

CS 61C: Great Ideas in Computer
Architecture (Machine Structures)
Lecture 30: Pipeline Parallelism 1

Instructor:

Dan Garcia


<http://inst.eecs.Berkeley.edu/~cs61c/sp13>

CS61C in the News

Is organic computing finally here?

SPOTLIGHT | NOVEMBER 2, 2011

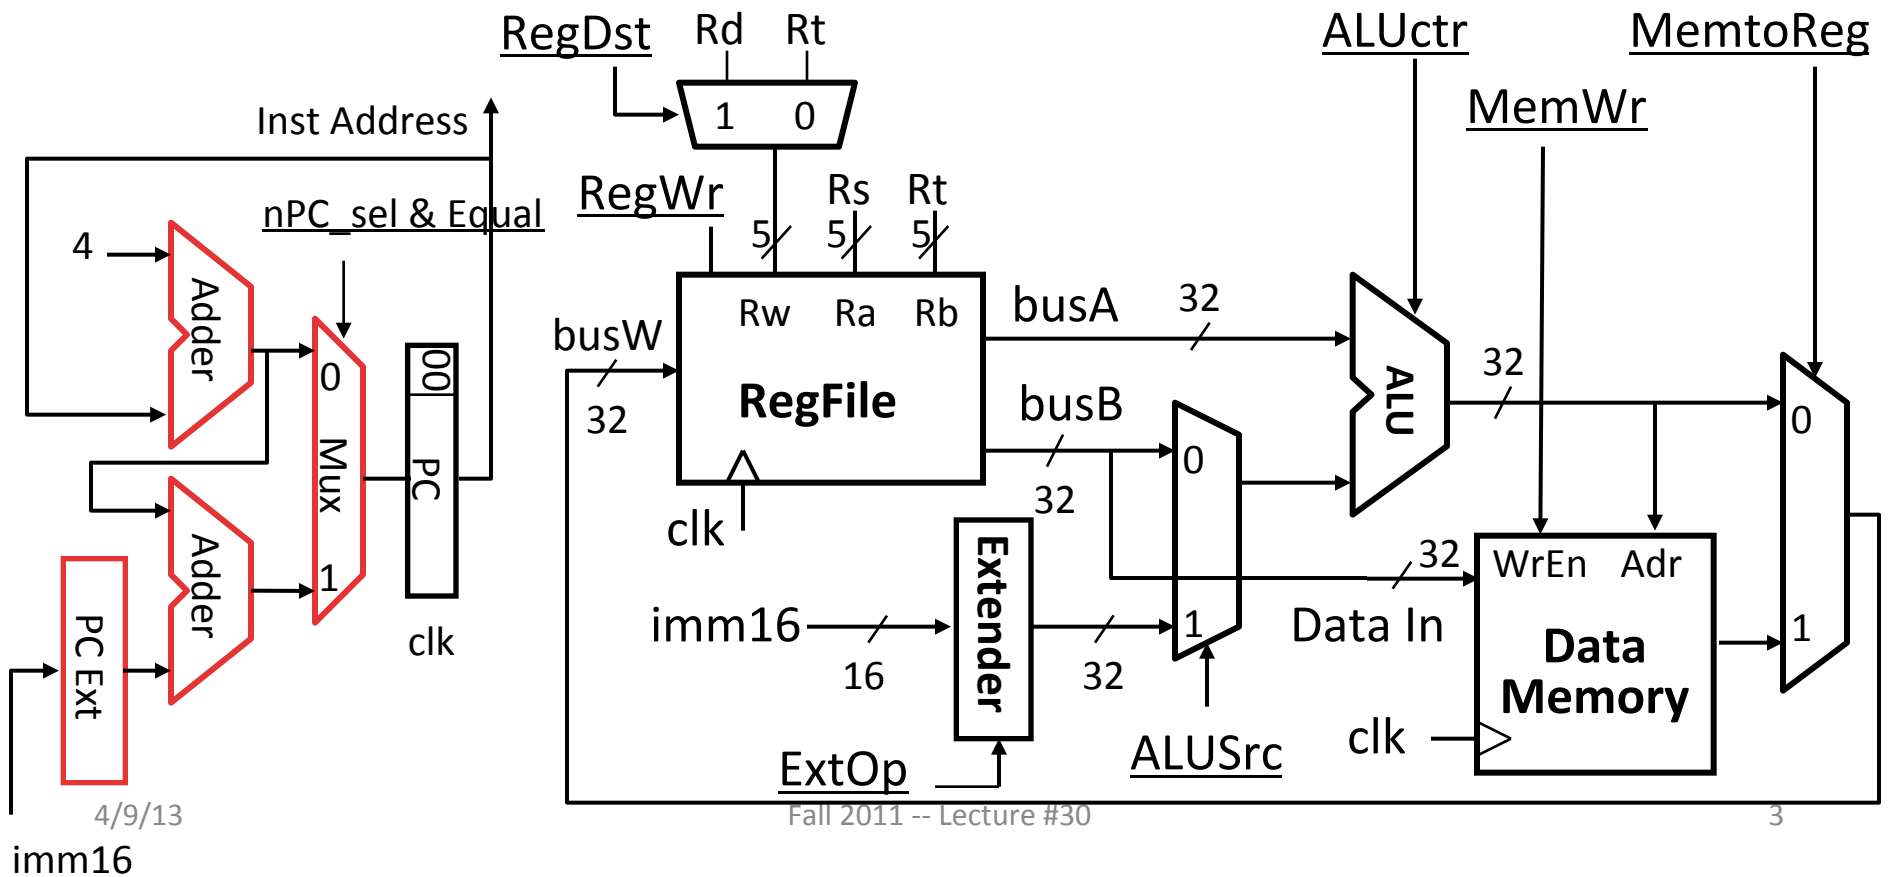
Japanese scientists have made organic molecules perform parallel computations like neurons in the human brain. They created this promising new approach with a ring-like molecule called 2,3-dichloro-5,6-dicyano-p-benzoquinone, or DDQ.



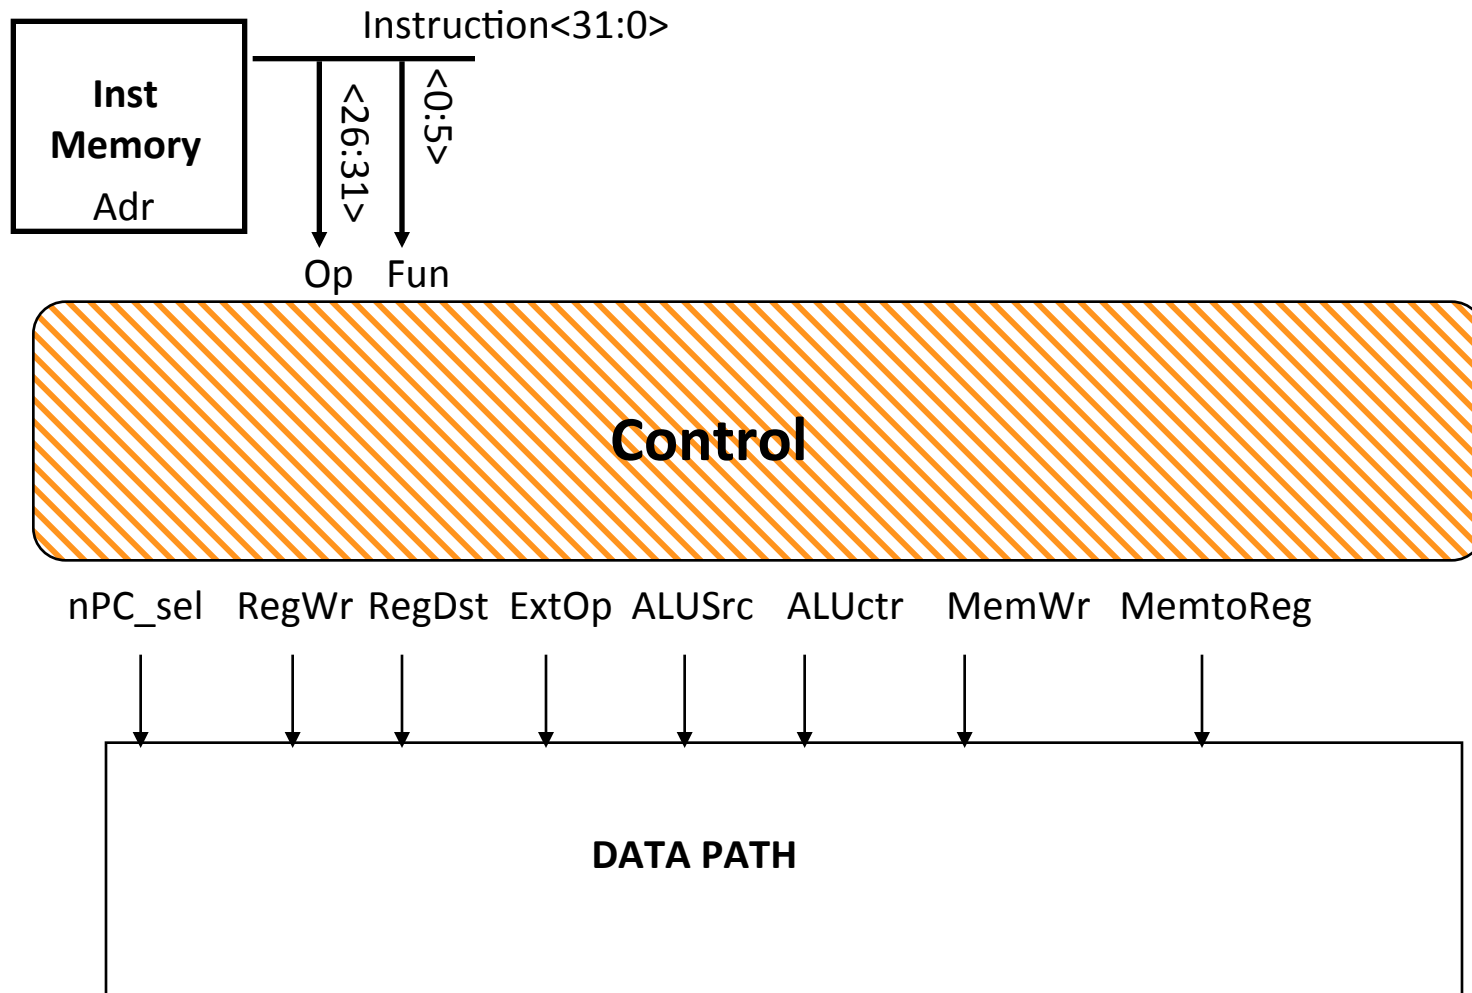
Today, computer chips can process data at 10 trillion (10^{13}) bits per second. But, even though neurons in the human brain fire at a rate of 100 times per second, the brain still outperforms the best computers at various tasks. **The main reason being that calculations done by computer chips happen in isolated pipelines one at a time.**

