
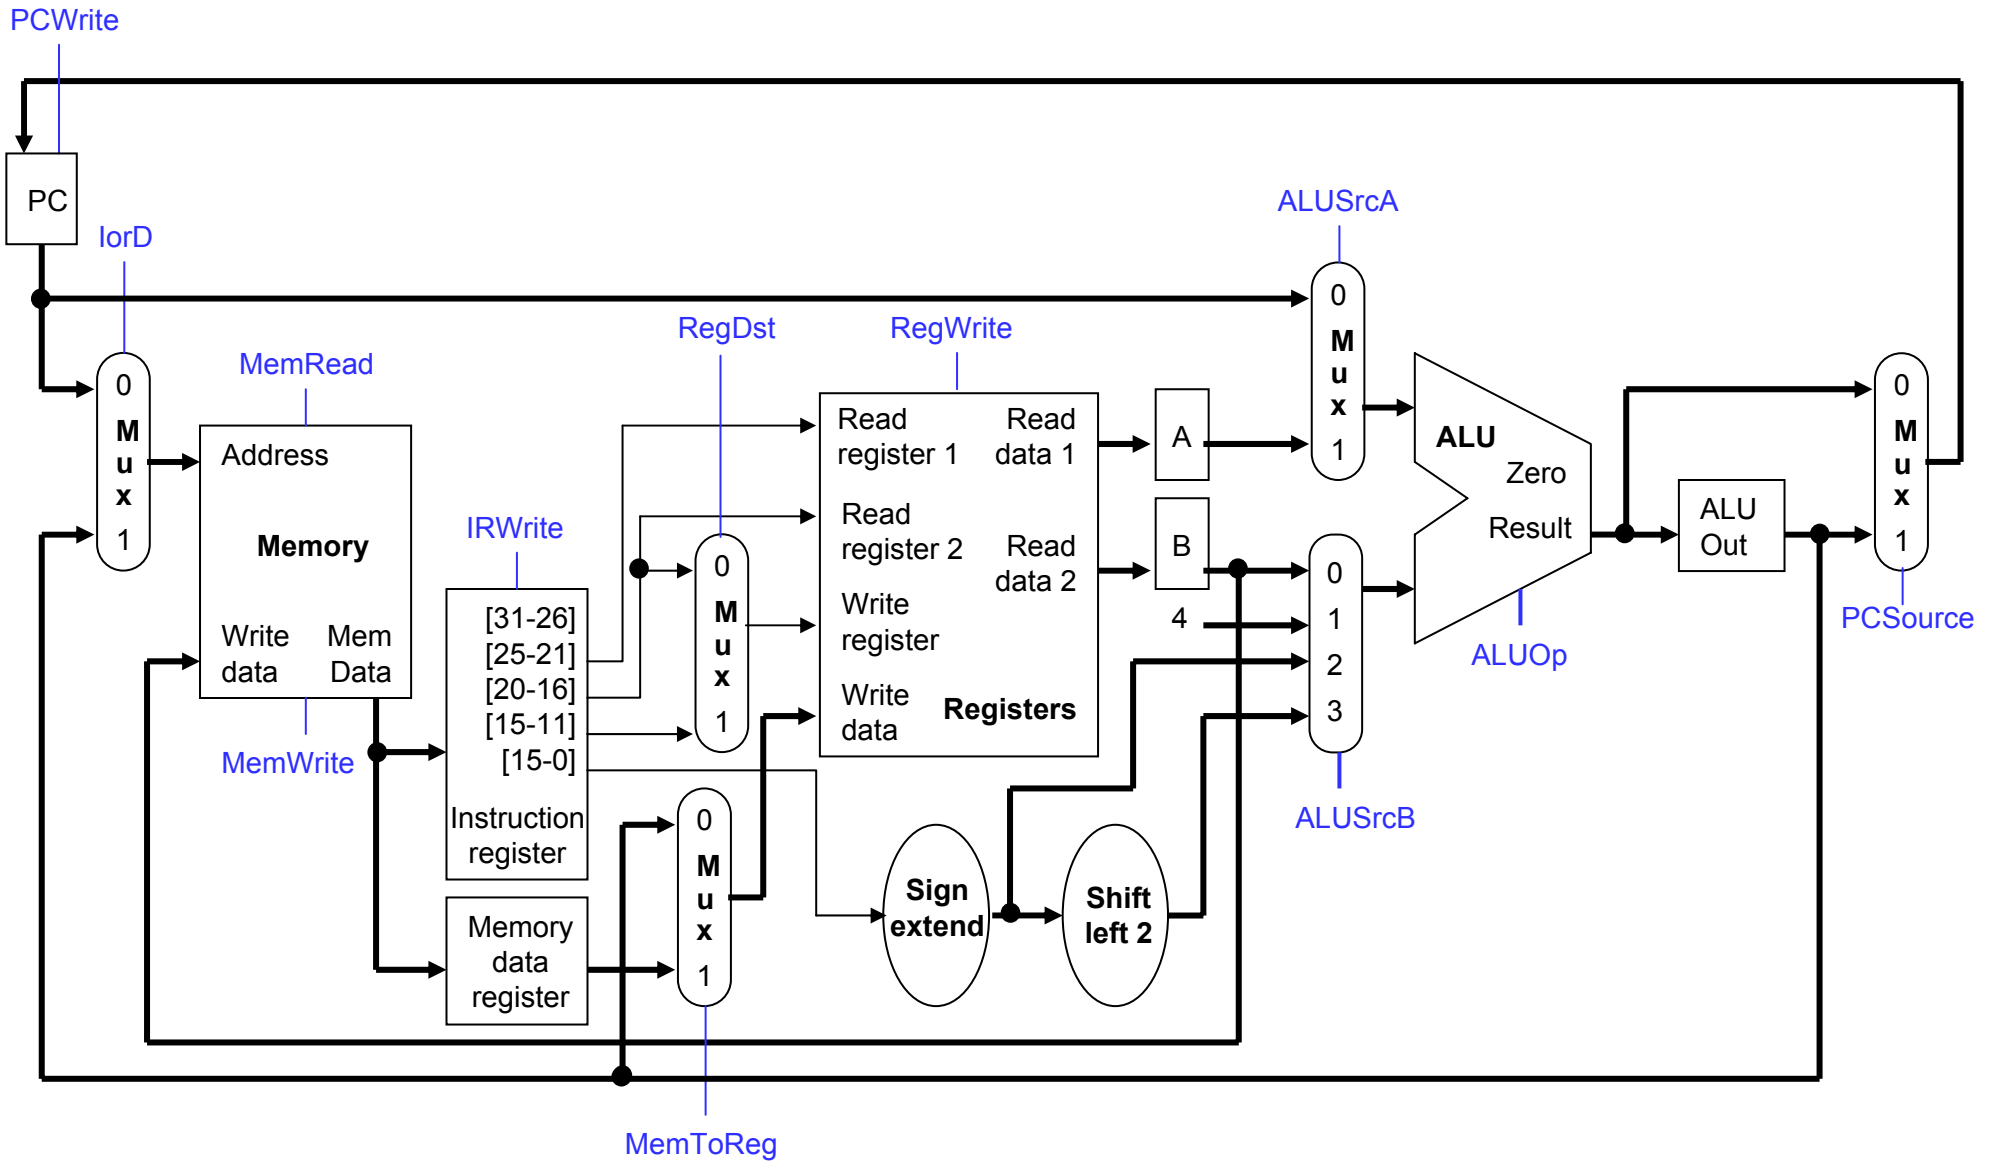


Multicycle conclusion

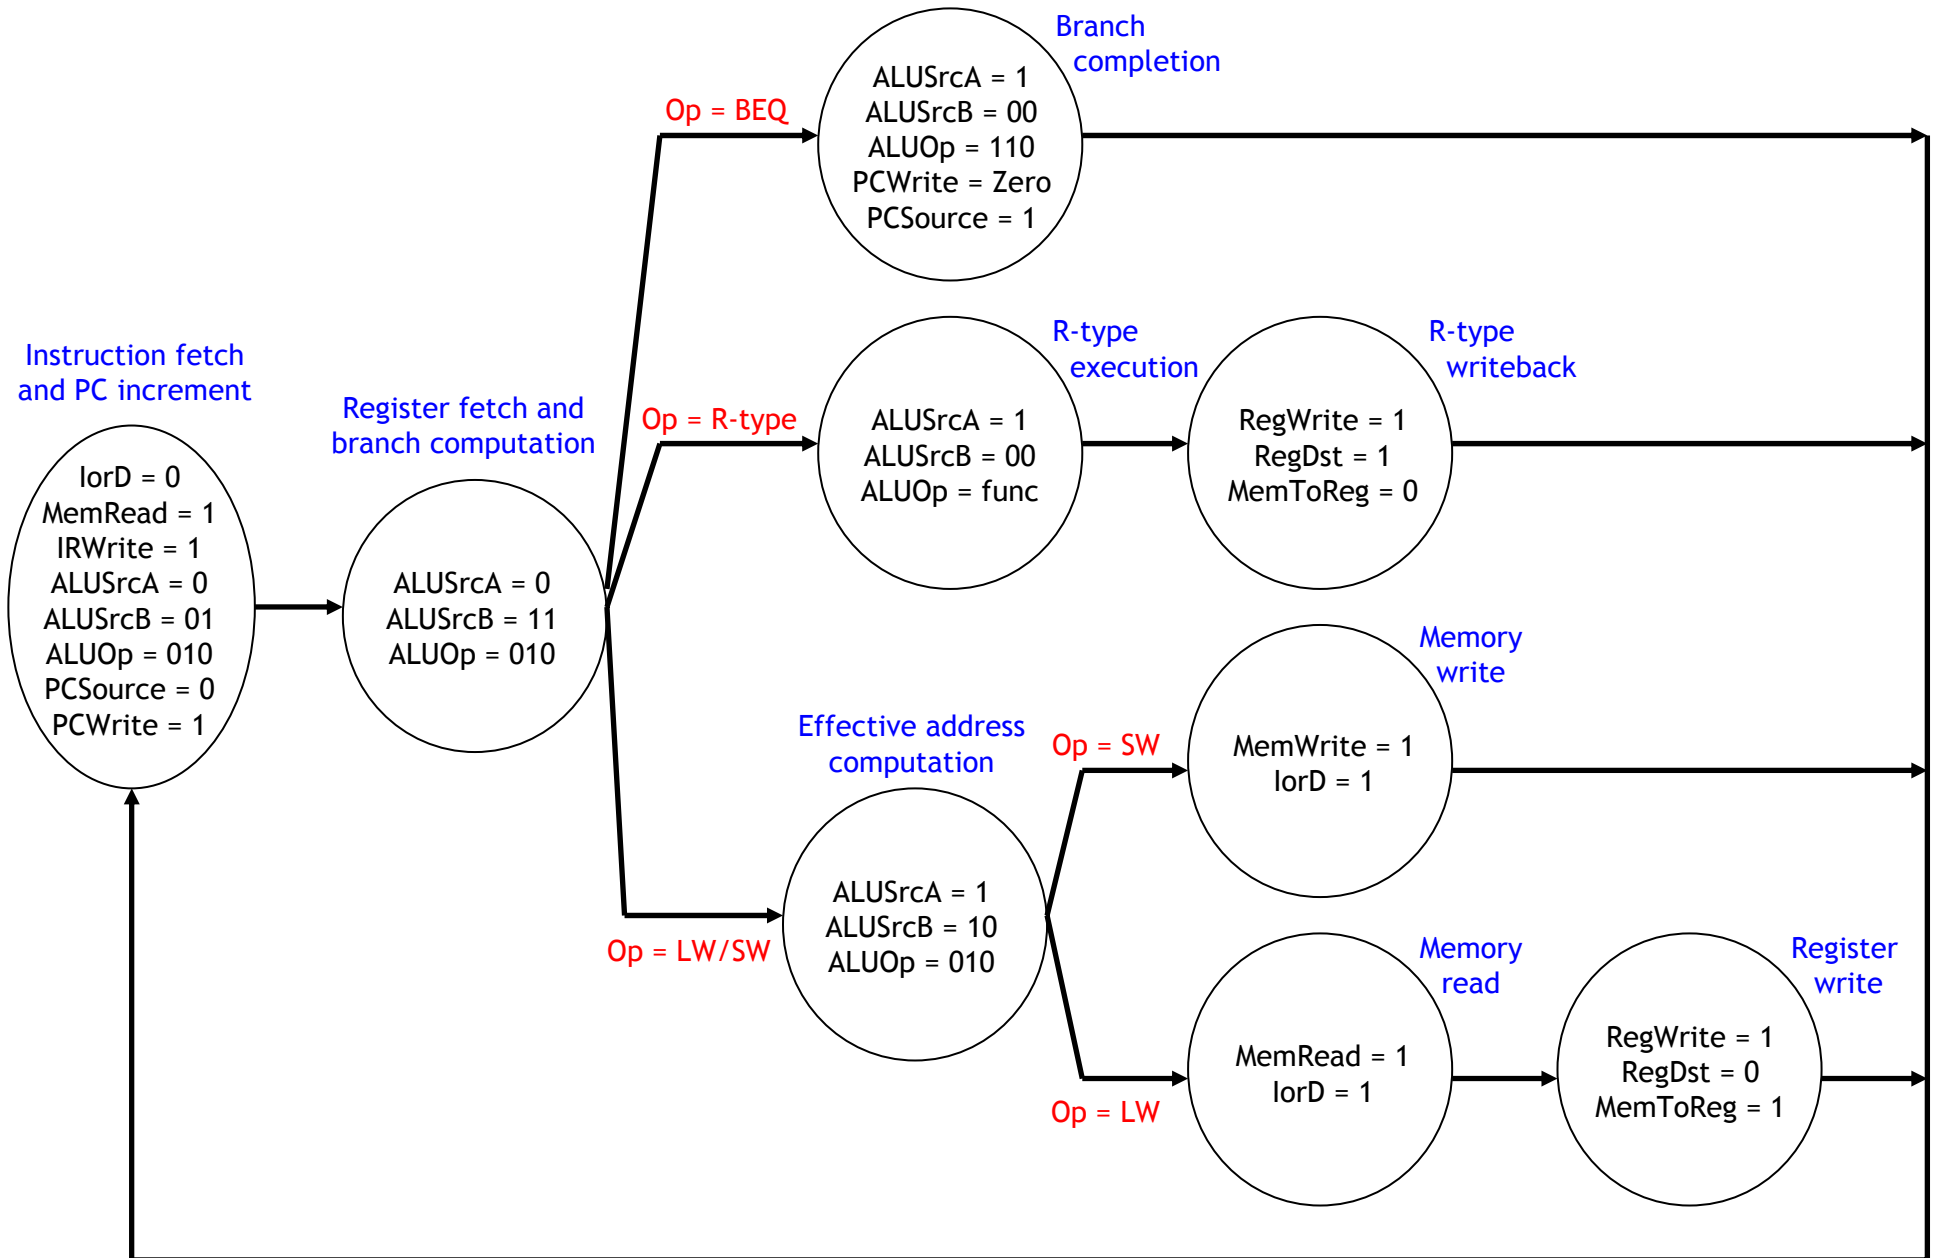


- The last few lectures covered a lot of material!
 - We introduced a **multicycle datapath**, where different instructions take different numbers of cycles to execute.
 - A **multicycle control unit** is a big state machine that generates the correct sequence of control signals for each type of instruction.
- Today we'll wrap up multicycle processors with two final topics.
 - **Microprogramming** is a different approach to building control units.
 - Comparing the performance of multicycle and single-cycle designs.

