CS 152
Computer Architecture and Engineering
Lecture 11

Multicycle Controller Design

Mar 8, 1999

John Kubiatowicz (http.cs.berkeley.edu/~kubitron)

lecture slides: http://www-inst.eecs.berkeley.edu/~cs152/

Overview of Control

° Control may be designed using one of several initial representations. The choice of sequence control, and how logic is represented, can then be determined independently; the control can then be implemented with one of several methods using a structured logic technique.

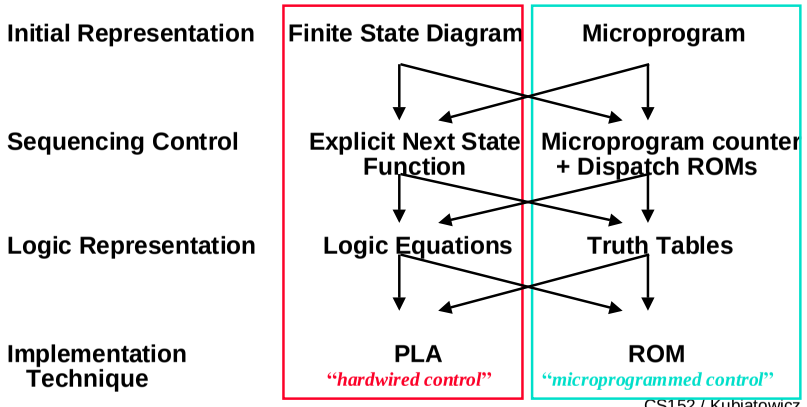

Recap: "Macroinstruction" Interpretation

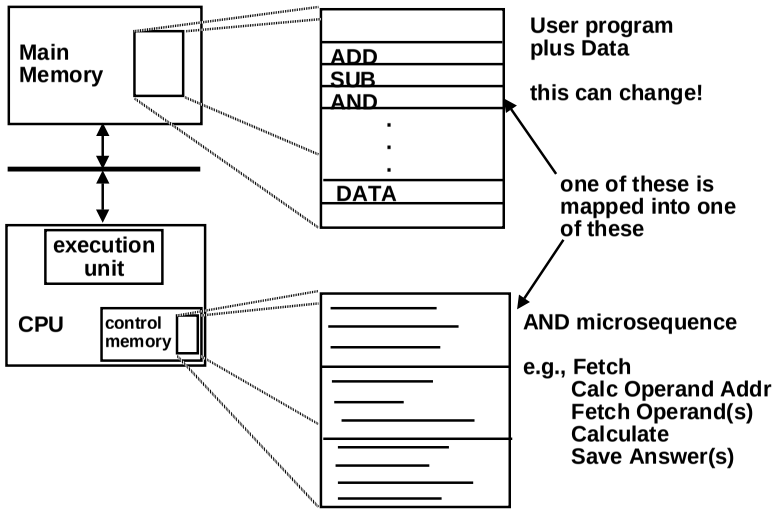

The Big Picture: Where are We Now?

° The Five Classic Components of a Computer

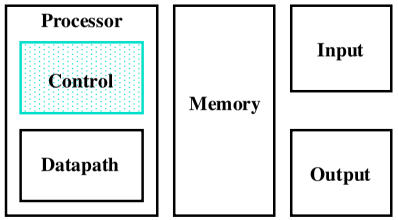

° Today's Topics:

- Microprogramed control
- Administrivia; Courses
- Microprogram it yourself
- Exceptions
- Intro to Pipelining (if time permits)

Recap: Horizontal vs. Vertical Microprogramming

NOTE: previous organization is not TRUE horizontal microprogramming; register decoders give flavor of *encoded* microoperations

Most microprogramming-based controllers vary between:

horizontal organization (1 control bit per control point)

vertical organization (fields encoded in the control memory and must be decoded to control something)

Horizontal

+ more control over the potential parallelism of operations in the datapath
- uses up lots of control store

Vertical

+ easier to program, not very different from programming a RISC machine in assembly language
- extra level of decoding may slow the machine down

Recap: Designing a Microinstruction Set

1) Start with list of control signals

2) Group signals together that make sense (vs. random): called "fields"

3) Places fields in some logical order (e.g., ALU operation & ALU operands first and microinstruction sequencing last)

4) Create a symbolic legend for the microinstruction format, showing name of field values and how they set the control signals
   - Use computers to design computers

5) To minimize the width, encode operations that will never be used at the same time

