

**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
- Patterson and John Hennessy
- Morgan kaufmann
- Assessment:
 

■ Assignments/Quizzes	20%
■ Lab	20%
■ Midterm	25%
■ Final	35%



**MK** Chapter 1 — Computer Abstractions and Technology — 2

---

---

---

---

---

---

---

---

### EECS2021 Computer Organization

- Monday and Wednesday 5:30-7:00pm
- CLH A
- 2 Lab sections
  - Monday 7-10pm LAS 1006/1004
  - Tuesday 7-10pm LAS 1006/1004/1002
- Labs start The week of Sept. 21
- 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
- Cache – 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.
- Cache memory (basic operation)

 Chapter 1 — Computer Abstractions and Technology — 6

---

---

---

---

---

---

---

---

### 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




---

---

---

---

---

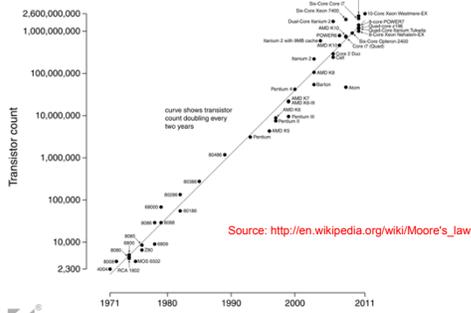
---

---

---

### Moore's law

Microprocessor Transistor Counts 1971-2011 & Moore's Law




---

---

---

---

---

---

---

---

### 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




---

---

---

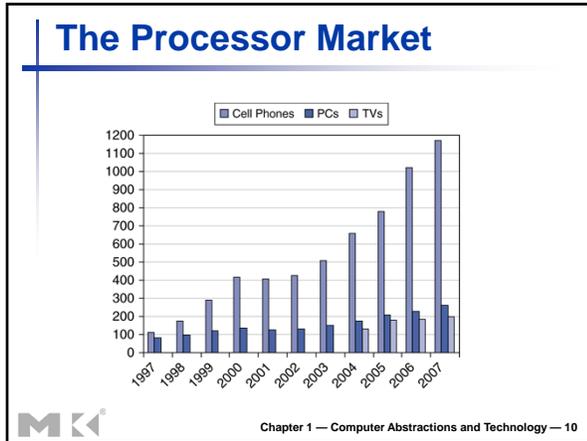
---

---

---

---

---



---

---

---

---

---

---

---

---

---

---

### 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

Morgan Kaufmann Publishers logo (MKP) and Chapter 1 — Computer Abstractions and Technology — 11

---

---

---

---

---

---

---

---

---

---

### 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

Morgan Kaufmann Publishers logo (MKP) and Chapter 1 — Computer Abstractions and Technology — 12

---

---

---

---

---

---

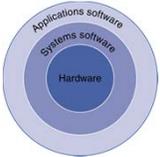
---

---

---

---

### 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 — 13

---

---

---

---

---

---

---

---

### 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)

Assembly language program (for MIPS)

Binary machine language program (for MIPS)

Compiler

Assembler

```

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

swap:
  muli $2, $5, 4
  add $2, $4, $2
  lw $16, 0($2)
  lw $16, 4($2)
  sw $16, 0($2)
  sw $16, 4($2)
  jr $31
          
```

```

00000010100010000000000000000000
0000000000110000001100000100001
10011001100100010000000000000000
10011001100100000000000000000000
10101001100100000000000000000000
10101000100010000000000000000000
000000111100000000000000000000
          
```

