18-447

Computer Architecture Lecture 5: Single-Cycle and Multi-Cycle Microarchitectures

> Prof. Onur Mutlu Carnegie Mellon University Spring 2014, 1/24/2014

A Single-Cycle Microarchitecture A Closer Look

Remember...

Single-cycle machine



Let's Start with the State Elements



For Now, We Will Assume

- Magic" memory and register file
- Combinational read
 - output of the read data port is a combinational function of the register file contents and the corresponding read select port
- Synchronous write
 - the selected register is updated on the positive edge clock transition when write enable is asserted
 - Cannot affect read output in between clock edges
- Single-cycle, synchronous memory
 - Contrast this with memory that tells when the data is ready
 - i.e., Ready bit: indicating the read or write is done

Instruction Processing

- 5 generic steps (P&H)
 - Instruction fetch (IF)
 - Instruction decode and register operand fetch (ID/RF)
 - Execute/Evaluate memory address (EX/AG)
 - Memory operand fetch (MEM)
 - Store/writeback result (WB)



What Is To Come: The Full MIPS Datapath



**Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

JAL, JR, JALR omitted

Single-Cycle Datapath for Arithmetic and Logical Instructions

R-Type ALU Instructions

- Assembly (e.g., register-register signed addition)
 ADD rd_{reg} rs_{reg} rt_{reg}
- Machine encoding

0	rs	rt	rd	0	ADD	R-type
6-bit	5-bit	5-bit	5-bit	5-bit	6-bit	

Semantics

if MEM[PC] == ADD rd rs rt $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$ $PC \leftarrow PC + 4$

ALU Datapath



if MEM[PC] == ADD rd rs rt $GPR[rd] \leftarrow GPR[rs] + GPR[rt]$ $PC \leftarrow PC + 4$ IF ID EX MEM WB Combinational state update logic

**Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

I-Type ALU Instructions

- Assembly (e.g., register-immediate signed additions)
 ADDI rt_{req} rs_{req} immediate₁₆
- Machine encoding

ADDI	rs	rt	immediate	I-type
6-bit	5-bit	5-bit	16-bit	

Semantics

if MEM[PC] == ADDI rt rs immediate $GPR[rt] \leftarrow GPR[rs] + sign-extend (immediate)$ $PC \leftarrow PC + 4$

Datapath for R and I-Type ALU Insts.



Single-Cycle Datapath for Data Movement Instructions

Load Instructions

- Assembly (e.g., load 4-byte word)
 LW rt_{reg} offset₁₆ (base_{reg})
- Machine encoding

LW	base	rt	offset	l-type
6-bit	5-bit	5-bit	16-bit	

Semantics
 if MEM[PC]==LW rt offset₁₆ (base)
 EA = sign-extend(offset) + GPR[base]
 GPR[rt] ← MEM[translate(EA)]
 PC ← PC + 4

LW Datapath



if MEM[PC]==LW rt offset₁₆ (base) EA = sign-extend(offset) + GPR[base] GPR[rt] \leftarrow MEM[translate(EA)] PC \leftarrow PC + 4 IF ID EX MEM[WB Combinational state update logic 15

Store Instructions

- Assembly (e.g., store 4-byte word) SW rt_{reg} offset₁₆ (base_{reg})
- Machine encoding

SW	base	rt	offset	l-type
6-bit	5-bit	5-bit	16-bit	

Semantics
 if MEM[PC]==SW rt offset₁₆ (base)
 EA = sign-extend(offset) + GPR[base]
 MEM[translate(EA)] ← GPR[rt]
 PC ← PC + 4

SW Datapath



if MEM[PC]==SW rt offset₁₆ (base) EA = sign-extend(offset) + GPR[base] MEM[translate(EA)] \leftarrow GPR[rt] PC \leftarrow PC + 4 IF ID EX MEM WB Combinational state update logic 17

Load-Store Datapath



Datapath for Non-Control-Flow Insts.



