

18-447 Lecture 9: Microcontrolled Multi-Cycle Implementations

James C. Hoe

Dept of ECE, CMU

February 18, 2009

Announcements: Read P&H Appendix D
Get started on Lab

Handouts: Handout #8: Project 1 (on Blackboard)

Single-Cycle Implementations

- ◆ Matches naturally with the sequential and atomic semantics inherent to most ISAs
 - instantiate programmer-visible state one-for-one
 - map instructions to combinational next-state logic
- ◆ But, contrived and inefficient
 - all instructions run as slow as the slowest instruction
 - must provide worst-case combinational resource in parallel as required by any instruction
- ◆ Not necessarily the simplest way to implement an ISA

Imagine trying to build a single-cycle implementation of a CISC ISA

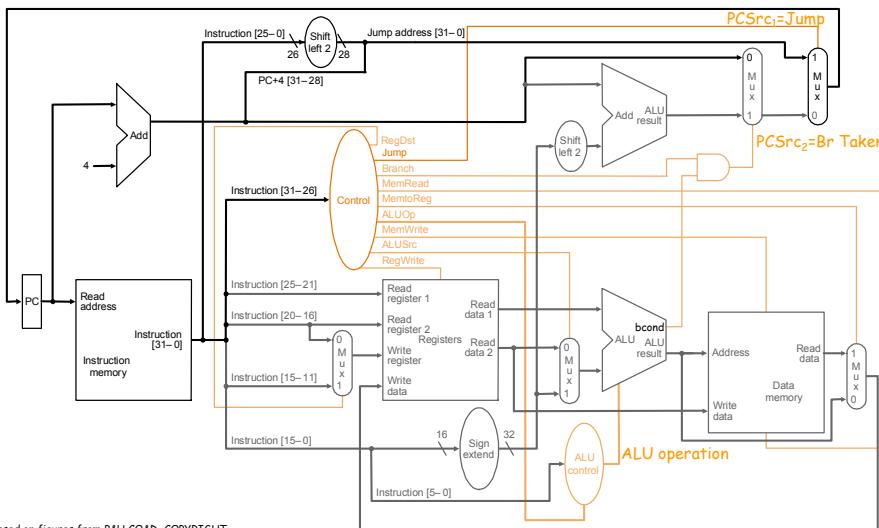
Single-Cycle Datapath Analysis

◆ 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	Delay
resources	mem	RF	ALU	mem	RF	
R-type	200	50	100		50	400
I-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

Worksheet ▶



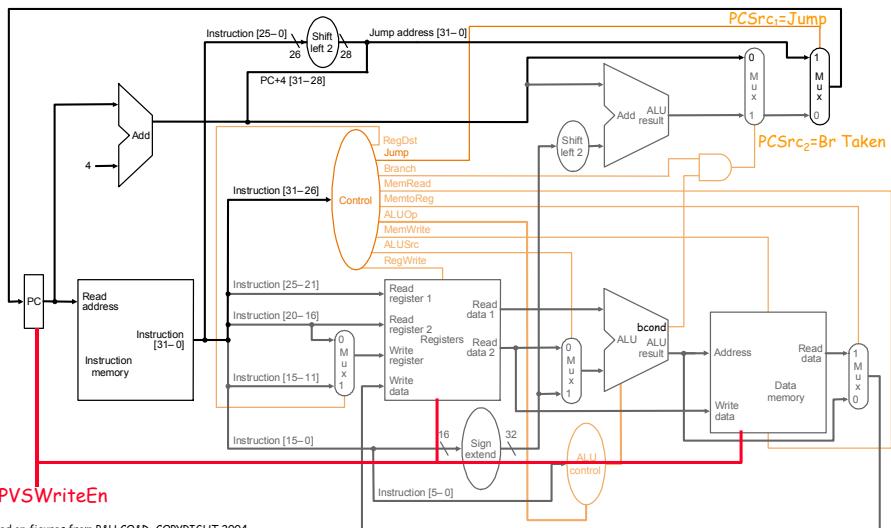
[Based on figures from P&H CO4D, COPYRIGHT 2004 Elsevier. ALL RIGHTS RESERVED.]