Chapter 1 — Computer Abstractions and Technology — 14

---

---

---

---

---

---

---

---

### 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 — 15

---

---

---

---

---

---

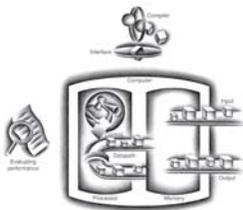
---

---

## Components of a Computer

1.3 Understanding Computers

**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

MKK Chapter 1 — Computer Abstractions and Technology — 16

---

---

---

---

---

---

---

---

---

---

## Anatomy of a Computer



MKK Chapter 1 — Computer Abstractions and Technology — 17

---

---

---

---

---

---

---

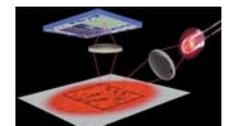
---

---

---

## 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

MKK Chapter 1 — Computer Abstractions and Technology — 18

---

---

---

---

---

---

---

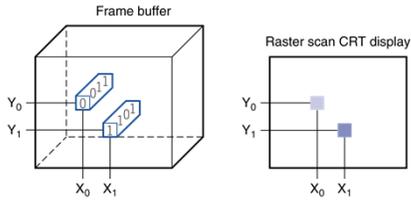
---

---

---

### Through the Looking Glass

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



Chapter 1 — Computer Abstractions and Technology — 19

---

---

---

---

---

---

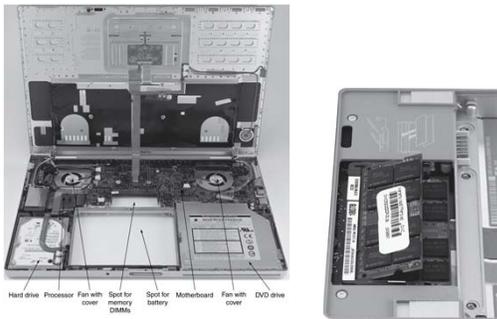
---

---

---

---

### Opening the Box



Chapter 1 — Computer Abstractions and Technology — 20

---

---

---

---

---

---

---

---

---

---

### 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)



Chapter 1 — Computer Abstractions and Technology — 21

---

---

---

---

---

---

---

---

---

---

### Networks

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



Chapter 1 — Computer Abstractions and Technology — 22

---

---

---

---

---

---

---

---

### Inside the Processor (CPU)

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



Chapter 1 — Computer Abstractions and Technology — 23

---

---

---

---

---

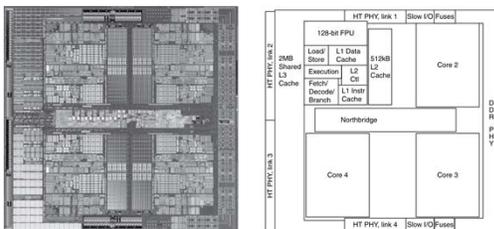
---

---

---

### Inside the Processor

- AMD Barcelona: 4 processor cores



Chapter 1 — Computer Abstractions and Technology — 24

---

---

---

---

---

---

---

---

## 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 — 25

---

---

---

---

---

---

---

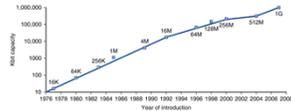
---

---

---

## 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 (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 — 26

---

---

---

---

---

---

---

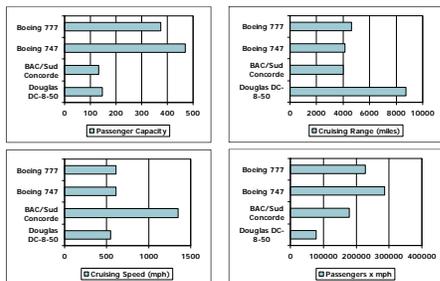
---

---

---

## Defining Performance

- Which airplane has the best performance?



Chapter 1 — Computer Abstractions and Technology — 27

---

---

---

---

---

---

---

---

---

---

### 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...



Chapter 1 — Computer Abstractions and Technology — 28

---

---

---

---

---

---

---

---

### Relative Performance

- Define Performance = 1/Execution Time
- "X is *n* time faster than Y"
 

$$\frac{\text{Performance}_x}{\text{Performance}_y} = \frac{\text{Execution time}_y}{\text{Execution time}_x} = n$$
- Example: time taken to run a program
  - 10s on A, 15s on B
  - $\text{Execution Time}_B / \text{Execution Time}_A = 15s / 10s = 1.5$
  - So A is 1.5 times faster than B



