

EEL-4713C Computer Architecture Lecture 1

Ann Gordon-Ross
Benton 319

EEL-4713C – Ann Gordon-Ross

Administrative matters

- Instructor:** Ann Gordon-Ross (Dr. Ann)
Benton 319; Office hours: By appointment
<http://www.ann.ece.ufl.edu>; ann@ece.ufl.edu
- TA: Shaon Yousuf <yousuf@hcs.ufl.edu>; Office hours: TBD
- Web Page:** Sakai and all files at
http://www.ann.ece.ufl.edu/courses/eel4713_13fal/
- Email:** Start subject with [EEL 4713] (don't send email via Sakai)
- Course files:** On Sakai and
http://www.ann.ece.ufl.edu/courses/eel4713_13fal/
- Schedule:** Pay special attention to the course schedule, linked off Sakai and http://www.ann.ece.ufl.edu/courses/eel4713_13fal/
- Text:** Computer Organization & Design
The Hardware / Software Interface (Revised 4th Edition – Green version)
by Patterson and Hennessy, Morgan Kaufman Publishers

EEL-4713C – Ann Gordon-Ross

Overview

- **Computer architecture is an exciting field**
 - Computer architects are always on the cutting edge
 - Designing several future generations of processors now
- **Exciting time to be in computer architecture!**
 - Paradigm shift from single-core to multi-core
 - But this class focuses on single-core
 - Multi-core architecture is just a collection of single cores, so must know single-core architecture first.
- **Computer architects have a different design philosophy as compared to software designers**

EEL-4713C – Ann Gordon-Ross

What is this class about?

- **Computer Architecture:**
 - **Instruction sets:** how are microprocessors programmed?
 - **Organization:** how does data flow in the microprocessor?
 - **Hardware design:** how are logic components implemented?

EEL-4713C – Ann Gordon-Ross

What is this class about?

- Computer Architecture:
 - **Instruction sets:** how are microprocessors programmed?
 - **Hardware/software interface:** How are instruction sets designed? How does it impact the design of microprocessors and the software running on them?
 - Example: Apple's move from PowerPC to "x86" (Intel)
 - Enabled greater choice in terms of processor configurations
 - Software migration was a major issue; addressed with "binary translation" software (Rosetta)

EEL-4713C – Ann Gordon-Ross

What is this class about?

- Computer Architecture:
 - **Instruction sets:** how are microprocessors programmed?
 - **Organization:** how does data flow in the microprocessor?
 - Instruction set defines the behavior for each and every instruction supported by a microprocessor; there are multiple organizations that can satisfy the functional behavior, and tradeoffs involved
 - How are the major components of the data path organized and controlled?
 - Example: Intel Pentium 4 vs. Core Duo
 - Additional CPU "core", plus changes in the pipeline design
 - "Wider" instruction issue (4 vs. 3), shorter pipeline
 - "Conroe is nothing like any previous Pentium 4 products. In fact, it's based on the mobile Core Duo design which is in itself based on Pentium M, which is based on the Pentium 3 architecture. So Intel has actually done a bit of a U-turn." (trustedreviews.com)

EEL-4713C – Ann Gordon-Ross

What is this class about?

- Computer Architecture:
 - **Instruction sets:** how are microprocessors programmed?
 - **Organization:** how does data flow in the microprocessor?
 - **Hardware design:** how are logic components implemented?
 - CMOS, transistor size scaling; power/performance tradeoffs
 - "The Core-based Intel Xeon is so power efficient, that Apple engineers were able to remove the liquid cooling system from the previous Power-PC based model" (apple.com)

EEL-4713C – Ann Gordon-Ross

What is this class about?

- Computer Architecture:
 - **Instruction sets:** how are microprocessors programmed?
 - **Organization:** how does data flow in the microprocessor?
 - **Hardware design:** how are logic components implemented?
- The process of designing complex digital logic systems
 - Based on knowledge of instruction sets and organization covered in class, you will design a micro-processor using VHDL

EEL-4713C – Ann Gordon-Ross

What should you expect to achieve in this class?

