

18-447

Computer Architecture
Lecture 6: Multi-Cycle and
Microprogrammed Microarchitectures

Prof. Onur Mutlu

Carnegie Mellon University

Spring 2014, 1/27/2014

Assignments

- Lab 2 due next Friday (start early)
- HW1 due this week
- HW0
 - Make sure you submitted this!

Extra Credit for Lab Assignment 2

- Complete your normal (single-cycle) implementation first, and get it checked off in lab.
- Then, implement the MIPS core using a microcoded approach similar to what we will discuss in class.
- We are not specifying any particular details of the microcode format or the microarchitecture; you can be creative.
- For the extra credit, the microcoded implementation should execute the same programs that your ordinary implementation does, and you should demo it by the normal lab deadline.
- You will get maximum 4% of course grade
- Document what you have done and demonstrate well

Readings for Today

- P&P, Revised Appendix C
 - Microarchitecture of the LC-3b
 - Appendix A (LC-3b ISA) will be useful in following this
- P&H, Appendix D
 - Mapping Control to Hardware
- Optional
 - Maurice Wilkes, “[The Best Way to Design an Automatic Calculating Machine](#),” Manchester Univ. Computer Inaugural Conf., 1951.

Readings for Next Lecture

- Pipelining
 - P&H Chapter 4.5-4.8
- Pipelined LC-3b Microarchitecture
 - <http://www.ece.cmu.edu/~ece447/s14/lib/exe/fetch.php?media=18447-lc3b-pipelining.pdf>

Quick Recap of Past Five Lectures

- Basics
 - Why Computer Architecture
 - Levels of Transformation
 - Memory Topics: DRAM Refresh and Memory Performance Attacks
- ISA Tradeoffs
- Single-Cycle Microarchitectures
- Multi-Cycle Microarchitectures
- Performance Analysis
 - Amdahl's Law
- Microarchitecture Design Principles

Review: 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?*

Review: 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


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

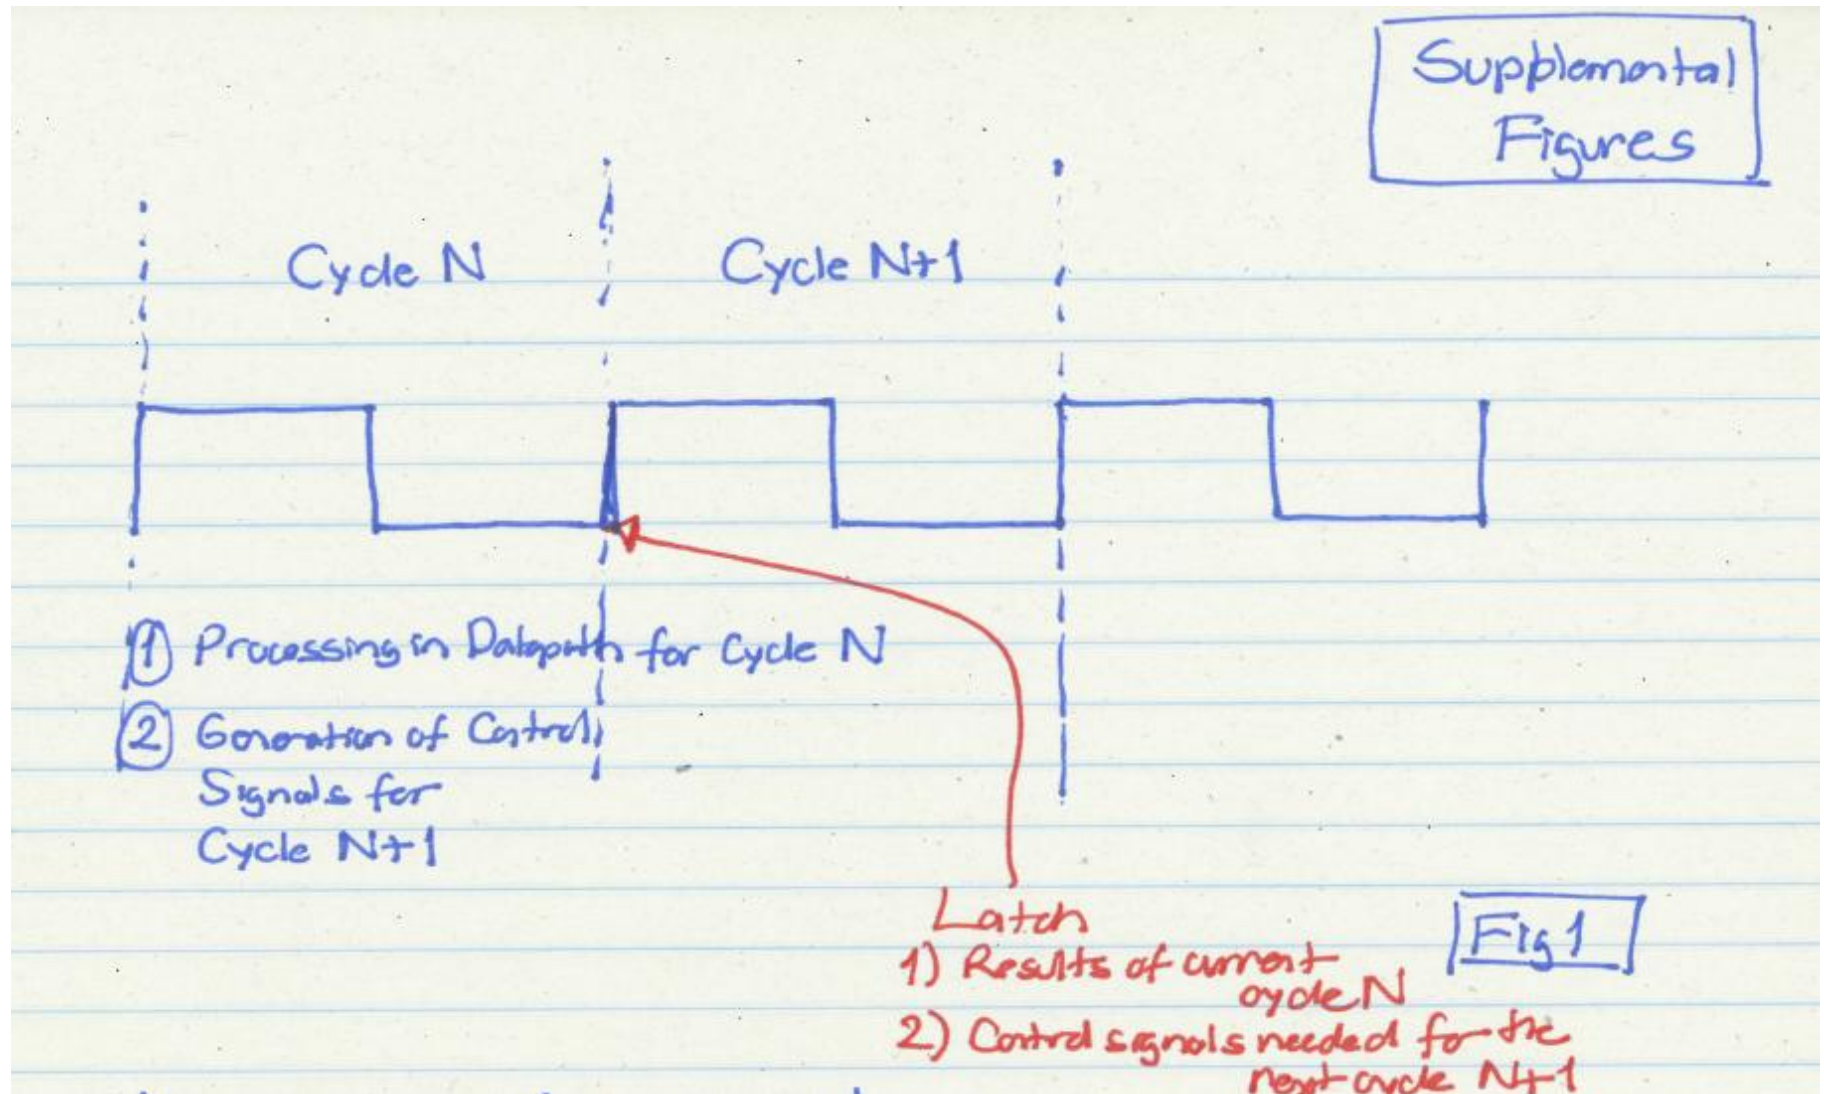
Microprogrammed Control Terminology

- 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?

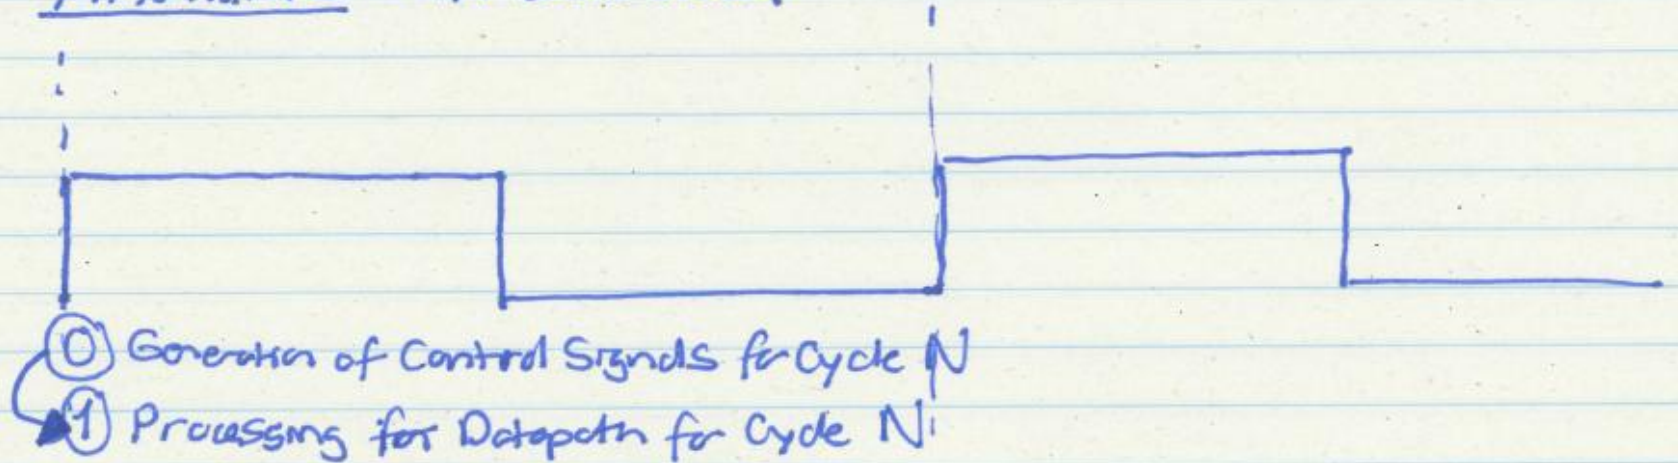
- The control signals (microinstruction) for the current state control
 - Processing in the data path
 - Generation of control signals (microinstruction) for the next cycle
 - See Supplemental Figure 1 (next slide)
- Datapath and microsequencer operate concurrently
- Question: why not generate control 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



A Bad Clock Cycle!

Alternative - A BAD ONE!



Step (1) is dependent on Step (0)

