

ECE-3056: Architecture, Concurrency, and Energy in Computation

Summary

This course introduces the basic organizational principles of the major components of a processor – the core, memory hierarchy, and the I/O subsystem. Students gain an understanding of the sources of performance and energy dissipation in modern processors and learn the multiple forms and levels of parallelism that have been employed to sustain performance improvements in the industry. Assignments using architecture-level simulators enable the students to explore the operation and tradeoffs in performance and energy and reinforce the concepts learned in the class room.

Prerequisites: ECE 2035, ECE 2031

Textbook(s):

Patterson & Hennessey, *Computer Organization & Design: The Hardware/Software Interface* (4th edition), Morgan Kaufmann, 2008. (required)

Course Objectives - As part of this course, students:

1. Understand the organizational principles of modern computer architecture related to cores, memory, and input/output subsystems
2. Acquire an in depth understanding of the concept of an instruction set architecture and its implications for software and hardware implementations
3. Understand fundamentals of concurrent execution and requirements for hardware and software implementation
4. Understand the sources, models, and impact of energy dissipation in modern computing systems
5. Learn common techniques for improving metrics such as performance and energy
6. Are exposed to case studies
7. Gain experience in simulating processor architectures and memory systems and learn to use such simulators to understand the operation of hardware, software, and their interactions

Course Outcomes - Upon successful completion of this course, students should be able to have the following capabilities.

1. Evaluate the cycles per instruction (CPI) for multi-cycle and pipelined data paths.
2. Schedule machine-instruction-level programs on a pipelined datapath with and without hazard handling.
3. Define different types of data and control dependences and describe multiple methods of handling them.
4. Understand the hardware organization of cache memories of different sizes and associativity.
5. Understand the behavior of a cache
6. Define and understand virtual memory, including the concepts of paging, TLB, and page table.
7. Define and distinguish between basic I/O mechanisms such as polling, DMA, interrupts.
8. Understand the operation of storage devices

9. Define and distinguish between various forms of parallelism: instruction level parallelism (ILP), thread level parallelism (TLP), and data level parallelism (DLP)
10. Perform a performance/energy analysis of data paths and the memory system

Topical Coverage

1. Architecture
 - Instruction set architectures
 - Multicycle and pipelined datapaths
 - Branch Prediction and concepts of speculation
 - Caches, Virtual Memory and the memory hierarchy
2. Concurrency
 - Instruction Level Parallelism (ILP), Data Level Parallelism (DLP), Thread Level Parallelism (TLP) basics and introduction to Multi-Core
3. Exceptions, Interrupts, and I/O
4. Energy & Power:
 - Basics of energy dissipation: dynamic and static
 - Microarchitecture power models

Topical Outline:

1. Instruction Set Architectures
 - a. Instructions, addressing modes, and sample ISAs
 - b. Multi-cycle data path and control
 - c. Controller implementation: state machine vs. microprogramming
2. Pipelining
 - a. Pipelining basics
 - b. Pipeline stages: fetch, decode, execute, memory write-back
 - c. Hazards and solutions
 - d. Branch prediction and delayed branches
 - e. Case Studies
3. Memory Systems
 - a. Basic organization of caches and main memory
 - b. Virtual memory basics, memory management
4. Concurrency
 - a. Evolution to multicore
 - b. Introduction to synchronization primitives and the concept of data coherence
 - c. Basics of message passing communication
5. Parallelism
 - a. ILP, DLP, TLP
 - b. Basic architectural support mechanisms
6. I/O Architectures
 - a. Buses and interconnects
 - b. Interrupts, DMA, polling
 - c. Disk structures, I/O scheduling
 - d. LANs, network interfaces, & basic interprocessor communication
 - e. Case Studies
7. Energy and Power dissipation
 - a. Dynamic and static energy dissipation fundamentals
 - b. Microarchitecture-level energy dissipation and power models
 - c. Power virus, kernel benchmarks and power

- d. Basics of voltage and frequency scaling
- e. Case studies

Laboratory Assignments

Laboratory assignments will use prepackaged architecture simulators and hardware description languages (HDLs). Three categories of assignments are conducted: i) datapath, ii) memory hierarchy, iii) power/energy