- **In-depth understanding of the inner-workings of modern computers, their evolution, and trade-offs present at the hardware/software boundary.**
 - Insight into fast/slow operations that are easy/hard to implement in hardware
 - Tradeoffs between these designs
 - Computer architecture design process
- **Hands-on experience with the design process in the context of a large, complex hardware system**
- From functional specification to control and datapath implementation and simulation
- Using modern CAD tools and methodologies (VHDL)

EEL-4713C – Ann Gordon-Ross

Course Structure

- **Lecture topics, order may change:**
 - Introduction and ISA/MIPS (Chapters 1 and 2)
 - Basic RISC datapath/control design
 - Pipelined processor design
 - Number systems and performance evaluation
 - Memory systems
 - Input/output
 - Parallelism and other advanced topics, time permitting
 - 4-5 extended lab period lectures or special topics
- **Slides and reading assignments posted on Sakai or off of course files repository linked off my webpage**
 - **Acknowledgement:**
 - The slides used in class, unless otherwise noted, are adapted from David Patterson's lecture slides

EEL-4713C – Ann Gordon-Ross

Course structure

- **Class syllabus:**
 - Also refer to policies document for information on academic honesty and late assignments
- **Book to be used as supplement for lectures**
 - When a topic is covered in class, not all details will be presented.
 - I expect you to read on your own to learn those details
- **Additional reading materials**
- **Key ingredient to success:**
 - Read material **before** lecture
- **Grading:**
 - Lab assignments – 55%
 - Homework questions from book – 10%
 - Exams (two midterms, second one is not cumulative) – 35%
 - Midterm 1 date tentative, Midterm 2 date fixed

EEL-4713C – Ann Gordon-Ross

Lab Assignments/Homework Questions

- **No late assignments/homework will be accepted, no matter what**
- **Homeworks and labs will essentially alternate**
- **Demo assignments in lab, turn in report via Sakai**
 - **Two sections:**
 - Setup section: Get started with tools used
 - Lab section: Hands-on design experience
- **Homework questions**
 - Helps you keep up with material for exams, reinforces concepts
 - You must use the **5th edition, the white one with the orange spine**
- **Dos and Donts**
 - While studying together in groups is encouraged to foster discussion and learning, all work submitted *must* be your own
 - **Not your neighbors, partners, past years' students, from the web, etc. not even with citation**
 - Plagiarism will result in an F in the course!

EEL-4713C – Ann Gordon-Ross

Lab Assignments

- **Lab assignments are a major component of this class**
 - Goal: expose you to the process of designing a microprocessor
 - Labs will upon each other
 - Challenging but rewarding
- **Throughout this class you will design a MIPS microprocessor:**
 - To the extent that it can be simulated within a VHDL-based hardware development framework
 - Starting with the major components of a MIPS datapath
 - Integrate the components and control logic into a processor implementing a subset of MIPS
- **Your tools:**
 - VHDL and Altera Quartus II
 - Proficiency with these is key to success

EEL-4713C – Ann Gordon-Ross

Next lectures

- **Homework #1 is posted, due next week**
- **All lab assignments and homeworks are available**
- **Reading for the next few lectures: chapters 1 and 2**
- **Computer Abstractions and Technology**
 - Textbook, chapter 1
- **Instruction set architectures**
 - Textbook, Chapter 2
 - Sections 2.1-2.8, 2.10, 2.12-2.13, 2.18-2.20

EEL-4713C – Ann Gordon-Ross

Internet companions

- **EEL-4713 Web site - Sakai:**
 - Lecture slides
 - Assignments
 - Announcements
 - Software documentation, tutorials
 - Discussion forum
 - Course schedule
 - All course files are linked off of my webpage, Sakai may simply refer you to that directory at times

EEL-4713C – Ann Gordon-Ross

What is “Computer Architecture”

Computer Architecture =

Instruction Set Architecture (ISA) + Machine Organization

Classic computer organization:

John von Neumann

Stored program computer

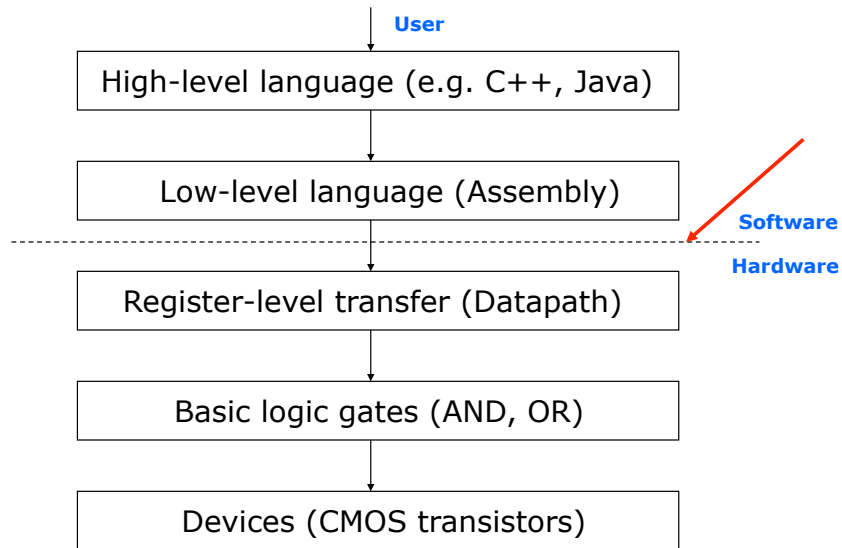
Read instruction and data from memory; decode and execute; write results back to memory

Five key components:

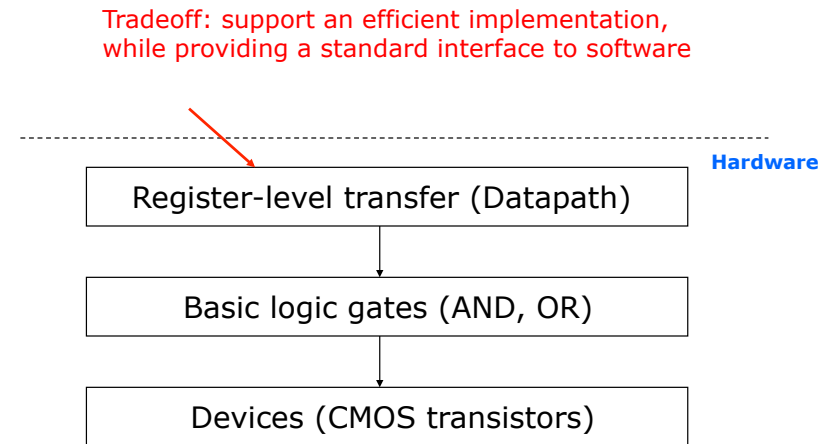
Input, Output, Memory, Datapath and Control

