

MK COMPUTER ORGANIZATION AND DESIGN
The Hardware/Software Interface

EECS2021


Computer Organization Fall 2015

The slides are based on the publisher slides and contribution from Profs Amir Asif and Peter Lian
The slides will be modified, annotated, explained on the board, and sometimes corrected in the class

EECS2021 Computer Organization

- Computer Organization and Design– The hardware/Software approach – ARM Edition
- Patterson and John Hennessy
- Morgan kaufmann
- Assessment:

■ Participation	5%
■ Quizzes	20%
■ Lab	20%
■ Midterm	20%
■ Final	35%



MK Chapter 1 — Computer Abstractions and Technology — 2


EECS2021 Computer Organization

- Monday and Wednesday 5:30-7:00pm
- LAS B
- 2 Lab sections
 - Monday, Tuesday 7-10pm LAS 1006
- Labs are posted on the course web page

MK Chapter 1 — Computer Abstractions and Technology — 3


EECS2021

- Instructor
 - Mokhtar Aboelaze
 - Office LAS2026 Phone ext: 40607
- Research interests
 - Computer Architecture
 - Low power architecture
 - Embedded systems
 - FPGA (in embedded applications)

 Chapter 1 — Computer Abstractions and Technology — 4


Topics

- Computer abstraction and technology – Ch 1
- Instruction language of the computer – Ch 2
- Verilog -- Notes
- Arithmetic for computers – Ch 3
- The processor – Ch 4
- Memory hierarchy and I/O – Ch 5

 Chapter 1 — Computer Abstractions and Technology — 5

What You Will Learn

- How programs are translated into the machine language
 - And how the hardware executes them
- The hardware/software interface
- What determines program performance
 - And how it can be improved
- How the ALU works and how to improve its performance using pipelining.
- Memory Hierarchy and I/O (basic operation)

 Chapter 1 — Computer Abstractions and Technology — 6

Why EECS2021

- Required
- You can not call yourself ECE or CS if you don't know pipelining, assembly, how to measure performance, or how to report performance.
- Knowledge of hardware helps you write better programs



Chapter 1 — Computer Abstractions and Technology — 7

The Computer Revolution

- Progress in computer technology
 - Underpinned by Moore's Law
- Makes novel applications feasible
 - Computers in automobiles
 - Cell phones
 - Human genome project
 - World Wide Web
 - Search Engines
- Computers are pervasive



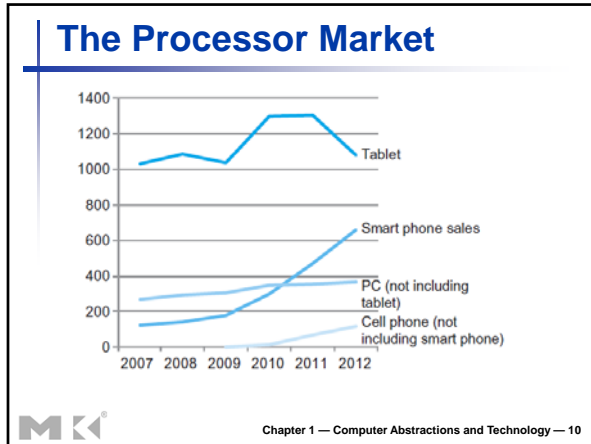
Chapter 1 — Computer Abstractions and Technology — 8