Alternative datapath (book): Multiple Cycle Datapath

° Miminizes Hardware: 1 memory, 1 adder

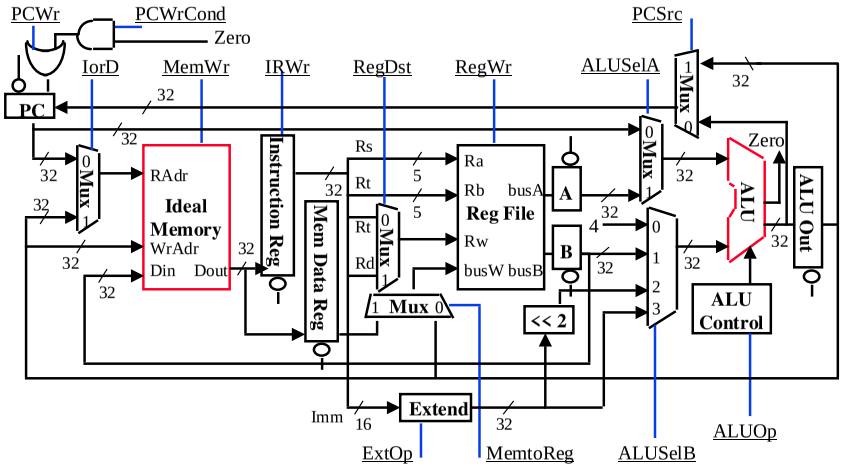

Finite State Machine (FSM) Spec

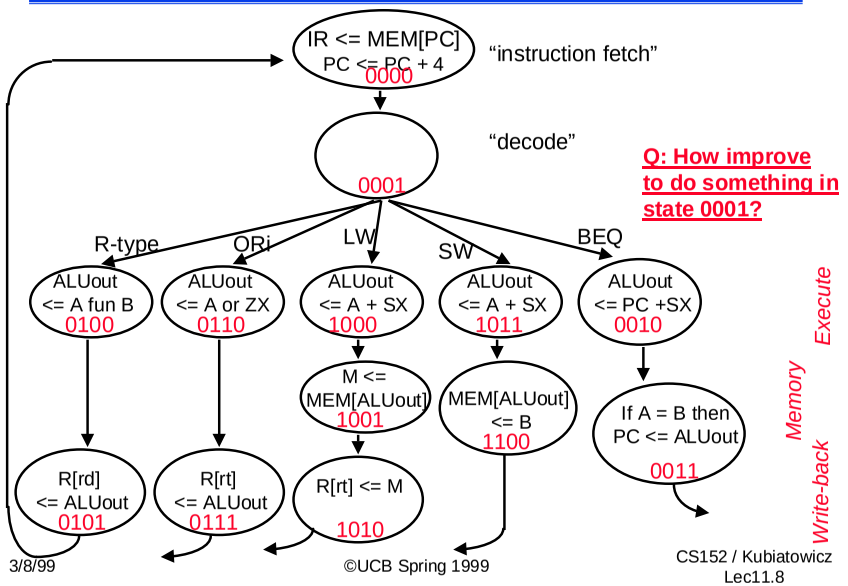

1&2) Start with list of control signals, grouped into fields

Single Bit Control

Signal name	Effect when deasserted	Effect when asserted
ALUSelA	1st ALU operand = PC	1st ALU operand = Reg[rs]
RegWrite	None	Reg. is written
MemtoReg	Reg. write data input = ALU	Reg. write data input = memory
RegDst	Reg. dest. no. = rt	Reg. dest. no. = rd
MemRead	None	Memory at address is read, MDR <= Mem[addr]
MemWrite	None	Memory at address is written
IorD	Memory address = PC	Memory address = S
IRWrite	None	IR <= Memory
PCWrite	None	PC <= PCSource
PCWriteCond	None	IF ALUzero then PC <= PCSource
PCSource	PCSource = ALU	PCSource = ALUout

Multiple Bit Control

Signal name	Value	Effect
ALUOp	00	ALU adds
	01	ALU subtracts
	10	ALU does function code
	11	ALU does logical OR
ALUSelB	000	2nd ALU input = Reg[rt]
	001	2nd ALU input = 4
	010	2nd ALU input = sign extended IR[15-0]
	011	2nd ALU input = sign extended, shift left 2 IR[15-0]
	100	2nd ALU input = zero extended IR[15-0]

Start with list of control signals, cont'd

° For next state function (next microinstruction address), use Sequencer-based control unit from last lecture

- Called "microPC" or "µPC" vs. state register

Signal	Value	Effect
Sequencing	00	Next µaddress = 0
	01	Next µaddress = dispatch ROM
	10	Next µaddress = µaddress + 1

° Could even include "branch" option which changes microPC by adding offset when certain control signals are true.

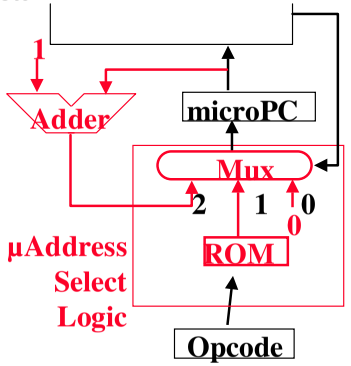

3) Microinstruction Format: unencoded vs. encoded fields

Field Name	Width		Control Signals Set
	wide	narrow	
ALU Control	4	2	ALUOp
SRC1	2	1	ALUSelA
SRC2	5	3	ALUSelB
ALU Destination	3	2	RegWrite, MemtoReg, RegDst
Memory	4	3	MemRead, MemWrite, IorD
Memory Register	1	1	IRWrite
PCWrite Control	4	3	PCWrite, PCWriteCond, PCSource
Sequencing	3	2	AddrCtl
Total width	26	17	bits

4) Legend of Fields and Symbolic Names

Field Name	Values for Field	Function of Field with Specific Value
ALU	Add	ALU adds
	Subt.	ALU subtracts
	Func code	ALU does function code
	Or	ALU does logical OR
SRC1	PC	1st ALU input = PC
	rs	1st ALU input = Reg[rs]
SRC2	4	2nd ALU input = 4
	Extend	2nd ALU input = sign ext. IR[15-0]
	Extend0	2nd ALU input = zero ext. IR[15-0]
	Extshft	2nd ALU input = sign ex., sl IR[15-0]
	rt	2nd ALU input = Reg[rt]
destination	rd ALU	Reg[rd] = ALUout
	rt ALU	Reg[rt] = ALUout
	rt Mem	Reg[rt] = Mem
Memory	Read PC	Read memory using PC
	Read ALU	Read memory using ALU output
	Write ALU	Write memory using ALU output
Memory register	IR	IR = Mem
PC write	ALU	PC = ALU
	ALUoutCond	IF ALU Zero then PC = ALUout
Sequencing	Seq	Go to sequential µinstruction
	Fetch	Go to the first microinstruction
	Dispatch	Dispatch using ROM.