Multi-cycle Implementation: Ver 1.0

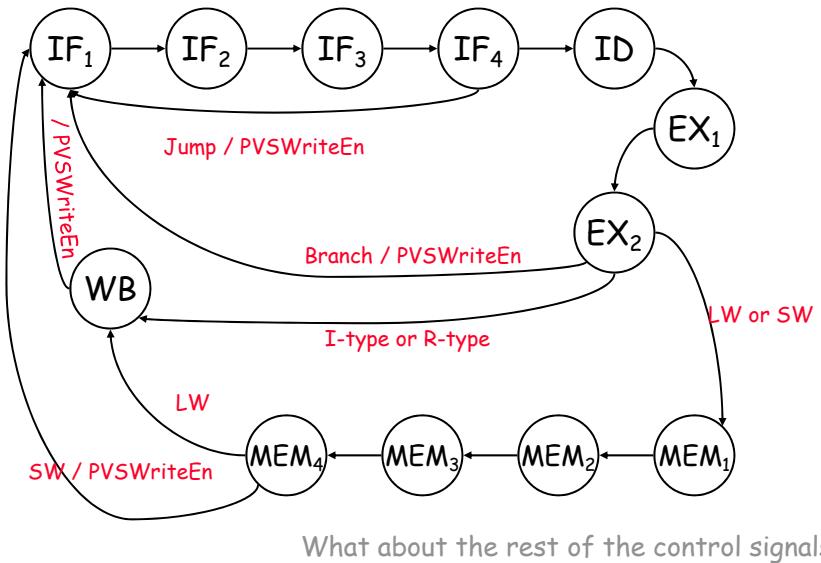
- ◆ Why not let each instruction type take only as much time as it needs
- ◆ Idea
 - run a 50 ps clock
 - let each instruction type take as many clock cycles as needed
 - programmer-visible state only updates at the end of an instruction's cycle-sequence
 - an instruction's effect is still purely combinational from PVS to PVS

A more realistic alternative to the "variable-length" clock design in the textbook

Multi-Cycle Datapath: Ver 1.0

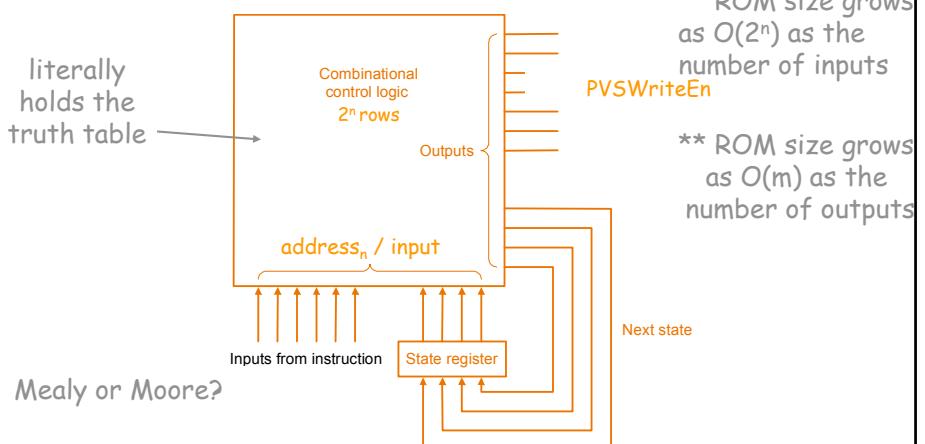


Sequential Control: Ver 1.0



MicroSequencer: Ver 1.0

- ◆ ROM as a combinational logic lookup table



Microcoding: Ver 1.0

state label	cntrl flow	conditional targets				
		R/I-type	LW	SW	Br	Jump
IF ₁	next	-	-	-	-	-
IF ₂	next	-	-	-	-	-
IF ₃	next	-	-	-	-	-
IF ₄	branch	ID	ID	ID	ID	IF ₁
ID	next	-	-	-	-	-
EX ₁	next	-	-	-	-	-
EX ₂	branch	WB	MEM ₁	MEM ₁	IF ₁	-
MEM ₁	next	-	-	-	-	-
MEM ₂	next	-	-	-	-	-
MEM ₃	next	-	-	-	-	-
MEM ₄	next	-	WB	IF ₁	-	-
WB	branch	IF ₁	IF ₁	IF ₁	IF ₁	IF ₁