If Step (0) takes non-zero time (it does!), clock cycle increases unnecessarily

→ Violates the "Critical Path Design" principle

Fig 2

A Simple LC-3b Control and Datapath

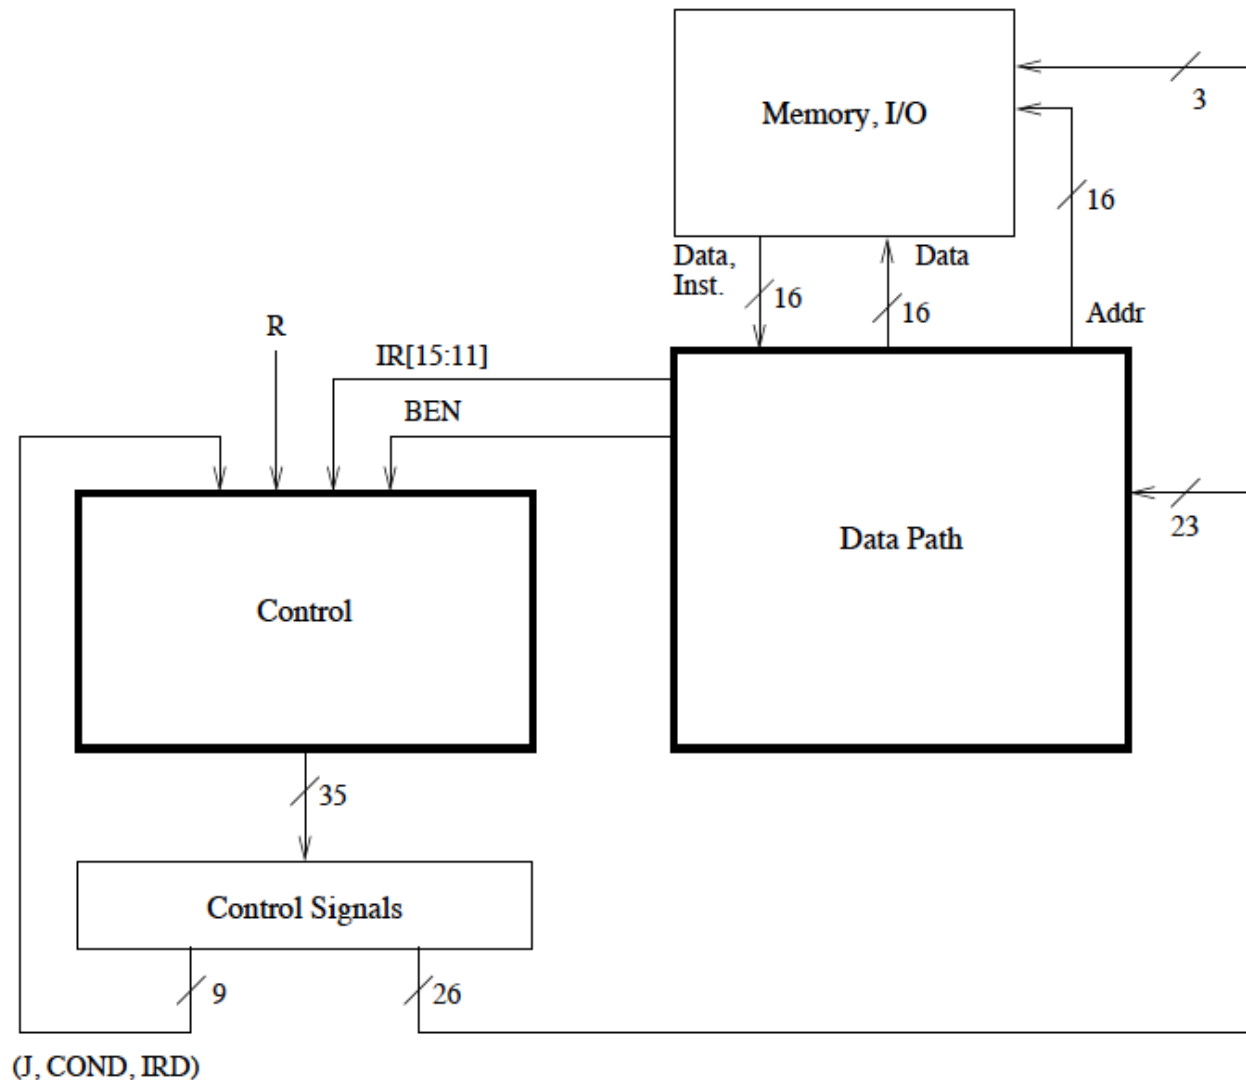


Figure C.1: Microarchitecture of the LC-3b, major components

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

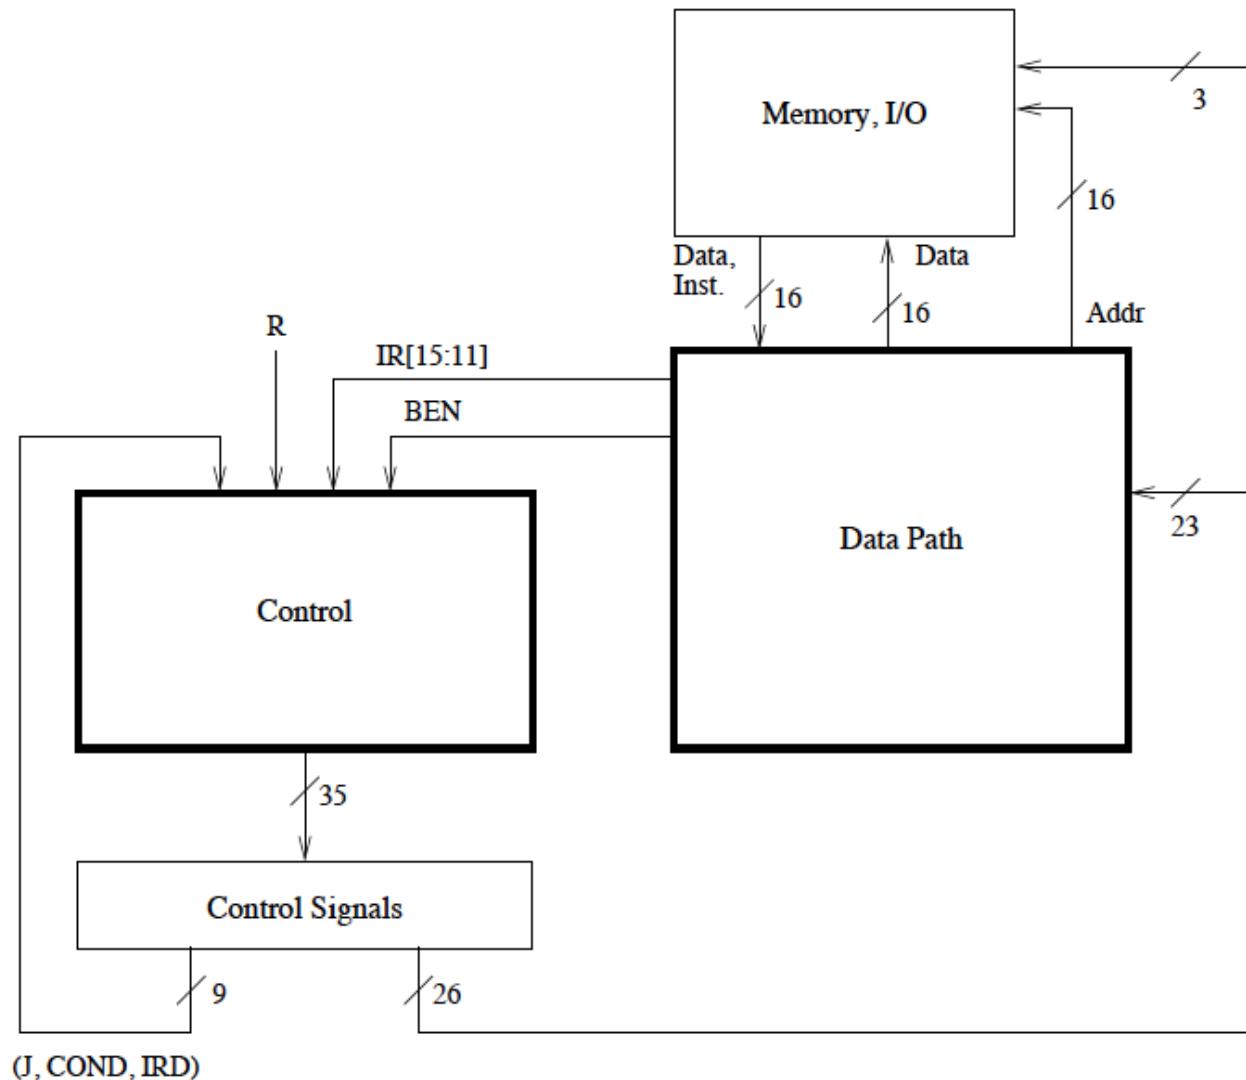


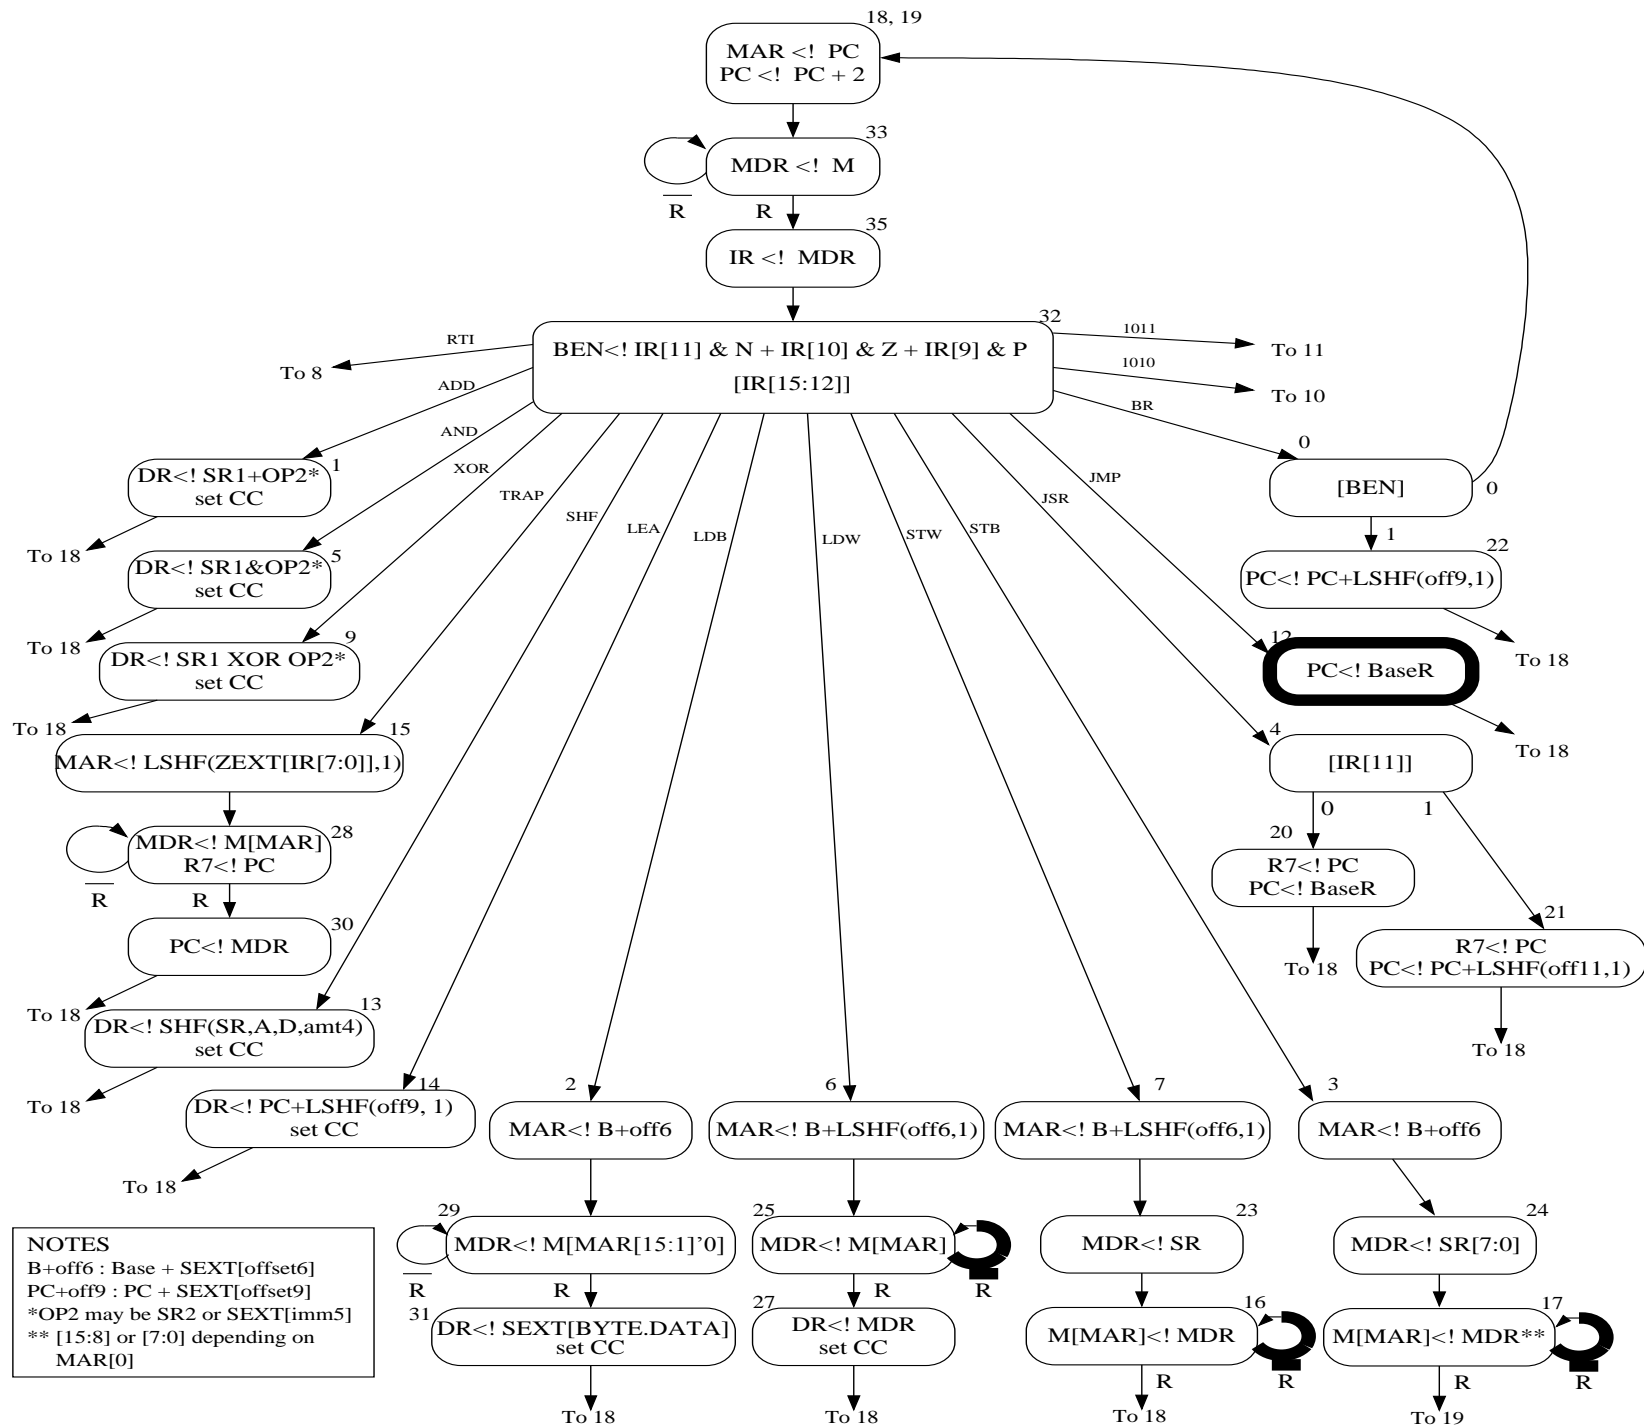
Figure C.1: Microarchitecture of the LC-3b, major components

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 → state 33 → state 35
 - Decode phase: state 32