The multicycle datapath



Finite-state machine for the control unit



Pitfalls of state machines

- As mentioned last time, we could translate this state diagram into a state table, and then make a logic circuit or stick it into a ROM.
- This works pretty well for our small example, but designing a finite-state machine for a larger instruction set is much harder.
 - There could be many **states** in the machine. For example, some MIPS instructions need 20 stages to execute in some implementations—each of which would be represented by a separate state.
 - There could be many **paths** in the machine. For example, the DEC VAX from 1978 had nearly 300 opcodes... that's a lot of branching!
 - There could be many **outputs**. For instance, the Pentium Pro's integer datapath has 120 control signals, and the floating-point datapath has 285 control signals. Dude!
- Implementing and maintaining the control unit for processors like these would be a nightmare. You'd have to work with large Boolean equations or a huge state table.

Motivation for microprogramming

- Think of the control unit's state diagram as a little program.
 - Each state represents a “command,” or a set of control signals that tells the datapath what to do.
 - Several commands are executed sequentially.
 - “Branches” may be taken depending on the instruction opcode.
 - The state machine “loops” by returning to the initial state.
- Why don't we invent a special language for making the control unit?
 - We could devise a more readable, higher-level notation rather than dealing directly with binary control signals and state transitions.
 - We would design control units by writing “programs” in this language.
 - We will depend on a hardware or software translator to convert our programs into a circuit for the control unit.