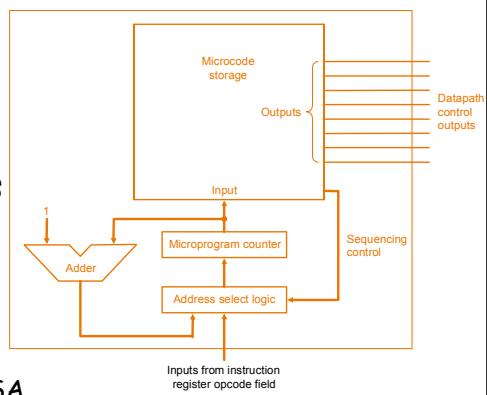
A systematic approach to FSM sequencing/control

Microcontroller

- ◆ A stripped-down processor for sequencing and control
 - control states are like µPC
 - µPC indexed into a µprogram ROM to select an µinstruction
 - well-formed control-flow architecture
 - fields in the µinstruction maps to control signals

Why not also support
µprogram-visible states
and ALU ops?

- ◆ Very elaborate microcontrollers have been built
 - some µcontrollers had full-fledged ISAs
 - µISAs for CISC machines in many ways inspired RISC ISA



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

Performance Analysis

- ◆ $T_{\text{wall-clock}} = \text{No.Instructions} \times CPI \times T_{\text{clk}}$
 - "No.Instruction" is fixed for a given ISA and application mix
- ◆ For a fixed ISA and application mix it is meaningful to compare

$$T_{\text{avg-inst}} = CPI \times T_{\text{clk}}$$

or

$$\text{MIPS} = IPC \times f_{\text{clk}}$$

million instructions
per second

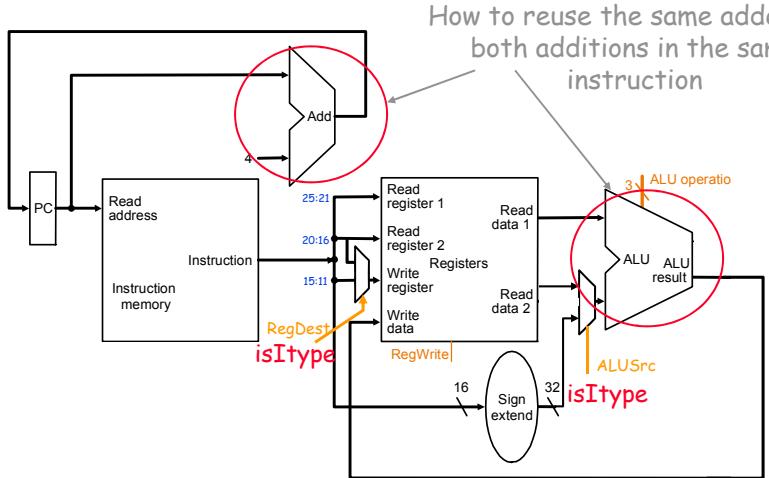
instructions per cycle

frequency in MHz

Performance Analysis

- ◆ Single-Cycle Implementation
 $1 \times 1,667 \text{MHz} = 1667 \text{ MIPS}$
- ◆ Multi-Cycle Implementation
 $IPC_{\text{avg}} \times 20,000 \text{ MHz} = 2235 \text{ MIPS}$
 what is average behavior?
- ◆ Assume: 25% LW, 10% SW, 45% ALU, 15% Branch and 5% Jumps
 - weighted arithmetic mean of CPI $\Rightarrow 8.95$
 - weighted harmonic mean of IPC $\Rightarrow 0.1117$