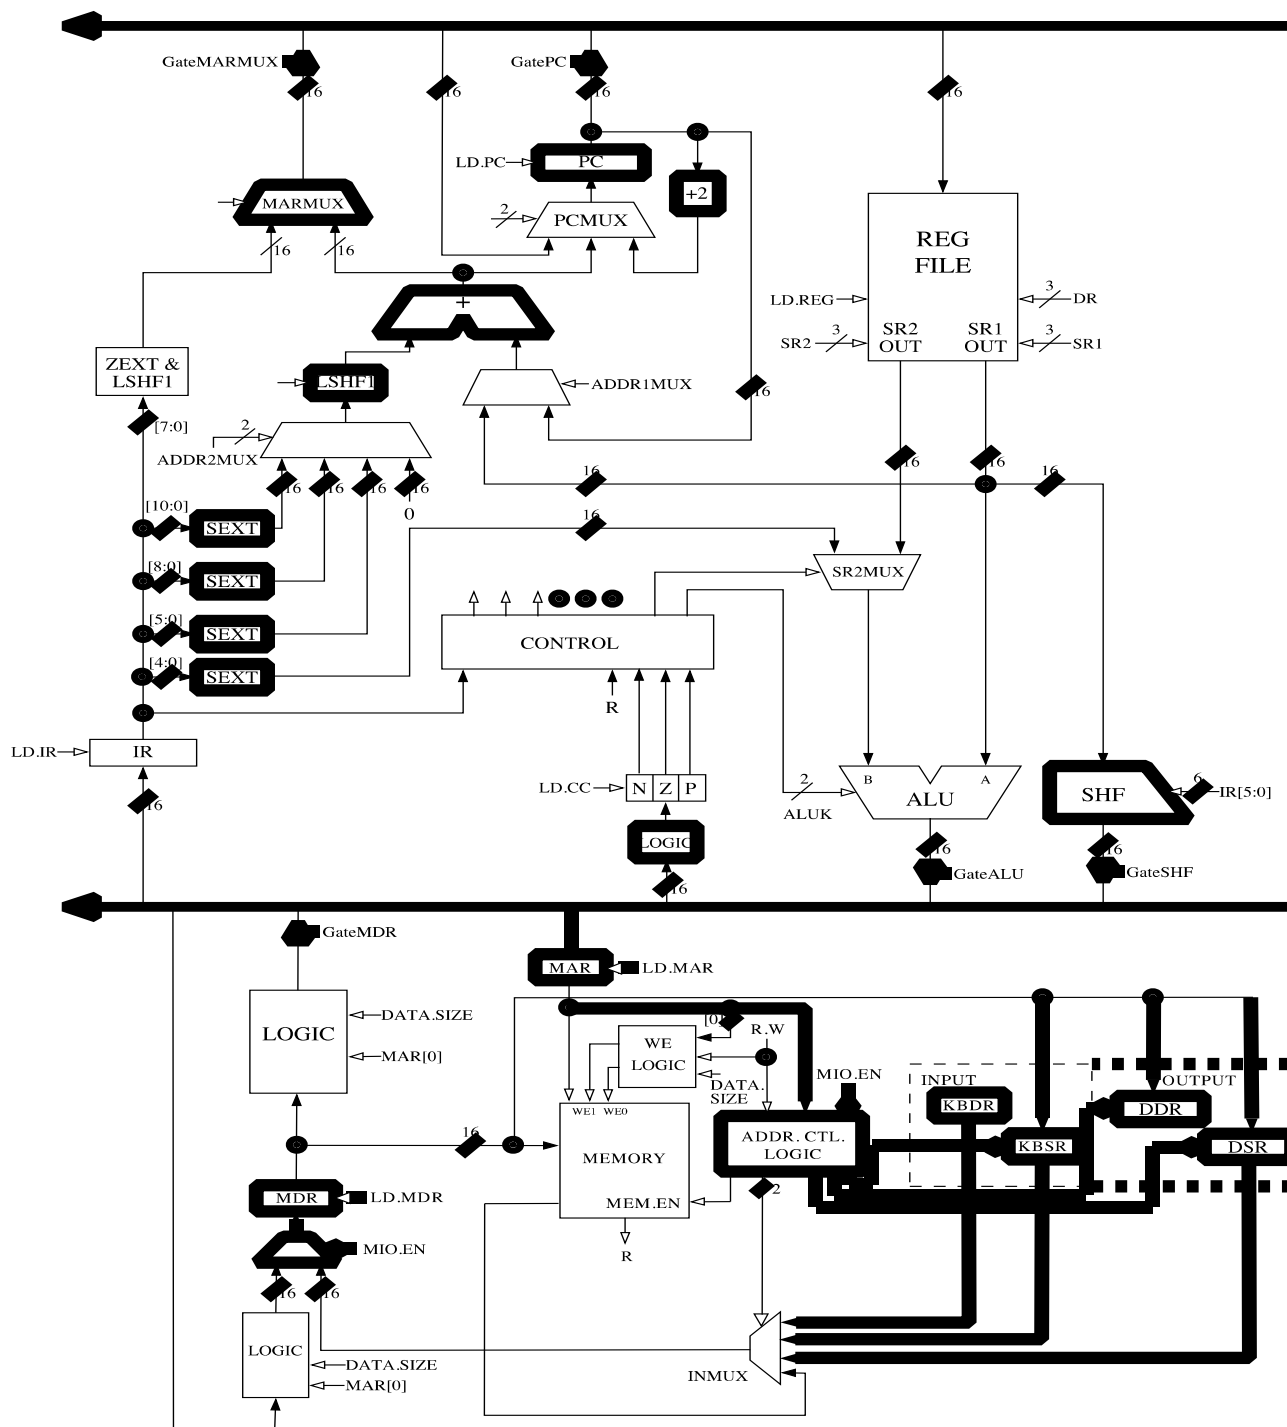


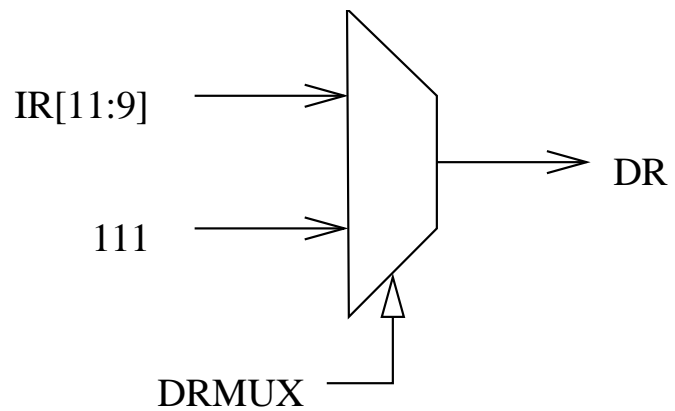
LC-3b State Machine: Some Questions

- 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 time?

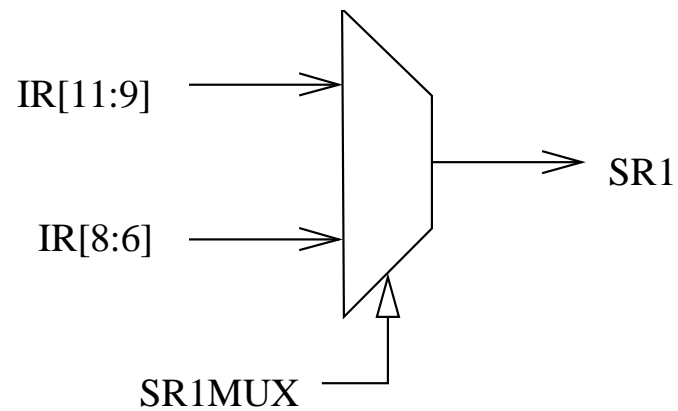
LC-3b Datapath

- 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






(a)



(b)



(c)

Signal Name	Signal Values	
LD.MAR/1:	NO, LOAD	
LD.MDR/1:	NO, LOAD	
LD.IR/1:	NO, LOAD	
LD.BEN/1:	NO, LOAD	
LD.REG/1:	NO, LOAD	
LD.CC/1:	NO, LOAD	
LD.PC/1:	NO, LOAD	
GatePC/1:	NO, YES	
GateMDR/1:	NO, YES	
GateALU/1:	NO, YES	
GateMARMUX/1:	NO, YES	
GateSHF/1:	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, XOR, PASSA	
MIO.EN/1:	NO, YES	
R.W/1:	RD, WR	
DATA.SIZE/1:	BYTE, WORD	
LSHF1/1:	NO, YES	