Microprogram it yourself!

Label	ALU	SRC1	SRC2	ALU Dest.	Memory	Mem. Reg.	PC Write	Sequencing
Fetch:	Add	PC	4		Read PC	IR	ALU	Seq

Administrivia

° Enjoyed meeting everyone after midterm
- Beer and pizza was a great way to say hello to everyone
- Lots of people heading for industry.

° Midterm graded, scores posted
- Average score: 70.0
- Std. Dev: 16.8

° Now, start reading Chapter 6

Midterm I distribution

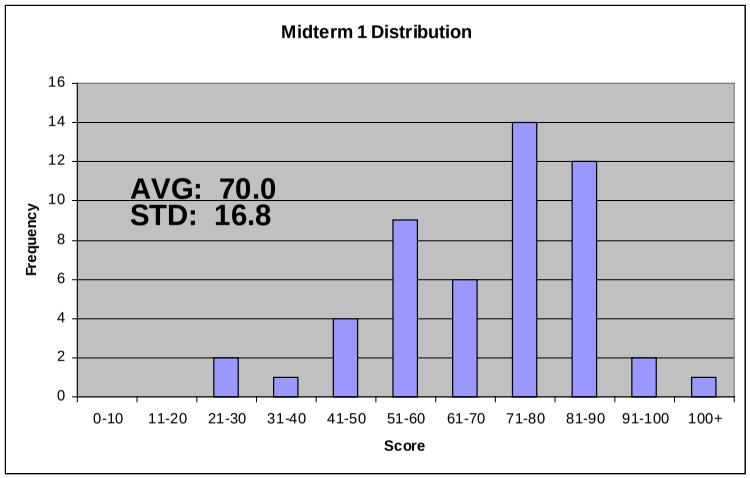

Multiplier

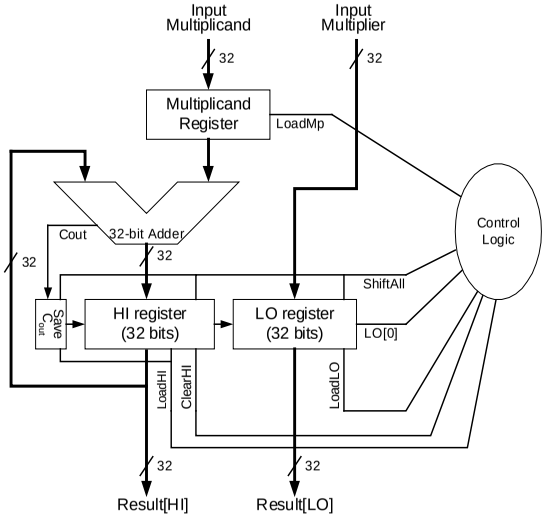

Single Bit Booth Multiplier

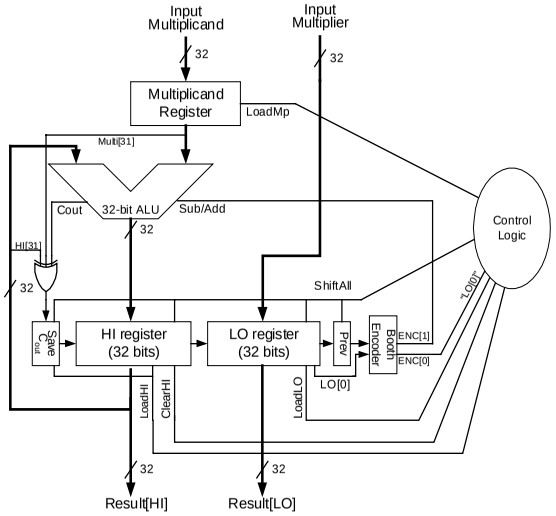

Double Bit Booth Multiplier

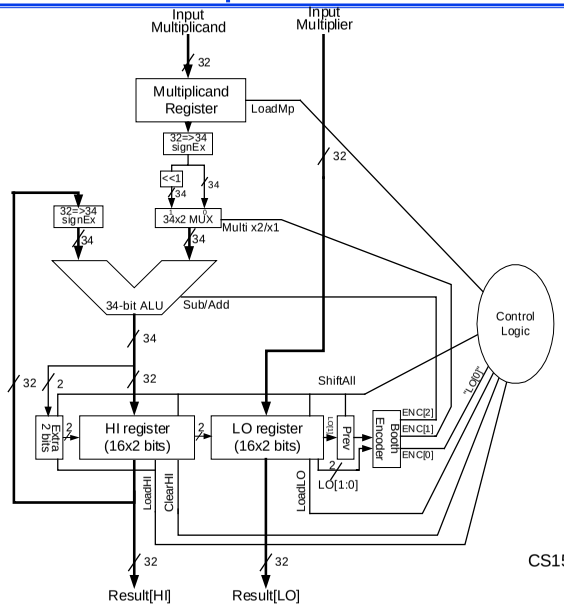

Administrivia: Courses to consider during Telebears

° **General Philosophy**
- **Take courses from great teachers (HKN ratings helps find them)**
  - **http://www-hkn.eecs.berkeley.edu/toplevel/coursesurveys.html**
- **Take variety of undergrad courses now to get introduction to areas; can learn advanced material on own later once know vocabulary**
- **Who knows what you will work on over a 40 year career?**