A good notation is very useful

- Instead of specifying the exact binary values for each control signal, we will define a symbolic notation that's easier to work with.
- As a simple example, we might replace $ALUSrcB = 01$ with $ALUSrcB = 4$.
- We can also create symbols that *combine* several control signals together. Instead of

lorD = 0
MemRead = 1
IRWrite = 1

it would be nicer to just say something like

Read PC

Microinstructions

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
-------	-------------	------	------	------------------	--------	-----------------	------

- We'll define **microinstructions** with eight fields.
 - These fields will be filled in symbolically, instead of in binary.
 - They determine all the control signals for the datapath. There are only 8 fields because some of them specify more than one of the 12 actual control signals.
 - A microinstruction corresponds to one execution stage, or one cycle.
- You can see that in each microinstruction, we can do something with the ALU, register file, memory, and program counter units.

Specifying ALU operations

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
-------	-------------	------	------	------------------	--------	-----------------	------

- **ALU control** selects the ALU operation.
 - **Add** indicates addition for memory offsets or PC increments.
 - **Sub** performs source register comparisons for “beq”.
 - **Func** denotes the execution of R-type instructions.
- **SRC1** is either **PC** or **A**, for the ALU’s first operand.
- **SRC2**, the second ALU operand, can be one of four different values.
 - **B** for R-type instructions and branch comparisons.
 - The constant **4** to increment the PC.
 - **Extend**, the sign-extended constant field for memory references.
 - **Extshift**, the sign-extended, shifted constant for branch targets.
- These correspond to the ALUOp, ALUSrcA and ALUSrcB control signals, except we use names like “Add” and not actual bits like “010.”

Specifying register and memory actions

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
-------	-------------	------	------	------------------	--------	-----------------	------

- **Register control** selects a register file action.
 - **Read** to read from registers “rs” and “rt” of the instruction word.
 - **Write ALU** writes ALUOut into destination register “rd”.
 - **Write MDR** saves MDR into destination register “rt”.
- **Memory** chooses the memory unit’s action.
 - **Read PC** reads an instruction from address PC into IR.
 - **Read ALU** reads data from address ALUOut into MDR.
 - **Write ALU** writes register B to address memory ALUOut.

Specifying PC actions

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
-------	-------------	------	------	------------------	--------	-----------------	------

- **PCWrite control** determines what happens to the PC.
 - **ALU** sets PC to ALUOut, used in incrementing the PC.
 - **ALU-Zero** writes ALUOut to PC only if the ALU's Zero condition is true. This is used to complete a branch instruction.
- **Next** determines the next microinstruction to be executed.
 - **Seq** causes the next microinstruction to be executed.
 - **Fetch** returns to the initial instruction fetch stage.
 - **Dispatch i** is similar to a "switch" or "case" statement; it branches depending on the actual instruction word.

The first stage, the microprogramming way

- This is all probably easier to understand with examples.
- Below are two lines of microcode to implement the first two multicycle execution stages, instruction fetch and register fetch.
- The first line, labelled **Fetch**, involves several actions.
 - Read from memory address PC.
 - Use the ALU to compute $PC + 4$, and store it back in the PC.
 - Continue on to the next sequential microinstruction.

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
Fetch	Add	PC	4		Read PC	ALU	Seq
	Add	PC	Extshift	Read			Dispatch 1

The second stage

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
Fetch	Add	PC	4		Read PC	ALU	Seq
	Add	PC	Extshift	Read			Dispatch 1

- The second line implements the register fetch stage.
 - Read registers *rs* and *rt* from the register file.
 - Pre-compute $PC + (\text{sign-extend}(\text{IR}[15-0]) \ll 2)$ for branches.
 - Determine the next microinstruction based on the opcode of the current MIPS program instruction.

```
switch (opcode) {  
    case 4:    goto BEQ1;  
    case 0:    goto Rtype1;  
    case 43:  
    case 35:   goto Mem1;  
}
```

Completing a beq instruction

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
BEQ1	Sub	A	B			ALU-Zero	Fetch

- Control would transfer to this microinstruction if the opcode was “beq”.
 - Compute A-B, to set the ALU’s Zero bit if A=B.
 - Update PC with ALUOut (which contains the branch target from the previous cycle) if Zero is set.
 - The beq is completed, so fetch the next instruction.
- The 1 in the label BEQ1 reminds us that we came here via the first branch point (“dispatch table 1”), from the second execution stage.

Completing an arithmetic instruction

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
Rtype1	func	A	B				Seq
				Write ALU			Fetch

- What if the opcode indicates an R-type instruction?
 - The first cycle here performs an operation on registers **A** and **B**, based on the MIPS instruction's **func** field.
 - The next stage writes the ALU output to register "rd" from the MIPS instruction word.
- We can then go back to the **Fetch** microinstruction, to fetch and execute the next MIPS instruction.