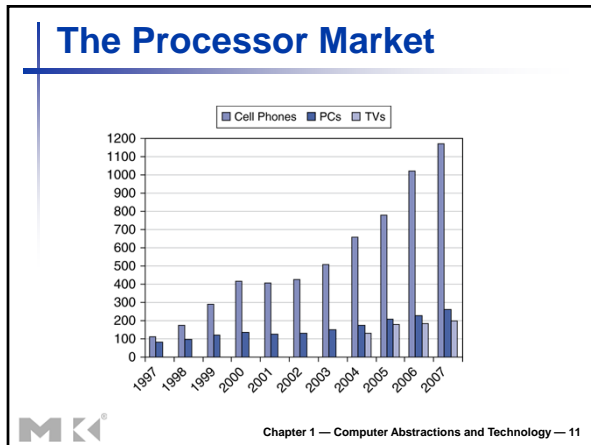
Classes of Computers

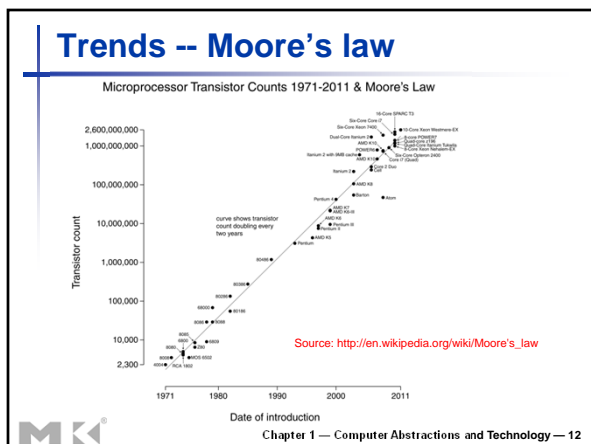
- Desktop computers
 - General purpose, variety of software
 - Subject to cost/performance tradeoff
- Server computers
 - Network based
 - High capacity, performance, reliability
 - Range from small servers to building sized
- Embedded computers
 - Hidden as components of systems
 - Stringent power/performance/cost constraints

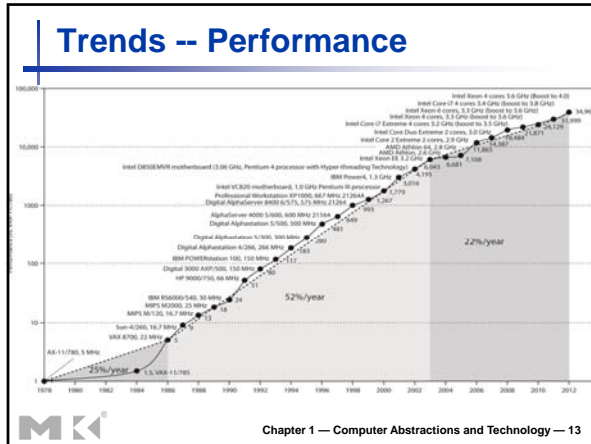


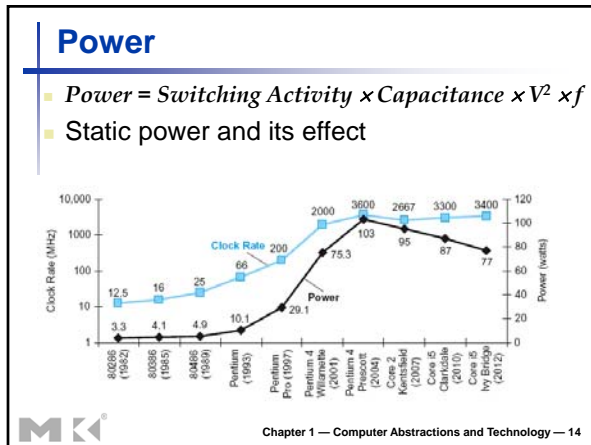
Chapter 1 — Computer Abstractions and Technology — 9





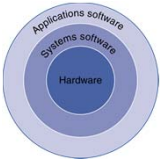






- ### Eight Great Ideas
1. Design for Moore's Law
 2. Use Abstraction to Simplify Design
 3. Make the Common case fast
 4. Performance via Parallelism
 5. Performance via Pipelining
 6. Performance via Prediction
 7. Hierarchy of Memories
 8. Dependability via Redundancy
- Chapter 1 — Computer Abstractions and Technology — 15

Below Your Program



- Application software
 - Written in high-level language
- System software
 - Compiler: translates HLL code to machine code
 - Operating System: service code
 - Handling input/output
 - Managing memory and storage
 - Scheduling tasks & sharing resources
- Hardware
 - Processor, memory, I/O controllers

Chapter 1 — Computer Abstractions and Technology — 16

How computers work?

- In your first year, you studies programming (java)
- Sequence of instructions
- Translated to machine language
- Instructions are fetched from the memory one after the other and executed

Chapter 1 — Computer Abstractions and Technology — 17

Levels of Program Code

- High-level language
 - Level of abstraction closer to problem domain
 - Provides for productivity and portability
- Assembly language
 - Textual representation of instructions
- Hardware representation
 - Binary digits (bits)
 - Encoded instructions and data

High-level language program (in C)

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

Assembly language program (for MIPS)

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

Binary machine language program (for MIPS)

```
000000010100001000000000000000010000
0000000000001100000011000000100001
1000110001100010000000000000000000
1000110011100100000000000000000000
1010110011110010000000000000000000
1010110001100010000000000000000100
00000011110000000000000000000000
```

Chapter 1 — Computer Abstractions and Technology — 18

Understanding Performance

- Algorithm
 - Determines number of operations executed
- Programming language, compiler, architecture
 - Determine number of machine instructions executed per operation
- Processor and memory system
 - Determine how fast instructions are executed
- I/O system (including OS)
 - Determines how fast I/O operations are executed