Chapter 1 — Computer Abstractions and Technology — 29

---

---

---

---

---

---

---

---

### 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 — 30

---

---

---

---

---

---

---

---

### 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 =  $4.0 \times 10^9$ Hz

Morgan Kaufmann Publishers  
Chapter 1 — Computer Abstractions and Technology — 31

---

---

---

---

---

---

---

---

---

---

### CPU Time

CPU Time = CPU Clock Cycles  $\times$  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

Morgan Kaufmann Publishers  
Chapter 1 — Computer Abstractions and Technology — 32

---

---

---

---

---

---

---

---

---

---

### 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  $\times$  clock cycles
- How fast must Computer B clock be?

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

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

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

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

Morgan Kaufmann Publishers  
Chapter 1 — Computer Abstractions and Technology — 33

---

---

---

---

---

---

---

---

---

---

### 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 — 34

---

---

---

---

---

---

---

---

### 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?

CPU Time<sub>A</sub> = Instruction Count × CPI<sub>A</sub> × Cycle Time<sub>A</sub>

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

CPU Time<sub>B</sub> = Instruction Count × CPI<sub>B</sub> × Cycle Time<sub>B</sub>

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

CPU Time<sub>B</sub> = I × 600ps

CPU Time<sub>A</sub> = I × 500ps

$$\frac{\text{CPU Time}_B}{\text{CPU Time}_A} = 1.2 \leftarrow \text{...by this much}$$


Chapter 1 — Computer Abstractions and Technology — 35

---

---

---

---

---

---

---

---

### 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 — 36

---

---

---

---

---

---

---

---

### 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 — 37

---

---

---

---

---

---

---

---

---

---

### Performance Summary

The BIG Picture

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

- 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 — 38

---

---

---

---

---

---

---

---

---

---

### 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 — 39

---

---

---

---

---

---

---

---

---

---

### Reporting Performance

- 2 Programs and 3 machines

	A	B	C
P1	1	10	20
P2	1000	100	20

MK<sup>®</sup> Chapter 1 — Computer Abstractions and Technology — 40

---

---

---

---

---

---

---

---

### Power Trends

In CMOS IC technology

Power = Capacitive load × Voltage<sup>2</sup> × Frequency

Annotations: ×30 (under Capacitive load), 5V → 1V (under Voltage<sup>2</sup>), ×1000 (under Frequency)

MK<sup>®</sup> Chapter 1 — Computer Abstractions and Technology — 41

---

---

---

---

---

---

---

---

### Reducing Power

- Suppose a new CPU has
  - 85% of capacitive load of old CPU
  - 15% voltage and 15% frequency reduction

$$\frac{P_{new}}{P_{old}} = \frac{C_{old} \times 0.85 \times (V_{old} \times 0.85)^2 \times F_{old} \times 0.85}{C_{old} \times V_{old}^2 \times F_{old}} = 0.85^4 = 0.52$$

- The power wall
  - We can't reduce voltage further
  - We can't remove more heat
- How else can we improve performance?

MK<sup>®</sup> Chapter 1 — Computer Abstractions and Technology — 42

---

---

---

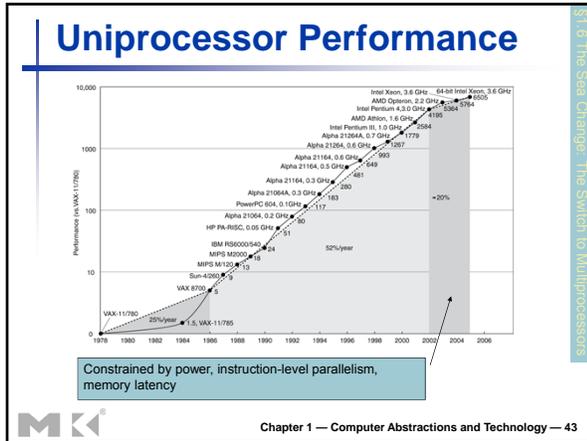
---

---

---

---

---




---

---

---

---

---

---

---

---

---

---

- ### Multiprocessors
- Multicore microprocessors
    - More than one processor per chip
  - Requires explicitly parallel programming
    - Compare with instruction level parallelism
      - Hardware executes multiple instructions at once
      - Hidden from the programmer
    - Hard to do
      - Programming for performance
      - Load balancing
      - Optimizing communication and synchronization
- Chapter 1 — Computer Abstractions and Technology — 44