Completing data transfer instructions

Label	ALU control	Src1	Src2	Register control	Memory	PCWrite control	Next
Mem1	Add	A	Extend				Dispatch 2
SW2					Write ALU		Fetch
LW2					Read ALU		Seq
				Write MDR			Fetch

- For both sw and lw instructions, we should first compute the effective memory address, $A + \text{sign-extend}(\text{IR}[15-0])$.
- Another dispatch or branch distinguishes between stores and loads.
 - For sw, we store data (from B) to the effective memory address.
 - For lw we copy data from the effective memory address to register rt.
- In either case, we continue on to **Fetch** after we're done.

Microprogramming vs. programming

- Microinstructions correspond to control signals.
 - They describe what is done in a single clock cycle.
 - These are the most basic operations available in a processor.
- Microprograms implement higher-level MIPS instructions.
 - MIPS assembly language instructions are comparatively complex, each possibly requiring multiple clock cycles to execute.
 - But each complex MIPS instruction can be implemented with several simpler microinstructions.



Similarities with assembly language

- Microcode is intended to make control unit design easier.
 - We defined symbols like `Read PC` to replace binary control signals.
 - A translator can convert microinstructions into a real control unit.
 - The translation is straightforward, because each microinstruction corresponds to one set of control values.
- Hey! This sounds similar to MIPS assembly language!
 - We use mnemonics like `lw` instead of binary opcodes like `100011`.
 - MIPS programs must be assembled to produce real machine code.
 - Each MIPS instruction corresponds to a 32-bit instruction word.

Managing complexity

- It looks like all we've done is devise a new notation that makes it easier to specify control signals.
- That's exactly right! It's all about managing complexity.
 - Control units are probably the most challenging part of CPU design.
 - Large instruction sets require large state machines with many states, branches and outputs.
 - Control units for multicycle processors are difficult to create and maintain.
- Applying programming ideas to hardware design is a useful technique.



Situations when microprogramming is bad

- One disadvantage of microprograms is that looking up control signals in a ROM can be *slower* than generating them from simplified circuits.
- Sometimes complex instructions implemented in hardware are *slower* than equivalent assembly programs written using simpler instructions
 - Complex instructions are usually very general, so they can be used more often. But this also means they can't be optimized for specific operands or situations.
 - Some microprograms just aren't written very efficiently. But since they're built into the CPU, people are stuck with them (at least until the next processor upgrade).

Fixing the bad situations

- Modern processors use a combination of hardwired logic and microcode to balance design effort with performance.
 - Control for many common instructions can be hardwired to “make the common case fast.”
 - Less-used or very complex instructions are microprogrammed to make the design easier and more flexible.
- For example, 8086 and VAX arithmetic instructions can directly access memory using a variety of exotic addressing modes.
 - A general “add” instruction must account for all possible addressing modes, and requires at least five cycles on a VAX 11/780.
 - The later VAX 8600 includes special hardwired circuitry for the most common cases—as an example, adding registers takes just one cycle.