Chapter 1 — Computer Abstractions and Technology — 19

Components of a Computer

The BIG Picture

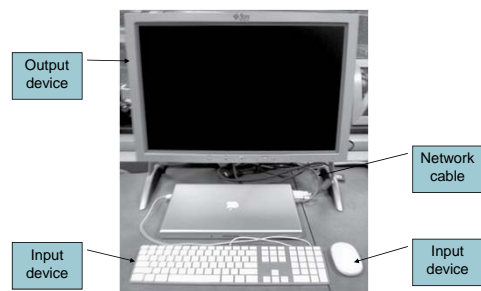


- Same components for all kinds of computer
 - Desktop, server, embedded
- Input/output includes
 - User-interface devices
 - Display, keyboard, mouse
 - Storage devices
 - Hard disk, CD/DVD, flash
 - Network adapters
 - For communicating with other computers



Chapter 1 — Computer Abstractions and Technology — 20

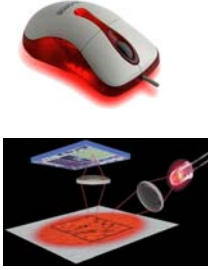
Anatomy of a Computer



Chapter 1 — Computer Abstractions and Technology — 21

Anatomy of a Mouse

- Optical mouse
 - LED illuminates desktop
 - Small low-res camera
 - Basic image processor
 - Looks for x, y movement
 - Buttons & wheel
- Supersedes roller-ball mechanical mouse

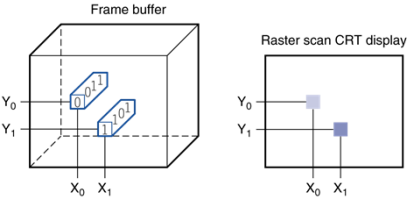


Morgan Kaufmann Publishers logo

Chapter 1 — Computer Abstractions and Technology — 22

Through the Looking Glass

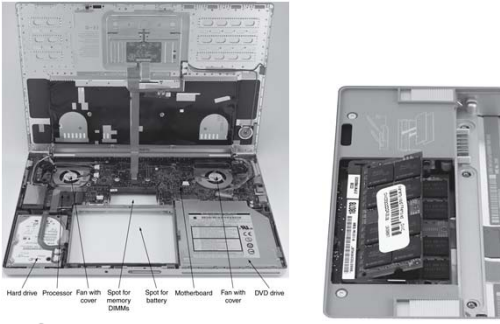
- LCD screen: picture elements (pixels)
 - Mirrors content of frame buffer memory



Morgan Kaufmann Publishers logo

Chapter 1 — Computer Abstractions and Technology — 23

Opening the Box



Hard drive Processor Fan with cover Spot for memory DIMMs Spot for battery Motherboard Fan with cover DVD drive

Morgan Kaufmann Publishers logo

Chapter 1 — Computer Abstractions and Technology — 24

A Safe Place for Data


- Volatile main memory
 - Loses instructions and data when power off
- Non-volatile secondary memory
 - Magnetic disk
 - Flash memory
 - Optical disk (CDROM, DVD)



Morgan Kaufmann Publishers logo (MKP) and Chapter 1 — Computer Abstractions and Technology — 25

Networks

- Communication and resource sharing
- Local area network (LAN): Ethernet
 - Within a building
- Wide area network (WAN: the Internet)
- Wireless network: WiFi, Bluetooth



Morgan Kaufmann Publishers logo (MKP) and Chapter 1 — Computer Abstractions and Technology — 26

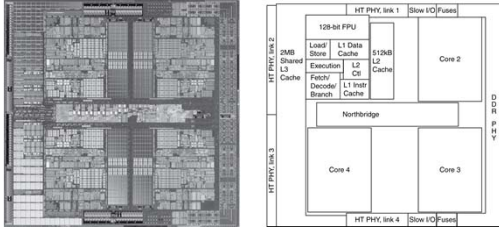
Inside the Processor (CPU)

- Datapath: performs operations on data
- Control: sequences datapath, memory, ...
- Cache memory
 - Small fast SRAM memory for immediate access to data

Morgan Kaufmann Publishers logo (MKP) and Chapter 1 — Computer Abstractions and Technology — 27

Inside the Processor

- AMD Barcelona: 4 processor cores



Chapter 1 — Computer Abstractions and Technology — 28

Abstractions

The BIG Picture

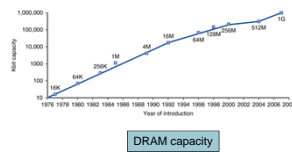
- Abstraction helps us deal with complexity
 - Hide lower-level detail
- Instruction set architecture (ISA)
 - The hardware/software interface
- Application binary interface
 - The ISA plus system software interface
- Implementation
 - The details underlying and interface



