





CS 211: Computer Architecture


Instructor: Prof. Bhagi Narahari
 Dept. of Computer Science
 Course URL: www.seas.gwu.edu/~narahari/cs211/



Summary: Architecture Trends ?

- **Moore's law: density doubles every 18-24 months**
 - smaller processors, faster clocks
 - Price drops due to volume and dev. costs what next?
- **Interconnect delays could dominate over feature delay**
 - Need for simpler architectures
 - Distributed logic and control
- **More functionality**
 - communicating processors
 - network of embedded processors
- **To extract max performance**
 - Thumb rules: Amdahl's law, Parallelism, Locality
 - Software and compiler support needed!!!


CS 211: Bhagi Narahari, CS, GWU



Next: Review Computer Organization in an hour!

- **Overview of Computer Organization**
 - Components
 - Sample processor design process


CS 211: Bhagi Narahari, CS, GWU



Review: Computer Organization Basics

- What are the components of a CPU
- What is the microarchitecture level ?
- What is an ISA - Instruction set architecture ?
- How does a sample processor design look ?
 - A simple processor architecture
- what is the basic concept of pipelining


CS 211: Bhagi Narahari, CS, GWU



Memory Unit

- An ordered sequence of storage cells, each capable of holding a piece of data.
- **Volatile Memory**
 - RAM – Random Access Memory
- **Non-volatile Memory**
 - ROM – Read Only Memory

CS 211: Bhagi Narahari, CS, GWU



Memory Hierarchy: The Tradeoff

	register reference	L1-cache reference	L2-cache reference	memory reference	disk memory reference
size:	608 B	128k B	512kB – 4MB	128 MB	27GB
speed:	1.4 ns	4.2 ns	16.8 ns	112 ns	9 ms
\$/Mbyte:			\$90/MB	\$2-6/MB	\$0.01/MB
block size:	4 B	4 B	16 B	4-8 KB	

larger, slower, cheaper →

(Numbers are for a 21264 at 700MHz)

CS 211: Bhagi Narahari, CS, GWU



Central Processing Unit (CPU)

- The CPU has two components:
 - **Arithmetic and Logic Unit (ALU)**
 - Performs arithmetic operations
 - Performs logical operations
 - **Control Unit**
 - Controls the action of the other computer components so that instructions are executed in the correct sequence
- **Note: we will get back to Memory design after covering processor design**

CS 211: Bhagi Narahari, CS, GWU



Input/Output (I/O) Devices

- I/O devices are the components of a computer system that accepts data to be processed and presents the results of the processing
- **Input Device Examples**
 - **Keyboard**
 - **Mouse**
- **Output Device Examples**
 - **Video display**
 - **Printer**
- **We won't touch on this in this course!**

CS 211: Bhagi Narahari, CS, GWU



The Operating System

- The Operating System (OS) is a set of programs that manages all of the computer's resources
- Unix, Linux, Windows 98, Me, NT, 2000, XP, and MacOS are all examples of modern operating systems
- **Not the focus of this course!**

CS 211: Bhagi Narahari, CS, GWU



Computer Architecture



- The computer architecture encompasses the user's view of the computer.
- This includes such things as the assembly language instruction set, the number and types of internal registers, the memory management system and the model for exception handling.

CS 211: Bhagi Narahari, CS, GWU

Architecture Models: Von Neumann architecture

- Memory holds data, instructions.
- Central processing unit (CPU) fetches instructions from memory.
 - Separate CPU and memory distinguishes programmable computer.
- CPU registers help out: program counter (PC), instruction register (IR), general-purpose registers, etc.

CS 211: Bhagi Narahari, CS, GWU



von Neumann vs. Harvard

- Harvard can't use self-modifying code.
- Harvard allows two simultaneous memory fetches.
- Most DSPs use Harvard architecture for streaming data:
 - greater memory bandwidth;
 - more predictable bandwidth.

CS 211: Bhagi Narahari, CS, GWU



Instruction Set Architecture

- The Instruction Set Architecture (ISA) describes a set of instructions whose syntactic and semantic characteristics are defined by the underlying computer architecture.

CS 211: Bhagi Narahari, CS, GWU



Programming model

- **Programming model:** registers visible to the programmer.
- Some registers are not visible (IR).