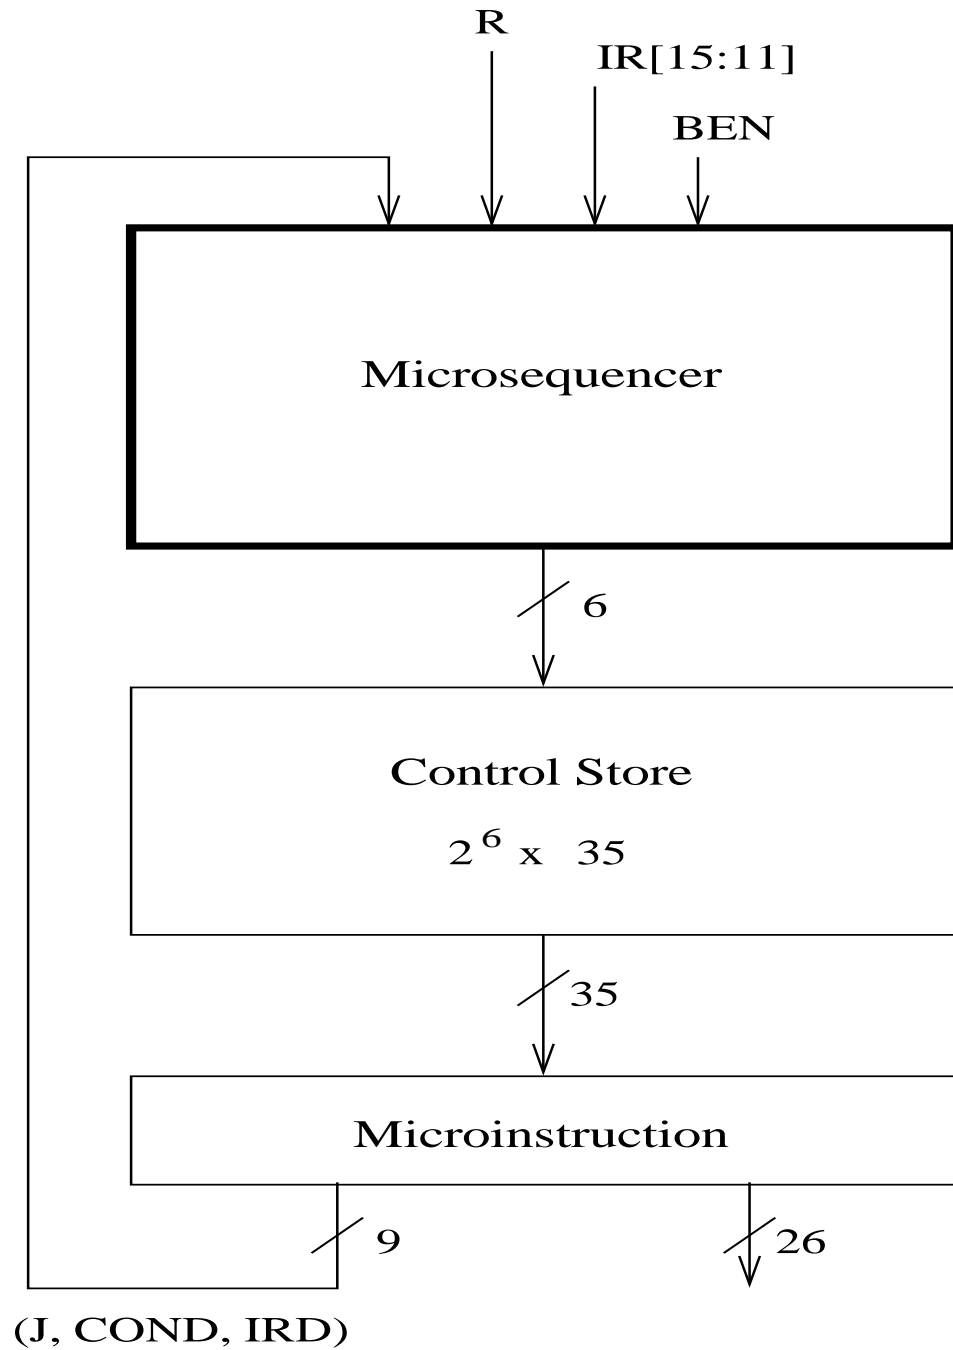
Table C.1: Data path control signals

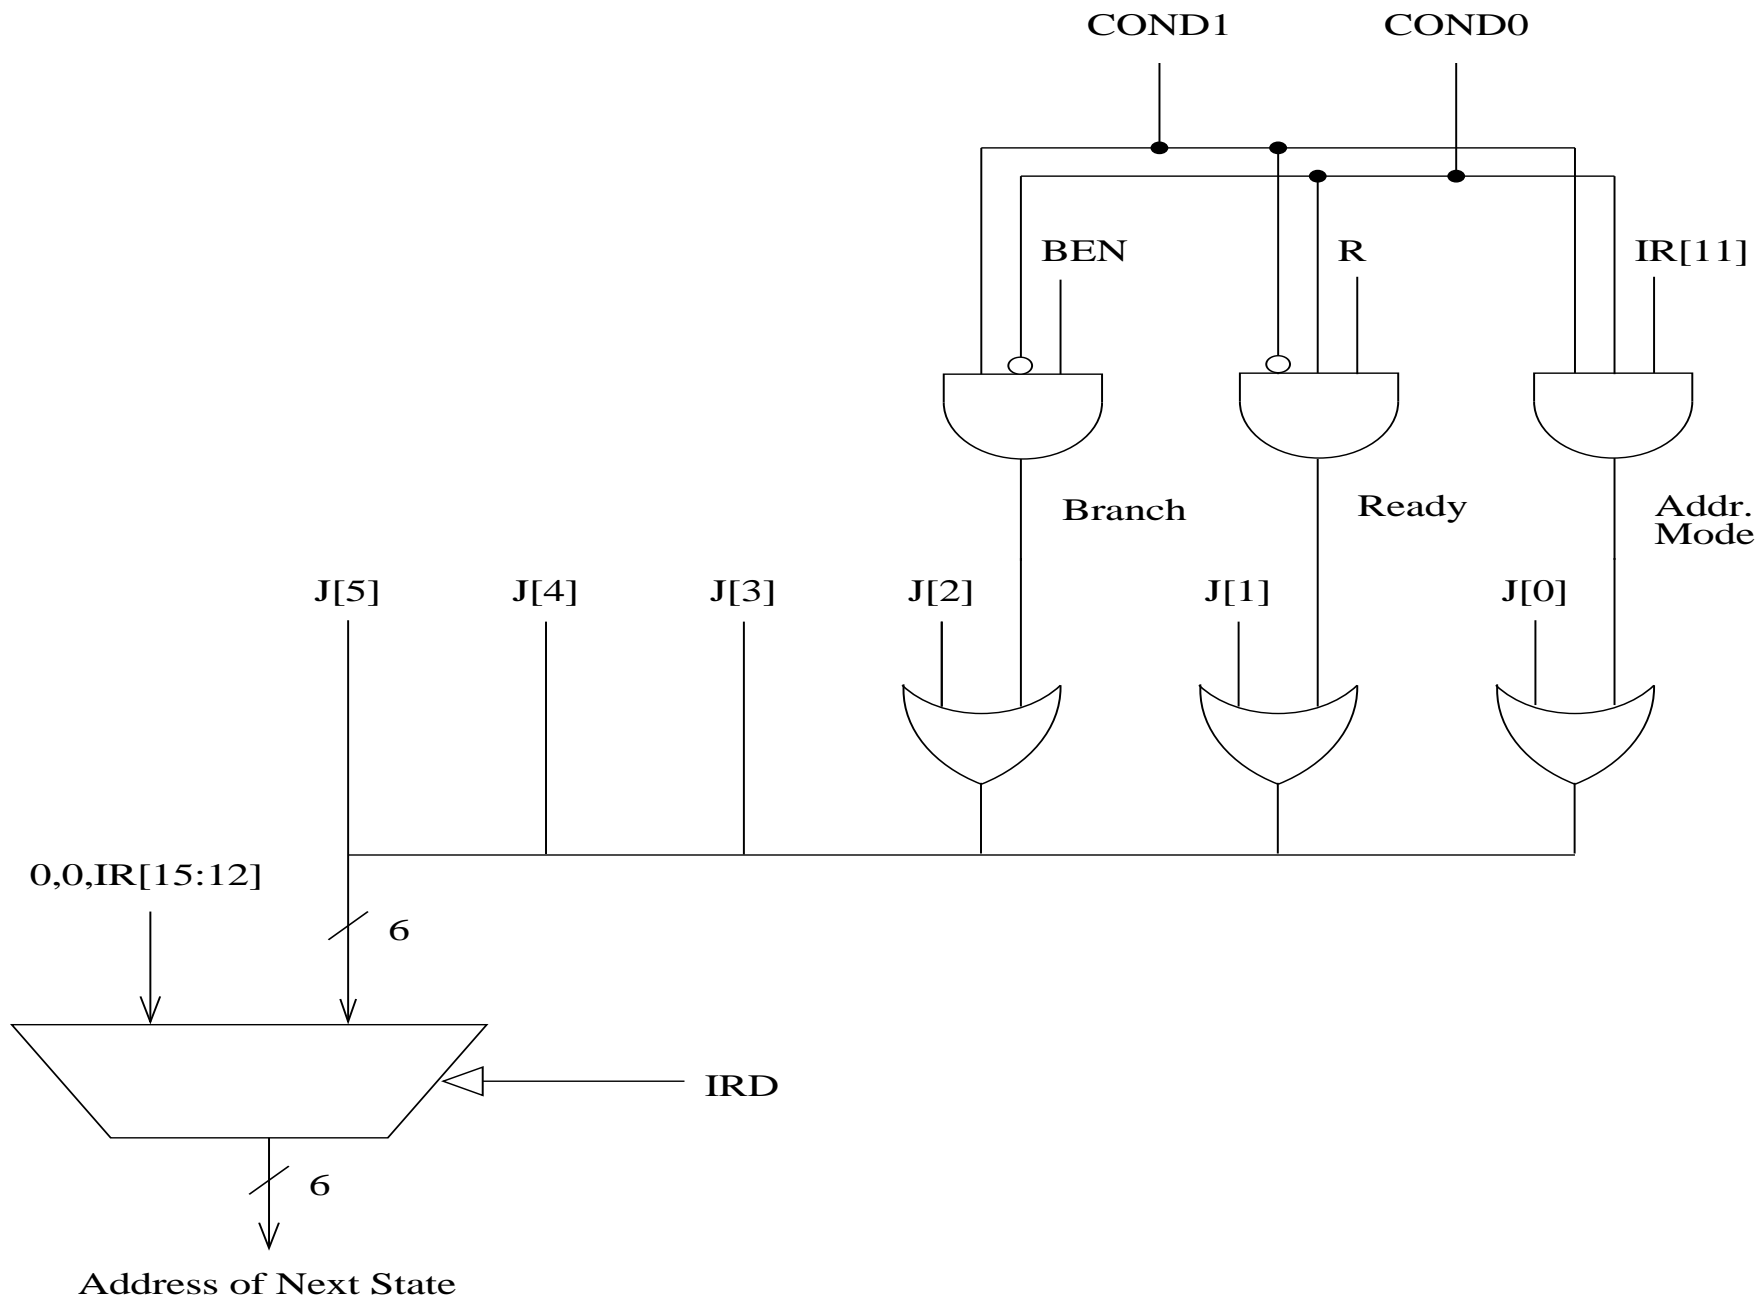
LC-3b Datapath: Some Questions

- 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 help 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)





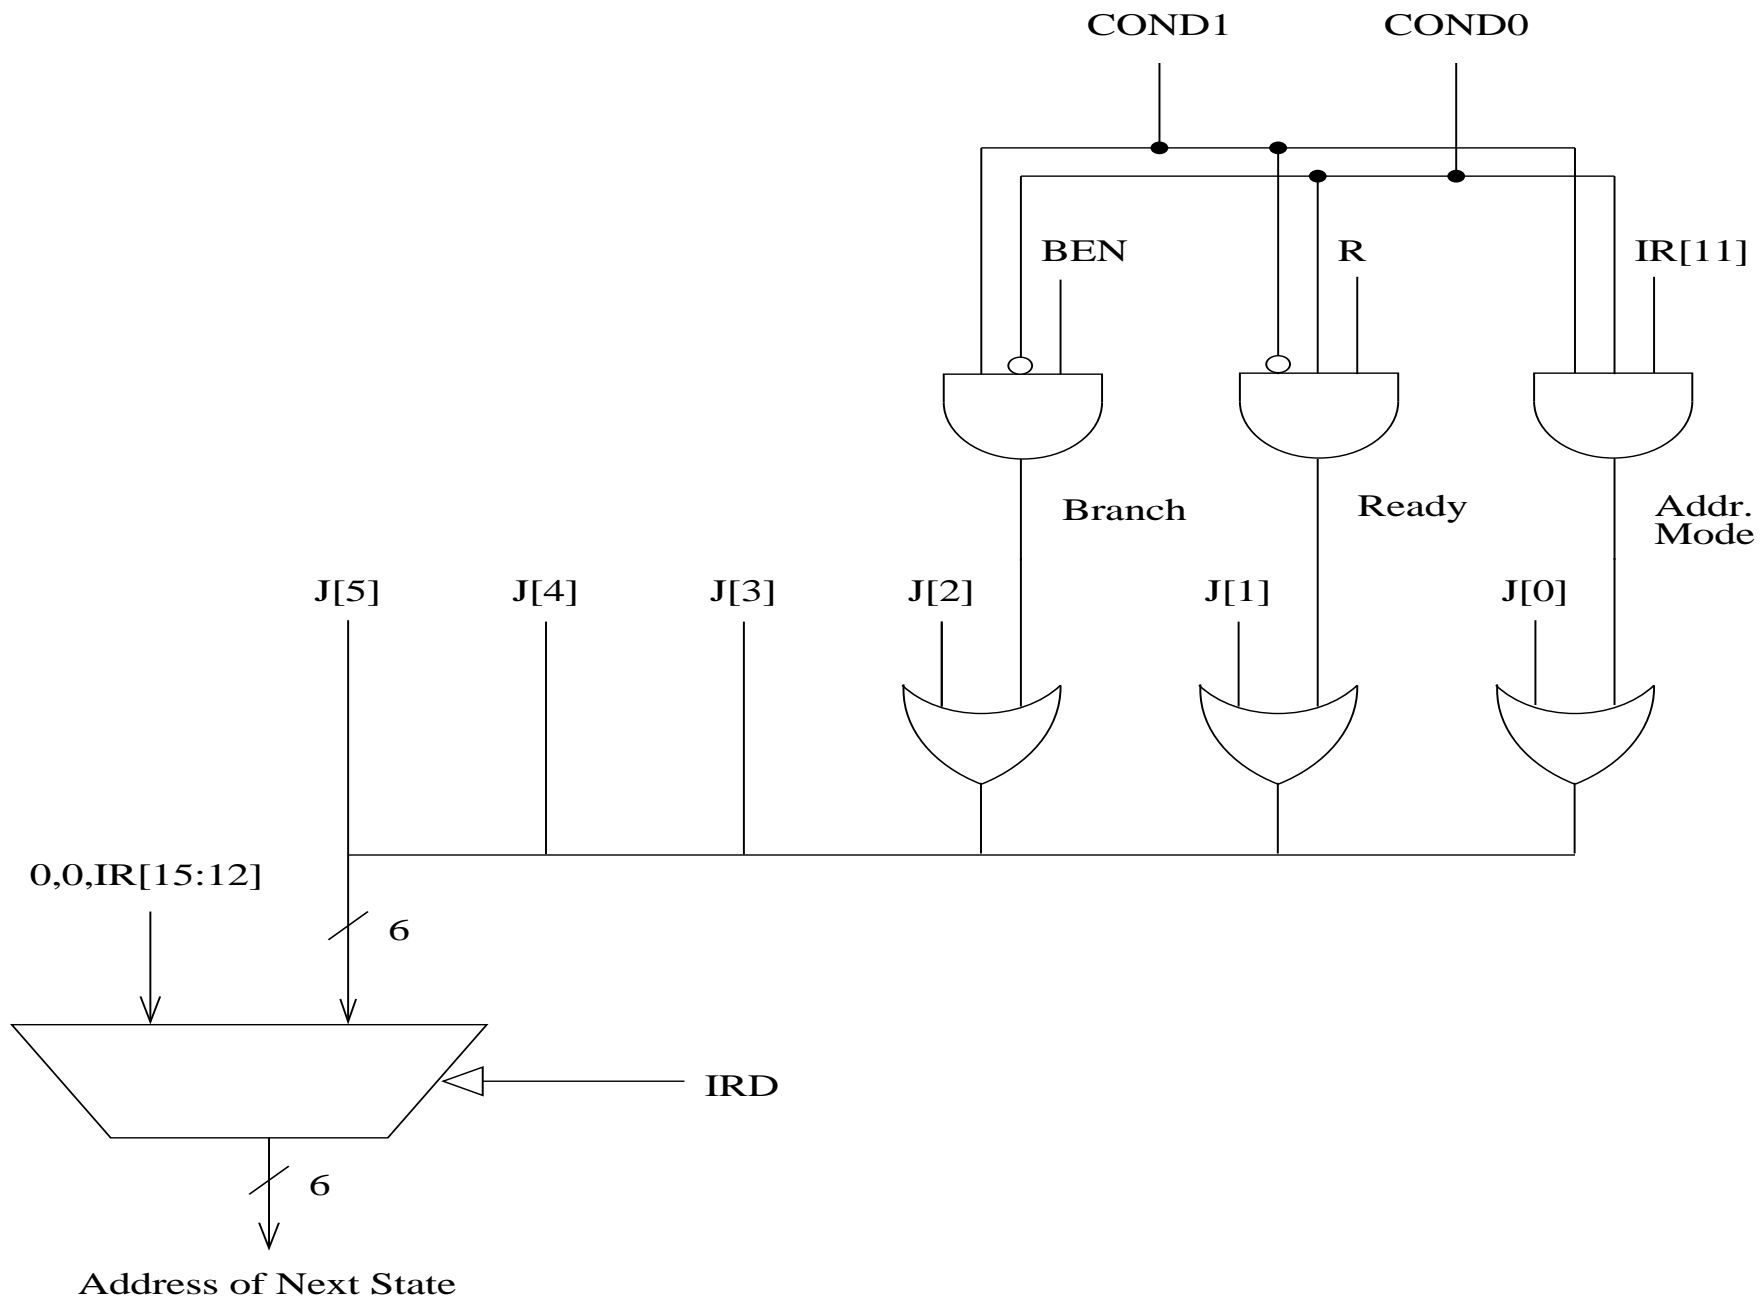
[illegible]

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 Values
J/6:	
COND/2:	COND ₀ ;Unconditional
	COND ₁ ;Memory Ready
	COND ₂ ;Branch
	COND ₃ ;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/s14/doku.php?id=techdocs>
- <http://www.ece.cmu.edu/~ece447/s14/lib/exe/fetch.php?media=lc3b-figures.pdf>