Datapath Control Signals

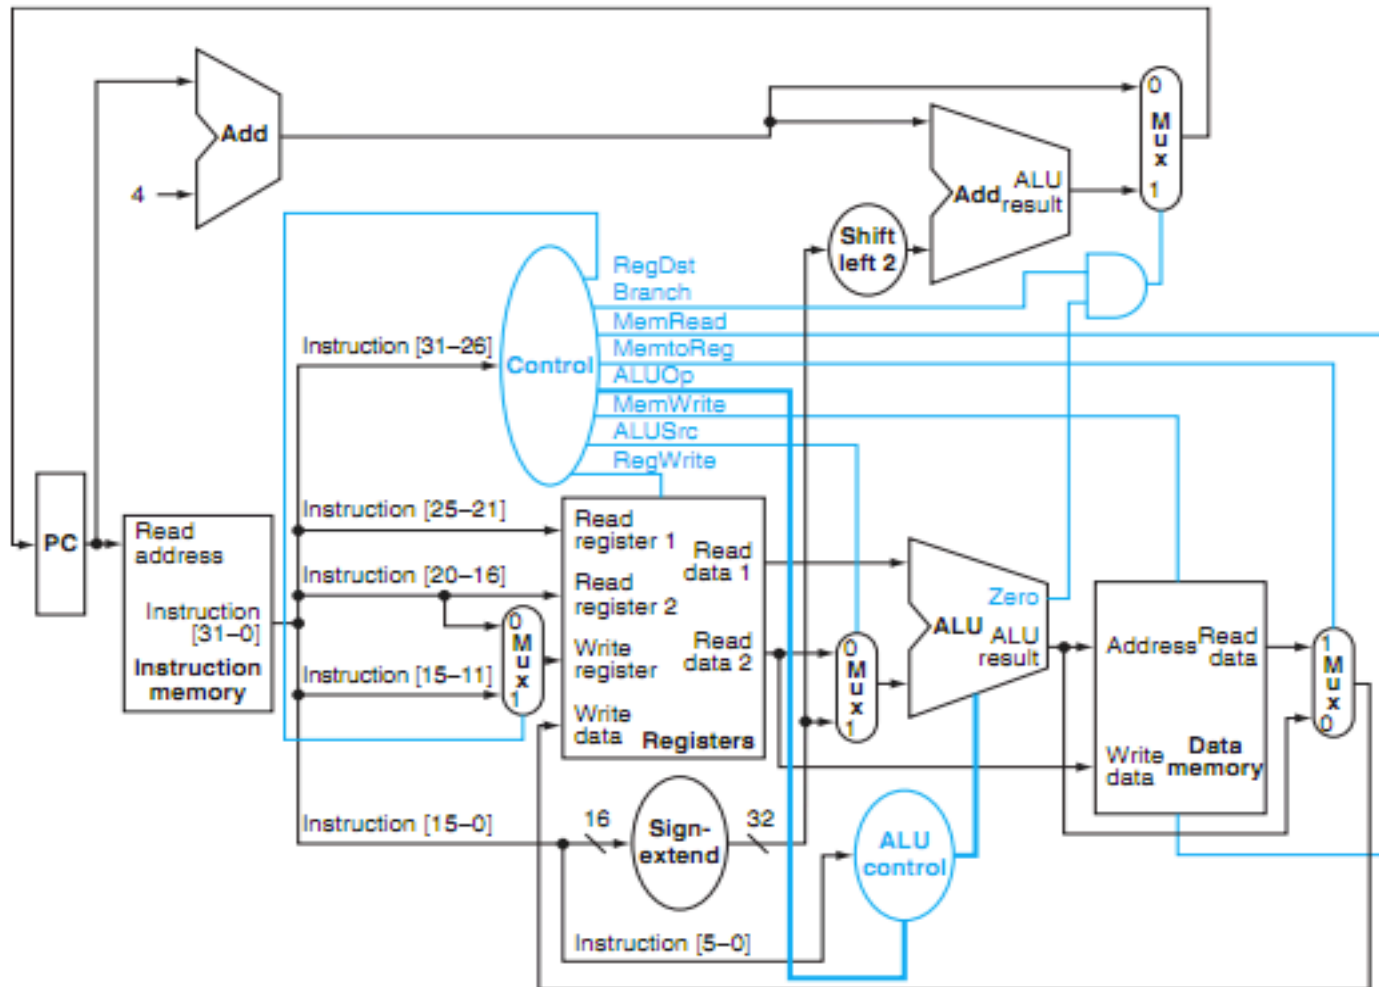
- ExtOp: “zero”, “sign”
- ALUSrc: 0 \Rightarrow regB;
1 \Rightarrow immed
- ALUctr: “ADD”, “SUB”, “OR”
- MemWr: 1 \Rightarrow write memory
- MemtoReg: 0 \Rightarrow ALU; 1 \Rightarrow Mem
- nPC_sel: 0 \Rightarrow “+4”; 1 \Rightarrow “br”
- RegDst: 0 \Rightarrow “rt”; 1 \Rightarrow “rd”
- RegWr: 1 \Rightarrow write register



Where Do Control Signals Come From?



P&H Figure 4.17



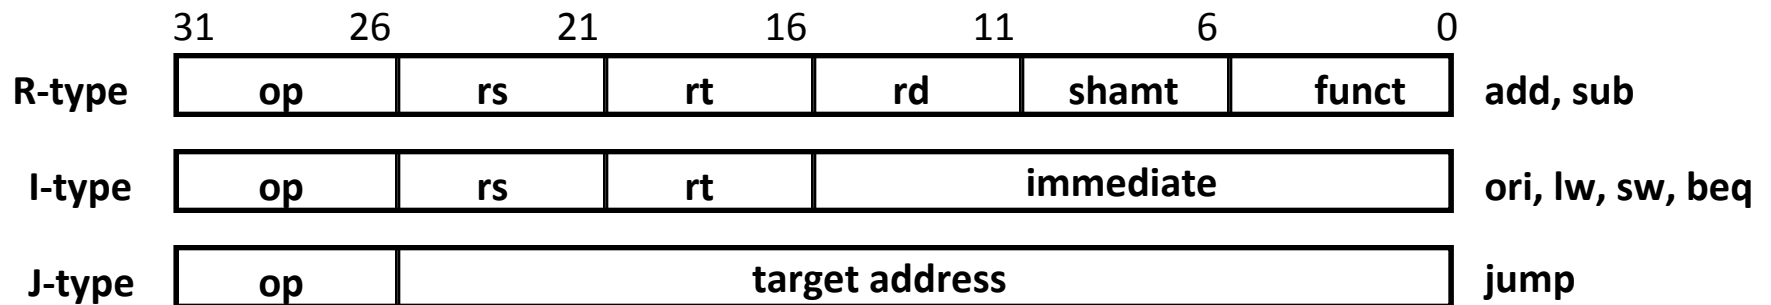
Summary of the Control Signals (1/2)

```
inst   Register Transfer
add     R[rd] ← R[rs] + R[rt]; PC ← PC + 4
        ALUSrc=RegB, ALUctr="ADD", RegDst=rd, RegWr, nPC_sel="+4"
sub     R[rd] ← R[rs] - R[rt]; PC ← PC + 4
        ALUSrc=RegB, ALUctr="SUB", RegDst=rd, RegWr, nPC_sel="+4"
ori     R[rt] ← R[rs] + zero_ext(Imm16); PC ← PC + 4
        ALUSrc=Im, Extop="Z", ALUctr="OR", RegDst=rt, RegWr, nPC_sel="+4"
lw      R[rt] ← MEM[ R[rs] + sign_ext(Imm16)]; PC ← PC + 4
        ALUSrc=Im, Extop="sn", ALUctr="ADD", MemtoReg, RegDst=rt, RegWr,
        nPC_sel = "+4"
sw      MEM[ R[rs] + sign_ext(Imm16)] ← R[rs]; PC ← PC + 4
        ALUSrc=Im, Extop="sn", ALUctr = "ADD", MemWr, nPC_sel = "+4"
beq     if (R[rs] == R[rt]) then PC ← PC + sign_ext(Imm16) || 00
        else PC ← PC + 4
        nPC_sel = "br", ALUctr = "SUB"
```