CS 211: Bhagi Narahari, CS, GWU



Multiple implementations

- Successful architectures have several implementations:
 - varying clock speeds;
 - different bus widths;
 - different cache sizes;
 - etc.





CS 211: Bhagi Narahari, CS, GWU



Assembly language

- One-to-one with instructions (more or less).
- Basic features:
 - One instruction per line.
 - Labels provide names for addresses (usually in first column).
 - Instructions often start in later columns.
 - Columns run to end of line.

CS 211: Bhagi Narahari, CS, GWU





Evolution of Instruction Sets

- Major advances in computer architecture are typically associated with landmark instruction set designs
 - Ex: Stack vs GPR (System 360)
- Design decisions must take into account:
 - technology
 - machine organization
 - programming languages
 - compiler technology
 - operating systems
 - applications
- And they in turn influence these

CS 211: Bhagi Narahari, CS, GWU



CISC vs. RISC

- Complex instruction set computer (**CISC**):
 - many addressing modes;
 - many operations.
- Reduced instruction set computer (**RISC**):
 - load/store;
 - pipelined instructions.

CS 211: Bhagi Narahari, CS, GWU



CISC Processors

- Instruction decoding is performed with large microcode ROMs
- Some instructions require more than a single instruction cycle to execute
- Many addressing modes supported
- Register set was designed to support specific functions

CS 211: Bhagi Narahari, CS, GWU



RISC Processors

- Instruction decoding is performed with static (hard-wired) logic for a much faster result
- Instructions are designed to execute in a single instruction cycle
- Data processing instructions operate only on registers. Load and store instructions were designated to access memory
- Register set is large and general purpose (in many cases)

CS 211: Bhagi Narahari, CS, GWU



IA - 32

- 1978: The Intel 8086 is announced (16 bit architecture)
- 1980: The 8087 floating point coprocessor is added
- 1982: The 80286 increases address space to 24 bits, +instructions
- 1985: The 80386 extends to 32 bits, new addressing modes
- 1989-1995: The 80486, Pentium, Pentium Pro add a few instructions (mostly designed for higher performance)
- 1997: 57 new "MMX" instructions are added, Pentium II
- 1999: The Pentium III added another 70 instructions (SSE)
- 2001: Another 144 instructions (SSE2)
- 2003: AMD extends the architecture to increase address space to 64 bits, widens all registers to 64 bits and other changes (AMD64)
- 2004: Intel capitulates and embraces AMD64 (calls it EM64T) and adds more media extensions

- *"This history illustrates the impact of the "golden handcuffs" of compatibility*
- "adding new features as someone might add clothing to a packed bag"*
- "an architecture that is difficult to explain and impossible to love"*

CS 211: Bhagi Narahari, CS, GWU



IA-32 Overview

- **Complexity:**
 - Instructions from 1 to 17 bytes long
 - one operand must act as both a source and destination
 - one operand can come from memory
 - complex addressing modes
e.g., "base or scaled index with 8 or 32 bit displacement"
- **Saving grace:**
 - the most frequently used instructions are not too difficult to build
 - compilers avoid the portions of the architecture that are slow

"what the 80x86 lacks in style is made up in quantity, making it beautiful from the right perspective"

CS 211: Bhagi Narahari, CS, GWU



Quick look at ISA

- Will use MIPS
 - Simple RISC ISA
 - Widely used

CS 211: Bhagi Narahari, CS, GWU



Instruction set characteristics

- Fixed vs. variable length.
- Addressing modes.
- Number of operands.
- Types of operands.

CS 211: Bhagi Narahari, CS, GWU



A "Typical" RISC - MIPS

- **fixed format instruction (3 formats I,R,J):**
 - 32 or 64 bits
- **32 General Purpose Registers (GPR)**
 - (R0 contains zero, DP take pair)
- **3-address, reg-reg arithmetic instruction**
- **Single address mode for load/store:**
base + displacement
 - no indirection
- **Simple branch conditions (based on register values)**

see: SPARC, MIPS, HP PA-Risc, DEC Alpha, IBM PowerPC,
CDC 6600, CDC 7600, Cray-1, Cray-2, Cray-3
- **Delayed branch**

CS 211: Bhagi Narahari, CS, GWU



Example: 32 bit MIPS

R-type: Register-Register

31	26	25	21	20	16	15	11	10	6	5	0
Op	Rs1	Rs2	Rd								

I-type: Register-Immediate (Load,Store)

31	26	25	21	20	16	15	0
Op	Rs1	Rd					

I-type: Branch

31	26	25	21	20	16	15	0
Op	Rs1	Rs2/Opx					

J-type: Jump / Call

31	26	25	0
Op			

CS 211: Bhagi Narahari, CS, GWU



Instruction Types

- **When is memory accessed ?**
 - How is address computed ?
- **When is control flow affected ?**
 - How is branch outcome computed
 - How is branch target address computed

CS 211: Bhagi Narahari, CS, GWU




Processor Design...