Single-Cycle Datapath for *Control Flow Instructions*

Unconditional Jump Instructions

- Assembly
 1 immo
 - J immediate₂₆
- Machine encoding

J	immediate	J-type
6-bit	26-bit	

Semantics
 if MEM[PC]==J immediate₂₆
 target = { PC[31:28], immediate₂₆, 2' b00 }
 PC ← target

Unconditional Jump Datapath



if MEM[PC]==J immediate26
 PC = { PC[31:28], immediate26, 2' b00 }

What about JR, JAL, JAL?

Conditional Branch Instructions

- Assembly (e.g., branch if equal)
 BEQ rs_{reg} rt_{reg} immediate₁₆
- Machine encoding

BEQ	rs	rt	immediate	l-type
6-bit	5-bit	5-bit	16-bit	

Semantics (assuming no branch delay slot) if MEM[PC]==BEQ rs rt immediate₁₆ target = PC + 4 + sign-extend(immediate) x 4 if GPR[rs]==GPR[rt] then PC ← target else PC ← PC + 4

Conditional Branch Datapath (For You to Fix)



How to uphold the delayed branch semanties?

Putting It All Together



**Based on original figure from [P&H CO&D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

JAL, JR, JALR omitted

Single-Cycle Control Logic

Single-Cycle Hardwired Control

As combinational function of Inst=MEM[PC]



- Consider
 - All R-type and I-type ALU instructions
 - LW and SW
 - BEQ, BNE, BLEZ, BGTZ
 - J, JR, JAL, JALR

Single-Bit Control Signals

	When De-asserted	When asserted	Equation
RegDest	GPR write select according to rt, i.e., inst[20:16]	GPR write select according to rd, i.e., inst[15:11]	opcode==0
ALUSrc	2 nd ALU input from 2 nd GPR read port	2 nd ALU input from sign- extended 16-bit immediate	(opcode!=0) && (opcode!=BEQ) && (opcode!=BNE)
MemtoReg	Steer ALU result to GPR write port	steer memory load to GPR wr. port	opcode==LW
RegWrite	GPR write disabled	GPR write enabled	(opcode!=SW) && (opcode!=Bxx) && (opcode!=J) && (opcode!=JR))

JAL and JALR require additional RegDest and MemtoReg options

	When De-asserted	When asserted	Equation
MemRead	Memory read disabled	Memory read port return load value	opcode==LW
MemWrite	Memory write disabled	Memory write enabled	opcode==SW
PCSrc ₁	According to PCSrc ₂	next PC is based on 26- bit immediate jump target	(opcode==J) (opcode==JAL)
PCSrc ₂	next PC = PC + 4	next PC is based on 16- bit immediate branch target	(opcode==Bxx) && "bcond is satisfied"

JR and JALR require additional PCSrc options

ALU Control

- case opcode
 - '0' \Rightarrow select operation according to funct
 - 'ALUi' \Rightarrow selection operation according to opcode
 - 'LW' \Rightarrow select addition
 - 'SW' \Rightarrow select addition
 - 'Bxx' \Rightarrow select bcond generation function
 - $_$ \Rightarrow don't care
- Example ALU operations
 - □ ADD, SUB, AND, OR, XOR, NOR, etc.
 - bcond on equal, not equal, LE zero, GT zero, etc.

R-Type ALU



I-Type ALU



LW



SW



Branch Not Taken



Branch Taken


Jump



2004 Elsevier. ALL RIGHTS RESERVED.]

What is in That Control Box?

- Combinational Logic → Hardwired Control
 - Idea: Control signals generated combinationally based on instruction
 - Necessary in a single-cycle microarchitecture...
- Sequential Logic → Sequential/Microprogrammed Control
 - Idea: A memory structure contains the control signals associated with an instruction
 - Control Store

Evaluating the Single-Cycle Microarchitecture

A Single-Cycle Microarchitecture

- Is this a good idea/design?
- When is this a good design?
- When is this a bad design?
- How can we design a better microarchitecture?