Summary of the Control Signals (2/2)

See Appendix A → **func**
 → **op**

	10 0000	10 0010	We Don't Care :-)				
	00 0000	00 0000	00 1101	10 0011	10 1011	00 0100	00 0010
	add	sub	ori	lw	sw	beq	jump
RegDst	1	1	0	0	x	x	x
ALUSrc	0	0	1	1	1	0	x
MemtoReg	0	0	0	1	x	x	x
RegWrite	1	1	1	1	0	0	0
MemWrite	0	0	0	0	1	0	0
nPCsel	0	0	0	0	0	1	?
Jump	0	0	0	0	0	0	1
ExtOp	x	x	0	1	1	x	x
ALUctr<2:0>	Add	Subtract	Or	Add	Add	Subtract	x



Boolean Exprs for Controller

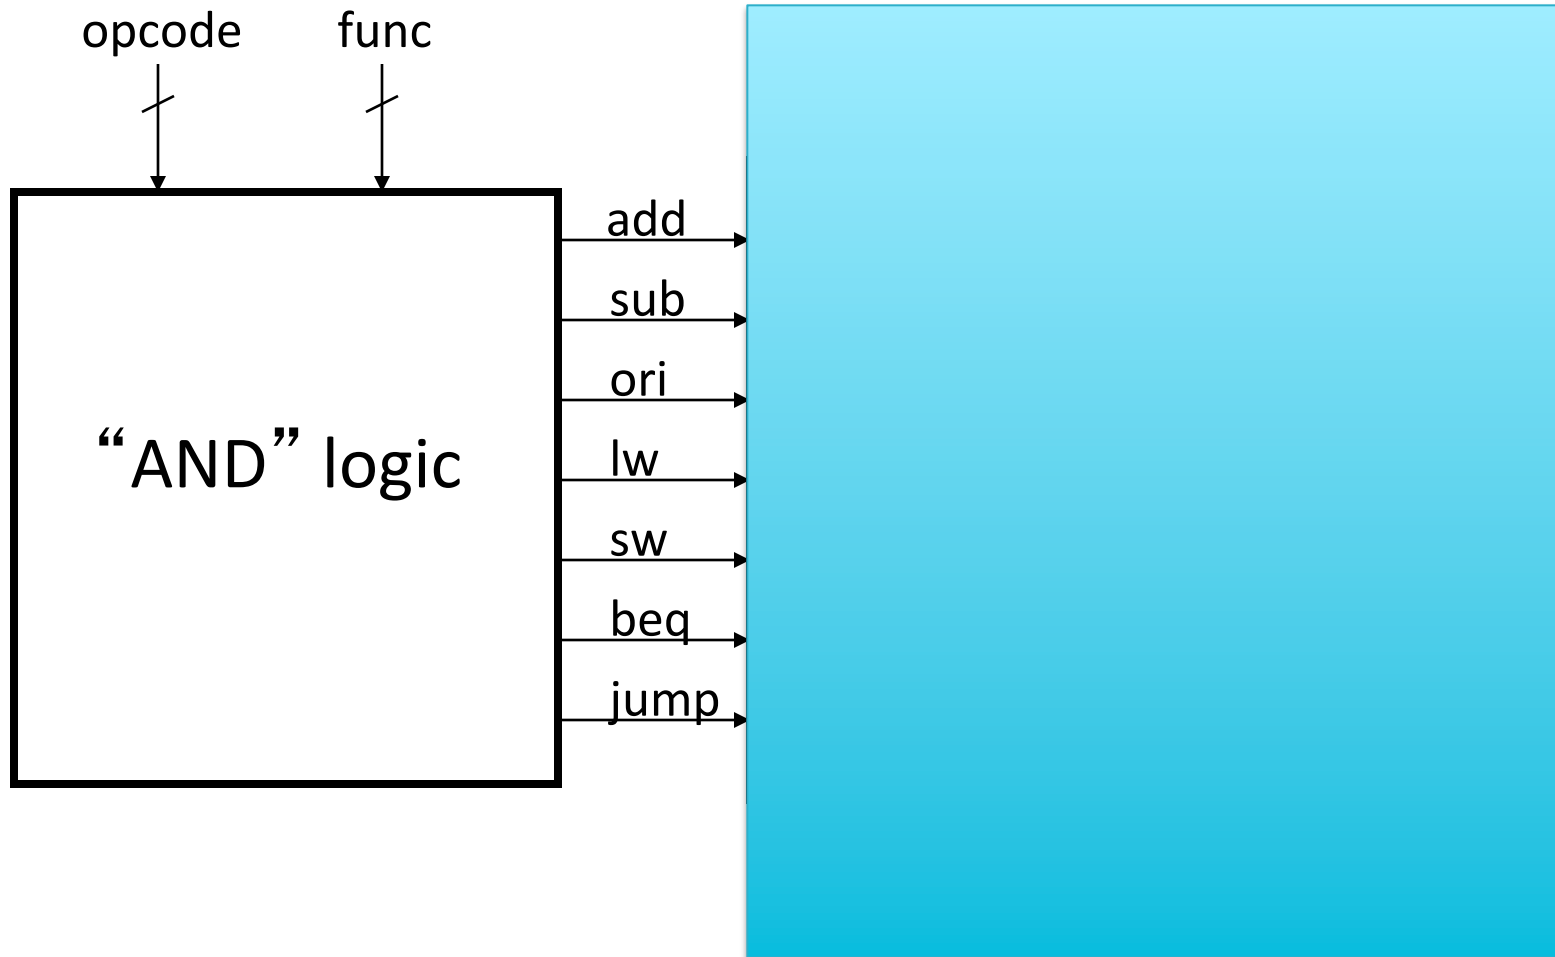


Op 0-5 are really Instruction bits 26-31
 Func 0-5 are really Instruction bits 0-5

$$\begin{aligned}
 \mathbf{rtype} &= \sim op_5 \cdot \sim op_4 \cdot \sim op_3 \cdot \sim op_2 \cdot \sim op_1 \cdot \sim op_0, \\
 ori &= \sim op_5 \cdot \sim op_4 \cdot op_3 \cdot op_2 \cdot \sim op_1 \cdot op_0 \\
 lw &= op_5 \cdot \sim op_4 \cdot \sim op_3 \cdot \sim op_2 \cdot op_1 \cdot op_0 \\
 sw &= op_5 \cdot \sim op_4 \cdot op_3 \cdot \sim op_2 \cdot op_1 \cdot op_0 \\
 beq &= \sim op_5 \cdot \sim op_4 \cdot \sim op_3 \cdot op_2 \cdot \sim op_1 \cdot \sim op_0 \\
 jump &= \sim op_5 \cdot \sim op_4 \cdot \sim op_3 \cdot \sim op_2 \cdot op_1 \cdot \sim op_0
 \end{aligned}$$

$$\begin{aligned}
 add &= \mathbf{rtype} \cdot func_5 \cdot \sim func_4 \cdot \sim func_3 \cdot \sim func_2 \cdot \sim func_1 \cdot \sim func_0 \\
 sub &= \mathbf{rtype} \cdot func_5 \cdot \sim func_4 \cdot \sim func_3 \cdot \sim func_2 \cdot func_1 \cdot \sim func_0
 \end{aligned}$$