- **Let's glance at processor and inst. Set design.**
 - This is review material...from a typical course on Computer Organization (pre-req)

CS 211: Bhagi Narahari, CS, GWU

The Big Picture: The Performance Perspective



- Performance of a machine is determined by:
 - Instruction count
 - Clock cycle time
 - Clock cycles per instruction
- Processor design (datapath and control) will determine:
 - Clock cycle time
 - Clock cycles per instruction

CS 211: Bhagi Narahari, CS, GWU

Microarchitecture Design: How ?

- Any design must attempt to meet the requirements
 - Where do the requirements come from ?
 - Ex: need to represent numbers in binary; integers, text, floating point
- How to proceed with design ?

CS 211: Bhagi Narahari, CS, GWU

Some History...

- The Indiana Legislature once introduced legislation declaring that the value of π was exactly 3.2

CS 211: Bhagi Narahari, CS, GWU


How to Design a Processor: step-by-step

- Analyze instruction set => datapath requirements
 - the meaning of each instruction is given by the *register transfers*
 - datapath must include storage element for ISA registers
 - possibly more
 - datapath must support each register transfer
- Select set of datapath components and establish clocking methodology
- Assemble datapath meeting the requirements
- Analyze implementation of each instruction to determine setting of control points that effects the register transfer.
- Assemble the control logic
- Let's look at a single cycle ISA...

CS 211: Bhagi Narahari, CS, GWU

The MIPS Instruction Formats

- All MIPS instructions are 32 bits long. The three instruction formats:



- The different fields are:
 - op: operation of the instruction
 - rs, rt, rd: the source and destination register specifiers
 - shamt: shift amount
 - funct: selects the variant of the operation in the "op" field
 - address / immediate: address offset or immediate value
 - target address: target address of the jump instruction

CS 211: Bhagi Narahari, CS, GWU

Step 1a: The MIPS-Inst Set (eg.)

- ADD and SUB

- addU rd, rs, rt
- subU rd, rs, rt

- OR Immediate:

- ori rt, rs, imm16

- LOAD and STORE Word

- lw rt, rs, imm16
- sw rt, rs, imm16

- BRANCH:

- beq rs, rt, imm16

- Register rs and rt are the source registers.
- If the instruction has three operand register, then rd is the destination register
- If the instruction has two operand register, then rt is the destination register

CS 211: Bhagi Narahari, CS, GWU

Logical Register Transfers

- RTL gives the meaning of the instructions
- All start by fetching the instruction

op | rs | rt | rd | shamt | funct = MEM[PC]

op | rs | rt | Imm16 = MEM[PC]

Register Transfers

ADDU	R[rd] ← R[rs] + R[rt];	PC ← PC + 4
SUBU	R[rd] ← R[rs] - R[rt];	PC ← PC + 4
ORI	R[rt] ← R[rs] zero_ext(Imm16);	PC ← PC + 4
LOAD	R[rt] ← MEM[R[rs] + sign_ext(Imm16)];	PC ← PC + 4
STORE	MEM[R[rs] + sign_ext(Imm16)] ← R[rt];	PC ← PC + 4
BEQ	if (R[rs] == R[rt]) then PC ← PC + 4 + sign_ext(Imm16) 00	
	else PC ← PC + 4	




CS 211: Bhagi Narahari, CS, GWU

Step 2: Components of the Datapath

- Combinational Elements
- Storage Elements
 - Clocking methodology

CS 211: Bhagi Narahari, CS, GWU

Combinational Logic Elements (Basic Building Blocks)

- **Adder**

- **MUX**

- **ALU**


CS 211: Bhagi Narahari, CS, GWU

Storage Element: Register (Basic Building Block)