A Single-Cycle Microarchitecture: Analysis

- Every instruction takes 1 cycle to execute
 - □ CPI (Cycles per instruction) is strictly 1
- How long each instruction takes is determined by how long the slowest instruction takes to execute
 - Even though many instructions do not need that long to execute
- Clock cycle time of the microarchitecture is determined by how long it takes to complete the slowest instruction
 - Critical path of the design is determined by the processing time of the slowest instruction

What is the Slowest Instruction to Process?

- Let's go back to the basics
- All six phases of the instruction processing cycle take a single machine clock cycle to complete
- Fetch
- Decode
- Evaluate Address
- Fetch Operands
- Execute
- Store Result

- 1. Instruction fetch (IF)
- 2. Instruction decode and
 - register operand fetch (ID/RF)
- 3. Execute/Evaluate memory address (EX/AG)
- 4. Memory operand fetch (MEM)
- 5. Store/writeback result (WB)

Do each of the above phases take the same time (latency) for all instructions?

Single-Cycle Datapath Analysis

Assume

- memory units (read or write): 200 ps
- ALU and adders: 100 ps
- register file (read or write): 50 ps
- other combinational logic: 0 ps

steps	IF	ID	EX	MEM	WB	
resources	mem	RF	ALU	mem	RF	Delay
R-type	200	50	100		50	400
l-type	200	50	100		50	400
LW	200	50	100	200	50	600
SW	200	50	100	200		550
Branch	200	50	100			350
Jump	200					200 4

Let's Find the Critical Path



R-Type and I-Type ALU



LW



SW



Branch Taken



2004 Elsevier. ALL RIGHTS RESERVED.]

ump



2004 Elsevier. ALL RIGHTS RESERVED.]

What About Control Logic?

- How does that affect the critical path?
- Food for thought for you:
 - Can control logic be on the critical path?
 - □ A note on CDC 5600: control store access too long...

What is the Slowest Instruction to Process?

- Memory is not magic
- What if memory *sometimes* takes 100ms to access?
- Does it make sense to have a simple register to register add or jump to take {100ms+all else to do a memory operation}?
- And, what if you need to access memory more than once to process an instruction?
 - Which instructions need this?
 - Do you provide multiple ports to memory?

Single Cycle uArch: Complexity

Contrived

□ All instructions run as slow as the slowest instruction

Inefficient

- All instructions run as slow as the slowest instruction
- Must provide worst-case combinational resources in parallel as required by any instruction
- Need to replicate a resource if it is needed more than once by an instruction during different parts of the instruction processing cycle
- Not necessarily the simplest way to implement an ISA
 - Single-cycle implementation of REP MOVS, INDEX, POLY?
- Not easy to optimize/improve performance
 - Optimizing the common case does not work (e.g. common instructions)
 - Need to optimize the worst case all the time

Microarchitecture Design Principles

Critical path design

□ Find the maximum combinational logic delay and decrease it

- Bread and butter (common case) design
 - Spend time and resources on where it matters
 - i.e., improve what the machine is really designed to do
 - Common case vs. uncommon case

Balanced design

- Balance instruction/data flow through hardware components
- Balance the hardware needed to accomplish the work

How does a single-cycle microarchitecture fare in light of these principles?

Multi-Cycle Microarchitectures

Multi-Cycle Microarchitectures

- Goal: Let each instruction take (close to) only as much time it really needs
 - Idea
 - Determine clock cycle time independently of instruction processing time
 - Each instruction takes as many clock cycles as it needs to take
 - Multiple state transitions per instruction
 - The states followed by each instruction is different

Remember: The "Process instruction" Step

- ISA specifies abstractly what A' should be, given an instruction and A
 - □ It defines an abstract finite state machine where
 - State = programmer-visible state
 - Next-state logic = instruction execution specification
 - From ISA point of view, there are no "intermediate states" between A and A' during instruction execution
 - One state transition per instruction
- Microarchitecture implements how A is transformed to A'
 - There are many choices in implementation
 - We can have programmer-invisible state to optimize the speed of instruction execution: multiple state transitions per instruction
 - Choice 1: $AS \rightarrow AS'$ (transform A to A' in a single clock cycle)
 - Choice 2: AS → AS+MS1 → AS+MS2 → AS+MS3 → AS' (take multiple clock cycles to transform AS to AS')