EEL-4713C – Ann Gordon-Ross

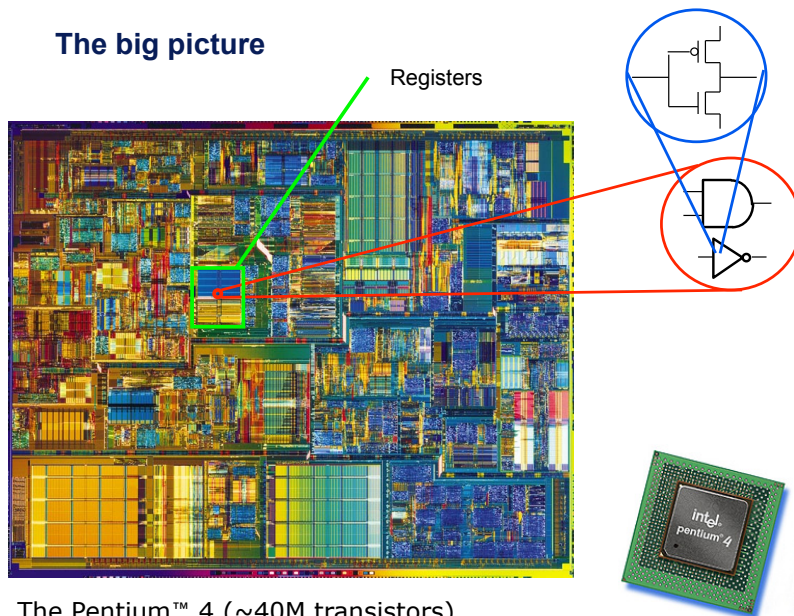
Abstraction layers



Hardware organization

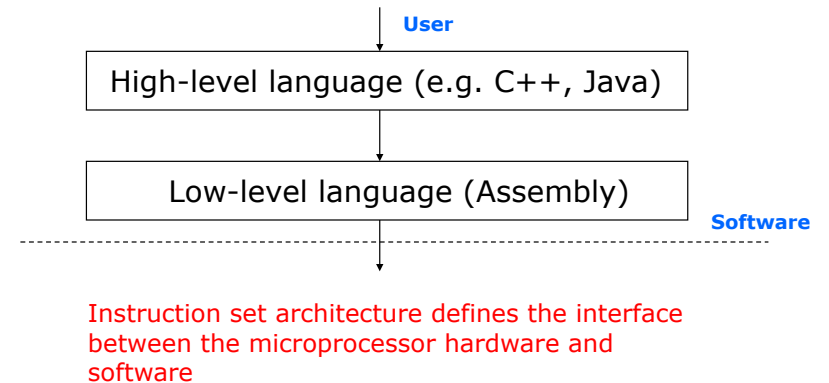


The big picture




The Pentium™ 4 (~40M transistors)

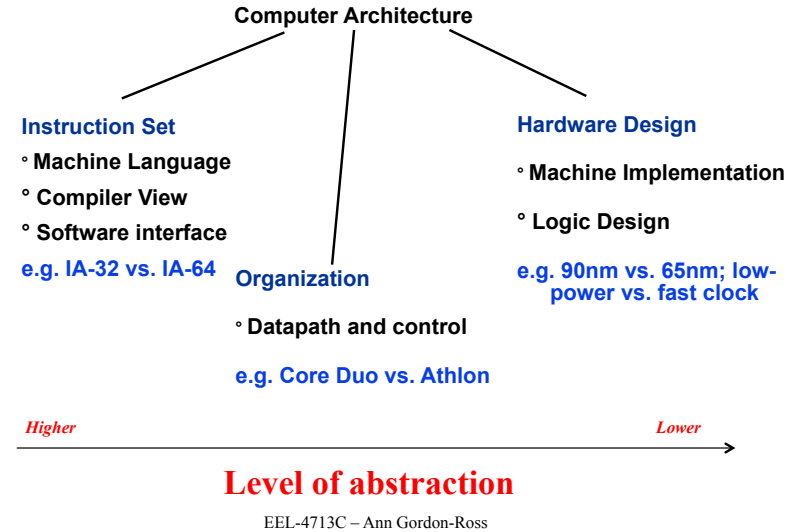
Software interface



The big picture (2)



