Computer Architecture

Unmesh D. Bordoloi





The first electronic computers



ENIAC (Electronic Numerical Integrator and Calculator) – The world's first electronic computer



- ENIAC provided conditional jumps and was programmable, clearly distinguishing it from earlier calculators
- Programming was done manually by plugging cables and setting switches, and data was entered on punched cards.
 Programming for typical calculations required from half an hour to a whole day
- ENIAC was a general-purpose machine, limited primarily by a small amount of storage and tedious programming

ENIAC

- J. Presper Eckert and John Mauchly at the Moore School of the University of Pennsylvania
- Funded by the United States Army
- Became operational during World War II but was not publicly disclosed until 1946

Von Neumann

- In 1944, John von Neumann was attracted to the ENIAC project. The group wanted to improve the way programs were entered and discussed storing programs as numbers; von Neumann helped crystallize the ideas and wrote a memo proposing a stored-program computer called EDVAC (Electronic Discrete Variable Automatic Computer).
- Herman Goldstine distributed the memo and put von Neumann's name on it, much to the dismay of Eckert and Mauchly, whose names were omitted. This memo has served as the basis for the commonly used term von Neumann computer. Several early pioneers in the computer field believe that this term gives too much credit to von Neumann, who wrote up the ideas, and too little to the engineers, Eckert and Mauchly, who worked on the machines.
- For this reason, the term does not appear elsewhere in your textbook book.

Why study Computer Architecture ?

Dilbert by Scott Adams I'M SO YOUR PROJECT IS T HAVEN! BORFO. MY TOP PRIORITY. L L STARTED LL ME EVERYTHING CAN T. TE TALKING GETS STAY. THAT I NEED TO YET. ALIAKE: WORSE??! KNOW. O-O-OREY

e ora inc.

Main goals of this course

- Understand and explain the main components of a computer
- Make connections between programs and the way they are executed on hardware infrastructure
- Evaluate how the different design choices influence the performance of computer
- Understand the interaction between the different components and their impact on performance

Will be evaluated by ...

Labs

- Written exam: Ability to solve problems
- Written exam: Ability to explain the fundamental concepts

In the classroom

- Sometimes, you will solve problems
 - I will interact with you, but this will not be graded
- For this, please be seated in a small group of 2 or 3
 - You may use your lab group

An experiment with Twitter

- Optional
- Use #tdts10
- Use to ask doubts, or share interesting results, when you do not feel like spamming the entire group
- Use also in classroom only when solving problems

Examination

- The course will have three obligatory exams/homework
 - Exam
 - Lab
- ■<u>Lab</u>
 - This is obligatory and you can get a pass/fail grade
- Final Exam

 There will be a final written exam. More details will be shared later

Personnel

- Unmesh Bordoloi- Kursledare (examinator)
 - Epost: <u>unmesh.bordoloi@liu.se</u>
- Dimitar Nikolov (Teaching Assistant)
 - Epost: <u>dimitar.nikolov@liu.se</u>
- Madeleine Dahlqvist (Kurssekreterare)
 - Epost: <u>madeline.hager.dahlqvist@liu.se</u>
- Patrick Lambrix (Director of Studies)
 - Epost: <u>patrick.lambrix@liu.se</u>



LiU



Textbook

- The recommended textbook is:
 - David A. Patterson and John L. Hennessy: Computer Organization and Design - The Hardware / Software Interface, 4th Edition, Morgan Kaufmann
 - Available as e-book in the university library

Other literature

- William Stallings: Computer Organization and Architecture, Prentice Hall International, Inc..
- Sven Eklund, Avancerad datorarkitektur, Studentlitteratur, 1994.
- Andrew S. Tanenbaum, Structured Computer Organization, 4th edition, Prentice Hall International, Inc., 1999.
- V. C. Hamacher, et al.: Computer Organization, 4th edition, McGraw-Hill, 1996.



4 hours!

• 3 + 1 hour!



Technology

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

Yesterday's fiction is now reality

- Applications empowered by computers:
 - Human genome project
 - World Wide Web
 - Search Engines
 - Social Networks

Classes of Computers



















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

Below your program

- We constantly interact with these computers
 - E.g., via apps on iPhone or via a Word Processor
- How does an application interact with the hardware?