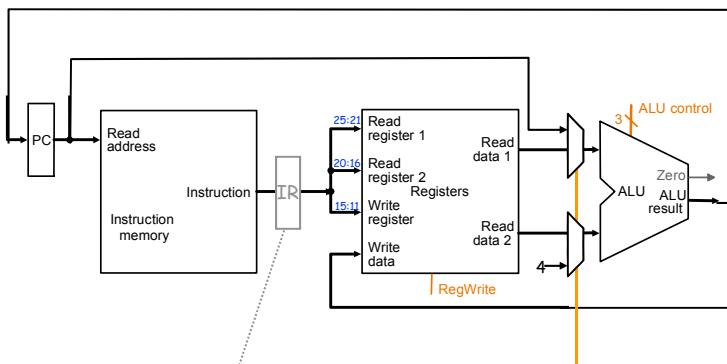
Reducing Datapath by Resource Reuse



Previous example of reuse by mutually exclusive conditions

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

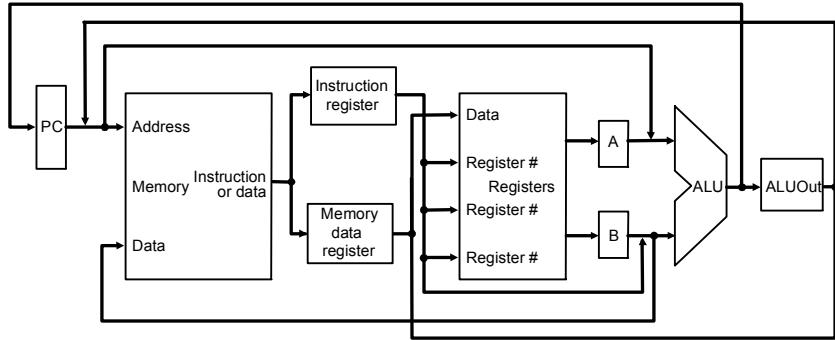
Reducing Datapath by Sequential Reuse



to IR or not to IR?

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

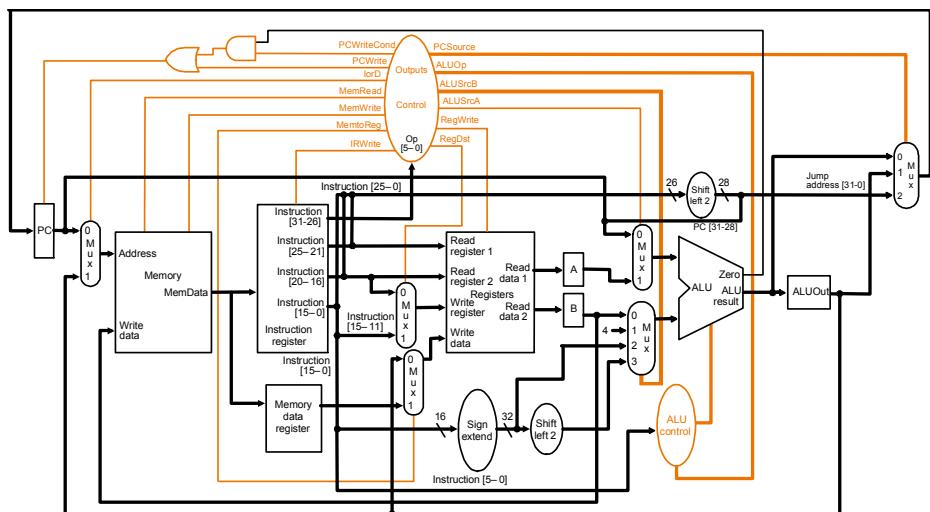
Removing Redundancies



- Could also reduce down to a single register read-write port!
- Do we want to remove all redundancies?

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

Control Points



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

New Sequential Control Signals

	When De-asserted	When asserted
ALUSrcA	1 st ALU input from PC	1 st ALU input from 1 st RF read port (latched in A)
IorD	PC supplies memory address	ALUOut supplies memory address
IRWrite	IR latching disabled	IR latching enabled
PCWrite	no effect	PC latching enabled unconditionally
PCWriteCond	no effect	PC latching enabled only if branch condition is satisfied

When both PCWrite and PCWriteCond are de-asserted, PC latching is disabled

New Sequential Control Signals

signal		effect
ALUSrcB[1:0]	00	2 nd ALU input from 2 nd RF read port (latched in B)
	01	2 nd ALU input is 4 (for PC increment)
	10	2 nd ALU input is sign-extended "IR[15:0]"
	11	2 nd ALU input is sign-extended "IR[15:0],00"
PCSource[1:0]	00	next PC from ALU
	01	next PC from ALUOut
	10	next PC from IR (jump target)

Old Control Signals (similar to single-cycle)