Course Overview

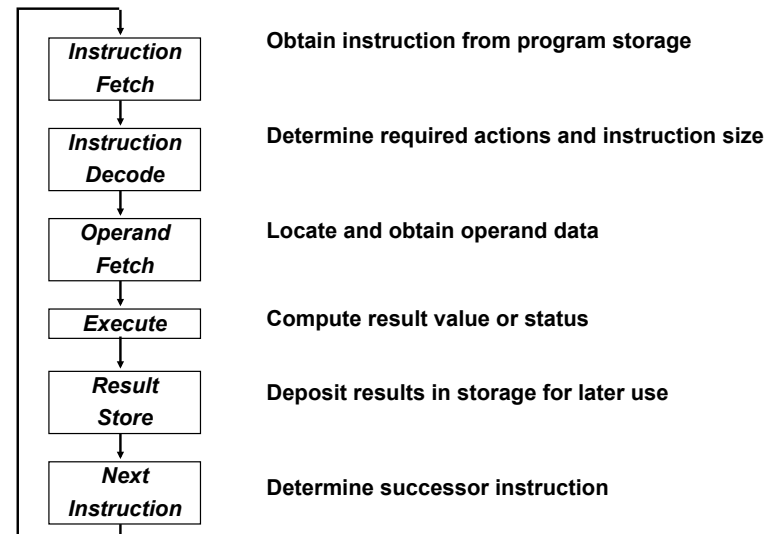


Topics addressed in this course

- How are programs written in a high-level language translated into the hardware language?
- What is the interface between the software and the hardware? What are the design criteria used in defining it?
- What determines the performance of a program? How can a programmer improve performance?
- What is the design process starting from the definition of a microprocessor's behavior and finishing with a functional implementation?
- What are techniques that a microprocessor designer can employ to improve performance while maintaining software compatibility?
- Focus on the architecture and organization aspects

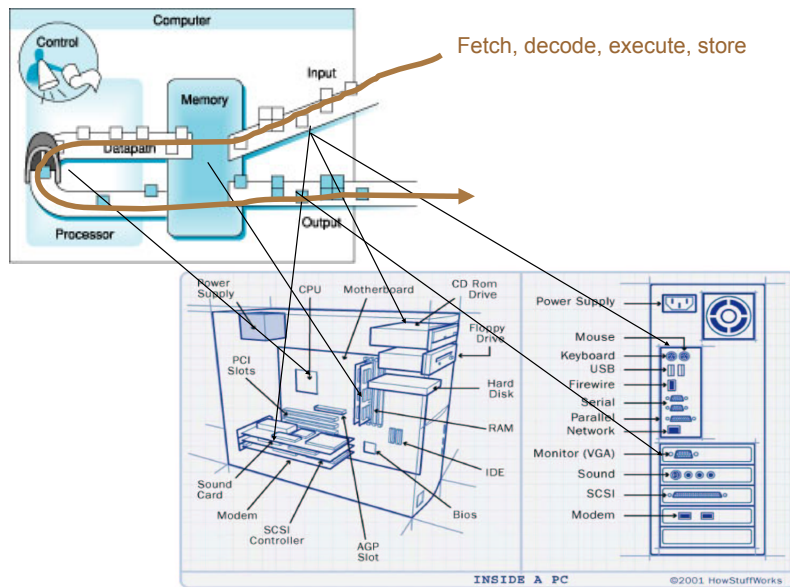
EEL-4713C – Ann Gordon-Ross

Execution cycle (control)



EEL-4713C – Ann Gordon-Ross

Five “classic” components of a computer organization



Understanding program performance




- **Algorithms and data structures**
 - Time/space complexity – e.g. naïve/bubble sort $O(n^2)$ vs. quick sort $O(n \log n)$ determines number of source-level statements executed
 - Not covered in this class
- **Programming language, compiler, architecture**
 - Determines number of machine-level instructions for each source-level statement
- **Processor and memory system**
 - Determines how fast instructions go through a fetch/execute/store cycle
- **I/O subsystem (hardware and software)**
 - How fast instructions which read from/write to I/O devices are executed

EEL-4713C – Ann Gordon-Ross


Before and during a program execution

- **Before** - Applications written in high-level language (e.g. C++) need to be translated to the machine language microprocessors recognize before they execute
 - Compilers
- **During** - At runtime, applications use services from an operating system to facilitate interaction with the hardware and sharing by multiple entities
 - E.g. Linux, Mac OS, Windows
 - Basic I/O operations on files, network sockets, ...
 - Memory allocation
 - Scheduling of CPU cycles across multiple processes

Application classes and characteristics


	Price of system	Price of micro-processor module	Critical system design issues
Desktop 	\$500-\$5,000	\$50-\$500	<ul style="list-style-type: none"> • Tradeoff price/performance • High graphics performance
Server 	\$5,000-\$5,000,000	\$200-\$10,000	<ul style="list-style-type: none"> • High throughput • High availability/dependability • High scalability
Embedded 	Free-\$100,000	\$0.01-\$100	<ul style="list-style-type: none"> • Low price • Low power consumption • Application-specific performance