A Simple LC-3b Control and Datapath

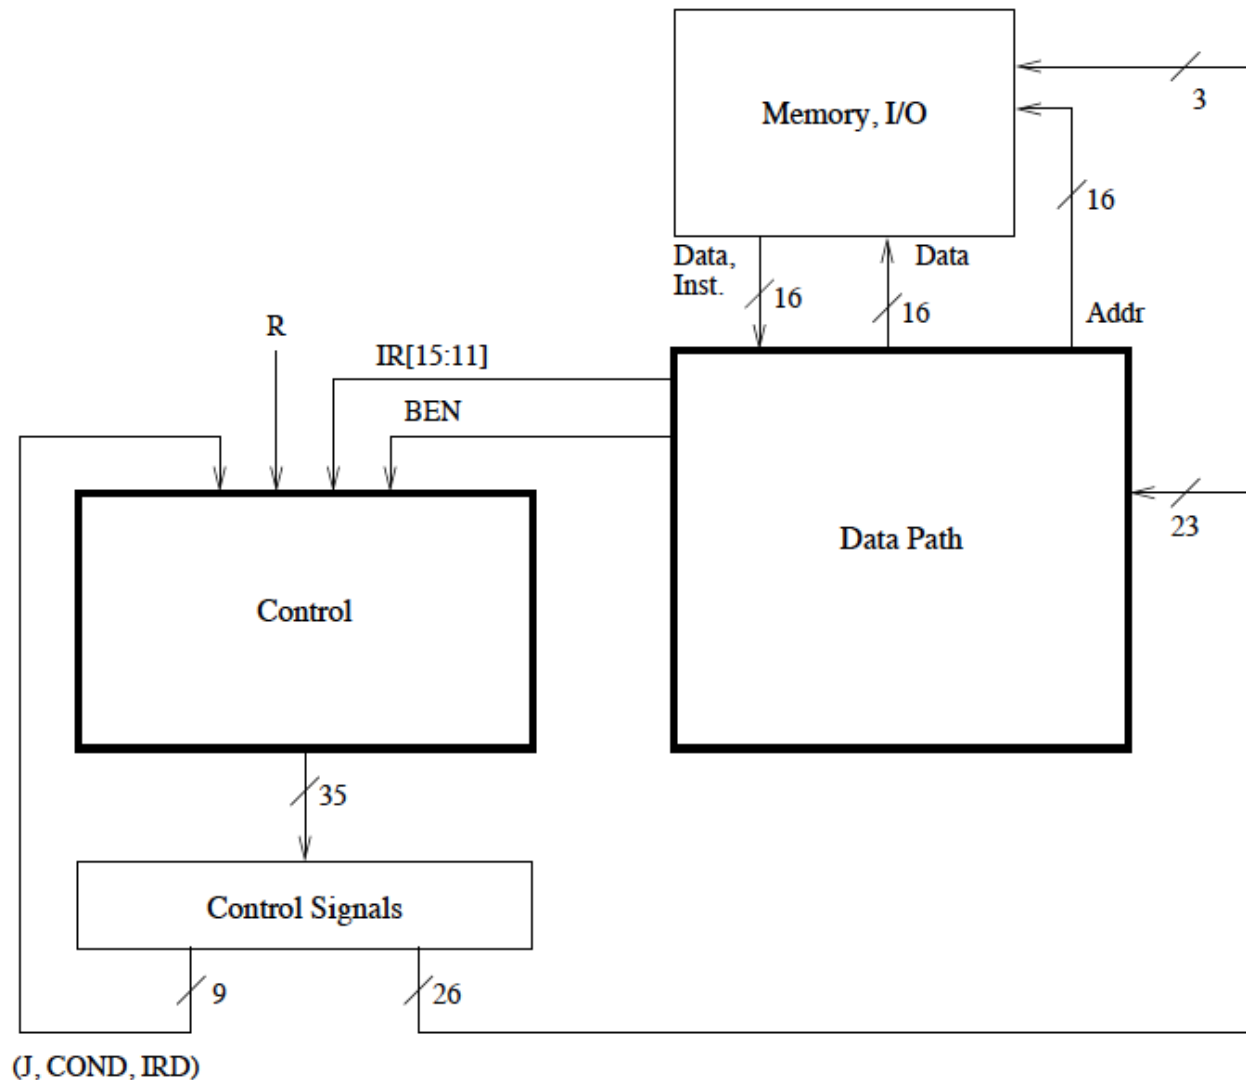
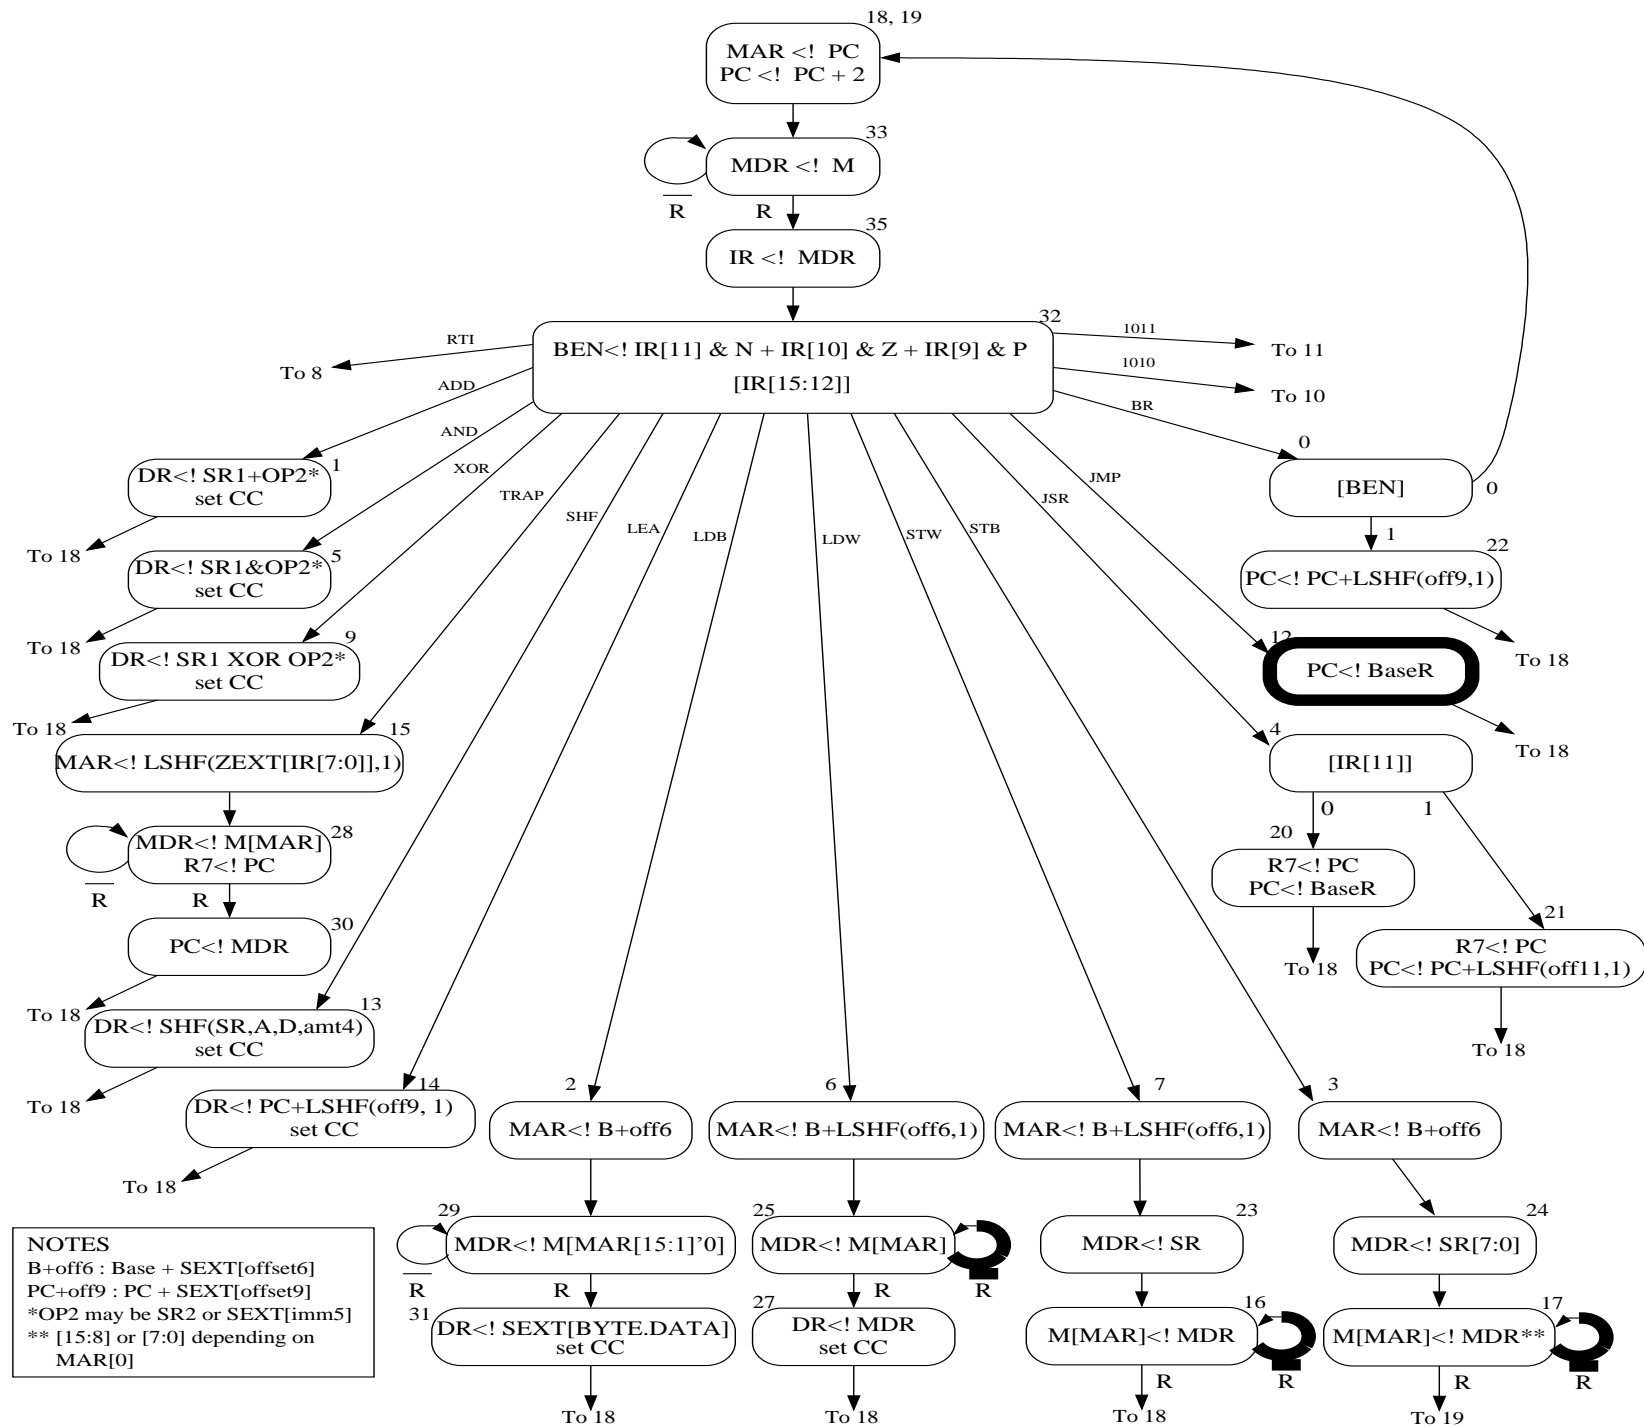
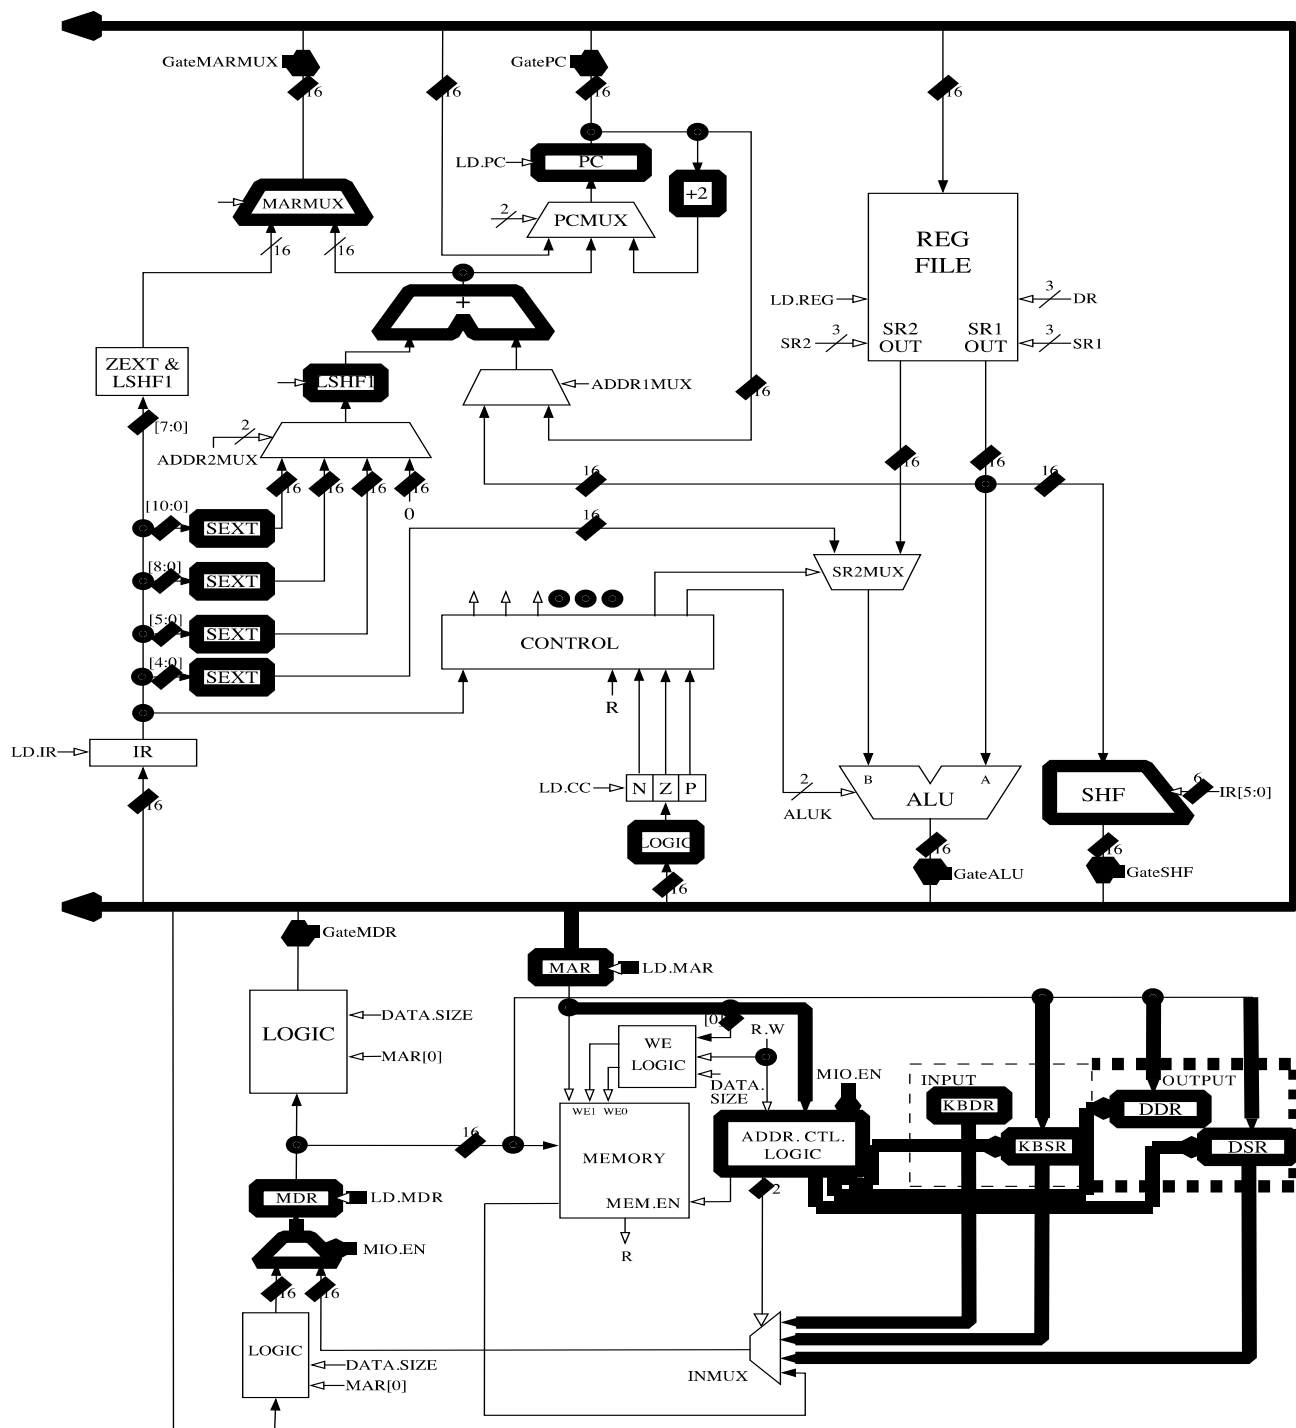
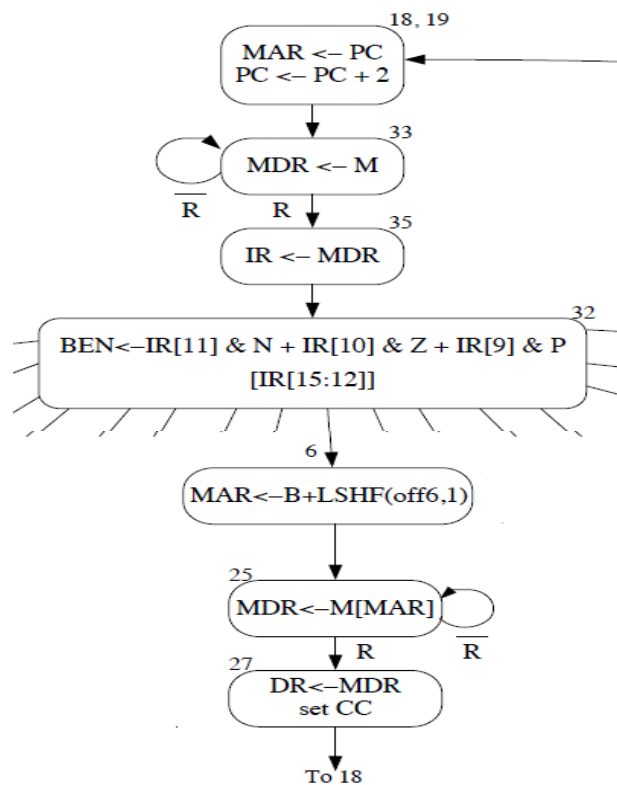