	When De-asserted	When asserted
RegDest	RF write select according to IR[20:16]	RF write select according to IR[15:11]
RegWrite	RF write disabled	RF write enabled
MemRead	Memory read disabled	Memory read port return load value
MemWrite	Memory write disabled	Memory write enabled
MemtoReg	Steer ALU result (latched in ALUOut) to RF write port	steer memory load result (latched in MDR) to RF write port

Synchronous Register Transfers

- ◆ Synchronous state with latch enables
 - PC, IR, RF, MEM
- ◆ Synchronous state that always latch
 - A, B, ALUOut, MDR
- ◆ One can enumerate all possible combinational "register transfers" in the datapath
- ◆ For example starting from PC
 - $IR \leftarrow MEM[PC]$
requires IRWrite=1, IorD=0, MemRead=1
 - $PC \leftarrow PC, JumpTarget$
requires PCWrite=1, PCSource=2'b10
 - $PC \leftarrow PC + ALUSrcB$
 - $MDR \leftarrow MEM[PC]$
 - $ALUOut \leftarrow PC + ALUSrcB$
 -

Useful Register Transfers

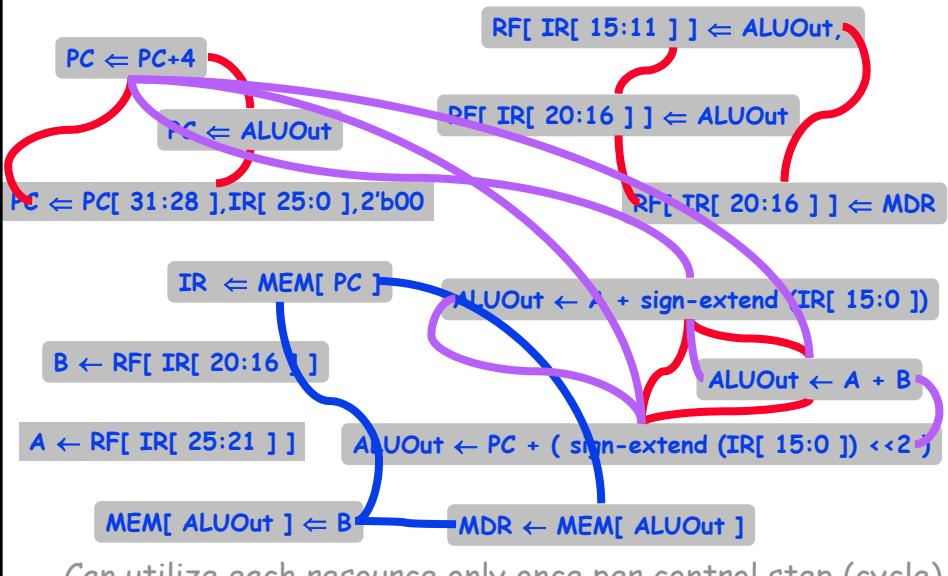
- ◆ $PC \leftarrow PC + 4$
- ◆ $PC \leftarrow ALUOut$ (if PCWriteCond is asserted)
- ◆ $PC \leftarrow PC[31:28], IR[25:0], 2'b00$
- ◆ $IR \leftarrow MEM[PC]$
- ◆ $A \leftarrow RF[IR[25:21]]$
- ◆ $B \leftarrow RF[IR[20:16]]$
- ◆ $ALUOut \leftarrow A + B$
- ◆ $ALUOut \leftarrow A + \text{sign-extend}(IR[15:0])$
- ◆ $ALUOut \leftarrow PC + (\text{sign-extend}(IR[15:0]) \ll 2)$
- ◆ $MDR \leftarrow MEM[ALUOut]$
- ◆ $MEM[ALUOut] \leftarrow B$
- ◆ $RF[IR[15:11]] \leftarrow ALUOut$,
- ◆ $RF[IR[20:16]] \leftarrow ALUOut$
- ◆ $RF[IR[20:16]] \leftarrow MDR$