Microprocessor markets



EEL-4713C – Ann Gordon-Ross


Microprocessor market



* No TV data available prior to 2004

EEL-4713C – Ann Gordon-Ross

Course Overview



EEL-4713C – Ann Gordon-Ross

Instruction Set Architecture

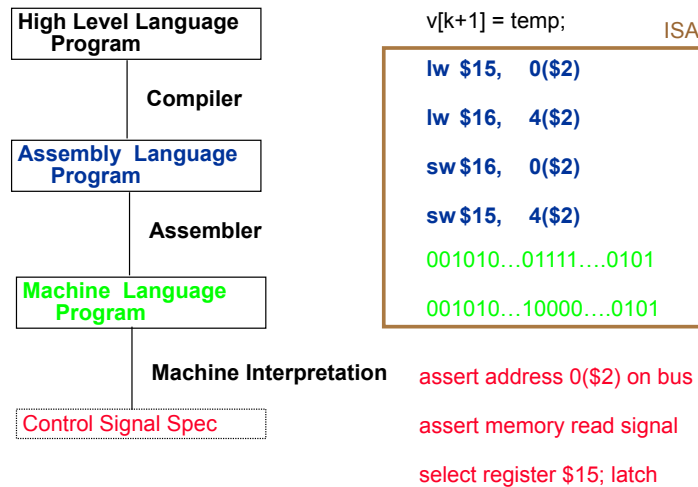
... the attributes of a [computing] system as seen by the programmer, i.e. the conceptual structure and functional behavior, as distinct from the organization of the data flows and controls of the logic design, and the physical implementation.

Amdahl, Blaaw, and Brooks, 1964

- Organization of programmable storage
- Data types & data structures: encodings & representations
- Instruction formats
- Instruction (or operation code) set
- Modes of addressing and accessing data items and instructions
- Exceptional conditions

EEL-4713C – Ann Gordon-Ross

Levels of Representation




EEL-4713C – Ann Gordon-Ross

Example Desktop/server Instruction Set Architectures

Same ISA	Different Hardware Implementations
° Digital Alpha	(v1, v3)
° HP PA-RISC	(v1.1, v2.0)
° Sun Sparc	(v8, v9)
° SGI MIPS	(MIPS I, II, III, IV, V)
° “x86” (IA-32)	(Intel 8086, 80286, 80386, 80486, Pentium, MMX, AMD Athlon,...)
° HP/Intel EPIC/IA-64	(Itanium)

EEL-4713C – Ann Gordon-Ross

Microprocessor sales by ISA



32- and 64-bit
ARM: 80% sales for cell phones
Other: application-specific or customized architectures

EEL-4713C – Ann Gordon-Ross

Example Instruction Set Architecture (ISA): MIPS R3000

Instruction Categories

- Load/Store
- Integer computation
- Jump and Branch
- Floating Point
- Memory Management
- System

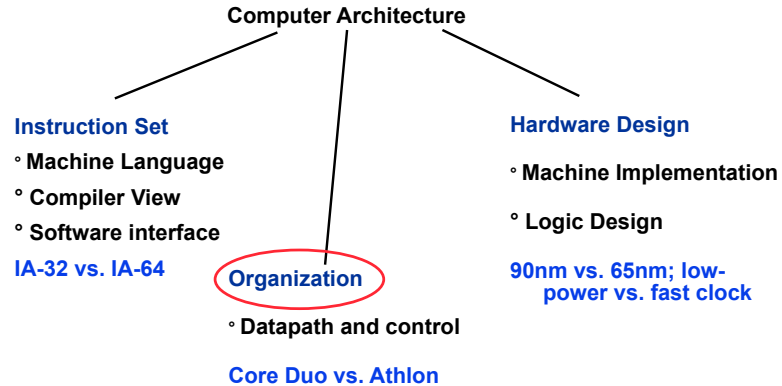
R0 - R31 Special range designations
PC
HI
LO

Instruction Format

OP	rs	rt	rd	shamt	funct
OP	rs	rt	immediate		
OP	target				

EEL-4713C – Ann Gordon-Ross

Course Overview



EEL-4713C – Ann Gordon-Ross

Organization

Logic Designer's View


- capabilities & performance characteristics of principal functional units
(e.g., registers, ALU, shifters, etc.)
- ways in which these components are interconnected
- nature of information flows between components
- logic and means by which such information flow is controlled.

