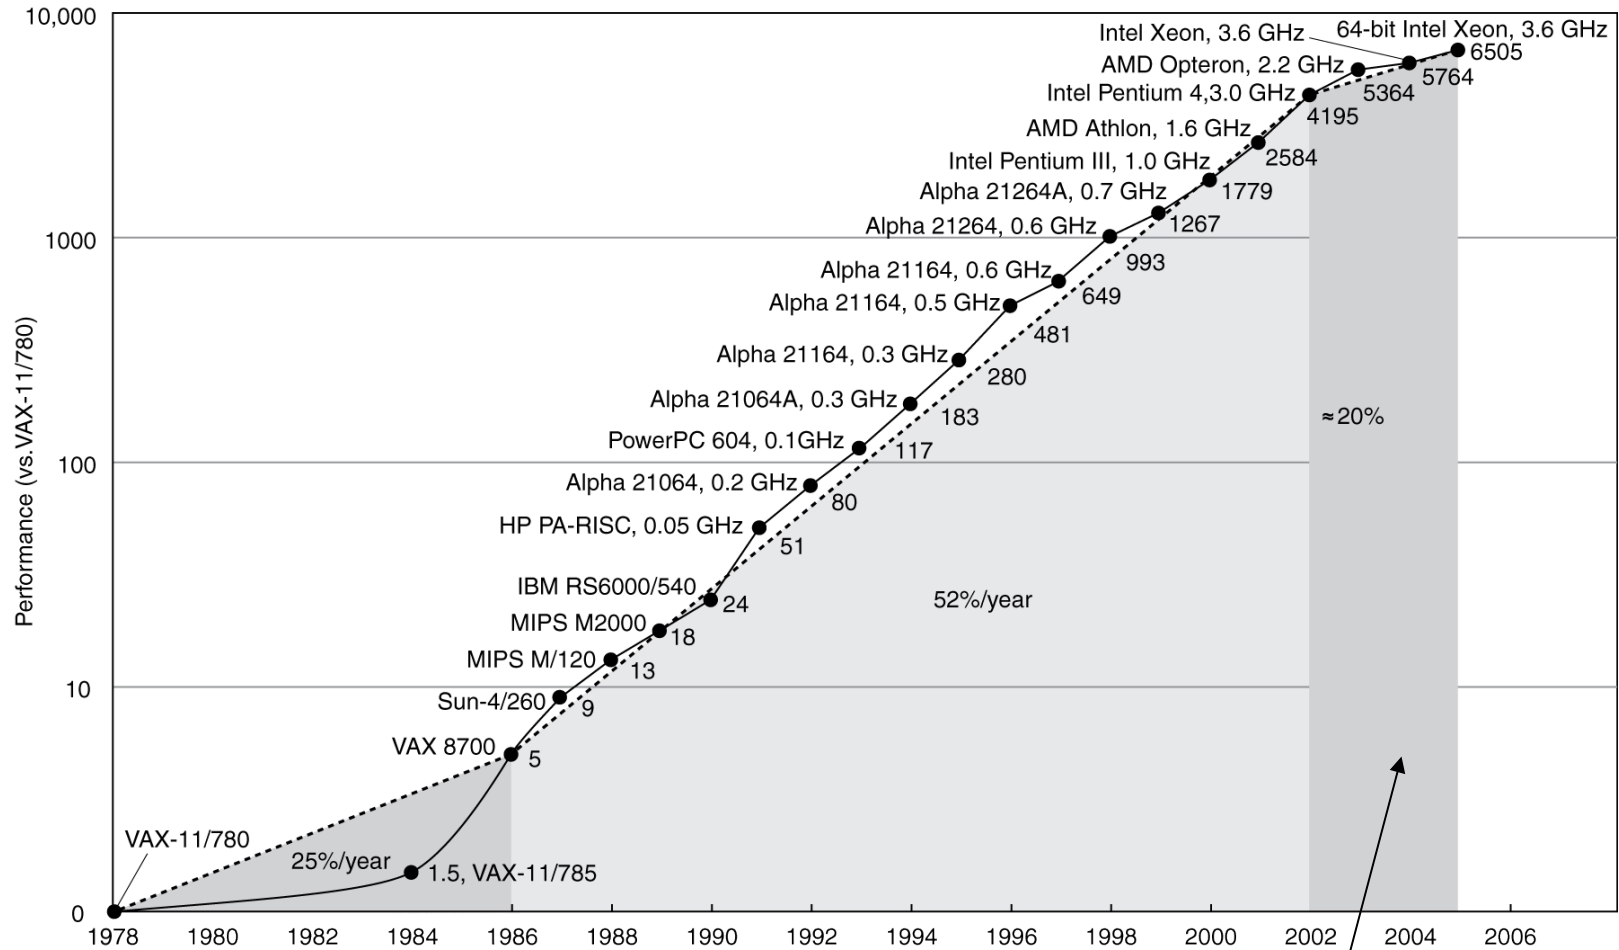
Computer architecture: a quantitative approach
By John L. Hennessy, David A. Patterson, Andrea C. Arpaci-Dusseau

Power Wall



- power = capacitive load x voltage² x frequency
 - cannot reduce voltage further (path length)
 - cannot remove more heat

Uniprocessor Performance



Constrained by power, instruction-level parallelism, memory latency

Multiprocessors

- multicore microprocessors
 - more than one processor per chip
- requires explicitly parallel programming
 - compare with instruction level parallelism (hidden)
- hard to do
 - programming for performance
 - load balancing
 - optimizing communication and synchronization