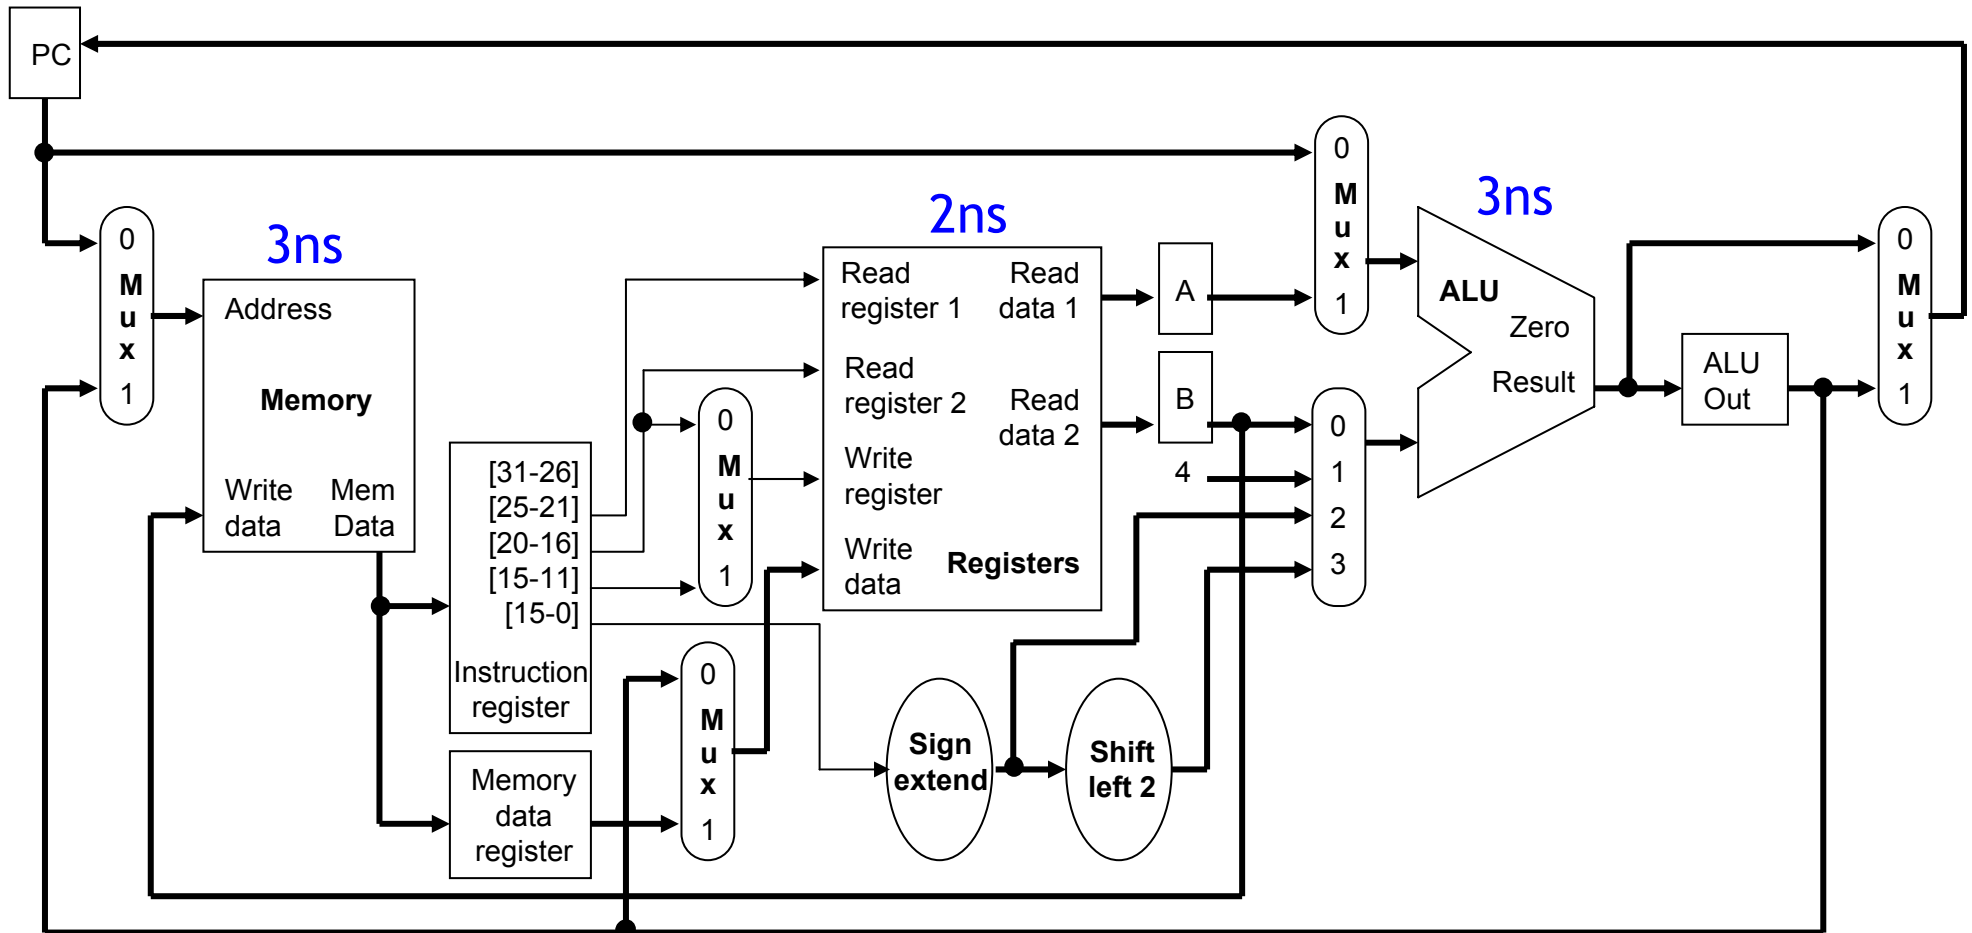
DEC VAX 11/780

- The VAX was designed in 1978 by Digital Equipment Corporation.
- It has one of the most complex instruction sets ever. (Compiler technology wasn't very good back then, and they wanted to make assembly programming easier.)
- VMS, the VAX multiuser, cluster-based operating system, was designed by Dave Cutler, who was also in charge of Windows NT.
- The VAX had a 32-bit processor, seven years before Intel's 80386.
- The cycle time was 200ns. 5MHz!
- All of this cost \$200,000.



Performance of a multicycle implementation

- Let's assume the following delays for the major functional units.



Comparing cycle times

- The clock period has to be long enough to allow all of the required work to complete within the cycle.
- In the single-cycle datapath, the “required work” was just the complete execution of any instruction.
 - The longest instruction, `lw`, requires 13ns ($3 + 2 + 3 + 3 + 2$).
 - So the clock cycle time has to be 13ns, for a 77MHz clock rate.
- For the multicycle datapath, the “required work” is only a single stage.
 - The longest delay is 3ns, for both the ALU and the memory.
 - So our cycle time has to be 3ns, or a clock rate of 333MHz.
 - The register file needs only 2ns, but it must wait an extra 1ns to stay synchronized with the other functional units.
- The single-cycle cycle time is limited by the slowest *instruction*, whereas the multicycle cycle time is limited by the slowest *functional unit*.

Comparing instruction execution times

- In the single-cycle datapath, each instruction needs an entire clock cycle, or 13ns, to execute.
- With the multicycle CPU, different instructions need different numbers of clock cycles, and hence different amounts of time.
 - A branch needs 3 cycles, or $3 \times 3\text{ns} = 9\text{ns}$.
 - Arithmetic and sw instructions each require 4 cycles, or 12ns.
 - Finally, a lw takes 5 stages, or 15ns.
- We can make some observations about performance already.
 - Loads take *longer* with this multicycle implementation, while all other instructions are faster than before.
 - So if our program doesn't have too many loads, then we should see an increase in performance.