° **CS169 Software Engineering**
- **Everyone writes programs, even hardware designers**
- **Often programs are written in groups => learn skill in school**

° **EE122 Introduction to Communication Networks**
- **World is getting connected; communications must play major role**

° **CS162 Operating Systems**
- **All special-purpose hardware will run a layer of software that uses processes and concurrent programming; CS162 is the closest thing**

Lab4: start using test benches

° **Idea: wrap testing infrastructure around devices under test.**

° **Include test vectors that are supposed to detect errors in implementation. Even strange ones…**

° **Can (and probably should in later labs) include assert statements to check for "things that should never happen"**

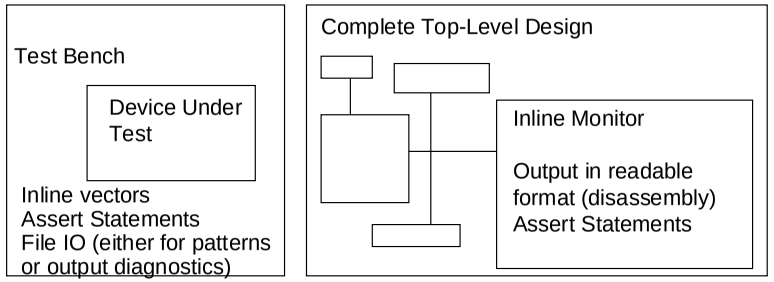

An Alternative MultiCycle DataPath

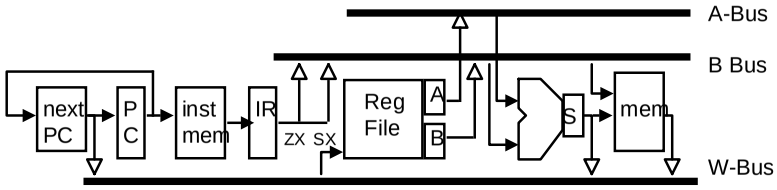

° In each clock cycle, each Bus can be used to transfer from one source

° µ-instruction can simply contain B-Bus and W-Dst fields

What about a 2-Bus Microarchitecture (datapath)?

Instruction Fetch

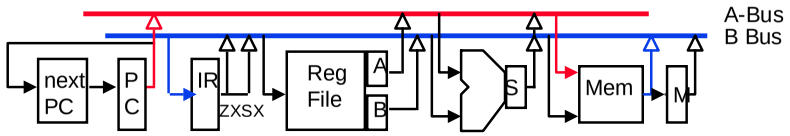

Decode / Operand Fetch

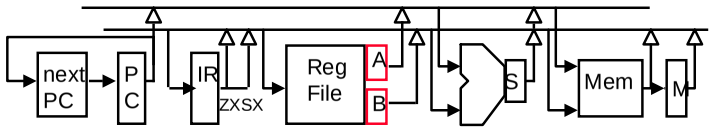

Load

Execute

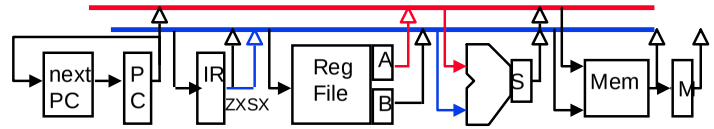

Mem

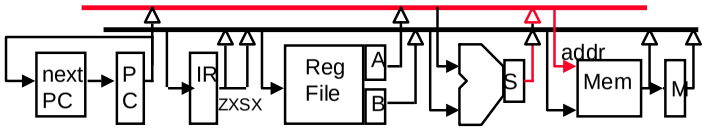

Write-back

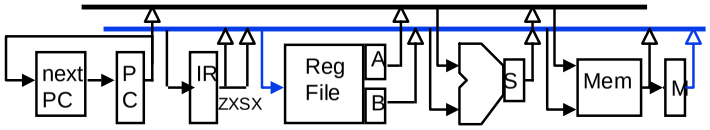

° What about 1 bus ? 1 adder? 1 Register port?

Legacy Software and Microprogramming

° IBM bet company on 360 Instruction Set Architecture (ISA): single instruction set for many classes of machines
  - (8-bit to 64-bit)

° Stewart Tucker stuck with job of what to do about software compatability

° If microprogramming could easily do same instruction set on many different microarchitectures, then why couldn't multiple microprograms do multiple instruction sets on the same microarchitecture?

