

Roadmap

C:

```
car *c = malloc(sizeof(car));
c->miles = 100;
c->gals = 17;
float mpg = get_mpg(c);
free(c);
```

Java:

```
Car c = new Car();
c.setMiles(100);
c.setGals(17);
float mpg =
    c.getMPG();
```

Memory & data
Arrays & structs
Integers & floats
RISC V assembly
Procedures & stacks
Executables
Memory & caches
Processor Pipeline
Performance
Parallelism

Assembly
language:

```
get_mpg(car*):
    lw    a5,0(a0)
    lw    a4,4(a0)
    divw  a5,a5,a4
    fcvt.s.w    fa0,a5
    ret
```



Machine
code:

```
0111010000011000
100011010000010000000010
1000100111000010
110000011111101000011111
```

OS:



Computer
system:



ACKNOWLEDGEMENT: These slides have been created by your CMPT 295 instructor and RISC-V ISA creators. However, please report all mistakes to your instructor.

RISC-V *CPU Datapath, Control Intro*

Putting it together!

High level languages (ex. C) become machine language through compilation, assembly, and linking.

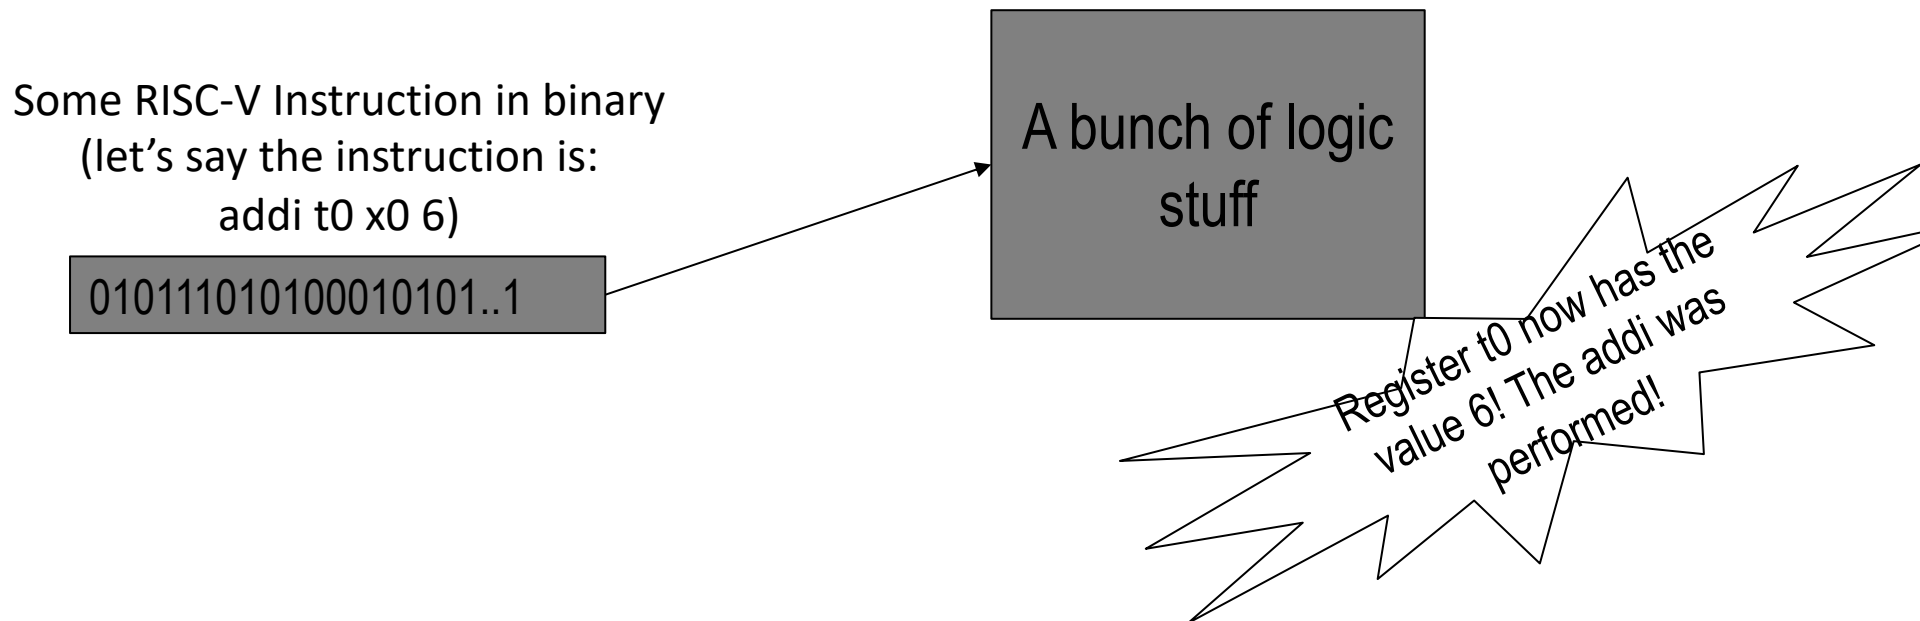
Registers, clock circuits, gates, and other logic devices are the fundamental building blocks of digital decision-making