Multi-Cycle Microarchitecture



Benefits of Multi-Cycle Design

Critical path design

- Can keep reducing the critical path independently of the worstcase processing time of any instruction
- Bread and butter (common case) design
 - Can optimize the number of states it takes to execute "important" instructions that make up much of the execution time

Balanced design

- No need to provide more capability or resources than really needed
 - An instruction that needs resource X multiple times does not require multiple X's to be implemented
 - Leads to more efficient hardware: Can reuse hardware components needed multiple times for an instruction

Performance Analysis

- Execution time of an instruction
 - □ {CPI} x {clock cycle time}
- Execution time of a program
 - Sum over all instructions [{CPI} x {clock cycle time}]
 - □ {# of instructions} x {Average CPI} x {clock cycle time}
- Single cycle microarchitecture performance
 - □ CPI = 1
 - Clock cycle time = long
- Multi-cycle microarchitecture performance
 - CPI = different for each instruction
 - Average CPI \rightarrow hopefully small
 - Clock cycle time = short

Now, we have two degrees of freedom to optimize independently

An Aside: CPI vs. Frequency

- CPI vs. Clock cycle time
- At odds with each other
 - Reducing one increases the other for a single instruction

Why?

- Average CPI can be amortized/reduced via concurrent processing of multiple instructions
 - □ The same cycle is devoted to processing multiple instructions
 - Example: Pipelining, superscalar execution

A Multi-Cycle Microarchitecture A Closer Look

How Do We Implement This?

- Maurice Wilkes, "The Best Way to Design an Automatic Calculating Machine," Manchester Univ. Computer Inaugural Conf., 1951.
- The concept of microcoded/microprogrammed machines
- Realization
 - One can implement the "process instruction" step as a finite state machine that sequences between states and eventually returns back to the "fetch instruction" state
 - □ A state is defined by the control signals asserted in it
 - Control signals for the next state determined in current state

The Instruction Processing Cycle



A Basic Multi-Cycle Microarchitecture

- Instruction processing cycle divided into "states"
 - A stage in the instruction processing cycle can take multiple states
- A multi-cycle microarchitecture sequences from state to state to process an instruction
 - The behavior of the machine in a state is completely determined by control signals in that state
- The behavior of the entire processor is specified fully by a finite state machine
- In a state (clock cycle), control signals control
 - How the datapath should process the data
 - How to generate the control signals for the next clock cycle

Microprogrammed Control Terminology

- Control signals associated with the current state
 Microinstruction
- Act of transitioning from one state to another
 - Determining the next state and the microinstruction for the next state
 - Microsequencing
- Control store stores control signals for every possible state
 Store for microinstructions for the entire FSM
- Microsequencer determines which set of control signals will be used in the next clock cycle (i.e., next state)

What Happens In A Clock Cycle?

- The control signals (microinstruction) for the current state control
 - Processing in the data path
 - Generation of control signals (microinstruction) for the next cycle
 - See Supplemental Figure 1 (next slide)
- Datapath and microsequencer operate concurrently
- Question: why not generate control signals for the current cycle in the current cycle?
 - This will lengthen the clock cycle
 - Why would it lengthen the clock cycle?
 - See Supplemental Figure 2

A Clock Cycle

	•	E	Supplemental Figures
Cyde N	Cycle N+1		
		4	
1) Processing in Palapath for	Cycle N		
(2) Gonoration of Control !			
Signals for			
Cycle N+1	and the state of the		and the second
	1) Results	of coment-	[F151]
	2) Controls	ignols needed f	NH1

A Bad Clock Cycle!