Language of the hardware

- The hardware understands on or off. Hence, the language of the hardware needs two symbols 0 and 1.
- Commonly referred to as binary digits or bits.
- 2 letters do not limit what computers can do (just like the finite letters in our languages)

Language of the hardware

- Instructions are collections of bits that the computer understands
- What our programs instruct the computer to do are the <u>instructions</u>:
- Hundreds/thousands/millions of lines of code (instructions) are hidden beneath our apps and programs

Hierarchical Layers 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



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

Under the Covers

- Understanding the underlying hardware (the computer!) is the main focus of this course
- So, what are the main components of the computer?

Components of a Computer



- Same components for all kinds of computer
 - Desktop, server, embedded

Anatomy of a Computer



Opening the Box



- Motherboard
 - I/O connections
 - Memory (DRAM)
 - CPU or central processing unit

AMD Barcelona: 4 processor cores





Inside the Processor (CPU)

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

Abstractions

- Abstraction helps us deal with complexity
 - Note abstraction in both hardware and software
 - Hide lower-level detail
- Instruction set architecture (ISA)
 - The hardware/software interface
- Implementation
 - The details underlying and interface

Our Topics in this Course

- Instructions
- Arithmetic
- The Processor
- Memory
- Input/Output
- Multi-cores and GPUs

Performance

- Why is performance important?
 - For purchasers: to choose between computers
 - For designers: to make the sales pitch
- Defining performance is not straightforward!
 - An analogy with airplanes shows the difficulty

Defining Performance

Which airplane has the best performance?



Response Time and Throughput

- Response time
 - As a user of a smart phone (embedded computer), or laptop, the one that responds faster is the better!
 - Response time = How long it takes to do a task?
 - Response time = the total time required for the computer to complete a task, including disk accesses, memory accesses, I/O activities, operating system overheads, CPU execution time ...

Response Time and Throughput

- Throughput
 - If I am running a data center with several servers, faster computer is the one that completes several tasks in one day!
 - 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?
 - Look in the textbook for a discussion (Page 28)

Relative Performance

- Define Performance = 1/Execution Time
- Performance_x > Performance_y
- 1/Execution Time x > 1/Execution Time y
- Execution Time y > Execution Time x

Performance_x/Performance_y = Execution time_y/Execution time_x = n

Relative Performance

- Define Performance = I/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

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
 - Different programs are affected differently by CPU and system performance

CPU Clocking

- Operation of digital hardware governed by a constant-rate clock
 - Clock period: duration of a clock cycle
 - e.g., 250ps = 0.25ns = 250×10⁻¹²s
 - Clock frequency (rate): cycles per second
 - e.g., 4.0 GHz = 4000 MHz = 4.0×10⁹Hz

CPU Time

CPU Time = CPU Clock Cycles × Clock Cycle Time

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

Clock Rate

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

- 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 compared to A
- How fast must Computer B clock be i.e., what is the clock rate for Computer B?

$$CPU Time_{B} = \frac{CPU Clock Cycles_{B}}{Clock Rate_{B}}$$

- 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 compared to A
- How fast must Computer B clock be?

$$CPU Time_{B} = \frac{CPU Clock Cycles_{B}}{Clock Rate_{B}}$$

$$\operatorname{Clock} \operatorname{Rate}_{B} = \frac{\operatorname{Clock} \operatorname{Cycles}_{B}}{\operatorname{CPU} \operatorname{Time}_{B}} = \frac{1.2 \times \operatorname{Clock} \operatorname{Cycles}_{A}}{6s}$$

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

$$CPU Time_{A} = \frac{CPU Clock Cycles_{A}}{Clock Rate_{A}}$$

 $Clock Cycles_{A} = CPU Time_{A} \times Clock Rate_{A}$ $= 10s \times 2GHz = 20 \times 10^{9}$

- 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?

$$CPU Time_{B} = \frac{CPU Clock Cycles_{B}}{Clock Rate_{B}}$$

$$Clock Rate_{B} = \frac{Clock Cycles_{B}}{CPU Time_{B}} = \frac{1.2 \times Clock Cycles_{A}}{6s}$$

$$Clock Rate_{B} = \frac{1.2 \times 20 \times 10^{9}}{6s} = \frac{24 \times 10^{9}}{6s} = 4GHz$$

Instruction Count and CPI

Clock Cycles = Instruction Count × Cycles per Instruction

CPU Time = Instruction Count × CPI × Clock Cycle Time

Instruction Count × CPI

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

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?

 $CPUTime_{A} = Instruction Count \times CPI_{A} \times Cycle Time_{A}$

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?

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