Register

- **Similar to the D Flip Flop except**

- N-bit input and output
- Write Enable input

Write Enable:

- negated (0): Data Out will not change
- asserted (1): Data Out will become Data In




CS 211: Bhagi Narahari, CS, GWU

Storage Element: Register File

- Register File consists of 32 registers:
 - Two 32-bit output busses: busA and busB
 - One 32-bit input bus: busW
- Register is selected by:
 - RA (number) selects the register to put on busA (data)
 - RB (number) selects the register to put on busB (data)
 - RW (number) selects the register to be written via busW (data) when Write Enable is 1
- Clock input (CLK)
 - The CLK input is a factor ONLY during write operation
 - During read operation, behaves as a combinational logic block:
 - RA or RB valid => busA or busB valid after "access time."

CS 211: Bhagi Narahari, CS, GWU



Storage Element: Idealized Memory

Memory (idealized)

- One input bus: Data In
- One output bus: Data Out


Memory word is selected by:

- Address selects the word to put on Data Out
- Write Enable = 1: address selects the memory word to be written via the Data In bus


Clock input (CLK)

- The CLK input is a factor ONLY during write operation
- During read operation, behaves as a combinational logic block:
 - Address valid => Data Out valid after "access time."

CS 211: Bhagi Narahari, CS, GWU



Clocking Methodology



- Clocks needed in sequential logic to decide when an element that contains state should be updated.
- A clock is a free-running circuit with a fixed cycle time or clock period. The clock frequency is the inverse of the cycle time.
- The clock cycle time or clock period is divided into two portions: when the clock is high and when the clock is low.
- Edge-triggered clocking: all state changes occur on a clock edge.

CS 211: Bhagi Narahari, CS, GWU

Step 3: Assemble DataPath meeting our requirements

- Register Transfer Requirements
⇒ Datapath Assembly
- Instruction Fetch
- Read Operands and Execute Operation

The common RTL operations for all instructions are:

- Fetch the instruction using the Program Counter (PC) at the beginning of an instruction's execution (PC → Instruction Memory → Instruction Word).
- Then at the end of the instruction's execution, you need to update the Program Counter (PC → Next Address Logic → PC).

More specifically, you need to increment the PC by 4 if you are executing sequential code. For Branch and Jump instructions, you need to update the program counter to "something else" other than plus 4.

The Next Address Logic block:

- Add 4 (number of bytes in an instruction) or
- Branch and Jump instructions

CS 211: Bhagi Narahari, CS, GWU


3a: Overview of the Instruction Fetch Unit

- The common RTL operations

➤ Fetch the Instruction: $\text{mem}[\text{PC}]$

➤ Update the program counter:

- Sequential Code: $\text{PC} \leftarrow \text{PC} + 4$
- Branch and Jump: $\text{PC} \leftarrow \text{"something else"}$




CS 211: Bhagi Narahari, CS, GWU

3b: Add & Subtract

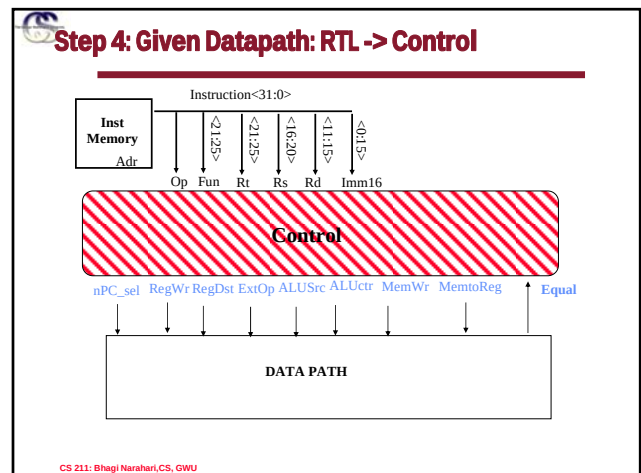



- $R[\text{rd}] \leftarrow R[\text{rs}] \text{ op } R[\text{rt}]$ Example: $\text{addU } \text{rd}, \text{rs}, \text{rt}$

