

Chapter 1: Fundamentals of Quantitative Design & Analysis (Part 1)

What is computer architecture?

Why study computer architecture?

Common principles

What is Computer Architecture?

*What is Computer Architecture?**

Instruction set architecture (ISA)

Interface between hardware and software

Instructions visible to programmer

e.g., Intel IA32 vs. Intel 64 (x86-64) vs. IA64; ARM v7 vs. ARM v8;
RISC-V

Organization or Microarchitecture

High-level aspects of the system

e.g., how many functional units, pipeline organization,
cache/memory hierarchy, cores, accelerators, interconnect, ...

e.g., AMD Ryzen 7 vs. Intel Core i9;

ARM Cortex-A76 vs. Cortex-A72

Implementation or hardware

Logic design, packaging, ...

e.g., AMD Ryzen 7 2700X vs. 2700 (3.7 vs. 3.2 GHz base clock)

Previously, Computer Architecture ~ ISA

Instruction set architectures

Most ISAs today are general-purpose register based

Operands may be registers or memory locations

Register-memory vs. load-store

Addressing modes

Register, immediate, displacement, ...

Operand sizes

8 bits, 16 bits, 32 bits, 64 bits, SP and DP FP

Operations: Arithmetic, memory, control flow, floating point

Encoding: fixed vs. variable length

Evolution of ISAs

Pre-1980s: lots of action → CISC vs. RISC wars → 2 to 3
decades of (almost) stability → new questions again

Our main focus: organization

Goals of the Computer Architect

Goals of the Computer Architect**

Depends on type of computer

Internet of things (IoT)/embedded

Personal mobile device

Desktop

Server

Cluster/warehouse-scale

Supercomputer

*Goals of the Computer Architect***

Functional goals

- Meet application area demands

- Compatibility with previous systems

- Standards (e.g., IEEE floating point)

- Last through trends

Performance: Latency, throughput, real-time constraints, scalability

Cost

Power, Energy, Temperature, ...