° Coined term "emulation": instruction set interpreter in microcode for non-native instruction set

° Very successful: in early years of IBM 360 it was hard to know whether old instruction set or new instruction set was more frequently used

Microprogramming Pros and Cons

- Ease of design
- Flexibility
  - Easy to adapt to changes in organization, timing, technology
  - Can make changes late in design cycle, or even in the field
- Can implement very powerful instruction sets (just more control memory)
- Generality
  - Can implement multiple instruction sets on same machine.
  - Can tailor instruction set to application.
- Compatibility
  - Many organizations, same instruction set
- Costly to implement
- Slow

Exceptions

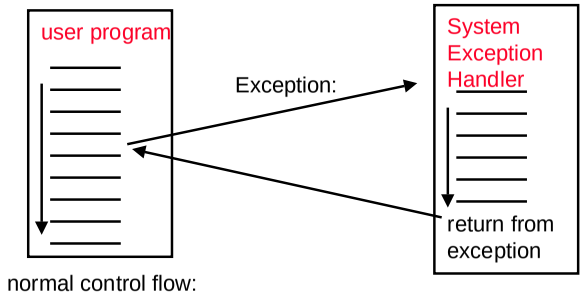

normal control flow:
sequential, jumps, branches, calls, returns

- Exception = unprogrammed control transfer
  - system takes action to handle the exception
    - must record the address of the offending instruction
    - record any other information necessary to return afterwards
  - returns control to user
  - must save & restore user state
- Allows constuction of a "user virtual machine"

Two Types of Exceptions

- Interrupts
  - caused by external events:
    - Network, Keyboard, Disk I/O, Timer
  - asynchronous to program execution
    - Most interrupts can be disabled for brief periods of time
    - Some (like "Power Failing") are non-maskable (NMI)
  - may be handled between instructions
  - simply suspend and resume user program
- Traps
  - caused by internal events
    - exceptional conditions (overflow)
    - errors (parity)
    - faults (non-resident page)
  - synchronous to program execution
  - condition must be remedied by the handler
  - instruction may be retried or simulated and program continued or program may be aborted

MIPS convention:

- exception means any unexpected change in control flow, without distinguishing internal or external; use the term interrupt only when the event is externally caused.

Type of event	*From where?*	*MIPS terminology*
I/O device request	External	Interrupt
Invoke OS from user program	Internal	Exception
Arithmetic overflow	Internal	Exception
Using an undefined instruction	Internal	Exception
Hardware malfunctions	Either	Exception or Interrupt

What happens to Instruction with Exception?

- MIPS architecture defines the instruction as having no effect if the instruction causes an exception.
- When get to virtual memory we will see that certain classes of exceptions must prevent the instruction from changing the machine state.
- This aspect of handling exceptions becomes complex and potentially limits performance => why it is hard

Precise Interrupts

- Precise ⇒ state of the machine is preserved as if program executed up to the offending instruction
  - All previous instructions completed
  - Offending instruction and all following instructions act as if they have not even started
  - Same system code will work on different implementations
  - Position clearly established by IBM
  - Difficult in the presence of pipelining, out-ot-order execution, ...
  - MIPS takes this position
- Imprecise ⇒ system software has to figure out what is where and put it all back together
- Performance goals often lead designers to forsake precise interrupts
  - system software developers, user, markets etc. usually wish they had not done this
- Modern techniques for out-of-order execution and branch prediction help implement precise interrupts

Big Picture: user / system modes

- By providing two modes of execution (user/system) it is possible for the computer to manage itself
  - operating system is a special program that runs in the privileged mode and has access to all of the resources of the computer
  - presents "virtual resources" to each user that are more convenient that the physical resources
    - files vs. disk sectors
    - virtual memory vs physical memory
  - protects each user program from others
  - protects system from malicious users.
  - OS is assumed to "know best", and is trusted code, so enter system mode on exception.
- Exceptions allow the system to taken action in response to events that occur while user program is executing:
  - Might provide supplemental behavior (dealing with denormal floating-point numbers for instance).
  - "Unimplemented instruction" used to emulate instructions that were not included in hardware (I.e. MicroVax)

Addressing the Exception Handler

- Traditional Approach: Interupt Vector
  - PC <- MEM[IV_base + cause || 00]
  - 370, 68000, Vax, 80x86, . . .
- RISC Handler Table
  - PC <– IT_base + cause || 0000
  - saves state and jumps
  - Sparc, PA, M88K, . . .
- MIPS Approach: fixed entry
  - PC <– EXC_addr
  - Actually very small table
    - RESET entry
    - TLB
    - other

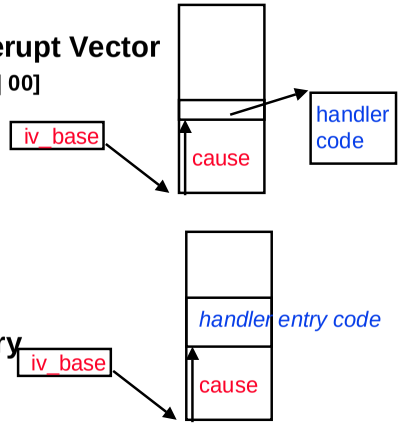

Saving State

- Push it onto the stack
  - Vax, 68k, 80x86
- Save it in special registers
  - MIPS EPC, BadVaddr, Status, Cause
- Shadow Registers
  - M88k
  - Save state in a shadow of the internal pipeline registers

Additions to MIPS ISA to support Exceptions?