➤ Ra, Rb, and Rw come from instruction's rs, rt, and rd fields

➤ ALUctr and RegWr: control logic after decoding the instruction



CS 211: Bhagi Narahari, CS, GWU



Control Signals

inst	Register Transfer
ADD	$R[rd] \leftarrow R[rs] + R[rt]; \quad PC \leftarrow PC + 4$ $ALUSrc = \text{RegB}, ALUctr = \text{"add"}, \text{RegDst} = rd, \text{RegWr}, nPC_sel = \text{"+4"}$
SUB	$R[rd] \leftarrow R[rs] - R[rt]; \quad PC \leftarrow PC + 4$ $ALUSrc = \text{RegB}, ALUctr = \text{"sub"}, \text{RegDst} = rd, \text{RegWr}, nPC_sel = \text{"+4"}$
ORI	$R[rt] \leftarrow R[rs] + \text{zero_ext}(Imm16); \quad PC \leftarrow PC + 4$ $ALUSrc = \text{Im}, \text{ExtOp} = \text{"Z"}, ALUctr = \text{"or"}, \text{RegDst} = rt, \text{RegWr}, nPC_sel = \text{"+4"}$
LOAD	$R[rt] \leftarrow \text{MEM}[R[rs] + \text{sign_ext}(Imm16)]; \quad PC \leftarrow PC + 4$ $ALUSrc = \text{Im}, \text{ExtOp} = \text{"Sn"}, ALUctr = \text{"add"}, \text{MemtoReg}, \text{RegDst} = rt, \text{RegWr}, nPC_sel = \text{"+4"}$
STORE	$\text{MEM}[R[rs] + \text{sign_ext}(Imm16)] \leftarrow R[rs]; \quad PC \leftarrow PC + 4$ $ALUSrc = \text{Im}, \text{ExtOp} = \text{"Sn"}, ALUctr = \text{"add"}, \text{MemWr}, nPC_sel = \text{"+4"}$
BEQ	$\text{if } (R[rs] == R[rt]) \text{ then } PC \leftarrow PC + \text{sign_ext}(Imm16) \parallel 00 \text{ else } PC \leftarrow PC + 4$ $nPC_sel = \text{EQUAL}, ALUctr = \text{"sub"}$


CS 211: Bhagi Narahari, CS, GWU

Step 5: Logic for each control signal

- $nPC_sel \leftarrow \text{if } (OP == BEQ) \text{ then EQUAL else } 0$
- $ALUSrc \leftarrow \text{if } (OP == \text{"000000"}) \text{ then "regB" else "immed"}$
- $ALUctr \leftarrow \text{if } (OP == \text{"000000"}) \text{ then funct}$
 $\quad \text{elseif } (OP == ORI) \text{ then "OR"}$
 $\quad \text{elseif } (OP == BEQ) \text{ then "sub"}$
 $\quad \text{else "add"}$
- $\text{ExtOp} \leftarrow \text{if } (OP == ORI) \text{ then "zero" else "sign"}$
- $\text{MemWr} \leftarrow (OP == \text{Store})$
- $\text{MemtoReg} \leftarrow (OP == \text{Load})$
- $\text{RegWr} \leftarrow \text{if } ((OP == \text{Store}) \parallel (OP == BEQ)) \text{ then } 0 \text{ else } 1$
- $\text{RegDst} \leftarrow \text{if } ((OP == \text{Load}) \parallel (OP == ORI)) \text{ then } 0 \text{ else } 1$


CS 211: Bhagi Narahari, CS, GWU

Example: Load Instruction



CS 211: Bhagi Narahari, CS, GWU

An Abstract View of the Implementation



- Logical vs. Physical Structure


CS 211: Bhagi Narahari, CS, GWU

Summary

- 5 steps to design a processor
 - 1. Analyze instruction set => datapath requirements
 - 2. Select set of datapath components & establish clock methodology
 - 3. Assemble datapath meeting the requirements
 - 4. Analyze implementation of each instruction to determine setting of control points that effects the register transfer.
 - 5. Assemble the control logic
- MIPS makes it easier
 - Instructions same size
 - Source registers always in same place
 - Immediates same size, location
 - Operations always on registers/immediates
- Single cycle datapath => CPI=1, CCT => long

CS 211: Bhagi Narahari, CS, GWU


Systematic Generation of Control



- In a single-cycle processor, each instruction is realized by exactly one control command or "microinstruction"
 - in general, the controller is a finite state machine
 - microinstruction can also control sequencing (see later)

CS 211: Bhagi Narahari, CS, GWU

What's wrong with our CPI=1 processor?