---

---

---

---

---

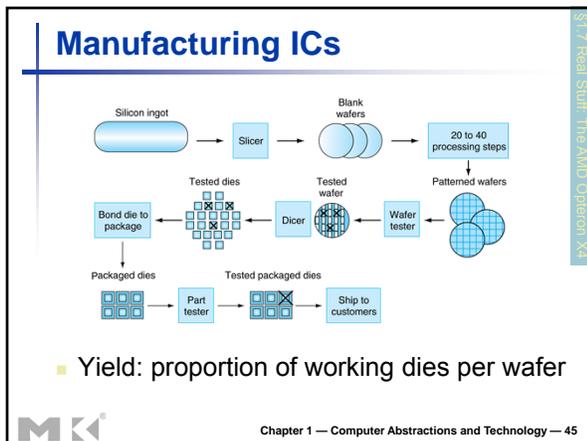
---

---

---

---

---




---

---

---

---

---

---

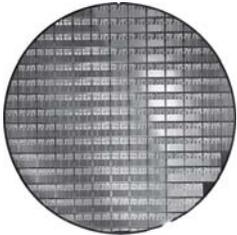
---

---

---

---

### AMD Opteron X2 Wafer



- X2: 300mm wafer, 117 chips, 90nm technology
- X4: 45nm technology

**MK** Chapter 1 — Computer Abstractions and Technology — 46

---

---

---

---

---

---

---

---

---

---

### Integrated Circuit Cost

$$\text{Cost per die} = \frac{\text{Cost per wafer}}{\text{Dies per wafer} \times \text{Yield}}$$

$$\text{Dies per wafer} \approx \frac{\text{Wafer area}}{\text{Die area}}$$

$$\text{Yield} = \frac{1}{(1 + (\text{Defects per area} \times \text{Die area}/2))^2}$$

- Nonlinear relation to area and defect rate
  - Wafer cost and area are fixed
  - Defect rate determined by manufacturing process
  - Die area determined by architecture and circuit design

**MK** Chapter 1 — Computer Abstractions and Technology — 47

---

---

---

---

---

---

---

---

---

---

### SPEC CPU Benchmark

- Programs used to measure performance
  - Supposedly typical of actual workload
- Standard Performance Evaluation Corp (SPEC)
  - Develops benchmarks for CPU, I/O, Web, ...
- SPEC CPU2006
  - Elapsed time to execute a selection of programs
    - Negligible I/O, so focuses on CPU performance
  - Normalize relative to reference machine
  - Summarize as geometric mean of performance ratios
    - CINT2006 (integer) and CFP2006 (floating-point)

$$\sqrt[n]{\prod_{i=1}^n \text{Execution time ratio}_i}$$

**MK** Chapter 1 — Computer Abstractions and Technology — 48

---

---

---

---

---

---

---

---

---

---





### Pitfall: MIPS as a Performance Metric

- MIPS: Millions of Instructions Per Second
  - Doesn't account for
    - Differences in ISAs between computers
    - Differences in complexity between instructions

$$\begin{aligned} \text{MIPS} &= \frac{\text{Instruction count}}{\text{Execution time} \times 10^6} \\ &= \frac{\text{Instruction count}}{\frac{\text{Instruction count} \times \text{CPI}}{\text{Clock rate}} \times 10^6} = \frac{\text{Clock rate}}{\text{CPI} \times 10^6} \end{aligned}$$

- CPI varies between programs on a given CPU




---

---

---

---

---

---

---

---

---

---

### Concluding Remarks

- Cost/performance is improving
  - Due to underlying technology development
- Hierarchical layers of abstraction
  - In both hardware and software
- Instruction set architecture
  - The hardware/software interface
- Execution time: the best performance measure
- Power is a limiting factor
  - Use parallelism to improve performance

Performance Metrics




---

---

---

---

---

---

---

---

---

---