The gcc example

- Let's assume the gcc instruction mix from p. 189, used in HW2, again.

Instruction	Frequency
Arithmetic	48%
Loads	22%
Stores	11%
Branches	19%

- In a single-cycle datapath, all instructions take 13ns to execute.
- The average execution time for an instruction on the multicycle processor works out to 12.09ns.

$$(48\% \times 12\text{ns}) + (22\% \times 15\text{ns}) + (11\% \times 12\text{ns}) + (19\% \times 9\text{ns}) = 12.09\text{ns}$$

- The multicycle implementation is faster in this case, but not by much. The speedup here is only 7.5%.

$$13\text{ns} / 12.09\text{ns} = 1.075$$

This CPU is too simple

- Our example instruction set is too simple to see large gains.
 - All of our instructions need about the same number of cycles (3-5).
 - The benefits would be much greater in a more complex CPU, where some instructions require many more stages than others.
- For example, the 80x86 and VAX both have instructions to push all the registers onto the stack in one shot (PUSHA and CALLS).
 - Pushing proceeds sequentially, register by register.
 - Implementing this in a single-cycle datapath would be foolish, since the instruction would need an obscene amount of time to store each register into memory.
 - But the 8086 and VAX are multicycle processors, so these complex instructions don't slow down the cycle time or other instructions.

Summary

- A **multicycle processor** splits instruction execution into several stages, each of which requires one clock cycle.
- Multicycle CPUs have two important advantages over single-cycle ones.
 - Each instruction can be executed in as few stages as necessary.
 - Redundant hardware units can be eliminated.
- However, multicycle control is more complex overall.
 - Extra multiplexers and temporary registers are needed.
 - The **control unit** must generate *sequences* of control signals.
- There are a few options in designing the control unit.
 - Generate control signals from MIPS instructions using a hardwired circuit or a programmable device like a ROM or PLA.
 - Design a **microprogramming** language, with microcode that can be automatically translated into the correct control signals.