Dependability

Security

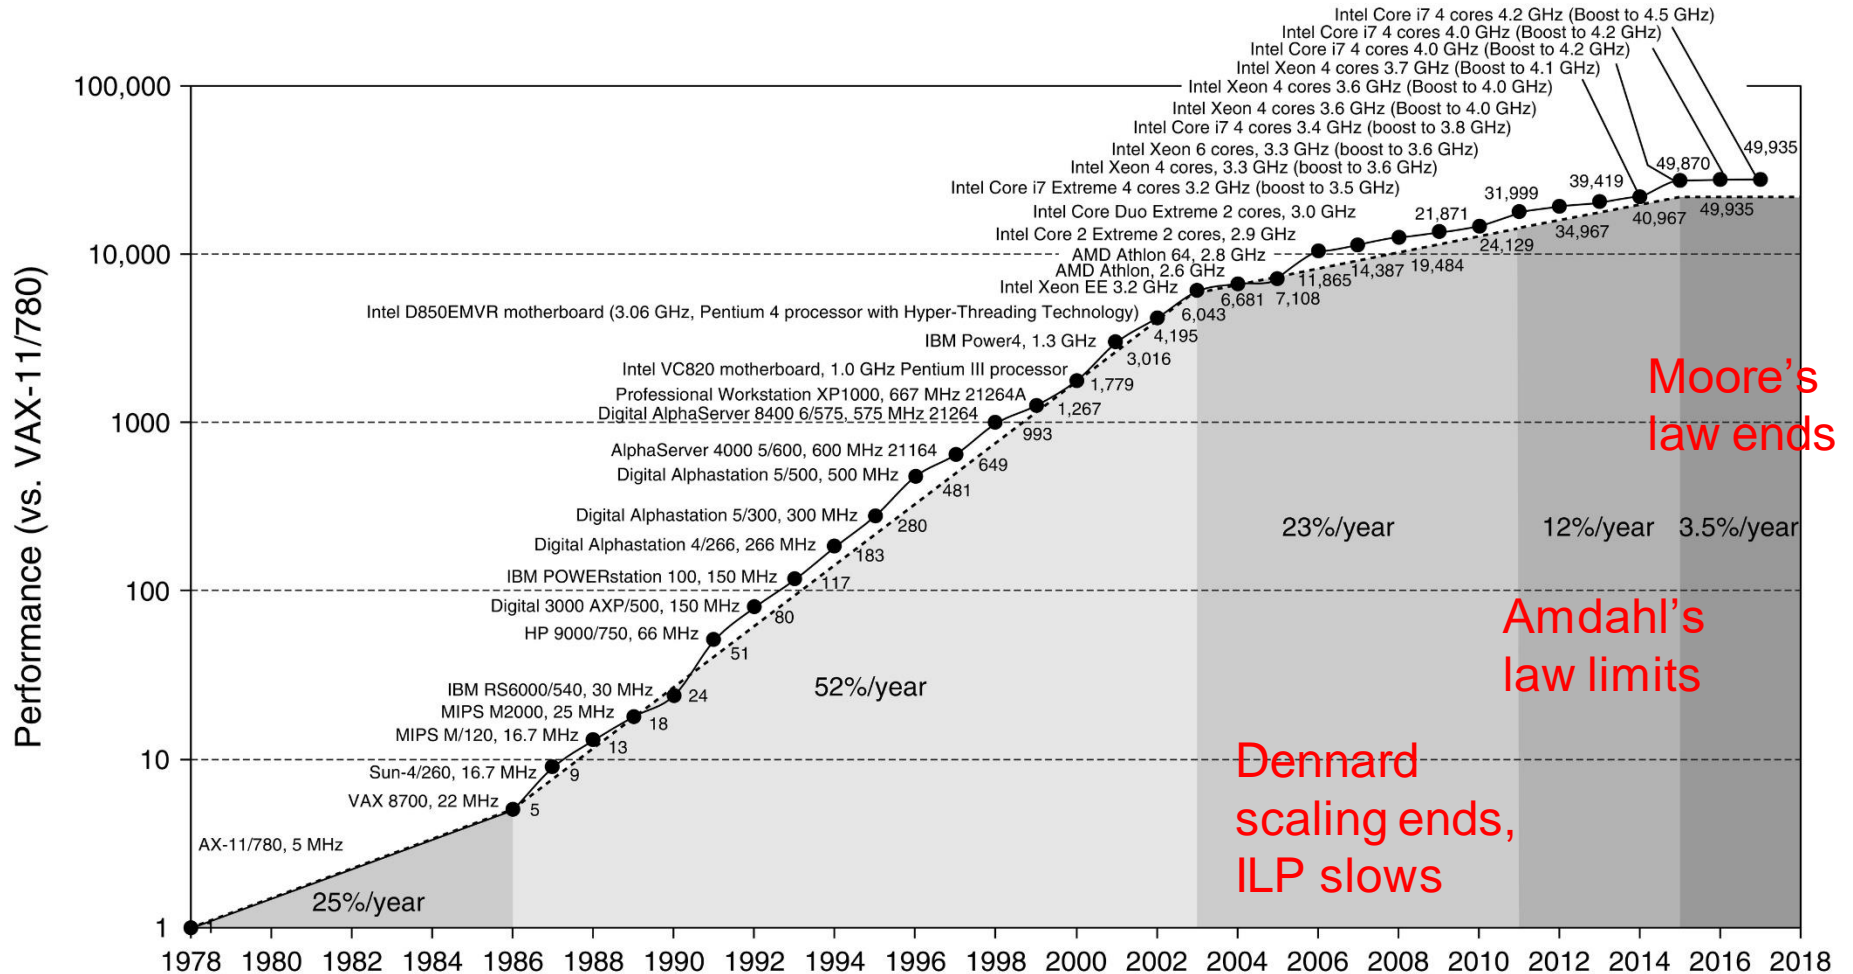
Maintainability, Verifiability, ...'ity...