- Exception state is kept in "coprocessor 0".
- EPC–a 32-bit register used to hold the address of the affected instruction (register 14 of coprocessor 0).
- Cause–a register used to record the cause of the exception. In the MIPS architecture this register is 32 bits, though some bits are currently unused. Assume that bits 5 to 2 of this register encodes the two possible exception sources mentioned above: undefined instruction=0 and arithmetic overflow=1 (register 13 of coprocessor 0).
- BadVAddr - register contained memory address at which memory reference occurred (register 8 of coprocessor 0)
- Status - interrupt mask and enable bits (register 12 of coprocessor 0)
- Control signals to write EPC , Cause, BadVAddr, and Status
- Be able to write exception address into PC, increase mux to add as input $01000000\ 00000000\ 00000000\ 01000000_{two}$ ($8000\ 0080_{hex}$)
- May have to undo PC = PC + 4, since want EPC to point to offending instruction (not its successor); PC = PC - 4

Recap: Details of Status register

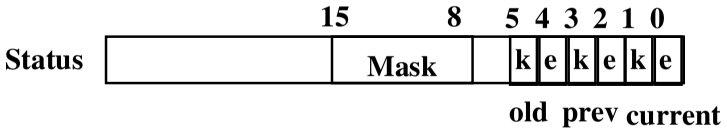

- Mask = 1 bit for each of 5 hardware and 3 software interrupt levels
  - 1 => enables interrupts
  - 0 => disables interrupts
- k = kernel/user
  - 0 => was in the kernel when interrupt occurred
  - 1 => was running user mode
- e = interrupt enable
  - 0 => interrupts were disabled
  - 1 => interrupts were enabled
- When interrupt occurs, 6 LSB shifted left 2 bits, setting 2 LSB to 0
  - run in kernel mode with interrupts disabled

Recap: Details of Cause register

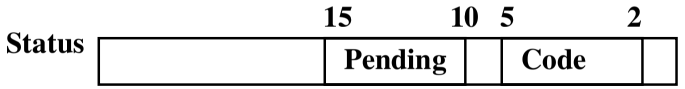

- Pending interrupt 5 hardware levels: bit set if interrupt occurs but not yet serviced
  - handles cases when more than one interrupt occurs at same time, or while records interrupt requests when interrupts disabled
- Exception Code encodes reasons for interrupt
  - 0 (INT) => external interrupt
  - 4 (ADDRL) => address error exception (load or instr fetch)
  - 5 (ADDRS) => address error exception (store)
  - 6 (IBUS) => bus error on instruction fetch
  - 7 (DBUS) => bus error on data fetch
  - 8 (Syscall) => Syscall exception
  - 9 (BKPT) => Breakpoint exception
  - 10 (RI) => Reserved Instruction exception
  - 12 (OVF) => Arithmetic overflow exception

How Control Detects Exceptions in our FSD

- Undefined Instruction–detected when no next state is defined from state 1 for the op value.
  - We handle this exception by defining the next state value for all op values other than lw, sw, 0 (R-type), jmp, beq, and ori as new state 12.
  - Shown symbolically using "other" to indicate that the op field does not match any of the opcodes that label arcs out of state 1.
- Arithmetic overflow–Chapter 4 included logic in the ALU to detect overflow, and a signal called Overflow is provided as an output from the ALU. This signal is used in the modified finite state machine to specify an additional possible next state
- Note: Challenge in designing control of a real machine is to handle different interactions between instructions and other exception-causing events such that control logic remains small and fast.
  - Complex interactions makes the control unit the most challenging aspect of hardware design

How add Exceptions for Overflow and Unimplmented?

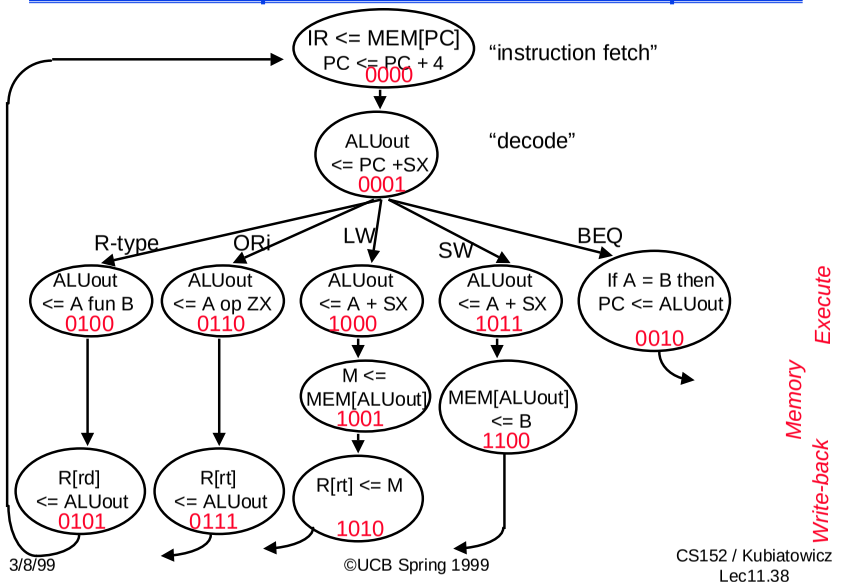

Modification to the Control Specification

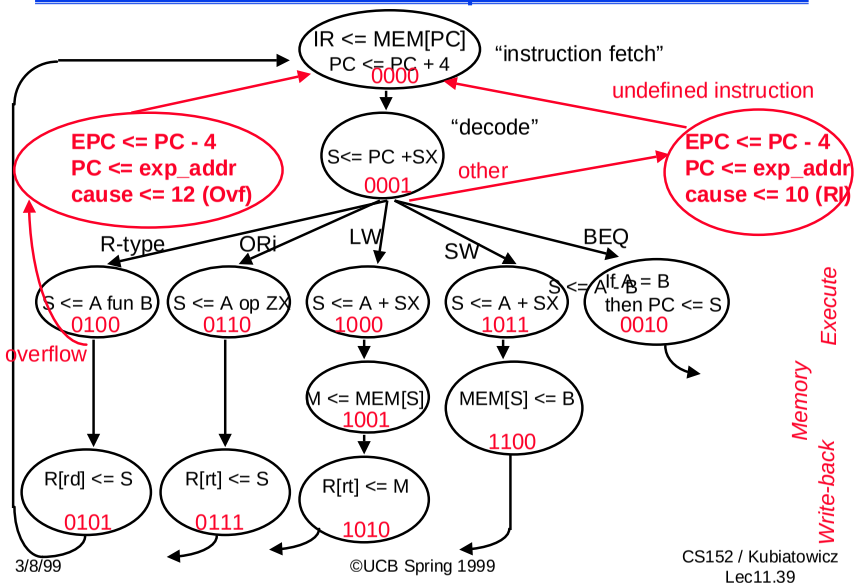

Pipelining is Natural!