Chapter 1 — Computer Abstractions and Technology — 29

Technology Trends

- Electronics technology continues to evolve
 - Increased capacity and performance
 - Reduced cost



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



Chapter 1 — Computer Abstractions and Technology — 30

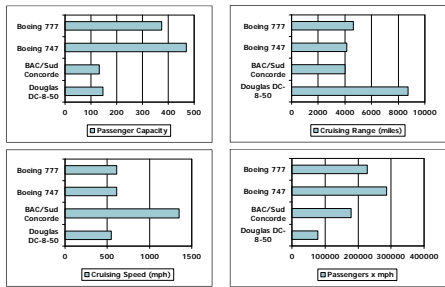
Performance

- Which computer is better? Even which computer is faster?
- We can not answer the question without defining what we mean by “better” or “faster”



Defining Performance

- Which airplane has the best performance?



Response Time and Throughput

- Response time
 - How long it takes to do a task
- Throughput
 - Total work done per unit time
 - e.g., tasks/transactions/... per hour
- How are response time and throughput affected by
 - Replacing the processor with a faster version?
 - Adding more processors?
- We'll focus on response time for now...



Relative Performance

- Define Performance = 1/Execution Time
- "X is *n* time faster than Y"
 - $Performance_x / Performance_y = Execution\ time_y / Execution\ time_x = n$
- Example: time taken to run a program
 - 10s on A, 15s on B
 - $Execution\ Time_B / Execution\ Time_A = 15s / 10s = 1.5$
 - So A is 1.5 times faster than B (*speedup*)



Chapter 1 — Computer Abstractions and Technology — 34

Measuring Execution Time

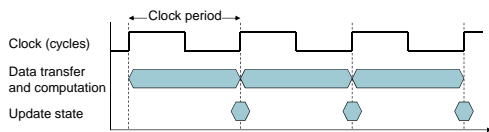
- Elapsed time
 - Total response time, including all aspects
 - Processing, I/O, OS overhead, idle time
 - Determines system performance
- CPU time
 - Time spent processing a given job
 - Discounts I/O time, other jobs' shares
 - Comprises user CPU time and system CPU time
 - Different programs are affected differently by CPU and system performance



Chapter 1 — Computer Abstractions and Technology — 35

CPU Clocking

- Operation of digital hardware governed by a constant-rate clock



- Clock period: duration of a clock cycle
 - e.g., 250ps = 0.25ns = $250 \times 10^{-12}s$
- Clock frequency (rate): cycles per second
 - e.g., 4.0GHz = 4000MHz = 4.0×10^9Hz




Chapter 1 — Computer Abstractions and Technology — 36

CPU Time

CPU Time = CPU Clock Cycles × Clock Cycle Time

$$= \frac{\text{CPU Clock Cycles}}{\text{Clock Rate}}$$

- Performance improved by
 - Reducing number of clock cycles
 - Increasing clock rate
 - Hardware designer must often trade off clock rate against cycle count



Chapter 1 — Computer Abstractions and Technology — 37


CPU Time Example

- Computer A: 2GHz clock, 10s CPU time
- Designing Computer B
 - Aim for 6s CPU time
 - Can do faster clock, but causes 1.2 × clock cycles
- How fast must Computer B clock be?

$$\text{Clock Rate}_B = \frac{\text{Clock Cycles}_B}{\text{CPU Time}_B} = \frac{1.2 \times \text{Clock Cycles}_A}{6s}$$

$$\text{Clock Cycles}_A = \text{CPU Time}_A \times \text{Clock Rate}_A$$

$$= 10s \times 2\text{GHz} = 20 \times 10^9$$

$$\text{Clock Rate}_B = \frac{1.2 \times 20 \times 10^9}{6s} = \frac{24 \times 10^9}{6s} = 4\text{GHz}$$


Chapter 1 — Computer Abstractions and Technology — 38


Instruction Count and CPI

Clock Cycles = Instruction Count × Cycles per Instruction

CPU Time = Instruction Count × CPI × Clock Cycle Time

$$= \frac{\text{Instruction Count} \times \text{CPI}}{\text{Clock Rate}}$$

- Instruction Count for a program
 - Determined by program, ISA and compiler
- Average cycles per instruction
 - Determined by CPU hardware
 - If different instructions have different CPI
 - Average CPI affected by instruction mix