Altratic - A BADONE!	the office inti
<u></u>	
O Governing of Control Stands for Cucle M	
A Part of control signals to cycle to	
TO Processing for Notopoth for Cycle N	
Step (1) is dependent on Step (0)	
If stop () takes non-zoro time (it does!), clock	cycle moreases
unnecessarly	
-> Viclales the "Critical Path Design" princip	ple
	Fig 21
	1-3-1

A Simple LC-3b Control and Datapath



Figure C.1: Microarchitecture of the LC-3b, major components

What Determines Next-State Control Signals?

- What is happening in the current clock cycle
 - □ See the 9 control signals coming from "Control" block
 - What are these for?
- The instruction that is being executed
 IR[15:11] coming from the Data Path
- Whether the condition of a branch is met, if the instruction being processed is a branch
 - BEN bit coming from the datapath
- Whether the memory operation is completing in the current cycle, if one is in progress
 - R bit coming from memory

A Simple LC-3b Control and Datapath



Figure C.1: Microarchitecture of the LC-3b, major components

The State Machine for Multi-Cycle Processing

- The behavior of the LC-3b uarch is completely determined by
 - the 35 control signals and
 - additional 7 bits that go into the control logic from the datapath
- 35 control signals completely describe the state of the control structure
- We can completely describe the behavior of the LC-3b as a state machine, i.e. a directed graph of
 - Nodes (one corresponding to each state)
 - Arcs (showing flow from each state to the next state(s))
An LC-3b State Machine

- Patt and Patel, App C, Figure C.2
- Each state must be uniquely specified
 Done by means of *state variables*
- 31 distinct states in this LC-3b state machine
 Encoded with 6 state variables
- Examples
 - State 18,19 correspond to the beginning of the instruction processing cycle
 - □ Fetch phase: state 18, 19 \rightarrow state 33 \rightarrow state 35
 - Decode phase: state 32



LC-3b State Machine: Some Questions

- How many cycles does the fastest instruction take?
- How many cycles does the slowest instruction take?
- Why does the BR take as long as it takes in the FSM?
- What determines the clock cycle?
- Is this a Mealy machine or a Moore machine?

LC-3b Datapath

- Patt and Patel, App C, Figure C.3
- Single-bus datapath design
 - At any point only one value can be "gated" on the bus (i.e., can be driving the bus)
 - Advantage: Low hardware cost: one bus
 - Disadvantage: Reduced concurrency if instruction needs the bus twice for two different things, these need to happen in different states
- Control signals (26 of them) determine what happens in the datapath in one clock cycle
 - □ Patt and Patel, App C, Table C.1







Signal Name	Signal Values	
LD.MAR/1: LD.MDR/1: LD.IR/1: LD.BEN/1: LD.REG/1: LD.CC/1: LD.PC/1:	NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD	
GatePC/1: GateMDR/1: GateALU/1: GateMARMUX/1: GateSHF/1:	NO, YES NO, YES NO, YES NO, YES NO, YES	
PCMUX/2:	PC+2 BUS ADDER	;select pc+2 ;select value from bus ;select output of address adder
DRMUX/1:	11.9 R7	;destination IR[11:9] ;destination R7
SR1MUX/1:	11.9 8.6	;source IR[11:9] ;source IR[8:6]
ADDR1MUX/1:	PC, BaseR	
ADDR2MUX/2:	ZERO offset6 PCoffset9 PCoffset11	;select the value zero ;select SEXT[IR[5:0]] ;select SEXT[IR[8:0]] ;select SEXT[IR[10:0]]
MARMUX/1:	7.0 ADDER	;select LSHF(ZEXT[IR[7:0]],1) ;select output of address adder
ALUK/2:	ADD, AND, X	COR, PASSA
MIO.EN/1: R.W/1: DATA.SIZE/1: LSHF1/1:	NO, YES RD, WR BYTE, WORL NO, YES)

Table C.1: Data path control signals

LC-3b Datapath: Some Questions

- How does instruction fetch happen in this datapath according to the state machine?
- What is the difference between gating and loading?
- Is this the smallest hardware you can design?