Controller Implementation



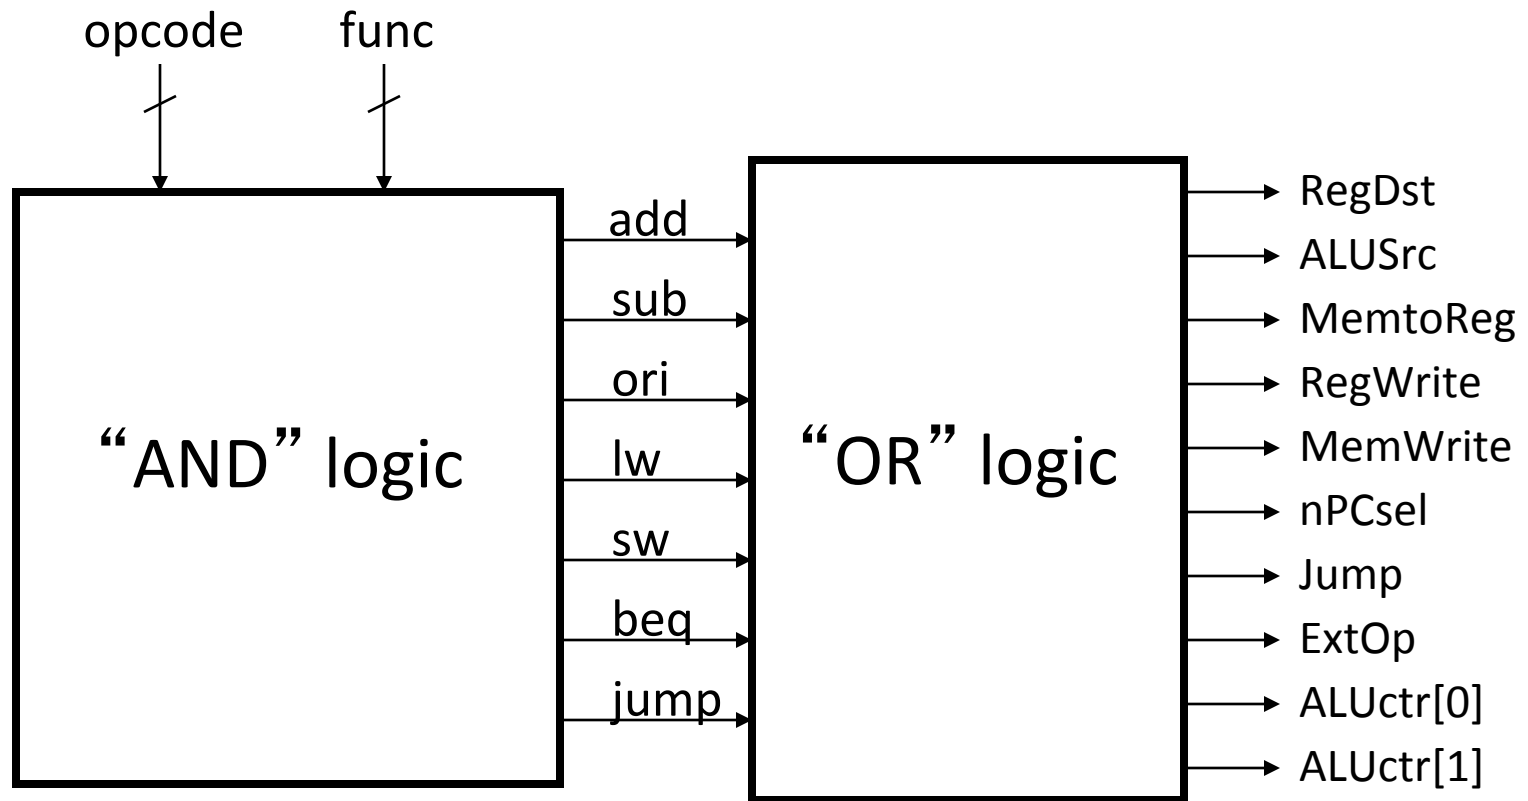
Boolean Exprs for Controller

```
RegDst      = add + sub
ALUSrc      = ori + lw + sw
MemtoReg    = lw
RegWrite     = add + sub + ori + lw
MemWrite    = sw
nPCsel      = beq
Jump        = jump
ExtOp       = lw + sw
ALUctr[0]   = sub + beq
ALUctr[1]   = ori
```

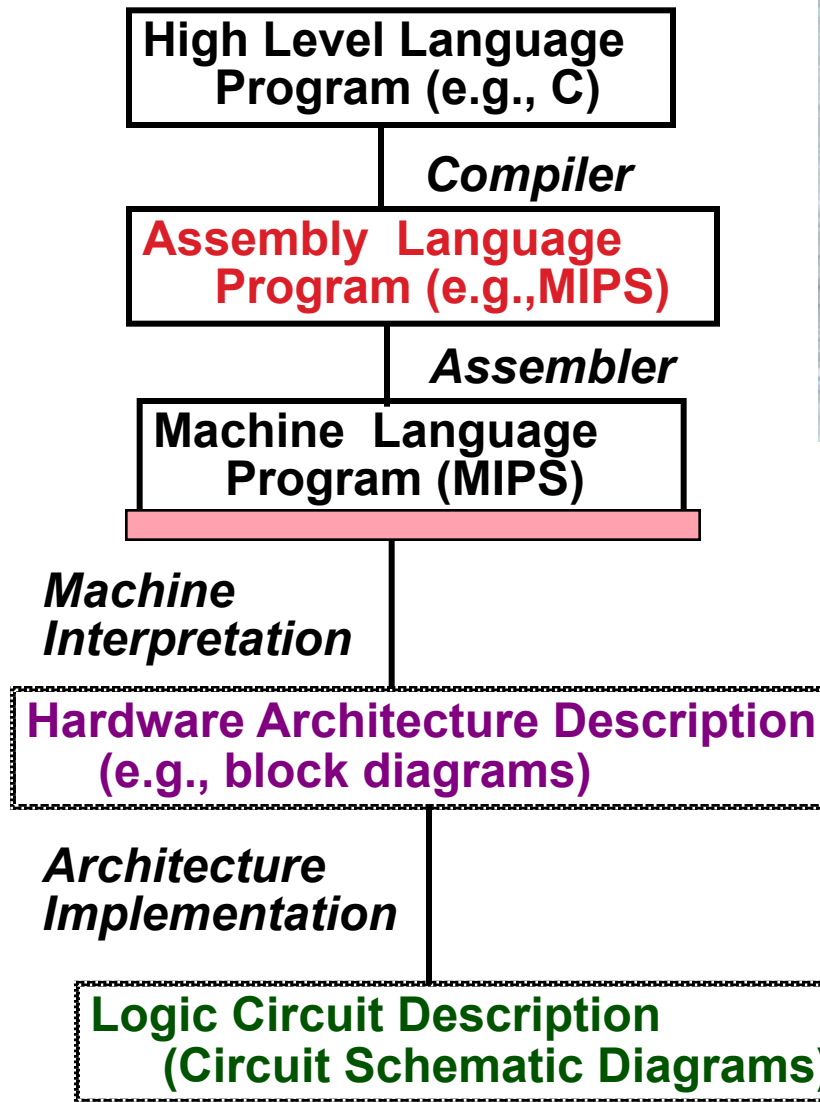
(assume ALUctr is 00 ADD, 01 SUB, 10 OR)

How do we implement this in gates?

Controller Implementation



Call home, we've made HW/SW contact!



Review: Single-cycle Processor

- Five steps to design a processor:

1. Analyze instruction set → datapath requirements


2. Select set of datapath components & establish clock methodology

3. Assemble datapath meeting the requirements

4. Analyze implementation of each instruction to determine setting of control points that effects the register transfer.

5. Assemble the control logic

- Formulate Logic Equations
- Design Circuits



Single Cycle Performance

- Assume time for actions are
 - 100ps for register read or write; 200ps for other events
- Clock rate is?

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps

- What can we do to improve clock rate?
- Will this improve performance as well?
 - Want increased clock rate to mean faster programs

Single Cycle Performance




- Assume time for actions are
 - 100ps for register read or write; 200ps for other events
- Clock rate is?

Instr	Instr fetch	Register read	ALU op	Memory access	Register write	Total time
lw	200ps	100 ps	200ps	200ps	100 ps	800ps
sw	200ps	100 ps	200ps	200ps		700ps
R-format	200ps	100 ps	200ps		100 ps	600ps
beq	200ps	100 ps	200ps			500ps


- What can we do to improve clock rate?
- Will this improve performance as well?
 - Want increased clock rate to mean faster programs

Gotta Do Laundry

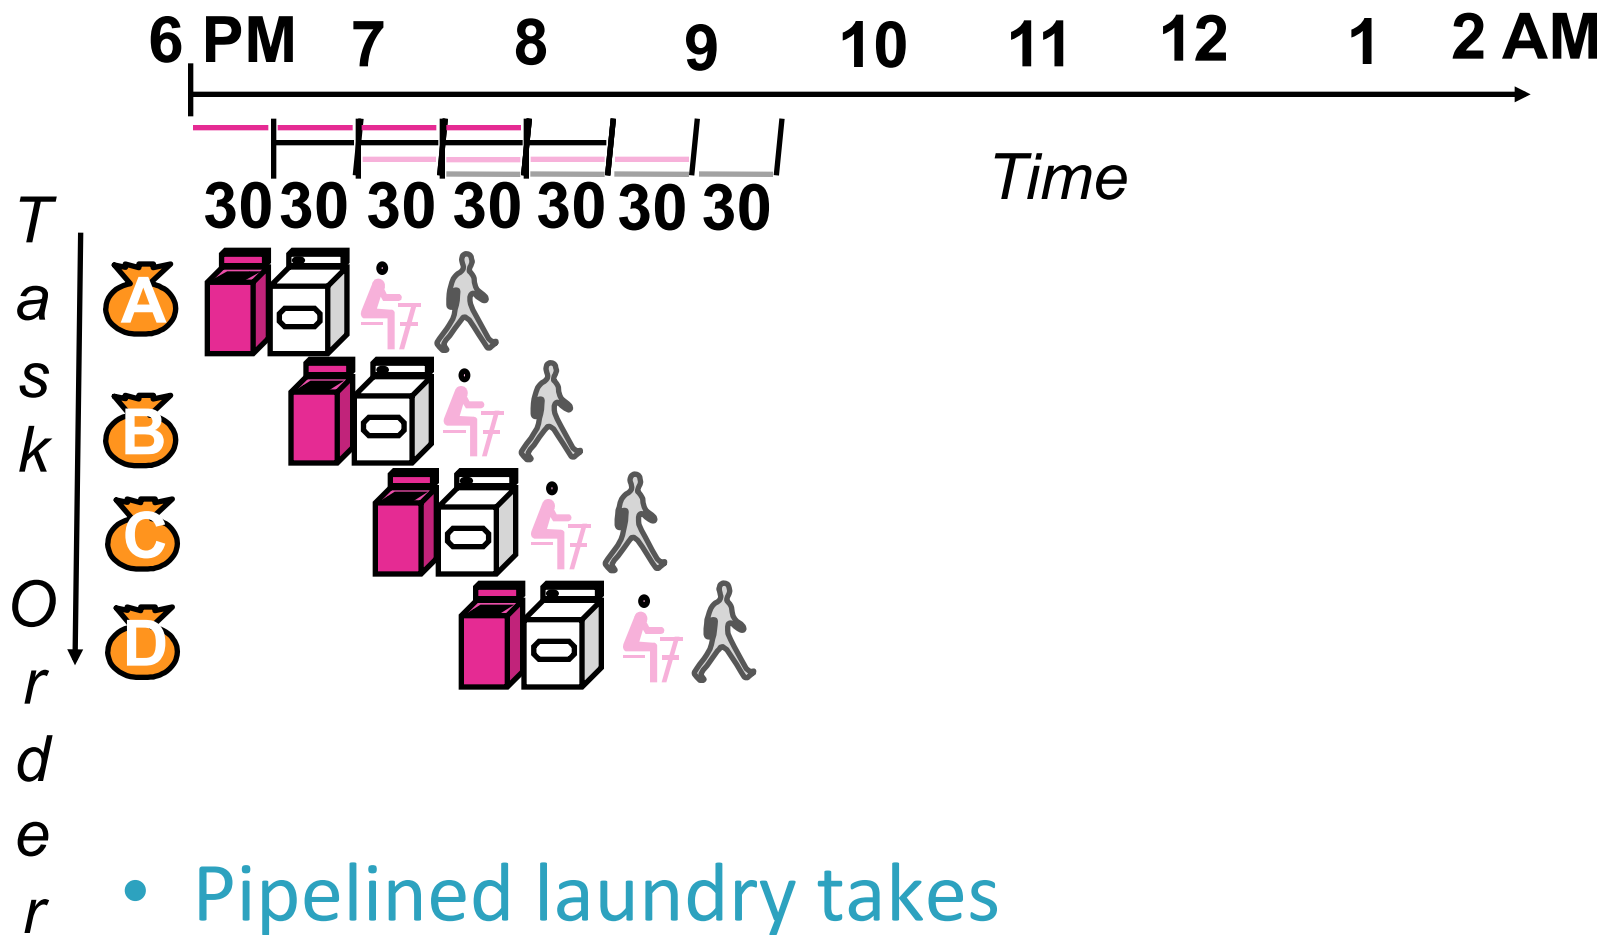
- Ann, Brian, Cathy, Dave each have one load of clothes to wash, dry, fold, and put away
 - Washer takes 30 minutes
 - Dryer takes 30 minutes
 - “Folder” takes 30 minutes
 - “Stasher” takes 30 minutes to put clothes into drawers