CPU

```
graph TD; A[High level languages (ex. C) become machine language through compilation, assembly, and linking.] --> C[CPU]; B[Registers, clock circuits, gates, and other logic devices are the fundamental building blocks of digital decision-making] --> C;
```

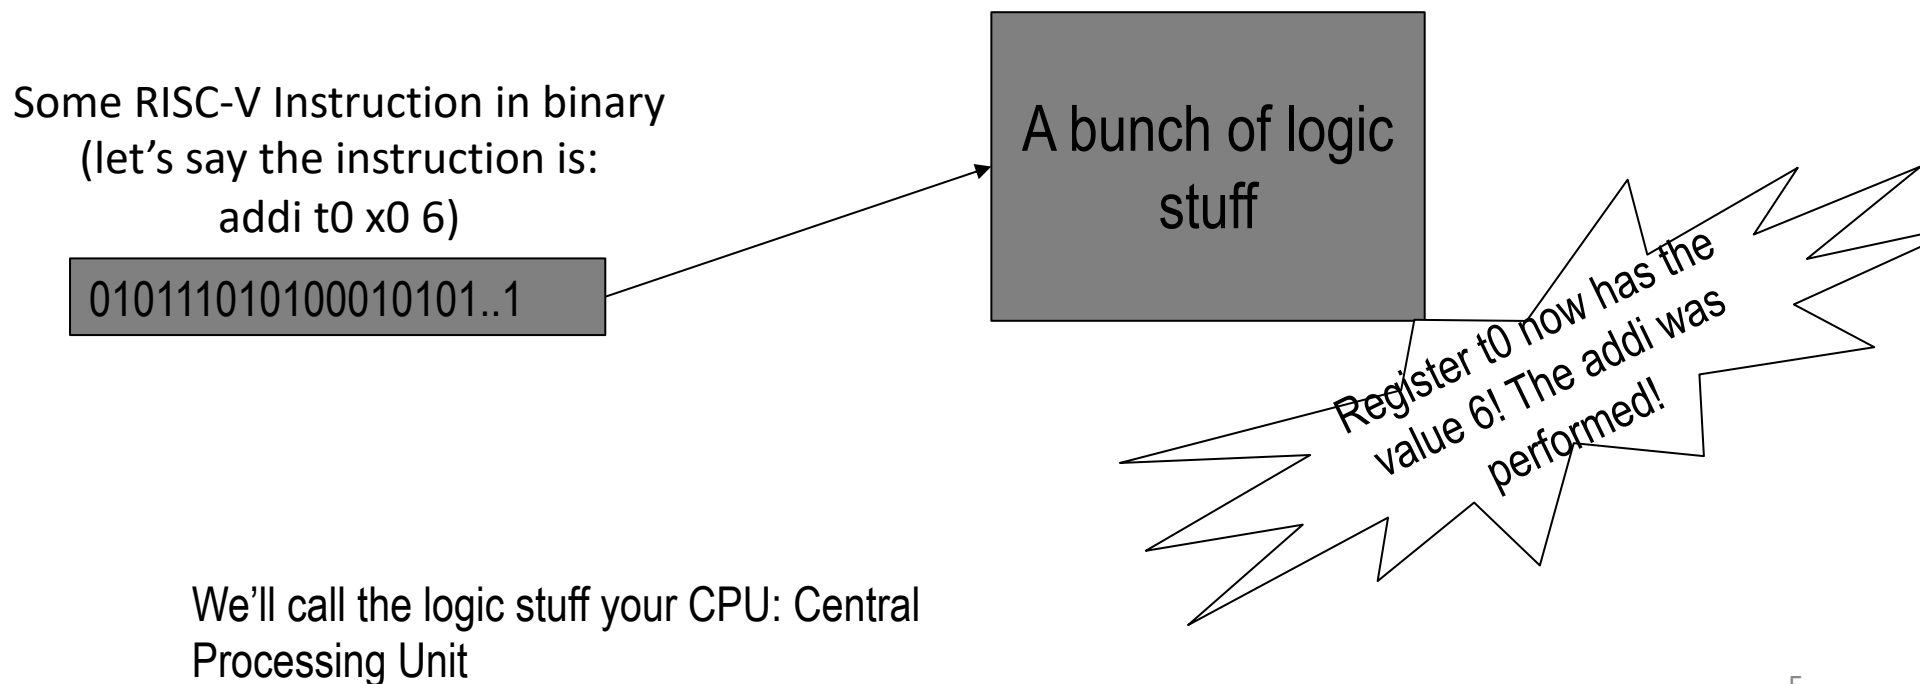
Today's goal:

Create a “circuit” of logic elements that, when given an assembly instruction, perform the action the instruction describes



Today's goal:

Create a “circuit” of logic elements that, when given an assembly instruction, perform the action the instruction describes



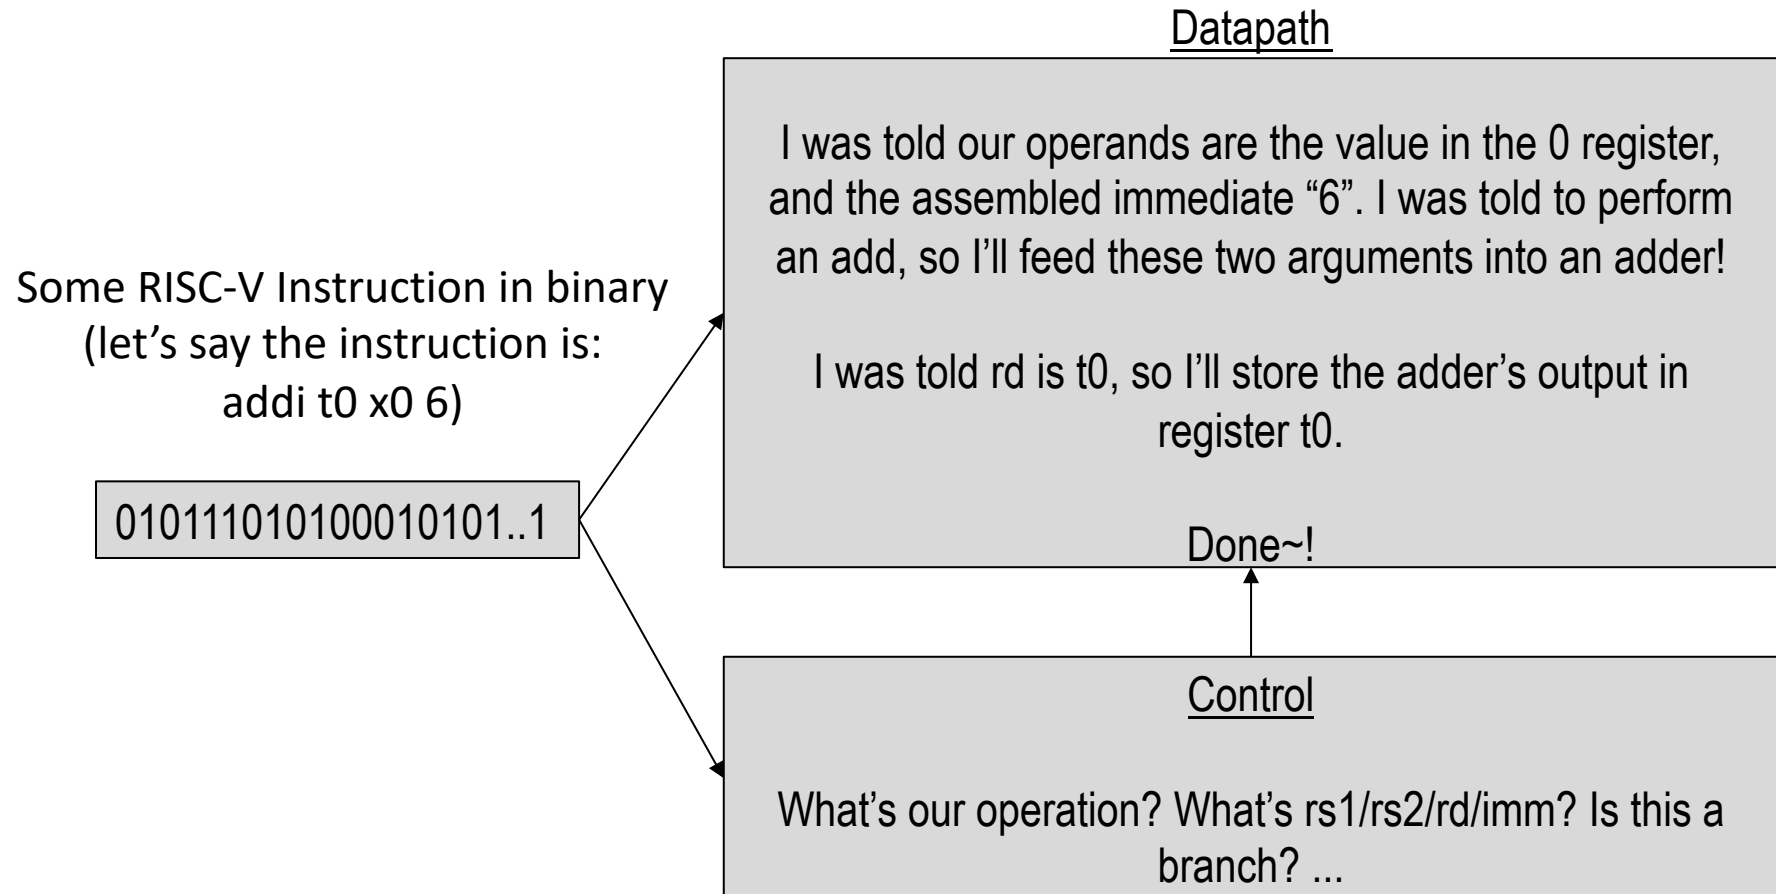
Agenda

- What's a CPU?
- Building from what we know
- Our CPU
- Processor Design Principles

Your CPU in two parts

- **Central Processing Unit (CPU):**
 - **Datapath:** contains the hardware necessary to perform operations required by the processor
 - Reacts to what the controller tells it! (ie. “I was told to do an add, so I’ll feed these arguments through an adder)
 - **Control:** decides what each piece of the datapath should do
 - What operation am I performing? Do I need to get info from memory? Should I write to a register? Which register?
 - Has to make decisions based on the input instruction only!

Your CPU in two parts



Designing our Datapath: Where to start?

- 📖 Let's start with a broad question:
 - 📖 What operations does our datapath need to be capable of performing?

- 📖 And also maybe a more specific one:
 - 📖 How can we ensure, when we build this, that all RISC-V instructions will be supported?

Designing our Datapath: Where to start?