Chapter 1 — Computer Abstractions and Technology — 39

CPI Example


- Computer A: Cycle Time = 250ps, CPI = 2.0
- Computer B: Cycle Time = 500ps, CPI = 1.2
- Same ISA
- Which is faster, and by how much?

$$\text{CPU Time}_A = \text{Instruction Count} \times \text{CPI}_A \times \text{Cycle Time}_A$$

$$= 1 \times 2.0 \times 250\text{ps} = 1 \times 500\text{ps} \quad \leftarrow \text{A is faster...}$$

$$\text{CPU Time}_B = \text{Instruction Count} \times \text{CPI}_B \times \text{Cycle Time}_B$$

$$= 1 \times 1.2 \times 500\text{ps} = 1 \times 600\text{ps}$$

$$\frac{\text{CPU Time}_B}{\text{CPU Time}_A} = \frac{1 \times 600\text{ps}}{1 \times 500\text{ps}} = 1.2 \quad \leftarrow \text{...by this much}$$



Chapter 1 — Computer Abstractions and Technology — 40

CPI in More Detail

- If different instruction classes take different numbers of cycles

$$\text{Clock Cycles} = \sum_{i=1}^n (\text{CPI}_i \times \text{Instruction Count}_i)$$

- Weighted average CPI

$$\text{CPI} = \frac{\text{Clock Cycles}}{\text{Instruction Count}} = \sum_{i=1}^n \left(\underbrace{\text{CPI}_i \times \frac{\text{Instruction Count}_i}{\text{Instruction Count}}}_{\text{Relative frequency}} \right)$$



Chapter 1 — Computer Abstractions and Technology — 41

CPI Example

- Alternative compiled code sequences using instructions in classes A, B, C

Class	A	B	C
CPI for class	1	2	3
IC in sequence 1	2	1	2
IC in sequence 2	4	1	1

- Sequence 1: IC = 5
- Sequence 2: IC = 6
- Clock Cycles = $2 \times 1 + 1 \times 2 + 2 \times 3 = 10$
- Clock Cycles = $4 \times 1 + 1 \times 2 + 1 \times 3 = 9$
- Avg. CPI = $10/5 = 2.0$
- Avg. CPI = $9/6 = 1.5$



Chapter 1 — Computer Abstractions and Technology — 42


Performance Summary

The BIG Picture

$$\text{CPU Time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Clock cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Clock cycle}}$$

IC CPI T_c


- Performance depends on
 - Algorithm: affects IC, possibly CPI
 - Programming language: affects IC, CPI
 - Compiler: affects IC, CPI
 - Instruction set architecture: affects IC, CPI, T_c



Chapter 1 — Computer Abstractions and Technology — 43

Benchmarks

- Performance is not determined by measuring the performance of a single program.
- A bunch of programs to measure performance
- SPEC CPU for example
- Spec rating is geometric mean of performance (1/execution_time)




Chapter 1 — Computer Abstractions and Technology — 44

Reporting Performance

- Assume 3 programs and 3 systems

	P1	P2	P3
A	10	8	25
B	12	9	20
C	8	8	30

- Arithmetic mean
- Geometric mean



Chapter 1 — Computer Abstractions and Technology — 45

Reporting Performance

- 2 Programs and 3 machines

	A	B	C
P1	1	10	20
P2	1000	100	20

MK Chapter 1 — Computer Abstractions and Technology — 46

SPEC Cfp2000

Benchmark	Ultra5 t	Optero n	EPECR atio	Itanium 2	SPECR atio	O/I time	O/ISpe cR
wupwise	1600	51.5	31.06	56.1	28.53	0.92	0.92
Swim	3100	125	24.73	70.7	43.85	1.77	1.77
Mgrid	1800	98	18.37	65.8	27.36	1.49	1.49
Applu	2100	94	22.43	50.9	41.25	1.85	1.85
Mesa	1400	64.6	21.69	108.0	12.99	0.6	0.6
Galgel	2900	86.4	33.57	40.0	72.47	2.16	2.16
Art	2600	92	28.13	21.0	123.67	4.40	4.40
Equake	1300	72.6	17.92	36.3	35.78	2.00	2.00
Facerec	1900	73.6	25.80	86.9	21.86	0.85	0.85
Amp	2200	136	16.14	132.0	16.63	1.03	1.03
Lucas	2000	88.8	22.52	107.0	18.76	0.83	0.83
Fma3d	2100	120	17.48	131.0	16.09	0.92	0.92
Sixtrack	1100	123	8.95	68.8	15.99	1.79	1.79
Aspl	2600	150	17.36	231.0	11.27	0.65	0.65
GM			20.86		27.12	1.30	1.30

M