RT Sequencing: R-Type ALU

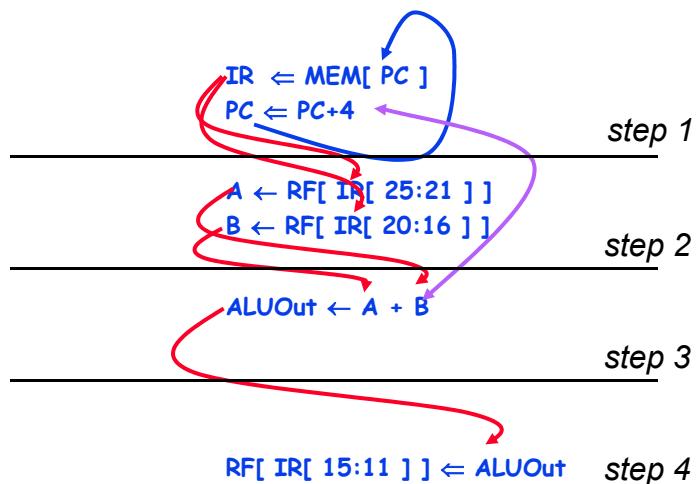
- ◆ IF
 - $IR \leftarrow MEM[PC]$
 - $PC \leftarrow PC + 4$
- ◆ ID
 - $A \leftarrow RF[IR[25:21]]$
 - $B \leftarrow RF[IR[20:16]]$
- ◆ EX
 - $ALUOut \leftarrow A + B$
- ◆ MEM
 -
- ◆ WB
 - $RF[IR[15:11]] \leftarrow ALUOut$

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

RT Datapath Conflicts



RT Sequencing: R-Type ALU



RT Sequencing: LW

- ◆ IF
 - $IR \leftarrow MEM[PC]$
 - $PC \leftarrow PC + 4$
- ◆ ID
 - $A \leftarrow RF[IR[25:21]]$
 - $B \leftarrow RF[IR[20:16]]$
- ◆ EX
 - $ALUOut \leftarrow A + \text{sign-extend}(IR[15:0])$
- ◆ MEM
 - $MDR \leftarrow MEM[ALUOut]$
- ◆ WB
 - $RF[IR[20:16]] \leftarrow MDR$

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

RT Sequencing: Branch

- ◆ IF
 - $IR \leftarrow MEM[PC]$
 - $PC \leftarrow PC + 4$
- ◆ ID
 - $A \leftarrow RF[IR[25:21]]$
 - $B \leftarrow RF[IR[20:16]]$
 - $ALUOut \leftarrow PC + (\text{sign-extend}(IR[15:0]) \ll 2)$
- ◆ EX
 - $PC \leftarrow ALUOut$ (only if condition is met)
- ◆ MEM
- ◆ WB

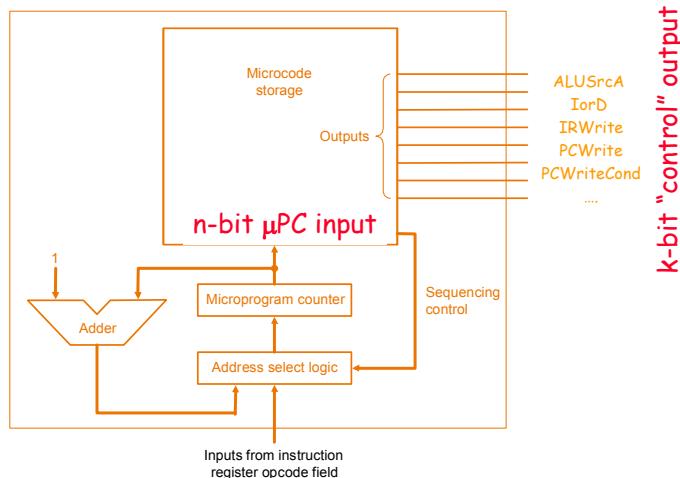
```
if MEM[PC]==BEQ rs rt immediate16
  target = PC + sign-extend(immediate) × 4 + 4
  if GPR[rs]==GPR[rt] then    PC ← target
                               else    PC ← PC + 4
```