- 📁 6 different formats: R, I, S, SB, U, UJ
 - 📁 Arithmetic, Immediate, Store, Branch, Upper-immediate, JAL
- 📁 Instructions are classified into these formats based on their behaviours, meaning each type does something a little different!
 - 📁 If we're building a CPU to run /all/ instructions, we'll need to figure out what functionalities each type needs
→ support them all!

Agenda

- What's a CPU?
- **Building from what we know**
- Our CPU
- Processor Design Principles

Working with an 'R-type'...

Work with the people around you. What needs to happen for an 'R-type' inst to execute?

 Wanna work with an example? Use:

add t0 t2 t3

Come up with a list of actions...

Working with an 'R-type'

What needs to be done before our R type is executed or evaluated?

- (1) Get the instruction
- (2) Parse instruction fields (rd, rs1, rs2, operation...)
- (3) Read data based on parsed operands
- (4) Perform operation
- (5) Write result to our destination register

Working with an 'R-type'

 (1) Get the instruction

 add t0 t2 t3

 (2) Parse instruction fields (rd, rs1, rs2, operation...)

 rd = t0 rs1 = t2 rs2 = t3

 (3) Read data based on parsed operands

 R[t2] R[t3]


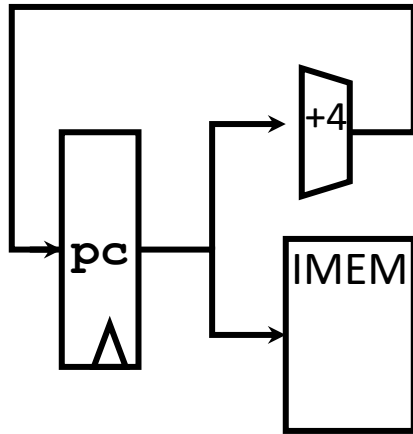
 (4) Perform operation

 R[t2] + R[t3]