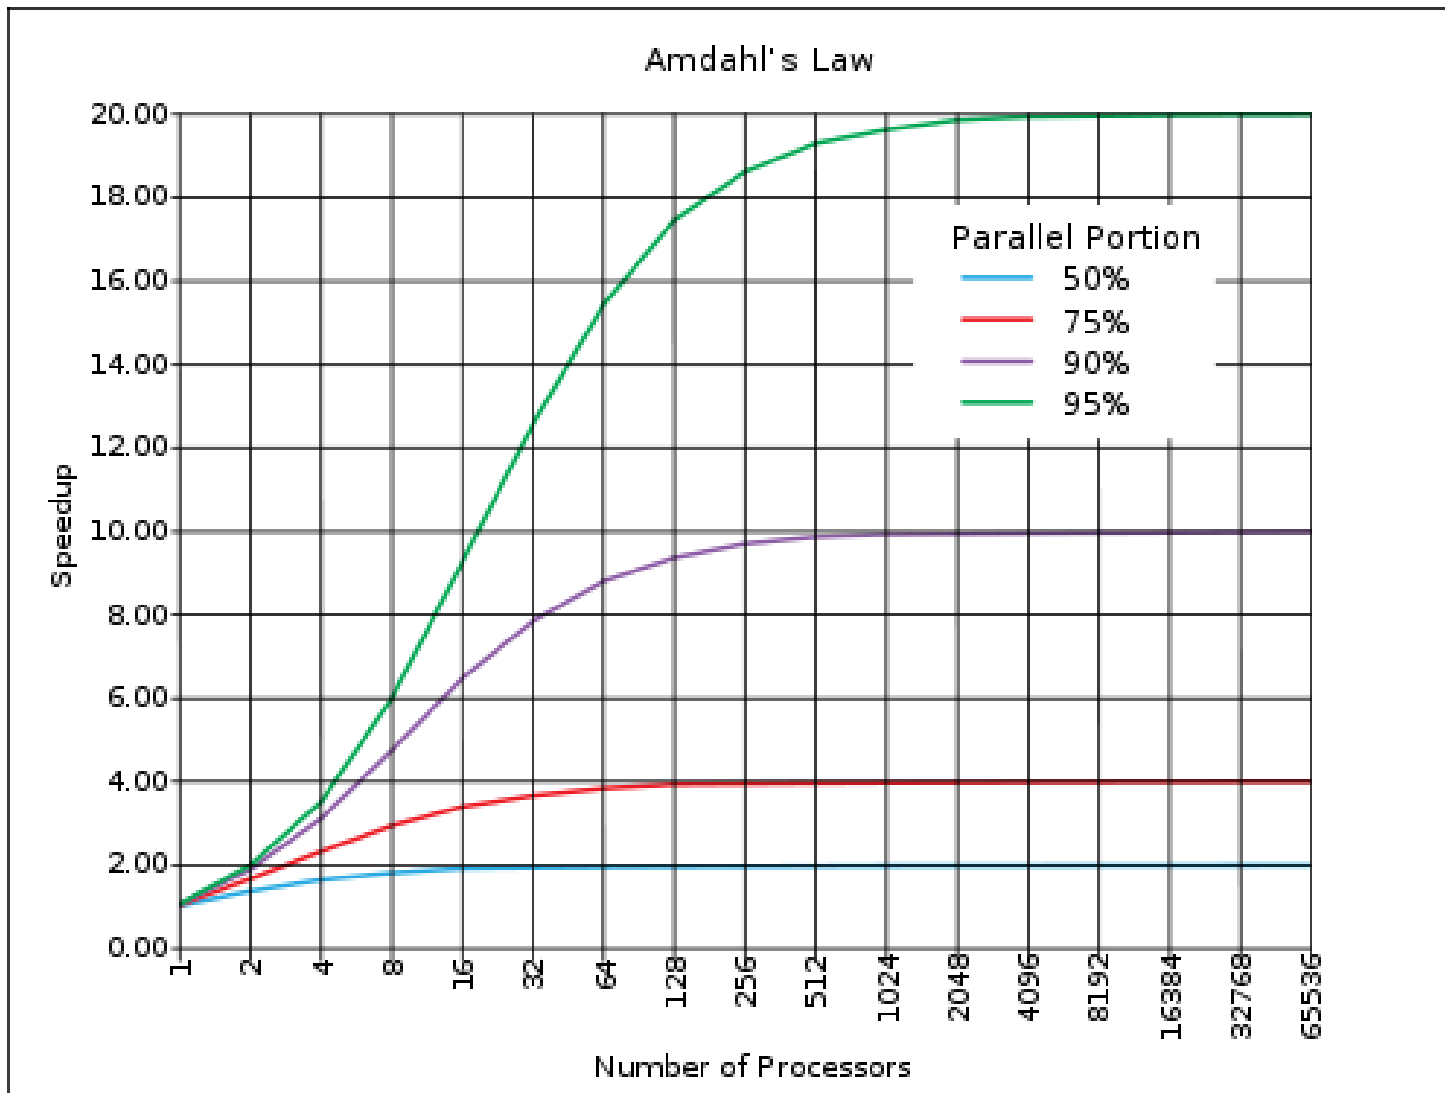
Amdahl's Law

- improve some part of a computer program
 - or it's execution speed (e.g., through parallelization)

$$T_{\text{improved}} = \frac{T_{\text{affected}}}{\text{improvement factor}} + T_{\text{unaffected}}$$

- limits overall performance improvement

Amdahl's Law



Source: Wikimedia Commons

Trade-Offs

- almost everything in CS is a trade-off
 - very few absolute truths
- “fast, good, or cheap – pick two”