Figure C.1: Microarchitecture of the LC-3b, major components

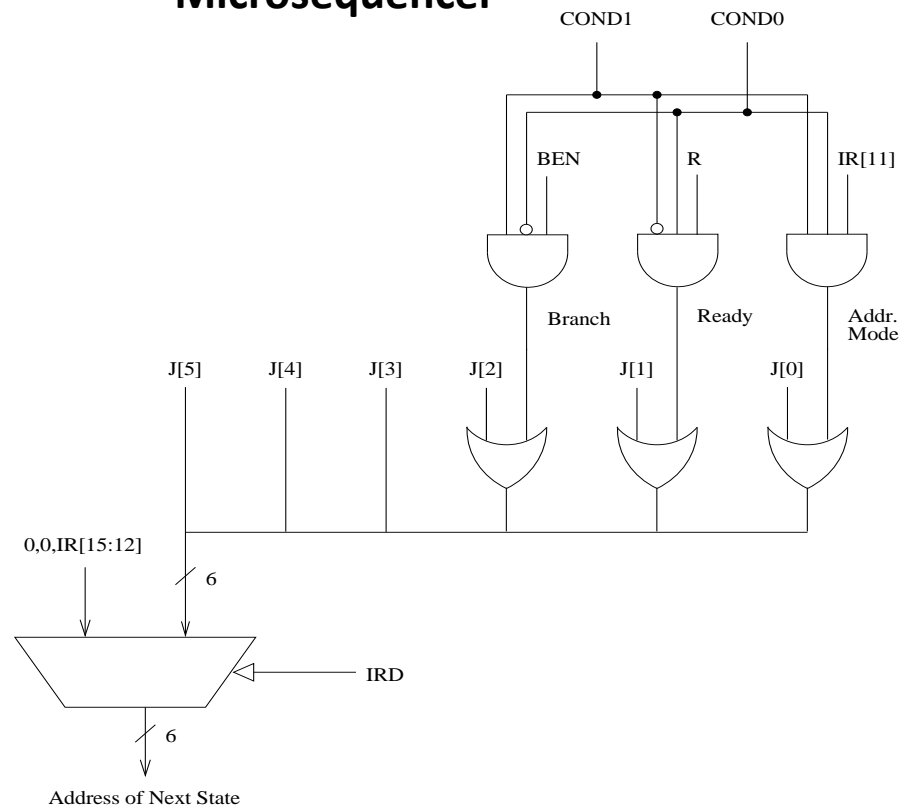





State Machine for LDW




Microsequencer

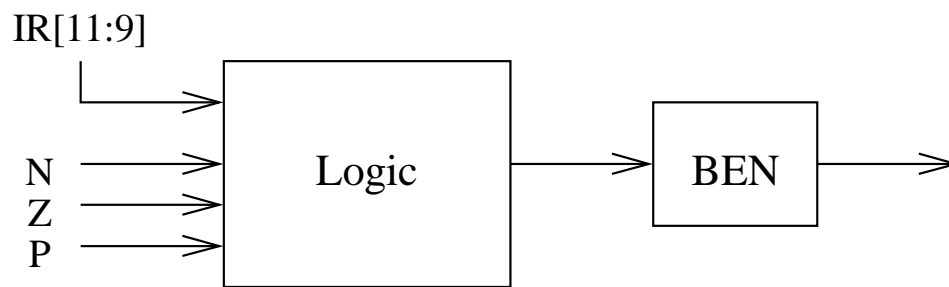
[illegible]



(a)



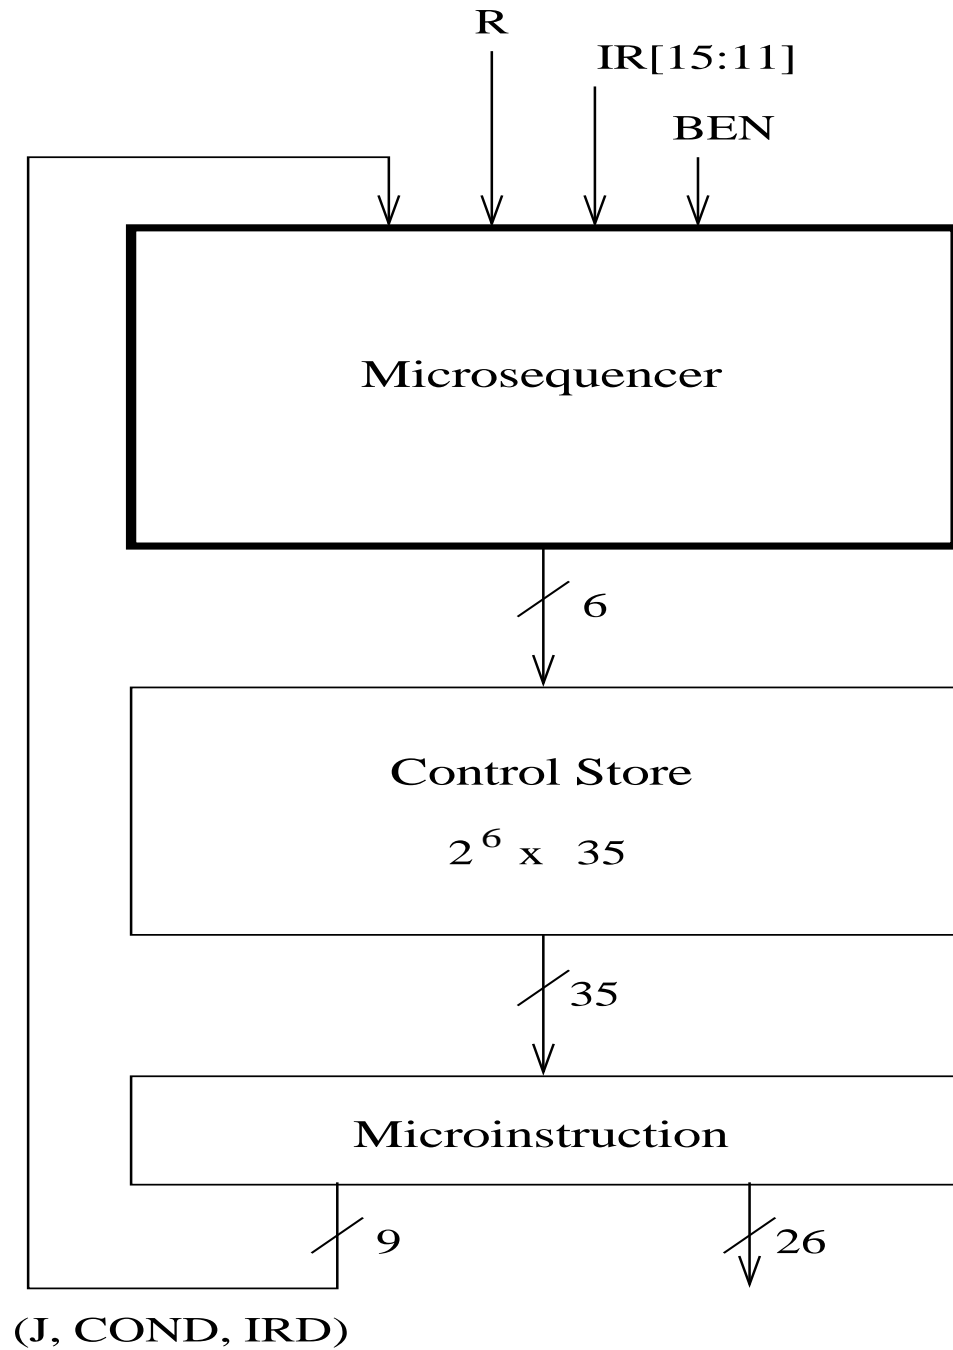
(b)