Sequential Laundry



Pipelined Laundry




Pipelining Lessons (1/2)



- Pipelining doesn't help latency of single task, it helps throughput of entire workload
- Multiple tasks operating simultaneously using different resources
- Potential speedup = Number pipe stages
- Time to "fill" pipeline and time to "drain" it reduces speedup: 2.3X v. 4X in this example

Pipelining Lessons (2/2)

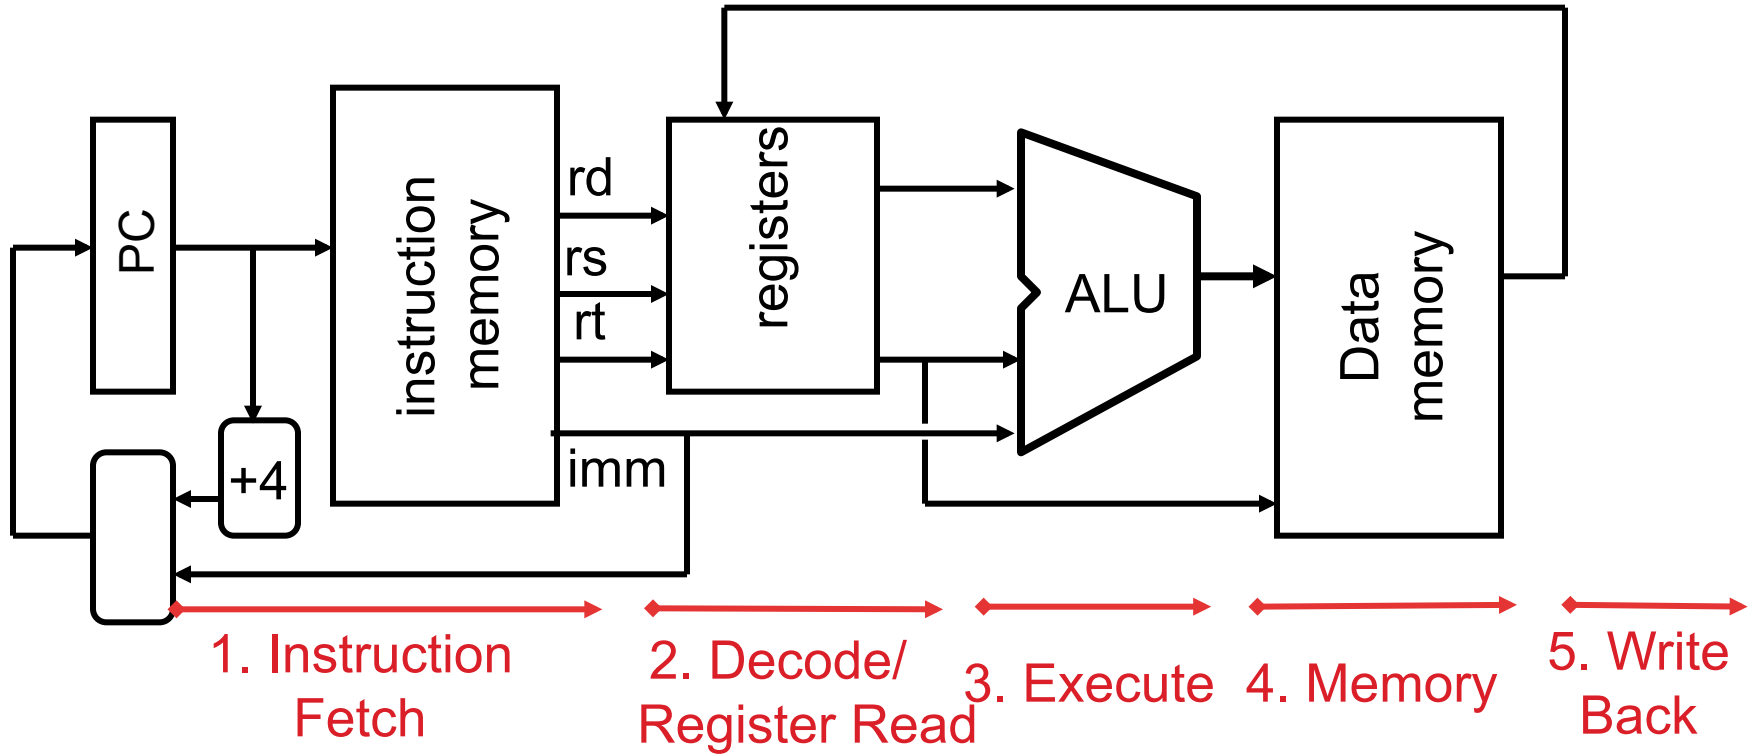


- Suppose new Washer takes 20 minutes, new Stasher takes 20 minutes. How much faster is pipeline?
- Pipeline rate limited by slowest pipeline stage
- Unbalanced lengths of pipe stages reduces speedup