 (5) Write result to our destination register

 R[t0] = R[t2] + R[t3]


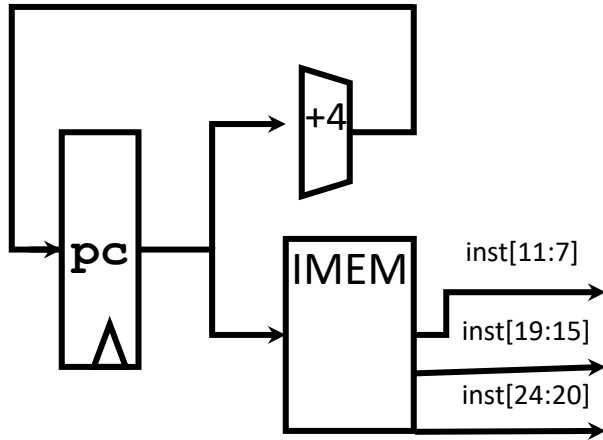
Implementing R-Types



(1) Get the instruction

- PC holds the address of our current instruction
- Where are the bits making up our inst stored?
- How does PC change after an R-Type is executed?

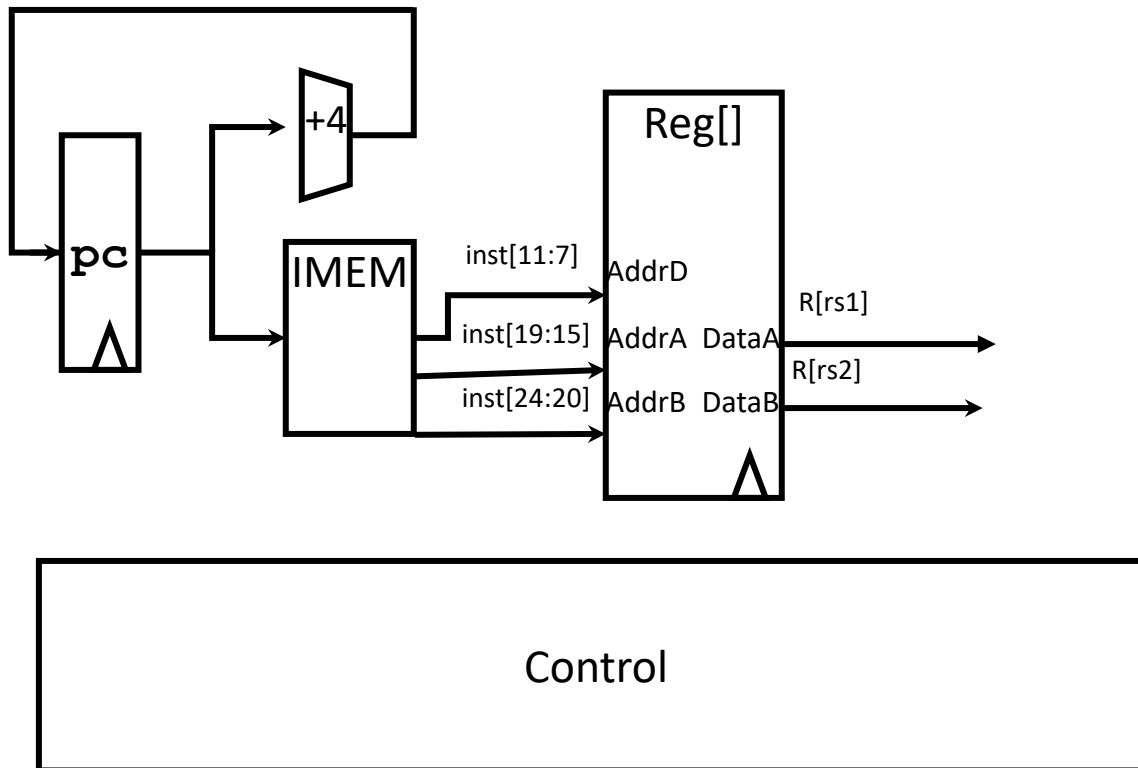
Implementing R-Types



(2) Parse instruction fields (rd, rs1, rs2, operation...)

- What registers are we operating on?
- Where do they lie in our instruction format?
- How big is each field?