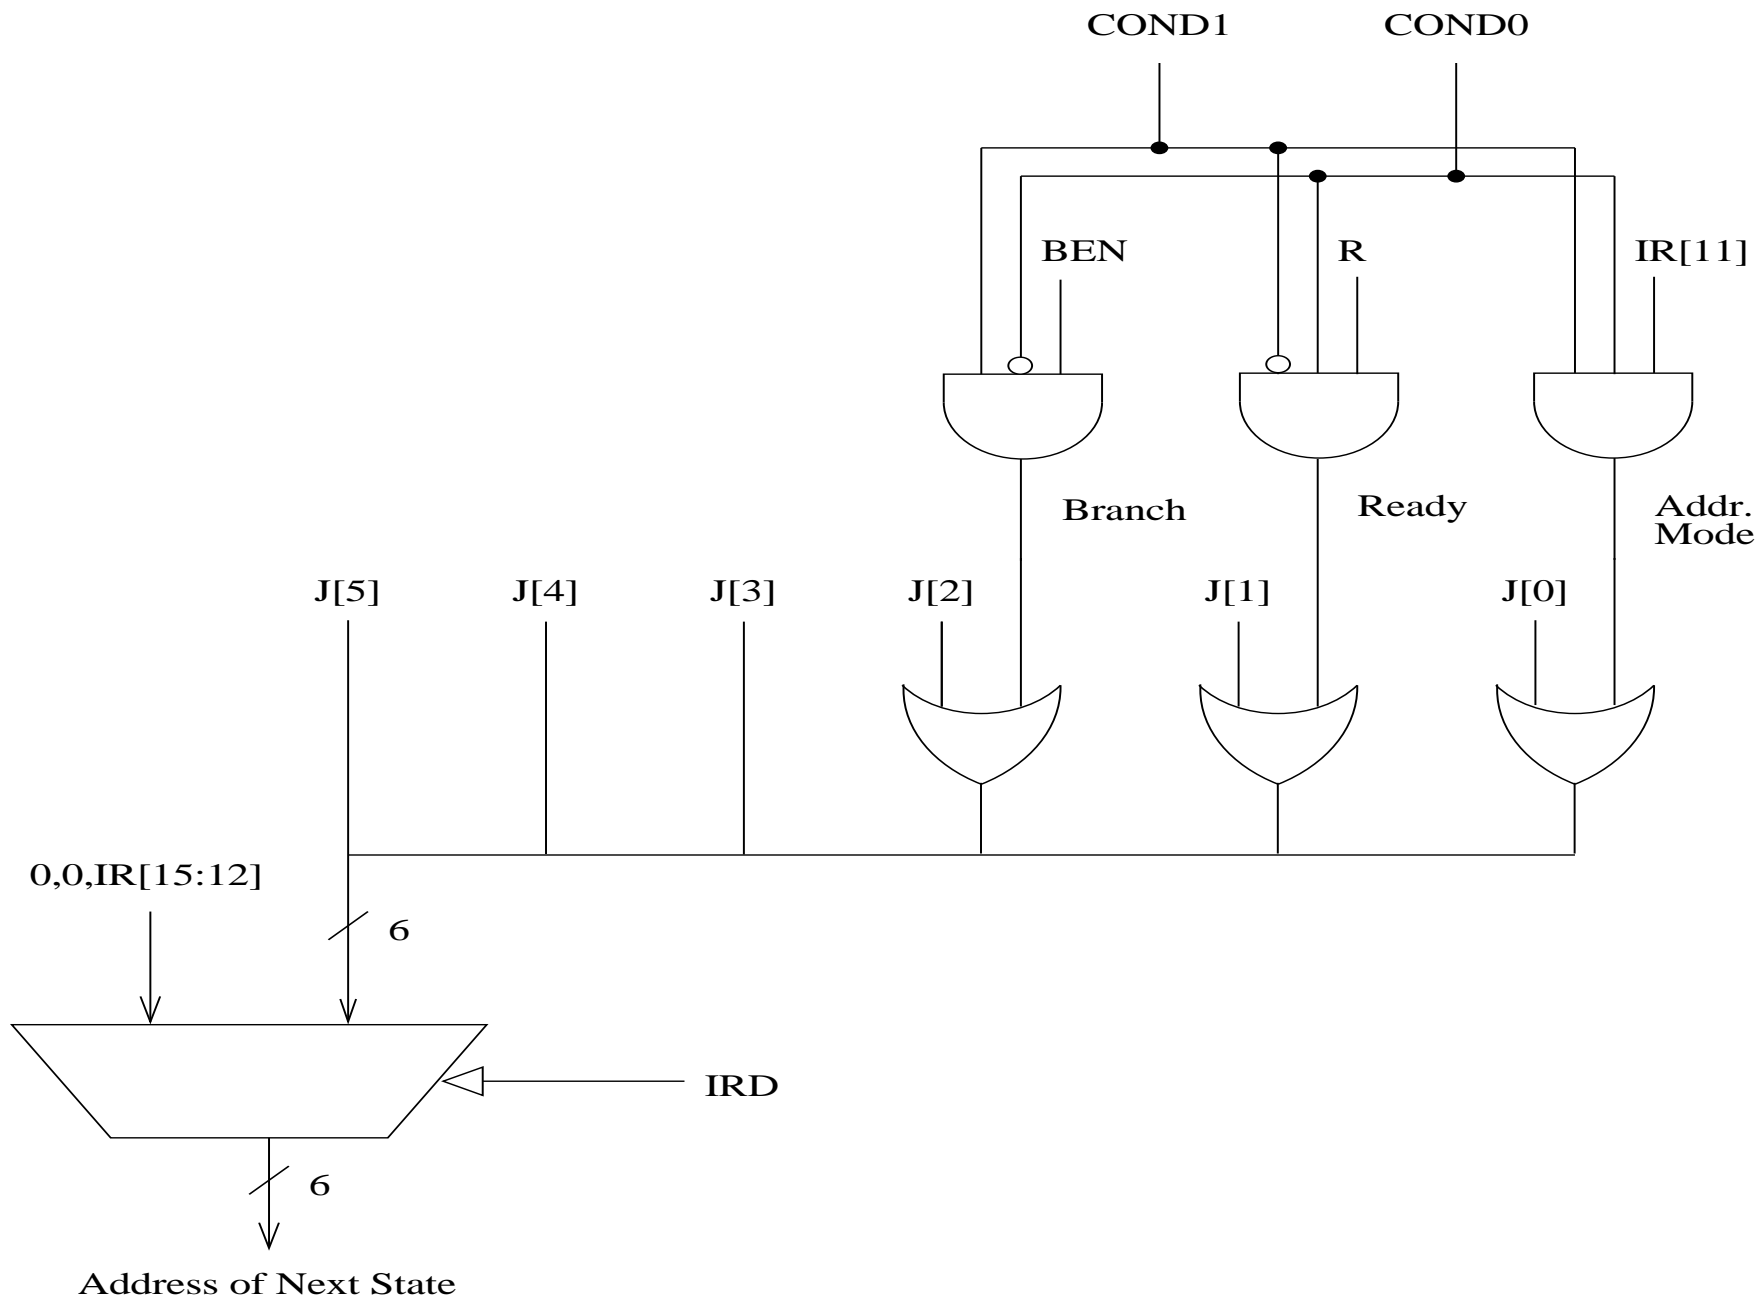


(c)

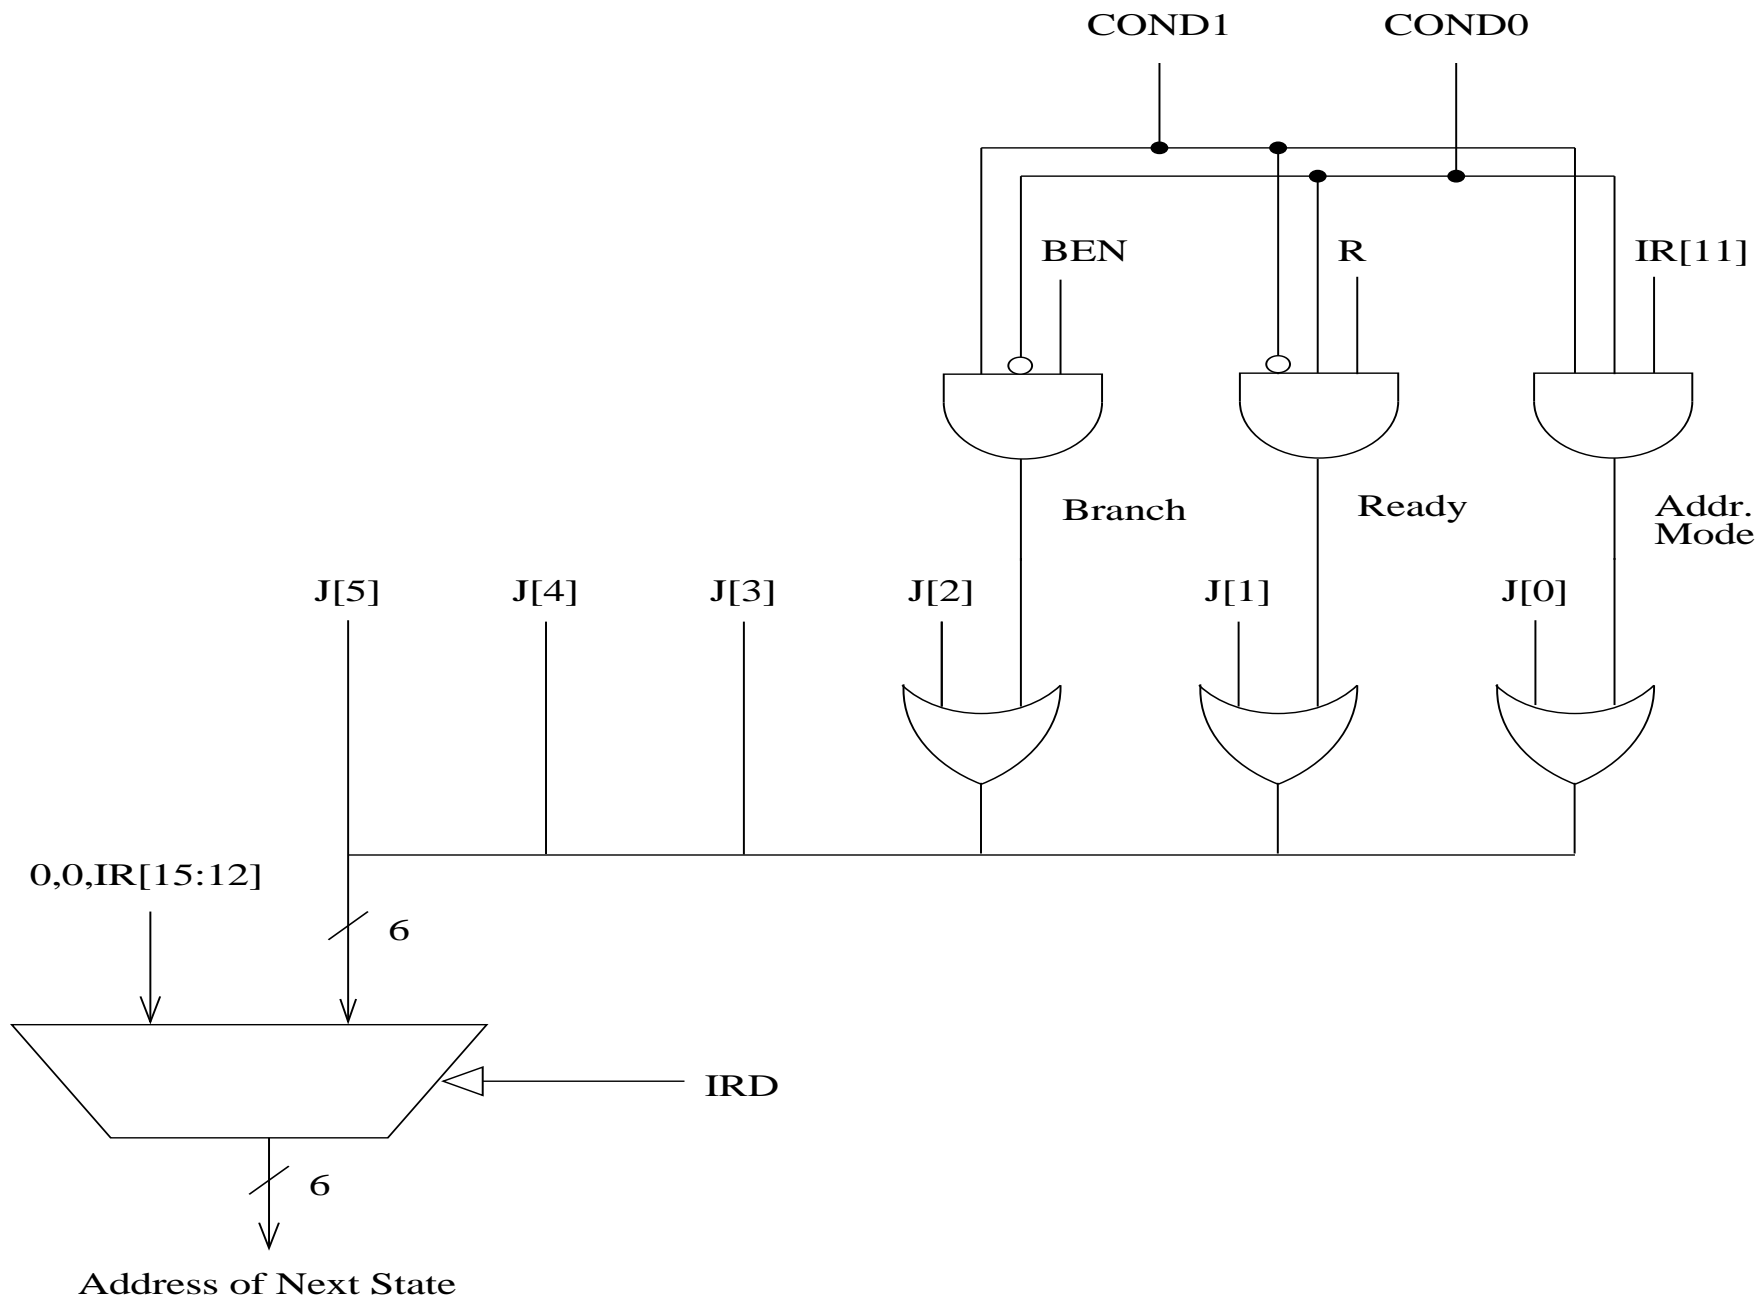
Signal Name	Signal Values	
LD.MAR/1:	NO, LOAD	
LD.MDR/1:	NO, LOAD	
LD.IR/1:	NO, LOAD	
LD.BEN/1:	NO, LOAD	
LD.REG/1:	NO, LOAD	
LD.CC/1:	NO, LOAD	
LD.PC/1:	NO, LOAD	
GatePC/1:	NO, YES	
GateMDR/1:	NO, YES	
GateALU/1:	NO, YES	
GateMARMUX/1:	NO, YES	
GateSHF/1:	NO, YES	
PCMUX/2:	PC+2 BUS ADDER	;select pc+2 ;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, XOR, PASSA	
MIO.EN/1:	NO, YES	
R.W/1:	RD, WR	
DATA.SIZE/1:	BYTE, WORD	
LSHF1/1:	NO, YES	

Table C.1: Data path control signals





[illegible]



End of the Exercise in Microprogramming

Homework 2

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

Lab 2 Extra Credit

- Microprogrammed ARM implementation
- Exercise your creativity!

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

The Control Store: Some Questions

- What control signals can be stored in the control store?

vs.

- What control signals have to be generated in hardwired logic?
 - i.e., what signal cannot be available without processing in the datapath?

Variable-Latency Memory

- The ready signal (R) enables memory read/write to execute correctly
 - Example: transition from state 33 to state 35 is controlled by the R bit asserted by memory when memory data is available
- Could we have done this in a single-cycle microarchitecture?

The Microsequencer: Advanced Questions

- What happens if the machine is interrupted?
- What if an instruction generates an exception?
- How can you implement a complex instruction using this control structure?
 - Think REP MOVS

The Power of Abstraction

- The concept of a control store of microinstructions enables the hardware designer with a new abstraction:
microprogramming
- The designer can translate any desired operation to a sequence microinstructions
- All the designer needs to provide is
 - The sequence of microinstructions needed to implement the desired operation
 - The ability for the control logic to correctly sequence through the microinstructions
 - Any additional datapath control signals needed (no need if the operation can be “translated” into existing control signals)

Let's Do Some More Microprogramming

- Implement REP MOVS in the LC-3b microarchitecture
- What changes, if any, do you make to the
 - state machine?
 - datapath?
 - control store?
 - microsequencer?
- Show all changes and microinstructions
- Coming up in Homework 3?