LC-3b Microprogrammed Control Structure

- Patt and Patel, App C, Figure C.4
- Three components:
 - Microinstruction, control store, microsequencer
- Microinstruction: control signals that control the datapath (26 of them) and determine the next state (9 of them)
- Each microinstruction is stored in a unique location in the control store (a special memory structure)
- Unique location: address of the state corresponding to the microinstruction
 - Remember each state corresponds to one microinstruction
- Microsequencer determines the address of the next microinstruction (i.e., next state)





Ð	Cond		~			40.	40 440	40 .07	(D. A.	19.07 19.07	9 9 9	10°.	Galen-	Gaten.	do to	Gaten.	Gate Chr.	b, the Od	ton at	50.42	40.04	-up with the	r man	MARY TOP	41, 44	the cf	R.H.	04. 04.	Lands	s. (41.
	'	1		1	1																				1					000000 (State 0)
	1				1													1							1					000001 (State 1)
	'																													000010 (State 2)
																		L .												000011 (State 3)
			· ·	·															_	_		-		_				-		000100 (State 4)
_	- 1				-													-		_	-			_					-	000101 (State 5)
-	- 1					-													-	-		-		_			-	-	-	000110 (State 6)
-					-	-													+	+	-	-		-			-	+	-	001000 (State 8)
	'	1			1																				1			\mathbf{T}		001000 (State 0) 001001 (State 9)
	'	1			1																							\top		001010 (State 10)
	1			1	1													1												001011 (State 11)
	1																													001100 (State 12)
																														001101 (State 13)
			· ·	·															_	_		-		_				-	_	001110 (State 14)
_	- 1		· · ·		-	-												-	+	+	-			-			_	+	-	001111 (State 15)
-	1	- 1			-													-	+	+	-			-	- 1		-	+		010000 (State 16) 010001 (State 17)
	1				-													-	+	+		1			- 1			+		010010 (State 18)
	1	- 1		- 1	1													-	+						- 1					010011 (State 19)
	'	1	1 1	1	1																				1					010100 (State 20)
	'	1			1													-							- 1					010101 (State 21)
	'	1	1 1	1																					1					010110 (State 22)
	1			1	1													T							1					010111 (State 23)
	'	-																		_				_				1		011000 (State 24)
			· ·		·	_												· ·	_	_				_			_	-	_	011001 (State 25)
					-	-													_	+				\rightarrow			_	+		011010 (State 26)
-			1 1			-													_	-	-	-		_			_	+	-	011011 (State 27)
-	- 1				-	-													-	-		-		_	- 1		-	-	-	011100 (State 28)
-	1				-	-													-	-	-	+		-			-	+	-	011101 (State 29)
	- 1	- 1	1 1		1														-	-		1			- 1			-		011110 (State 30) 011111 (State 31)
	'	- 1		1	1															+					1			+		100000 (State 32)
	'	1	1 1	1	1																				1					100001 (State 33)
	1		1 1	1																										100010 (State 34)
	'																													100011 (State 35)
					·													L ·												100100 (State 36)
			· ·																	_				_						100101 (State 37)
-						-													+	+	-	-		-			-	+	-	100110 (State 38)
_					-	-												-	_	-	-	-		_			-	-	-	100111 (State 39)
	1				1															+					1	-				101000 (State 40) 101001 (State 41)
	'			- 1	-													-	-	-								1		101010 (State 42)
				T	1	1												-	+		1				1			1		101011 (State 43)
																														101100 (State 44)
			<u> </u>																T	T										101101 (State 45)
																														101110 (State 46)
					-	-														\perp	1				- 1		_	1	\square	101111 (State 47)
-			· ·		· ·	1	-	<u> </u>											+	+	+	+		_		+	+	+	+	110000 (State 48)
-	-,					1	-	-	\vdash									-	+	+	+	+		_	-	+	+	+	+	110001 (State 49)
-	-,				-	-		-										-	_	+	+	+		_		-	_	+	+	110010 (State 50)
\vdash					1	1		-											+	+	+	\vdash			1	+	+	+	+	110100 (State 52)
				-	-	\vdash												-	+	+	+	\uparrow			T	+	+	+	+	110101 (State 53)
				1	-													-			1	1			1			1		110110 (State 54)
				I	1	1												-	+		1				Ţ			1		110111 (State 55)
				1	1																				1					111000 (State 56)
																														111001 (State 57)
						1														_	1						_	1	\perp	111010 (State 58)
						_														\perp	1							1	\square	111011 (State 59)
					· 	1		<u> </u>											+	_	+	\square					_	-	+	111100 (State 60)
_				· ·	-	-														_	1			_				-	+	111101 (State 61)
-	-,			- 1		1			\square									-	+	+	+	\vdash			- 1	-	+	+	+	111110 (State 62)
									-												-					_	_	-	1	1 11111 (State 63)