- 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

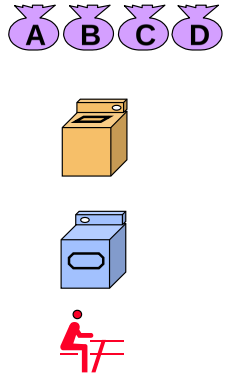

Sequential Laundry

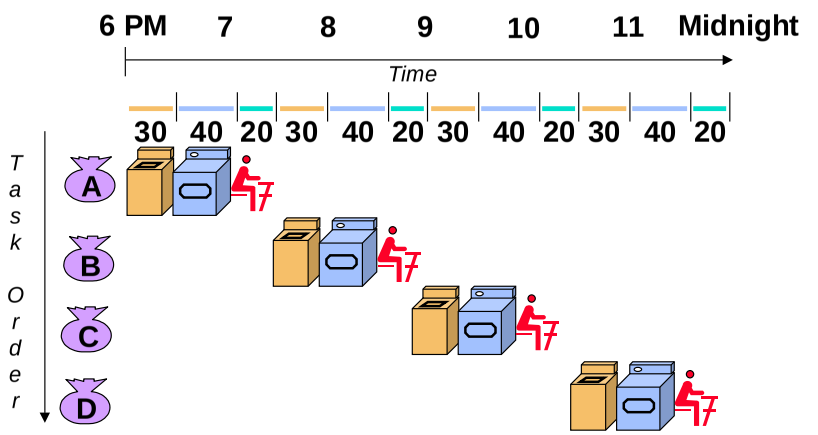

° Sequential laundry takes 6 hours for 4 loads

° If they learned pipelining, how long would laundry take?

Pipelined Laundry: Start work ASAP

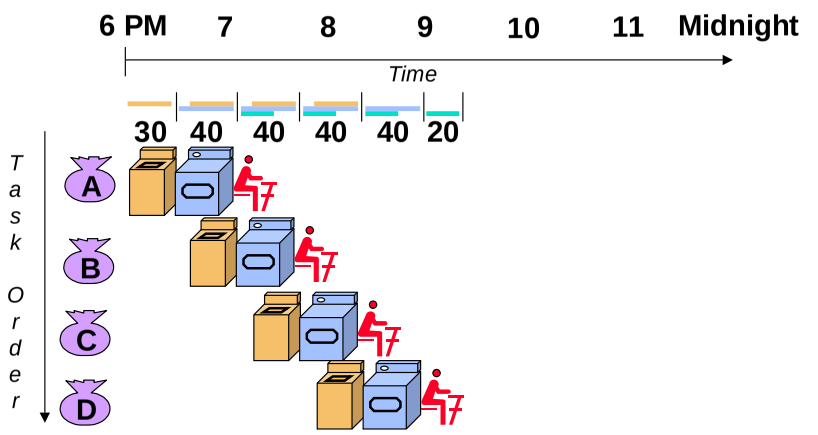

° Pipelined laundry takes 3.5 hours for 4 loads

Pipelining Lessons

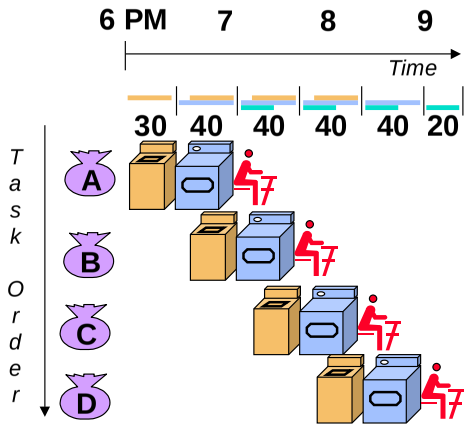

° Pipelining doesn't help latency of single task, it helps throughput of entire workload

° Pipeline rate limited by slowest pipeline stage

° Multiple tasks operating simultaneously using different resources

° Potential speedup = Number pipe stages

° Unbalanced lengths of pipe stages reduces speedup

° Time to "fill" pipeline and time to "drain" it reduces speedup

° Stall for Dependences

Pipelined Execution

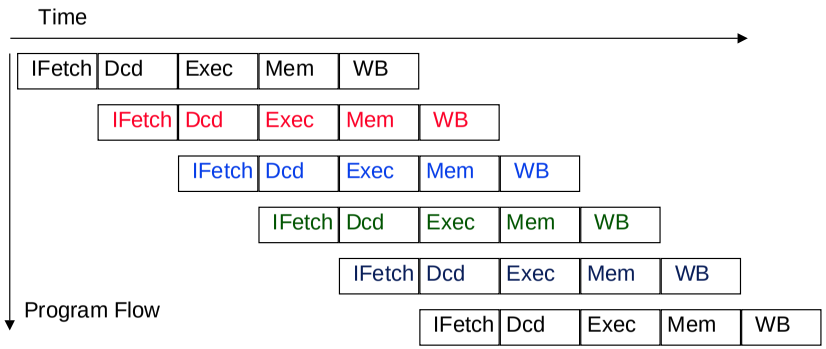

° Utilization?