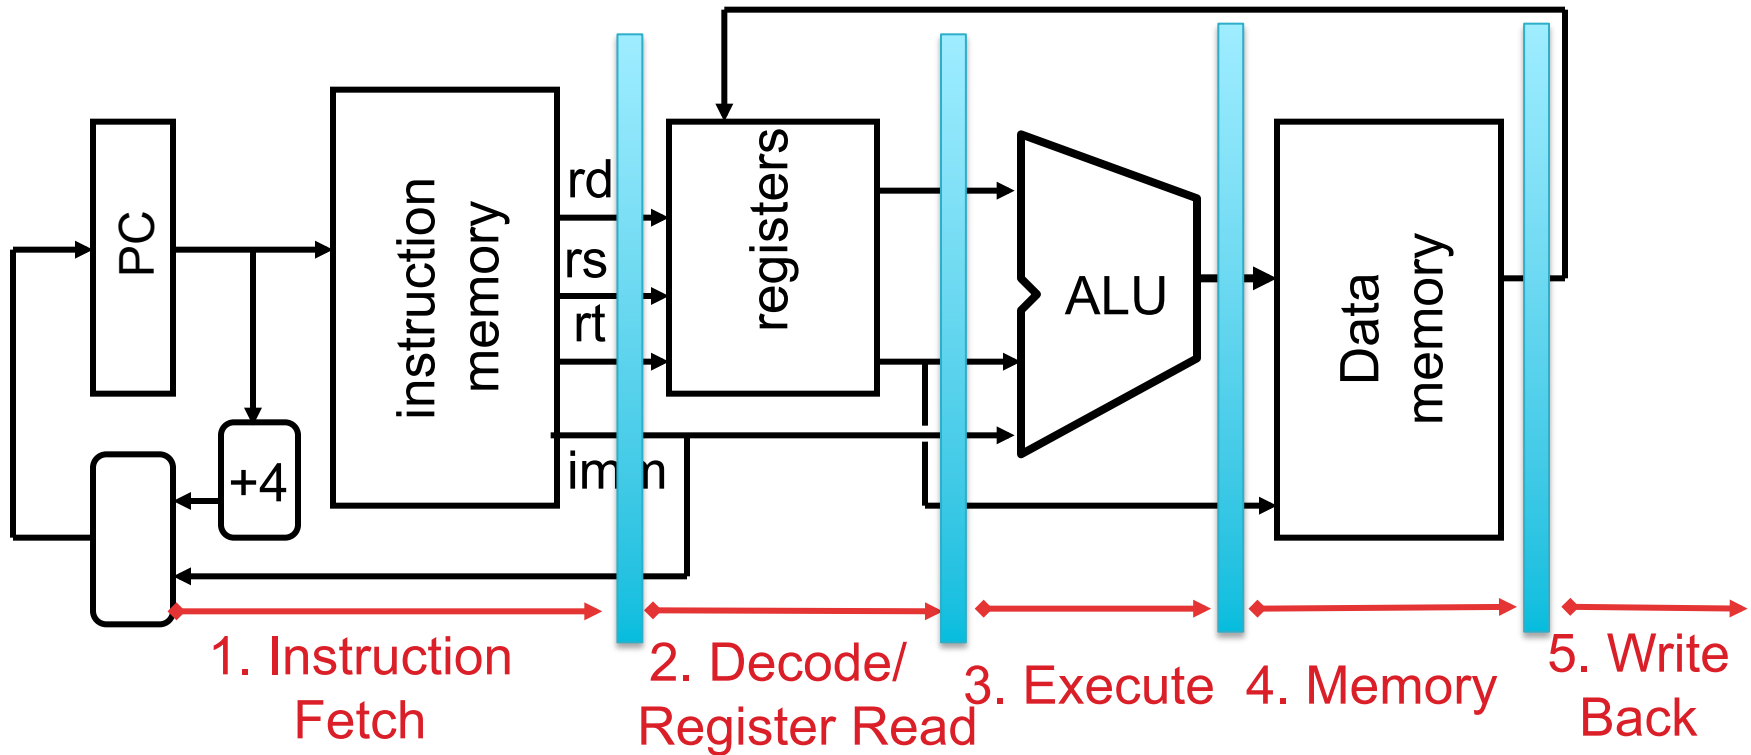
Steps in Executing MIPS

- 1) IFtch: Instruction Fetch, Increment PC
- 2) Dcd: Instruction Decode, Read Registers
- 3) Exec:
 - Mem-ref: Calculate Address
 - Arith-log: Perform Operation
- 4) Mem:
 - Load: Read Data from Memory
 - Store: Write Data to Memory
- 5) WB: Write Data Back to Register

Single Cycle Datapath

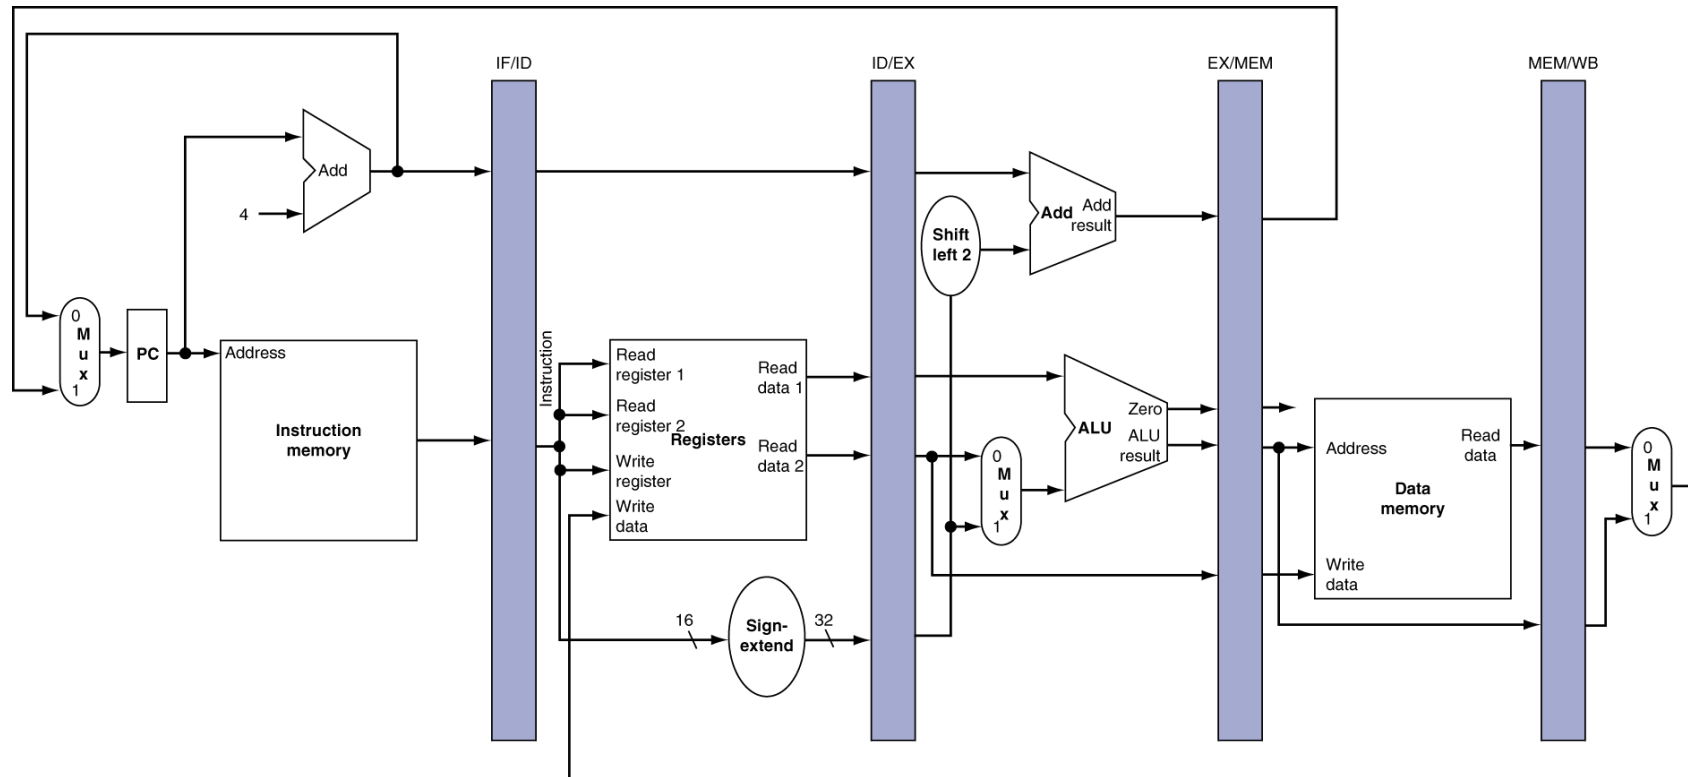


Pipeline registers




- Need registers between stages
 - To hold information produced in previous cycle