Combined RT Sequencing

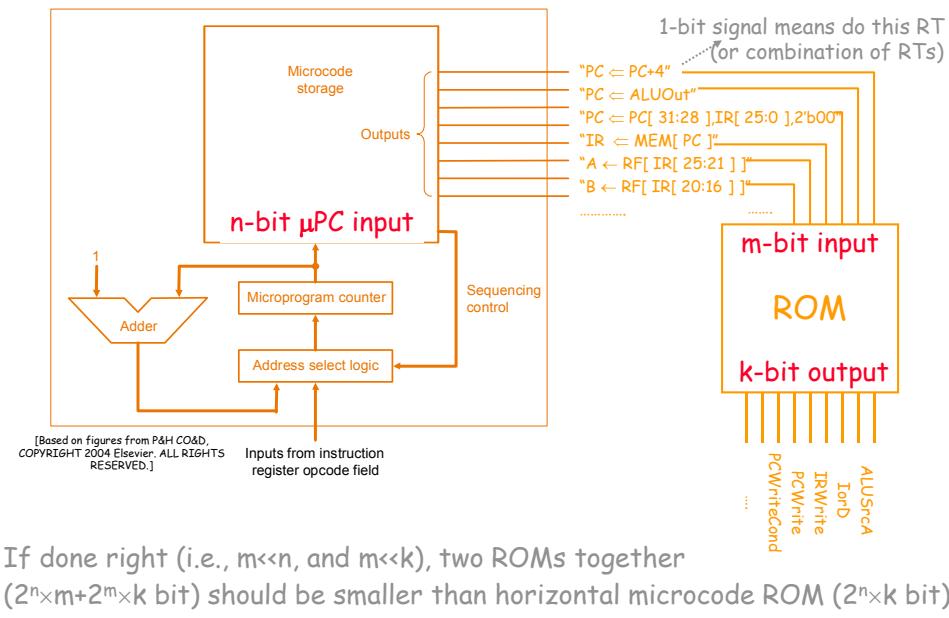
	R-Type	LW	SW	Branch	Jump
common steps	start: $IR \leftarrow MEM[PC]$ $PC \leftarrow PC+4$				
	$A \leftarrow RF[IR[25:21]]$ $B \leftarrow RF[IR[20:16]]$ $ALUOut \leftarrow PC + (\text{sign-extend}(IR[15:0]) \ll 2)$				
opcode dependent steps	$ALUOut \leftarrow A + B$ $RF[IR[15:11]] \leftarrow ALUOut$ $RF[IR[15:11]] \leftarrow MDR$	$ALUOut \leftarrow A + \text{sign-extend}(IR[15:0])$ $MDR \leftarrow MEM[ALUOut]$ $RF[IR[15:11]] \leftarrow MDR$	$ALUOut \leftarrow A + \text{sign-extend}(IR[15:0])$ $MEM[ALUOut] \leftarrow F$	$PC \leftarrow ALUOut$ $PC \leftarrow PC[31:28], goto IR[25:0], goto start$	$PC \leftarrow PC[31:28], goto IR[25:0], goto start$

RTs in each state corresponds to some setting of the control signals

Horizontal Microcode



Vertical Microcode



Microcoding for CISC

- ◆ Can we extend the μ controller and datapath
 - to support a new instruction I haven't thought of yet
 - to support a complex instruction, e.g. polyf
- ◆ Yes, and probably more
 - if I can sequence an arbitrary RISC instruction then I can sequence an arbitrary "RISC program" as a μ program sequence
 - will need some μ ISA state (e.g. loop counters) for more elaborate μ programs
 - more elaborate μ ISA features also make life easier
- ◆ μ coding allows very simple datapath do very powerful computation
 - a datapath as simple as a Turning machine is universal
 - μ code enables a minimal datapath to emulate any ISA you like (with a commensurate slow down)

Nanocode and Millicode

◆ Nanocode

- a level below 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
- e.g., the **polyf** sequence may be generated by a separate nanicontroller in the FPU

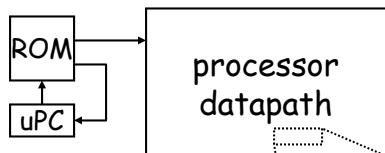
◆ Millicode

- a level above μ code
- ISA-level subroutines hardcoded into a ROM that can be called by the μ controller to handle really complicated operations
 - e.g., to add **polyf** to MIPS, one may code up **polyf** as a software routine that is called by the μ controller when the **polyf** opcode is decoded

- ◆ In both cases, we avoid complicating the main μ controller with **polyf** support **The power of abstractions!!**

Nanocode Concept Illustrated

a "mcoded" processor implementation



We refer to this
as "nanocode"
when a mcoded
subsystem is embedded
in a mcoded system

a "mcoded" FPU implementation