Need to be familiar with design alternatives and criteria for selecting among them

Why Study Computer Architecture? - Historical Trends



Why Study Computer Architecture? - Historical Trends



Why Study Computer Architecture? - Historical Trends**




Figure 1.1 Growth in processor performance over 40 years. This chart plots program performance relative to the VAX 11/780 as measured by the SPEC integer benchmarks (see Section 1.8). Prior to the mid-1980s, growth in processor performance was largely technology-driven and averaged about 22% per year, or doubling performance every 3.5 years. The increase in growth to about 52% starting in 1986, or doubling every 2 years, is attributable to more advanced architectural and organizational ideas typified in RISC architectures. By 2003 this growth led to a difference in performance of an approximate factor of 25 versus the performance that would have occurred if it had continued at the 22% rate. In 2003 the limits of power due to the end of Dennard scaling and the available instruction-level parallelism slowed uniprocessor performance to 23% per year until 2011, or doubling every 3.5 years. (The fastest SPECintbase performance since 2007 has had automatic parallelization turned on, so uniprocessor speed is harder to gauge. These results are limited to single-chip systems with usually four cores per chip.) From 2011 to 2015, the annual improvement was less than 12%, or doubling every 8 years in part due to the limits of parallelism of Amdahl's Law. Since 2015, with the end of Moore's Law, improvement has been just 3.5% per year, or doubling every 20 years! Performance for floating-point-oriented calculations follows the same trends, but typically has 1% to 2% higher annual growth in each shaded region. Figure 1.11 on page 27 shows the improvement in clock rates for these same eras. Because SPEC has changed over the years, performance of newer machines is estimated by a scaling factor that relates the performance for different versions of SPEC: SPEC89, SPEC92, SPEC95, SPEC2000, and SPEC2006. There are too few results for SPEC2017 to plot yet.

Why Study Computer Architecture?

Why Study Computer Arch? Technology Trends**

Technology changes fast and on different curves

Capacity: past history:

Transistors/chip: 1.5X/year (Moore)

DRAM: 1.4X/year

Disk: 1.3X to 2X/year

Flash: 1.5X/year

All of the above changing now

Performance trends (~ 20 years):

CMOS scaling trends

Transistor vs. wire speed

Voltage curve flat (Dennard)