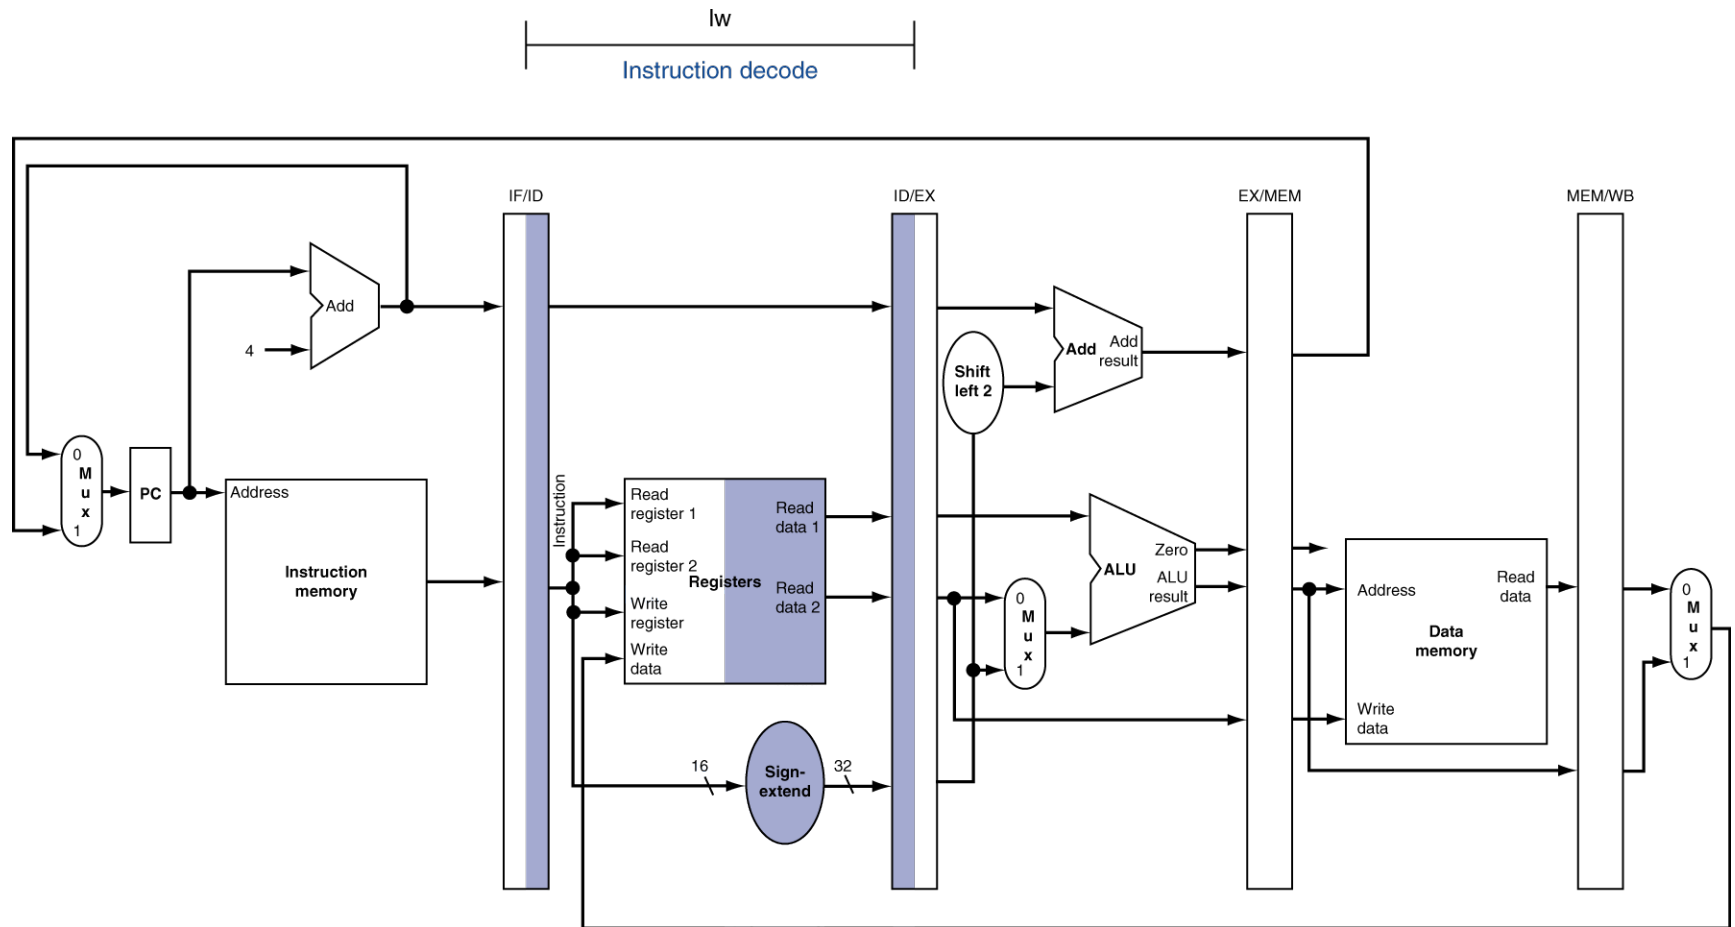
More Detailed Pipeline



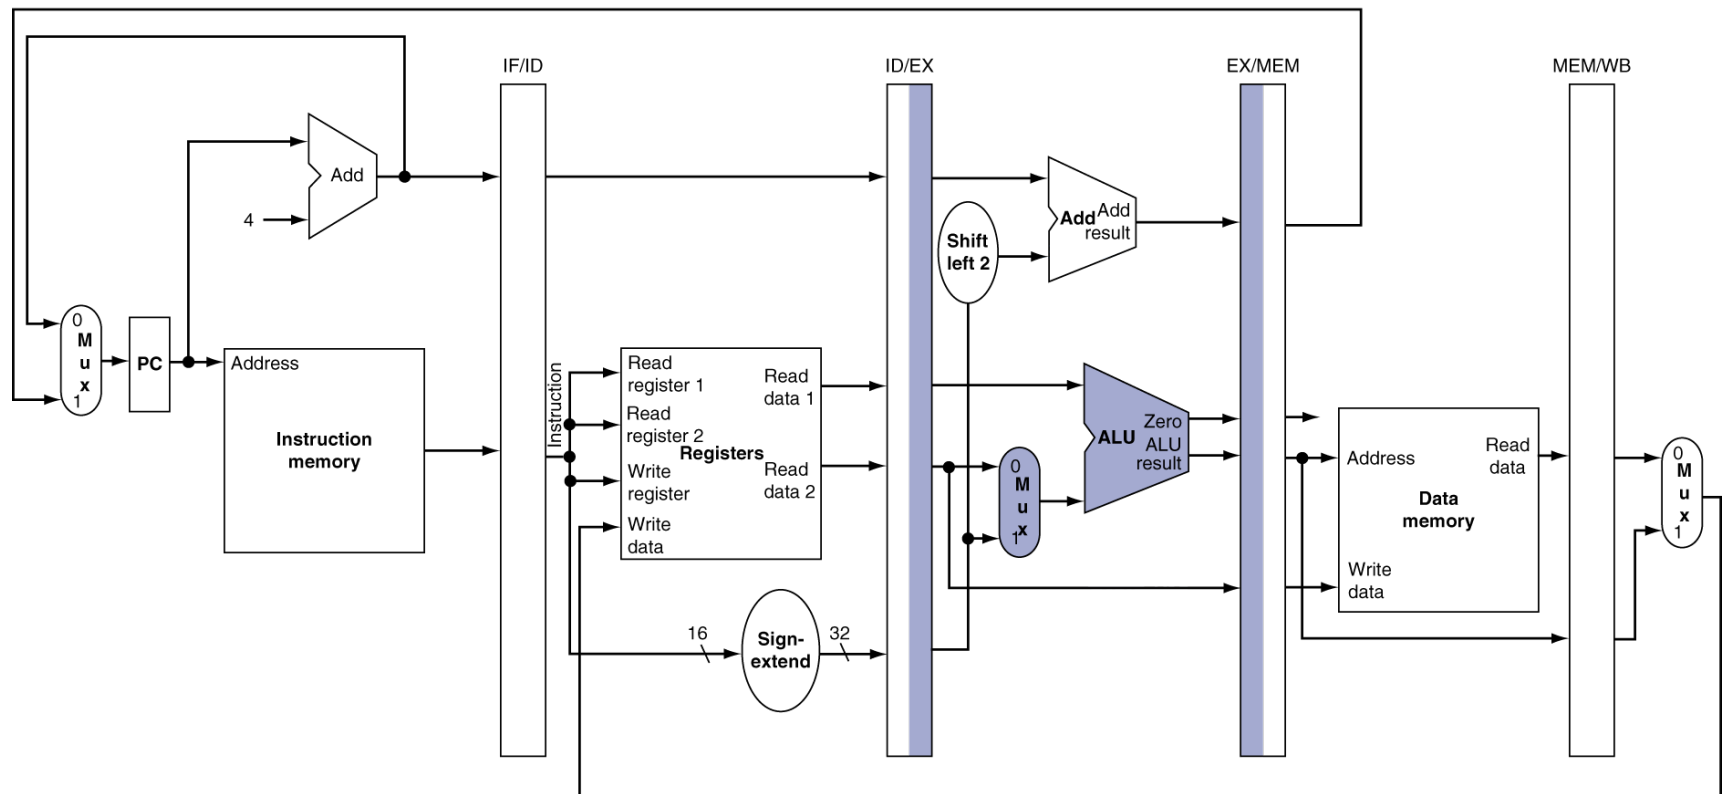
IF for Load, Store, ...