- Long Cycle Time
- All instructions take as much time as the slowest
- Real memory is not as nice as our idealized memory
 - cannot always get the job done in one (short) cycle

CS 211: Bhagi Narahari, CS, GWU

Partitioning the CPI=1 Datapath

- Add registers between smallest steps



CS 211: Bhagi Narahari, CS, GWU

Example Multicycle Datapath

The diagram illustrates a multicycle datapath with the following components and stages:

- Instruction Fetch:** The PC (Program Counter) is updated from `nextPC_sel` and feeds into the IR (Instruction Register).
- Operand Fetch:** The IR feeds into the Reg File (Register File), which provides operands A and B to the Ext ALU. The Ext ALU also receives an `Equal` signal and an `ExtOp` (extended operation) signal.
- ALU:** The Ext ALU performs operations on A and B, receiving `ALUSrc` and `ALUDetr` signals. It outputs a 5-bit `Op` (operation) signal.
- Mem Access:** The `Op` signal and the ALU result feed into the Mem Access unit. The Mem Access unit also receives `MemRd` (memory read) and `MemWr` (memory write) signals. It outputs a 5-bit `M` (memory access) signal.
- Result Store:** The Mem Access unit feeds into the Reg. File (Register File), which also receives `RegDst` (register destination) and `RegWr` (register write) signals. The Reg. File outputs a 5-bit `Result Store` signal.

The critical path is highlighted in red, showing the sequence of operations that determine the clock cycle time.

- **Critical Path ?**

CS 211: Bhagi Narahari, CS, GWU

Controller Design

- The state diagrams that arise define the controller for an instruction set processor are highly structured
- Use this structure to construct a simple “microsequencer”
- Control reduces to programming this very simple device

⇒ microprogramming

The diagram illustrates a microsequencer architecture. It consists of three main components: a **micro-PC** (Program Counter), a **sequencer**, and a **control unit**. The **control unit** is a large rectangle divided into four sections: **sequencer control** (top-left), **datapath control** (top-right), **microinstruction** (middle), and an empty bottom section. The **micro-PC** is a small rectangle on the left. The **sequencer** is a rounded rectangle at the bottom. Arrows indicate the flow of control: from **micro-PC** to **sequencer control**, from **sequencer control** to **sequencer**, and from **sequencer** back to **micro-PC**. Additionally, there are two downward arrows from the **control unit** (one from the **sequencer control** section and one from the empty bottom section) and one downward arrow from the **sequencer**.



Microprogramming

- Microprogramming is a convenient method for implementing *structured* control state diagrams:
 - Random logic replaced by microPC sequencer and ROM
 - Each line of ROM called a μ instruction: contains sequencer control + values for control points
 - limited state transitions:
 - branch to zero, next sequential,
 - branch to μ instruction address from dispatch ROM
- Horizontal μ Code: one control bit in μ instruction for every control line in datapath
- Vertical μ Code: groups of control-lines coded together in μ instruction (e.g. possible ALU dest)
- Control design reduces to Microprogramming
 - Part of the design process is to develop a "language" that describes control and is easy for humans to understand

CS 211: Bhagi Narahari, CS, GWU

Microprogramming

sequencer: datapath control

microinstruction (μ)

micro-PC

μ-sequencer: fetch, dispatch, sequential

Dispatch ROM

Opcode

Decode

Decode

To DataPath

μ-Code ROM

Decoders implement our μ-code language:

For instance:

rt-ALU

rd-ALU

mem-ALU

- **Microprogramming is a fundamental concept**
 - implement an instruction set by building a very simple processor and interpreting the instructions
 - essential for very complex instructions and when few register transfers are possible
 - overkill when ISA matches datapath 1-1

CS 211: Bhagii Narahari, CS, GWU



Microprogramming one inspiration for RISC

- If simple instruction could execute at very high clock rate...
 - you could even write compilers to produce microinstructions...
- If most programs use simple instructions and addressing modes...
- If microcode is kept in RAM instead of ROM so as to fix bugs ...
- Then why not skip instruction interpretation by a microprogram and simply compile directly into lowest language of machine? (microprogramming is overkill when ISA matches datapath 1-1)

CS 211: Bhagi Narahari, CS, GWU



How to improve performance?