LC-3b Microsequencer

- Patt and Patel, App C, Figure C.5
- The purpose of the microsequencer is to determine the address of the next microinstruction (i.e., next state)
- Next address depends on 9 control signals

Signal Name	Signal Val	ues
J/6: COND/2:	COND ₀ COND ₁ COND ₂ COND ₃	;Unconditional ;Memory Ready ;Branch ;Addressing Mode
IRD/1:	NO, YES	

Table C.2: Microsequencer control signals



The Microsequencer: Some Questions

- When is the IRD signal asserted?
- What happens if an illegal instruction is decoded?
- What are condition (COND) bits for?
- How is variable latency memory handled?
- How do you do the state encoding?
 - Minimize number of state variables
 - Start with the 16-way branch
 - Then determine constraint tables and states dependent on COND

An Exercise in Microprogramming

Handouts

- 7 pages of Microprogrammed LC-3b design
- http://www.ece.cmu.edu/~ece447/s13/doku.php?id=manu als
- http://www.ece.cmu.edu/~ece447/s13/lib/exe/fetch.php?m edia=lc3b-figures.pdf

A Simple LC-3b Control and Datapath



Figure C.1: Microarchitecture of the LC-3b, major components











Signal Name	Signal Values	
LD.MAR/1: LD.MDR/1: LD.IR/1: LD.BEN/1: LD.REG/1: LD.CC/1: LD.PC/1:	NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD NO, LOAD	
GatePC/1: GateMDR/1: GateALU/1: GateMARMUX/1: GateSHF/1:	NO, YES NO, YES NO, YES NO, YES NO, YES	
PCMUX/2:	PC+2 BUS ADDER	;select pc+2 ;select value from bus ;select output of address adder
DRMUX/1:	11.9 R7	;destination IR[11:9] ;destination R7
SR1MUX/1:	11.9 8.6	;source IR[11:9] ;source IR[8:6]
ADDR1MUX/1:	PC, BaseR	
ADDR2MUX/2:	ZERO offset6 PCoffset9 PCoffset11	;select the value zero ;select SEXT[IR[5:0]] ;select SEXT[IR[8:0]] ;select SEXT[IR[10:0]]
MARMUX/1:	7.0 ADDER	;select LSHF(ZEXT[IR[7:0]],1) ;select output of address adder
ALUK/2:	ADD, AND, X	COR, PASSA
MIO.EN/1: R.W/1: DATA.SIZE/1: LSHF1/1:	NO, YES RD, WR BYTE, WORL NO, YES)