° Now we just have to make it work

Single Cycle, Multiple Cycle, vs. Pipeline

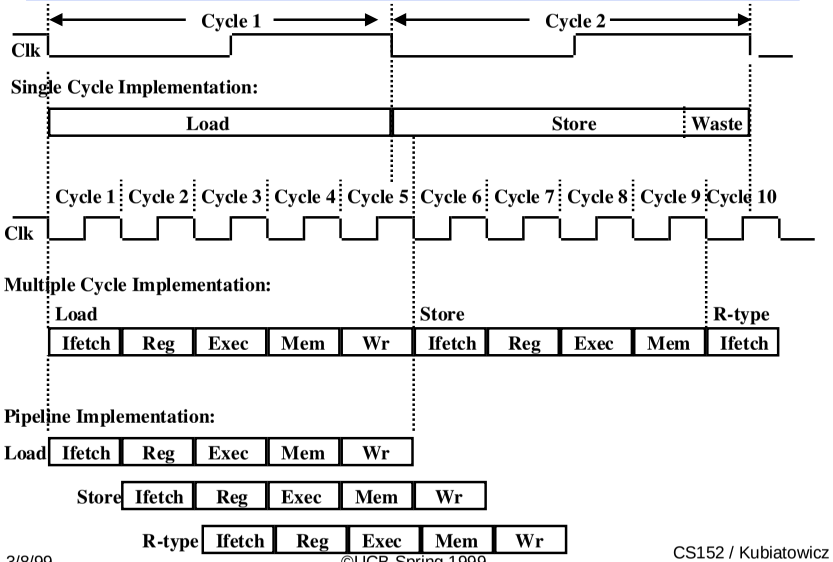

Why Pipeline?

- Suppose we execute 100 instructions
- Single Cycle Machine
  - 45 ns/cycle x 1 CPI x 100 inst = 4500 ns
- Multicycle Machine
  - 10 ns/cycle x 4.6 CPI (due to inst mix) x 100 inst = 4600 ns
- Ideal pipelined machine
  - 10 ns/cycle x (1 CPI x 100 inst + 4 cycle drain) = 1040 ns

Why Pipeline? Because the resources are there!

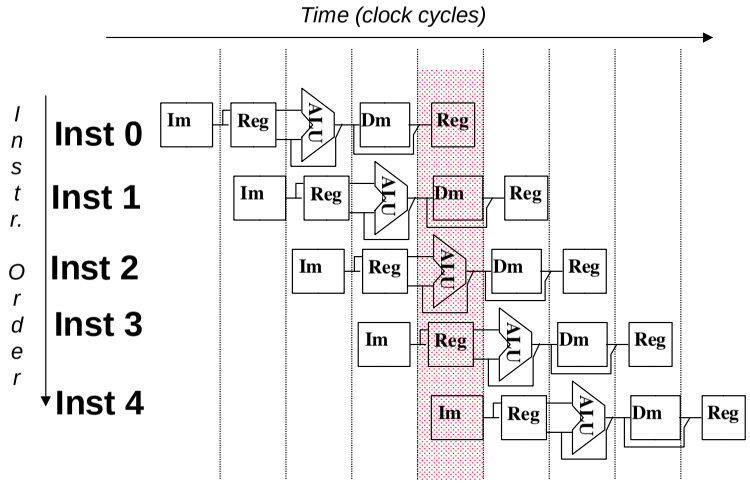

Can pipelining get us into trouble?

- Yes: Pipeline Hazards
  - structural hazards: attempt to use the same resource two different ways at the same time
    - E.g., combined washer/dryer would be a structural hazard or folder busy doing something else (watching TV)
  - data hazards: attempt to use item before it is ready
    - E.g., one sock of pair in dryer and one in washer; can't fold until get sock from washer through dryer
    - instruction depends on result of prior instruction still in the pipeline
  - control hazards: attempt to make a decision before condition is evaulated
    - E.g., washing football uniforms and need to get proper detergent level; need to see after dryer before next load in
    - branch instructions
- Can always resolve hazards by waiting
  - pipeline control must detect the hazard
  - take action (or delay action) to resolve hazards

Summary 1/3

- Specialize state-diagrams easily captured by microsequencer
  - simple increment & “branch” fields
  - datapath control fields
- Control design reduces to Microprogramming
- Exceptions are the hard part of control
- Need to find convenient place to detect exceptions and to branch to state or microinstruction that saves PC and invokes the operating system
- As we get pipelined CPUs that support page faults on memory accesses which means that the instruction cannot complete AND you must be able to restart the program at exactly the instruction with the exception, it gets even harder

Summary 2/3

- 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
- Pipelining is a fundamental concept
  - multiple steps using distinct resources
- Utilize capabilities of the Datapath by pipelined instruction processing
  - start next instruction while working on the current one
  - limited by length of longest stage (plus fill/flush)
  - detect and resolve hazards

Summary: Microprogramming one inspiration for RISC

- If simple instruction could execute at very high clock rate…
- If 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 …
- If same memory used for control memory could be used instead as cache for “macroinstructions”…
- 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)