- Recall performance is function of
 - CPI: cycles per instruction
 - Clock cycle
 - Instruction count
- Reducing any of the 3 factors will lead to improved performance

CS 211: Bhagi Narahari, CS, GWU



How to improve performance?


- First step is to apply concept of pipelining to the instruction execution process
 - Overlap computations
- What does this do?
 - Decrease clock cycle
 - Decrease effective CPI compared to original clock cycle

CS 211: Bhagi Narahari, CS, GWU

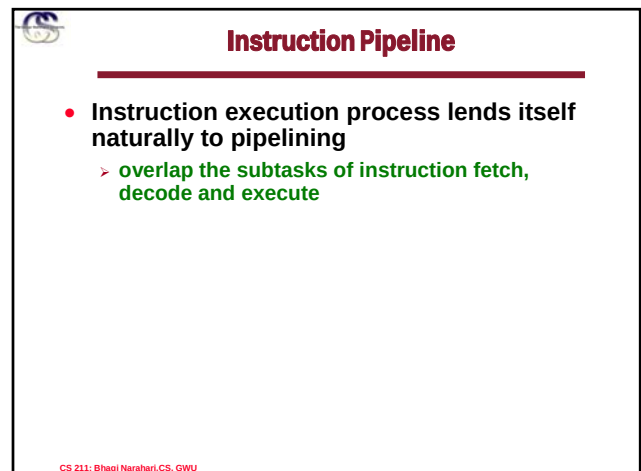


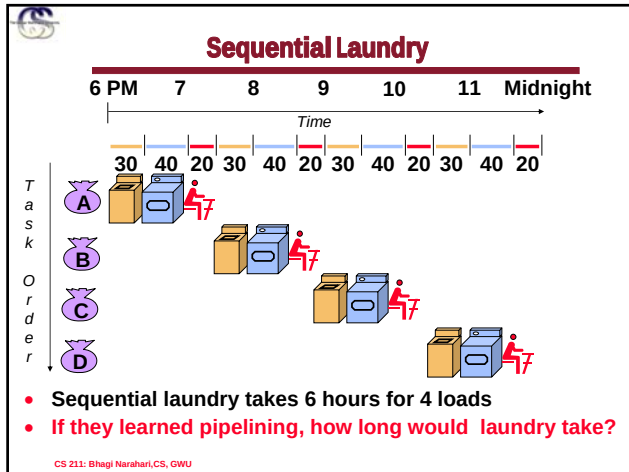


Pipelining: Its Natural!

- Laundry Example
- Ann, Brian, Cathy, Dave each have one load of clothes to wash, dry, and fold
- Washer takes 30 minutes
- Dryer takes 40 minutes
- "Folder" takes 20 minutes



CS 211: Bhagi Narahari, CS, GWU





How to improve performance?

- Recall performance is function of
 - CPI: cycles per instruction
 - Clock cycle
 - Instruction count
- Reducing any of the 3 factors will lead to improved performance

CS 211: Bhagi Narahari, CS, GWU



How to improve performance?

- First step is to apply concept of pipelining to the instruction execution process
 - Overlap computations
- What does this do?
 - Decrease clock cycle
 - Decrease effective CPU time compared to original clock cycle

CS 211: Bhagi Narahari, CS, GWU



Pipeline Approach to Improve System Performance

- Analogous to fluid flow in pipelines and assembly line in factories
- Divide process into “stages” and send tasks into a pipeline
 - Overlap computations of different tasks by operating on them concurrently in different stages

CS 211: Bhagi Narahari, CS, GWU



Instruction Level Parallel Processors (ILP)

- early ILP - one of two orthogonal concepts:
 - pipelining - vertical approach
 - multiple (non-pipelined) units - horizontal approach
- progression to multiple pipelined units
- instruction issue became bottleneck, led to
 - superscalar ILP processors
 - Very Large Instruction Word (VLIW)
- Note: key performance metric in all ILP processor classes is **IPC** (instructions per cycle)
 - this is the degree of parallelism achieved

CS 211: Bhagi Narahari, CS, GWU



Instruction Pipeline

- Instruction execution process lends itself naturally to pipelining
 - overlap the subtasks of instruction fetch, decode and execute