Aside: Alignment Correction in Memory

- Remember unaligned accesses
- LC-3b has byte load and byte store instructions that move data not aligned at the word-address boundary
 - Convenience to the programmer/compiler
- How does the hardware ensure this works correctly?
 - Take a look at state 29 for LDB
 - States 24 and 17 for STB
 - Additional logic to handle unaligned accesses

Aside: Memory Mapped I/O

- Address control logic determines whether the specified address of LDx and STx are to memory or I/O devices
- Correspondingly enables memory or I/O devices and sets up muxes
- Another instance where the final control signals (e.g., MEM.EN or INMUX/2) cannot be stored in the control store
 - Dependent on address

Advantages of Microprogrammed Control

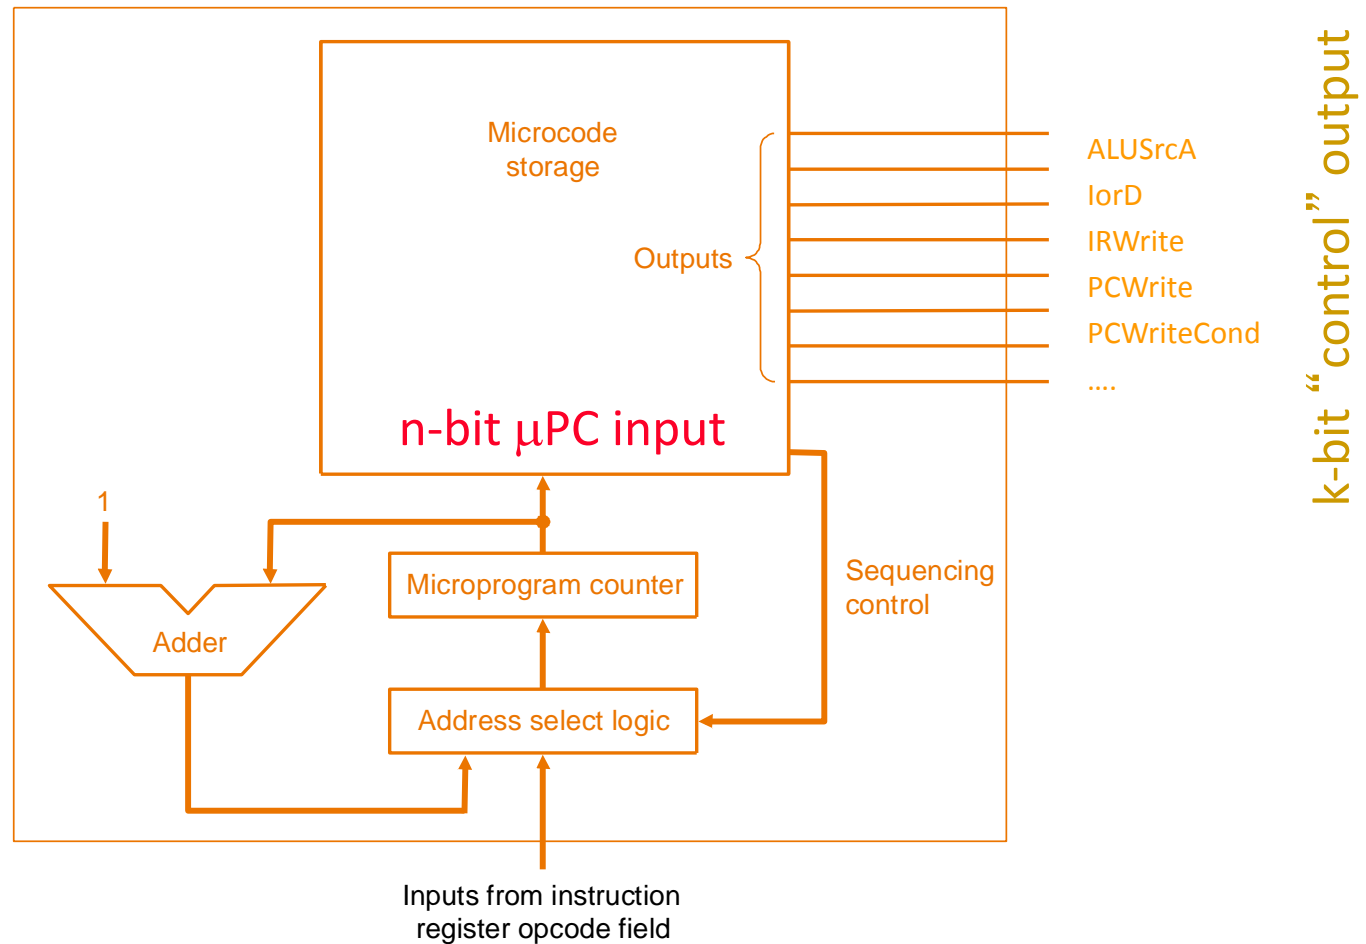
- Allows a very simple design to do powerful computation by controlling the datapath (using a sequencer)
 - High-level ISA translated into microcode (sequence of microinstructions)
 - Microcode (ucode) enables a minimal datapath to emulate an ISA
 - Microinstructions can be thought of a user-invisible ISA
- Enables easy extensibility of the ISA
 - Can support a new instruction by changing the ucode
 - Can support complex instructions as a sequence of simple microinstructions
- If I can sequence an arbitrary instruction then I can sequence an arbitrary “program” as a microprogram sequence
 - will need some new state (e.g. loop counters) in the microcode for sequencing more elaborate programs

Update of Machine Behavior

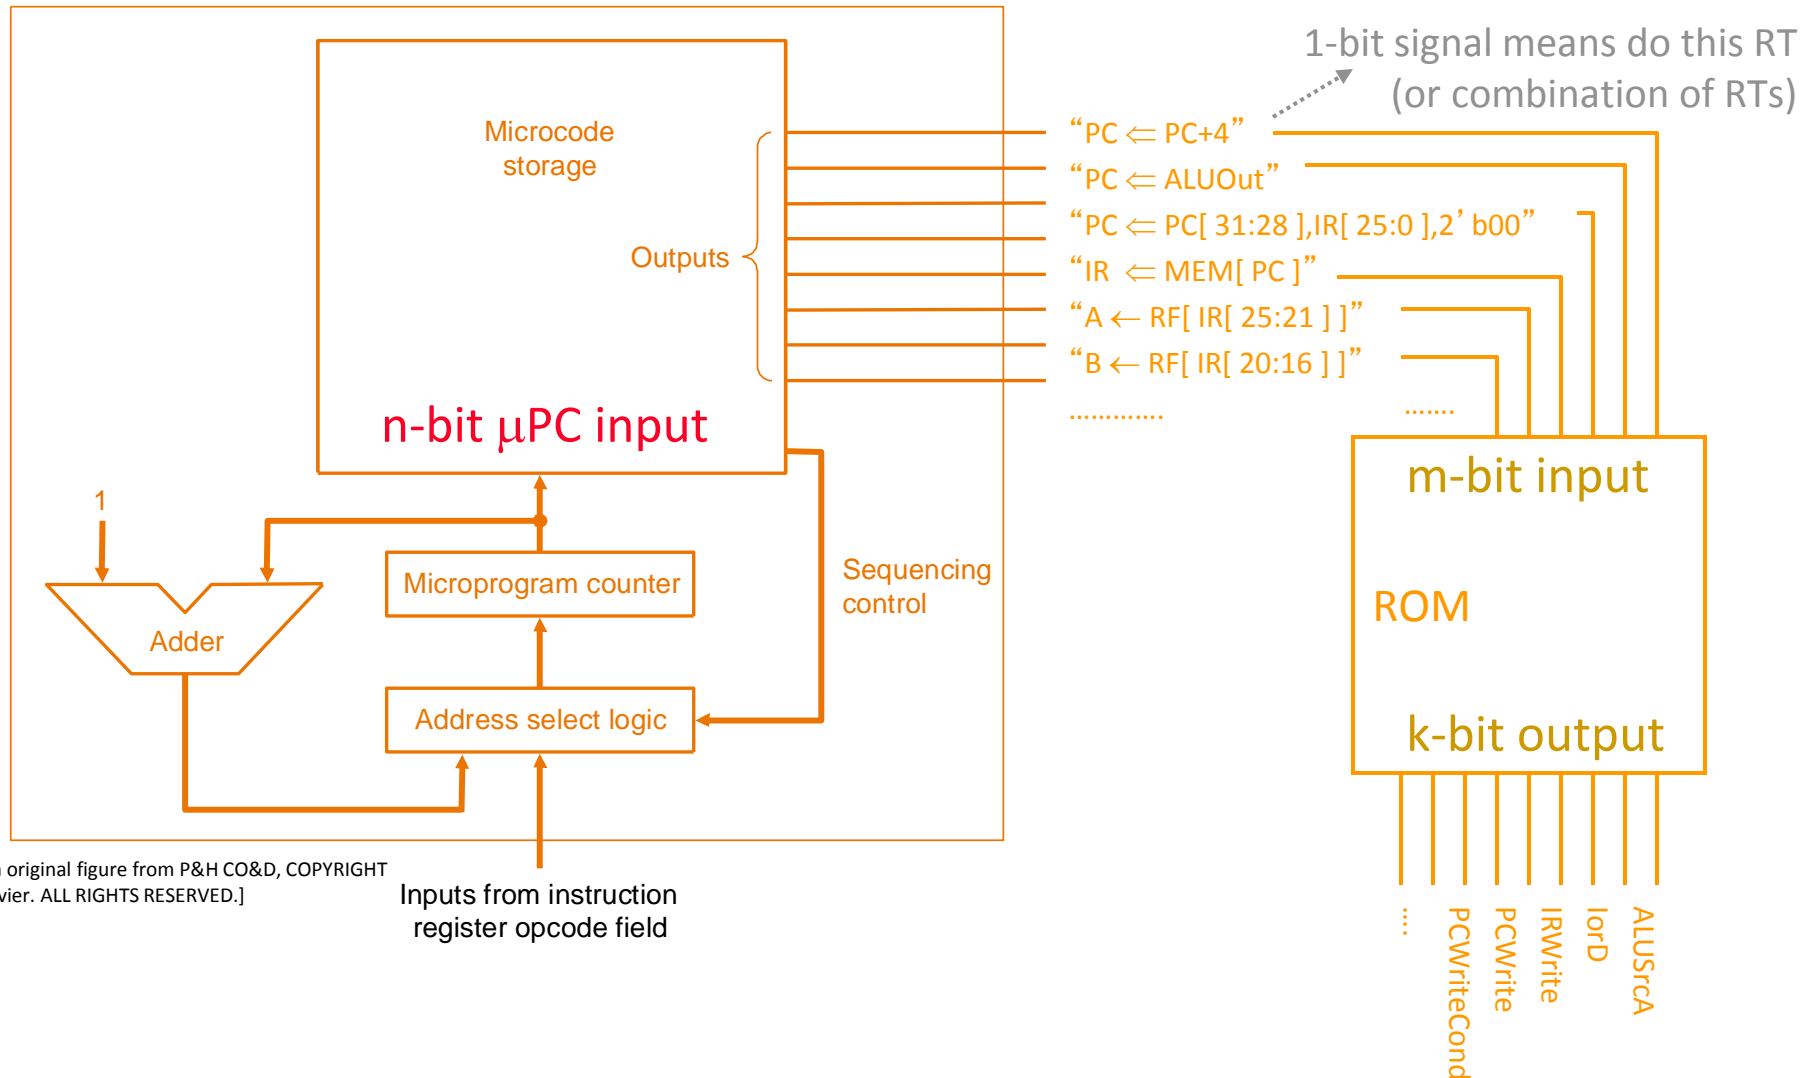
- The ability to update/patch microcode in the field (after a processor is shipped) enables
 - Ability to add new instructions without changing the processor!
 - Ability to “fix” buggy hardware implementations

- Examples
 - IBM 370 Model 145: microcode stored in main memory, can be updated after a reboot
 - IBM System z: Similar to 370/145.
 - Heller and Farrell, “[Millicode in an IBM zSeries processor](#),” IBM JR&D, May/Jul 2004.
 - B1700 microcode can be updated while the processor is running
 - User-microprogrammable machine!

Horizontal Microcode



Vertical Microcode



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

Inputs from instruction register opcode field


If done right (i.e., $m \ll n$, and $m \ll k$), two ROMs together ($2^n \times m + 2^m \times k$ bit) should be smaller than horizontal microcode ROM ($2^n \times k$ bit)

Nanocode and Millicode

- *Nanocode*: a level below traditional μ code
 - μ programmed control for sub-systems (e.g., a complicated floating-point module) that acts as a slave in a μ controlled datapath
- *Millicode*: a level above traditional μ code
 - ISA-level subroutines that can be called by the μ controller to handle complicated operations and system functions
 - E.g., Heller and Farrell, “[Millicode in an IBM zSeries processor](#),” IBM JR&D, May/Jul 2004.
- In both cases, we avoid complicating the main μ controller
- You can think of these as “microcode” at different levels of abstraction


Nanocode Concept Illustrated

a “ μ coded” processor implementation



We refer to this as “nanocode” when a μ coded subsystem is embedded in a μ coded system

a “ μ coded” FPU implementation



Multi-Cycle vs. Single-Cycle uArch

- Advantages
- Disadvantages
- You should be very familiar with this right now

Microprogrammed vs. Hardwired Control

- Advantages
- Disadvantages
- You should be very familiar with this right now

Can We Do Better?

- What limitations do you see with the multi-cycle design?
- Limited concurrency
 - Some hardware resources are idle during different phases of instruction processing cycle
 - “Fetch” logic is idle when an instruction is being “decoded” or “executed”
 - Most of the datapath is idle when a memory access is happening