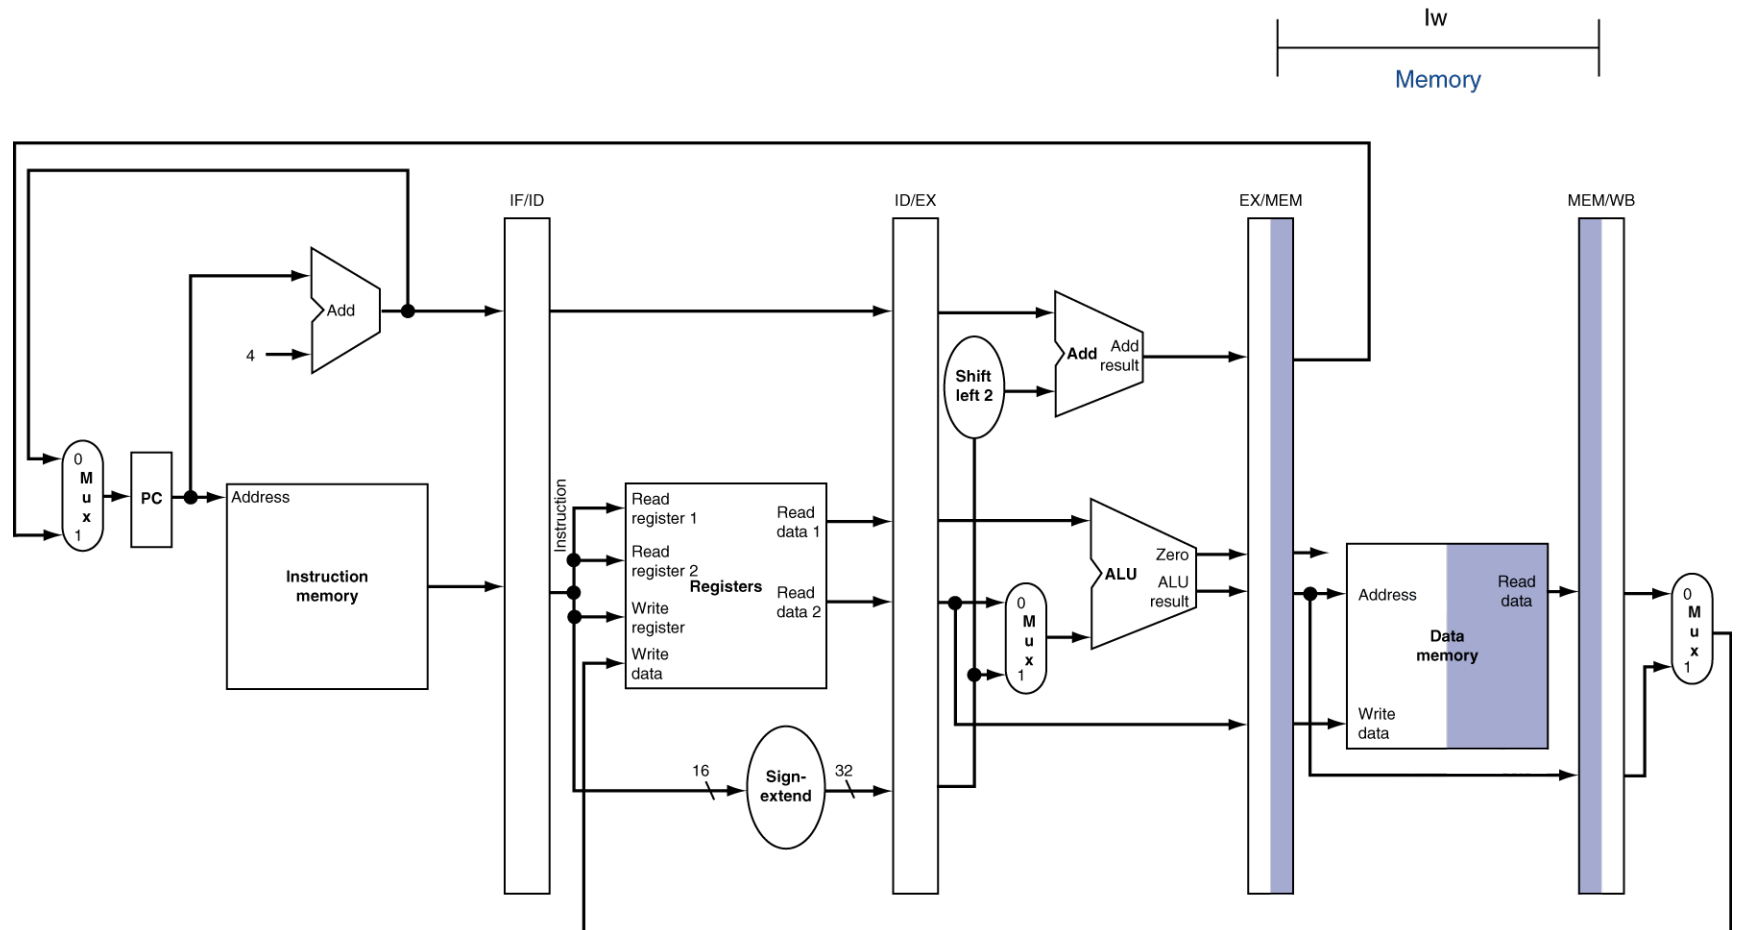
ID for Load, Store, ...



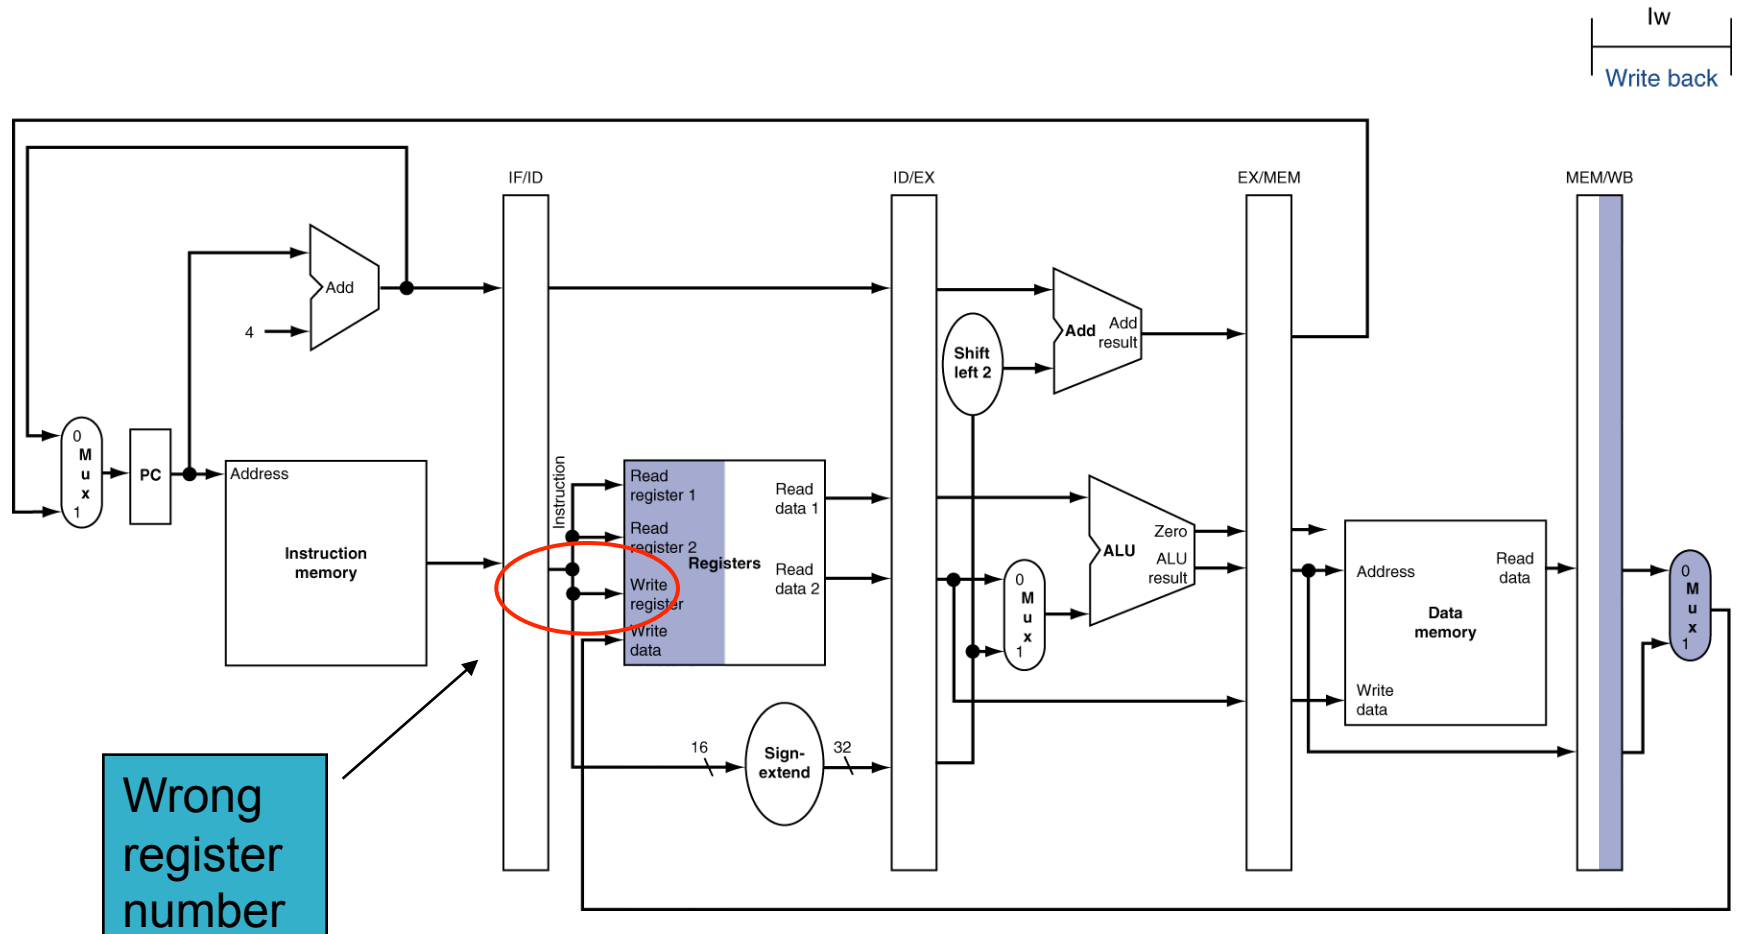
EX for Load




MEM for Load



WB for Load – Oops!



Corrected Datapath for Load



Peer Instruction

- 1) Thanks to pipelining, I have reduced the time it took me to wash my shirt.
- 2) Longer pipelines are always a win (since less work per stage & a faster clock).
- 3) We can rely on compilers to help us avoid data hazards by reordering instrs.

	1	2	3
a:	F	F	F
b:	F	F	T
b:	F	T	F
c:	F	T	T
c:	T	F	F
d:	T	F	T
d:	T	T	F
e:	T	T	T

So, in conclusion

- You now know how to implement the control logic for the single-cycle CPU.
 - (actually, you already knew it!)
- Pipelining improves performance by increasing instruction throughput: exploits ILP
 - Executes multiple instructions in parallel
 - Each instruction has the same latency
- Next: hazards in pipelining:
 - Structure, data, control