Choreography of units to realize the ISA


Register Transfer Level description

EEL-4713C – Ann Gordon-Ross

Example: Pentium III die




Example: Pentium III pipeline overview



© 2003 Elsevier Science (USA). All rights reserved.

EEL-4713C – Ann Gordon-Ross

Course Overview



EEL-4713C – Ann Gordon-Ross

Hardware design and implementation

- Impact performance, cost, and power consumption of architectures
- So far we have enjoyed exponential improvements over time in:
 - Microprocessor performance
 - Main memory capacity
 - Secondary storage capacity
- “Moore’s Law”
 - Not an actual physical law; observation of a technology trend
 - Microprocessor capacity doubles roughly every 18-24 months


EEL-4713C – Ann Gordon-Ross

Technology => dramatic change

- Processor
 - logic capacity: about 30% per year
 - clock rate: about 20% per year
- Memory
 - DRAM capacity: about 60% per year (4x every 3 years)
 - Memory speed: about 10% per year
 - Cost per bit: reduced by about 25% per year
- Disk
 - capacity: about 60% per year

EEL-4713C – Ann Gordon-Ross

DRAM capacity



EEL-4713C – Ann Gordon-Ross


Microprocessor performance

- Improvements also exponential
- Key technology driver: *device scaling*
- As transistors get smaller (e.g. 180nm to 90nm to 65nm feature sizes)
 - They tend to also get faster and consume less power
 - Faster clock rates
 - More transistors can be packed in the same area
 - Superscalar pipelines; multiple cores; larger caches
- Problems faced by scaling at current (nanoscale) technologies:
 - Fast transistors, but slow interconnect
 - Transient errors
 - Low power per device, but billions of them packed together

EEL-4713C – Ann Gordon-Ross


The power wall

- Dynamic power = capacitive load * Voltage² * Frequency
 - Load: function of transistor, wire technologies, fan-in/out
 - As frequency increases, voltage had to be dropped to maintain power at check => 5V down to 1V
 - At very low voltages, leakage and static power consumption become problems, approximately 40%
 - A “wall” blocking frequency scaling



EEL-4713C – Ann Gordon-Ross

Uniprocessor Performance



Constrained by power, instruction-level parallelism, memory latency

EEL-4713C – Ann Gordon-Ross

From uniprocessors to multiprocessors

- Clock frequency scaling limited
- Can get better performance by exploiting parallelism – multiple operations per cycle
- Instruction-level (superscalars): diminishing returns circa 2004
- Process/thread-level parallelism: multi-core processors

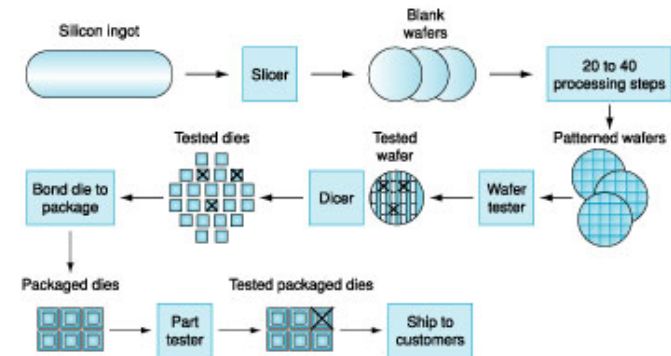
EEL-4713C – Ann Gordon-Ross

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

EEL-4713C – Ann Gordon-Ross

CMOS technology process



EEL-4713C – Ann Gordon-Ross

Next lectures

- **Sign up for the Google group, check for assignment #1**
- **Reading for the next few lectures: chapters 1 and 2**
- **Computer Abstractions and Technology**
 - Textbook, chapter 1
- **Instruction set architectures**
 - Textbook, Chapter 2
 - Sections 2.1-2.8, 2.10, 2.12-2.13, 2.18-2.20

EEL-4713C – Ann Gordon-Ross