Table C.1: Data path control signals





Ð	Cond		~			40.	40 440	40 .07	(D. A.	19.07 19.07	9 9 9	10°.	Galen-	Gaten.	do to	Gaten.	Galeo, Mr.	b, the Od	ton at	50.42	40.04	-up with the	r man	MARY TOP	41, 44	the cf	R.H.	04. 04.	Lands	s. (41.
	'	1		1	1																				1					000000 (State 0)
	1			-	1													1							1					000001 (State 1)
	'																													000010 (State 2)
																		L .												000011 (State 3)
			· ·																_	_		-		_				-		000100 (State 4)
_	- 1				-													-		_	-			_					-	000101 (State 5)
-	- 1					-													-	-		-		_			-	-	-	000110 (State 6)
-					-	-													+	+	-	-		-			-	+	-	001111 (State 7)
	'	1			1																				1			\mathbf{T}		001000 (State 0) 001001 (State 9)
	'	1			1																							\top		001010 (State 10)
	1			1	1													1												001011 (State 11)
	1																													001100 (State 12)
																														001101 (State 13)
			· ·	·															_	_		-		_				-	_	001110 (State 14)
_	-		· · ·		-	-												-	+	+	-			-			_	+	-	001111 (State 15)
-	1	- 1			-													-	+	+	-			-	- 1		-	+		010000 (State 16) 010001 (State 17)
	1				-													-	+	+		1		-	- 1			+		010010 (State 18)
	1	- 1		- 1	1													-	+						- 1					010011 (State 19)
	'	1	1 1	1	1																				1					010100 (State 20)
	'	1			1													-							- 1					010101 (State 21)
	'	1	1 1	1																					1					010110 (State 22)
	1			1	1													T							1					010111 (State 23)
	'	-																		_				_				1		011000 (State 24)
			· ·		·	_												· ·	_	_				_			_	-	_	011001 (State 25)
					-	-													_	+				\rightarrow			_	+		011010 (State 26)
-			1 1			-													_	-	-	-		_			_	+	-	011011 (State 27)
-	- 1				-	-													-	-		-		_	- 1		-	-		011100 (State 28)
-	1				-	-													-	-	-	+		-			-	+	-	011101 (State 29)
	- 1	- 1	1 1		1														-	-		1			- 1		-	-		011110 (State 30) 011111 (State 31)
	'	- 1		1	1															+					1			+		100000 (State 32)
	'	1	1 1	1	1																				1					100001 (State 33)
	1		1 1	1																										100010 (State 34)
	'																													100011 (State 35)
					·													L ·												100100 (State 36)
			· ·																	_				_						100101 (State 37)
-						-													+	+	-	-		_			-	+	-	100110 (State 38)
_					-	-												-	_	-	-	-		_			-	-	-	100111 (State 39)
	1				1															+					1	-				101000 (State 40) 101001 (State 41)
	'			- 1	-													-	-	-								1		101010 (State 42)
				T	1	1												-	+		1				T			1		101011 (State 43)
						L																								101100 (State 44)
			<u> </u>																T	T										101101 (State 45)
																														101110 (State 46)
					-	-														\perp	1				- 1		_	1	\square	101111 (State 47)
-			· ·		· ·	1	-	<u> </u>											+	+	+	+		_		+	+	+	+	110000 (State 48)
-	-,					1	-	-	\vdash									-	+	+	+	+		_	-	+	+	+	+	110001 (State 49)
-	-,				-	-		-										-	_	+	+	+		_		-	_	+	+	110010 (State 50)
\vdash					1	1		-											+	+	+	\vdash			1	+	+	+	+	110100 (State 52)
				-	-	\vdash												-	+	+	+	\uparrow			T	+	+	+	+	110101 (State 53)
				1	-													-			1	1			1			1		110110 (State 54)
				I	1	1												-	+		1				Ţ			1		110111 (State 55)
				1	1																				1					111000 (State 56)
																														111001 (State 57)
						1														_	1						_	1	\perp	111010 (State 58)
						_														\perp	1							1	\square	111011 (State 59)
					· 	1		<u> </u>											+	_	+	\square					_	-	+	111100 (State 60)
_				· ·	-	-														_	1			_				-	+	111101 (State 61)
-	-,			- 1		1			\square									-	+	+	+	\vdash			- 1	-	+	+	+	111110 (State 62)
									-												-					_	_	-	1	1 11111 (State 63)



End of the Exercise in Microprogramming

Homework 2

 You will write the microcode for the entire LC-3b as specified in Appendix C