Power steeper

Reliability worse

Physical limits and end of Moore's law

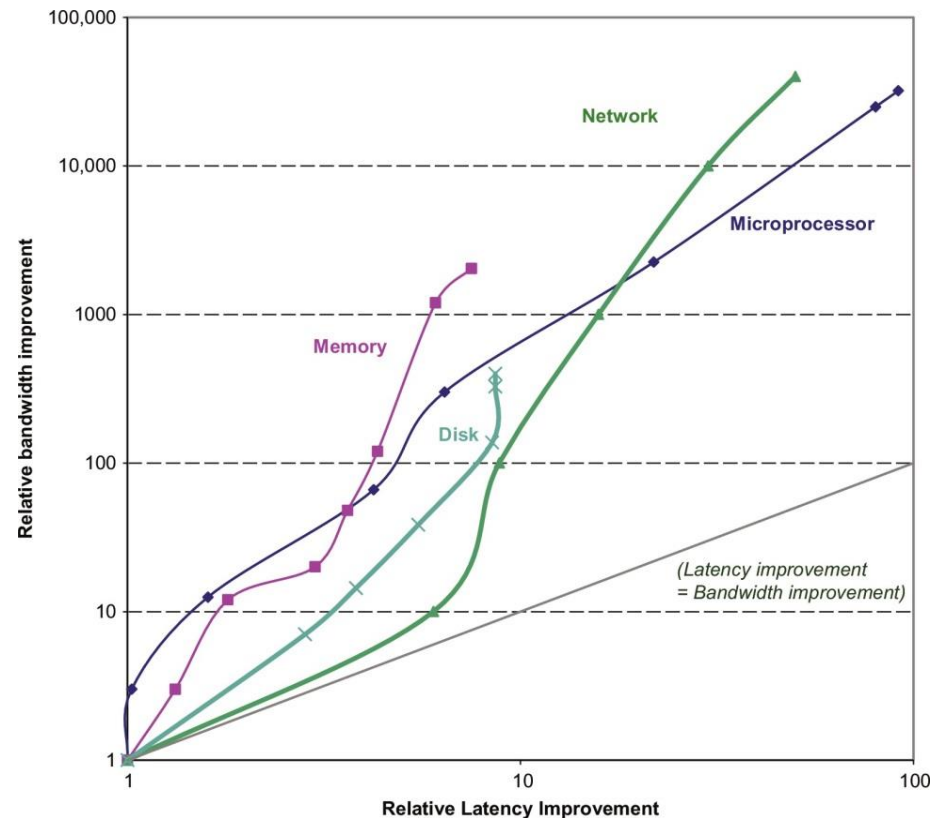
Major recent implications

Memory wall, ILP wall, power wall

Reliability/verifiability/maintainability/... walls

Led to multicore in spite of limited software base (Amdahl limit)

Recent: specialization/domain-specific architectures/heterogeneous systems



© 2019 Elsevier Inc. All rights reserved.

Why Study Computer Arch? Technology Trends**




Figure 1.9 Log-log plot of bandwidth and latency milestones in Figure 1.10 relative to the first milestone. Note that latency improved 8–91 x, while bandwidth improved about 400–32,000 x. Except for networking, we note that there were modest improvements in latency and bandwidth in the other three technologies in the six years since the last edition: 0%–23% in latency and 23%–70% in bandwidth. Updated from Patterson, D., 2004. Latency lags bandwidth. Commun. ACM 47 (10), 71–75.

Why Study Computer Architecture? (Cont.)**

Technology trends

Applications change

Scientific, business, personal computing, cloud, internet of things

Databases, graphics, multimedia, communications, machine learning

Software as a service

Next killer app??

New requirements: cost, availability, energy, maintainability, **security!**

New languages

E.g., shift from assembly to high-level languages

E.g., shift from C/C++ to Java/Python/Ruby

Compiler / hardware boundary shifts

[Revisit RISC vs. CISC]

Why Study Computer Architecture Today?



'Nobel Prize for Computing': Newly named Turing Award winners foretell a 'new golden age' for computer architecture at ISCA. .



John Hennessy and David Patterson Deliver Turing Lecture at ISCA 2018

2017 ACM A.M. Turing Award recipients John Hennessy and David Patterson delivered the Turing Lecture on June 4 at [ISCA 2018](#) in Los Angeles. The lecture took place from 5 to 6 p.m. PDT and was open to the public. A video of the lecture can be viewed below.

Titled "A New Golden Age for Computer Architecture: Domain-Specific Hardware/Software Co-Design, Enhanced Security, Open Instruction Sets, and Agile Chip Development," the talk covers recent developments and future directions in computer architecture.

Hennessy and Patterson were recognized with the Turing Award for "pioneering a systematic, quantitative approach to the design and evaluation of computer architectures with enduring impact on the microprocessor industry."

Golden Age of Computer Architecture!

See slides here:

<http://iscaconf.org/isca2018/docs/HennessyPattersonTuringLectureISCA4June2018.pdf>

Full video here:

<https://www.acm.org/hennessy-patterson-turing-lecture>

QnA: Why software community needs to learn about hardware now?

<https://youtu.be/3LVeEjsn8Ts?t=4268>

Relationship to Prerequisites

Prerequisite

How to design a computer?

This course

How to design a computer WELL?

Emphasis on Quantitative vs. Qualitative

Be sure to check the handout for details on the prerequisites