Implementing R-Types

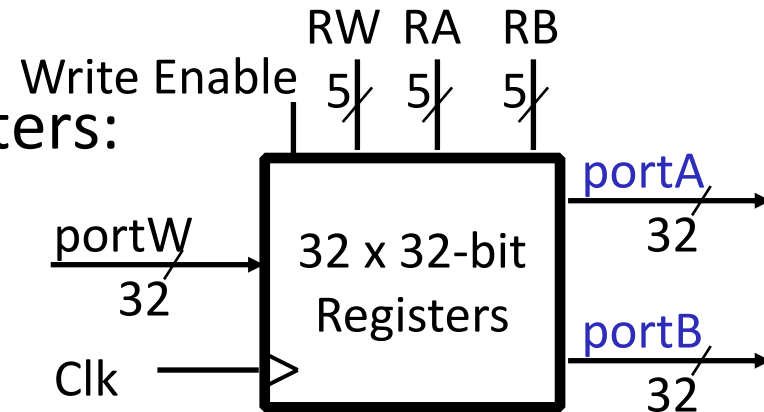


(3) Read data based on parsed operands

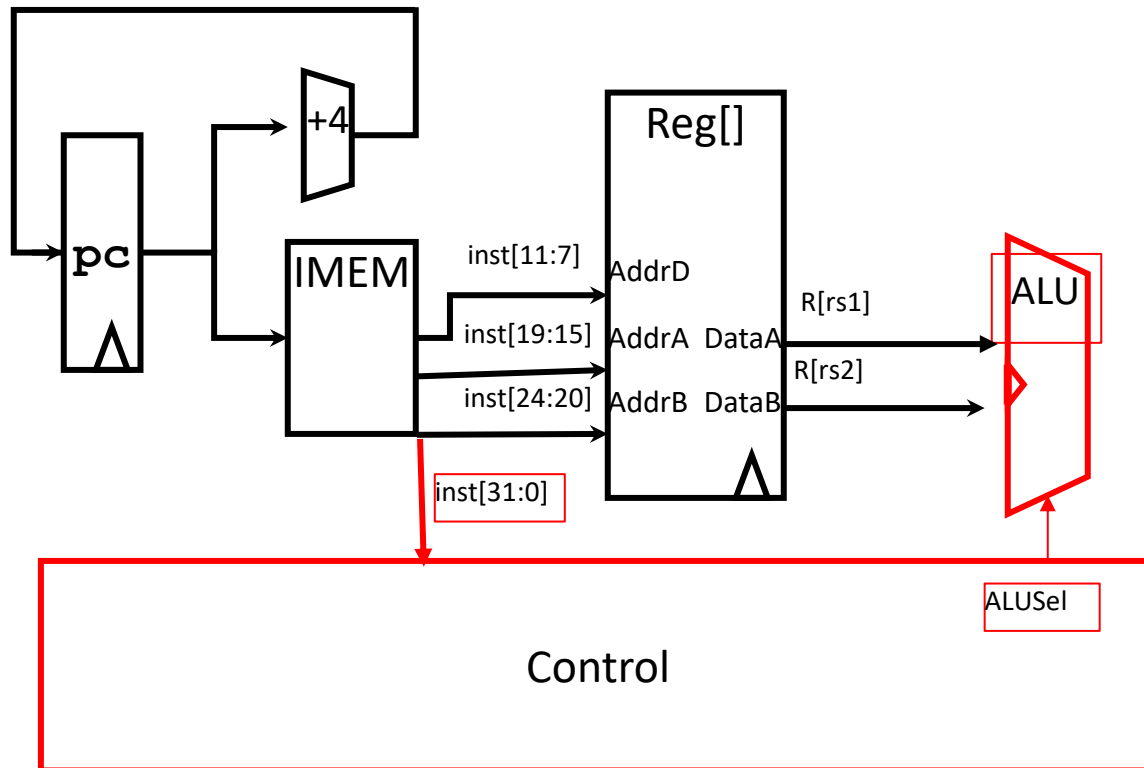
- New hardware: Register File
- Abstraction for all our registers (x0..x31 minus PC) and some mux'ing
- Reading, Writing happens here

Storage Element: Register File

- **Register File** consists of 31 registers:
 - Output ports **portA** and **portB**
 - Input port **portW**
- Register selection
 - Place data of register **RA** (number) onto **portA**
 - Place data of register **RB** (number) onto **portB**
 - Store data on **portW** into register **RW** (number) when **Write Enable** is 1
- Clock input (CLK)
 - CLK is passed to all internal registers so they can be written to if they match **RW** and **Write Enable** is 1



Implementing R-Types



(4) Perform operation

- New hardware: ALU (Arithmetic Logic Unit)
- Abstraction for adders, multipliers, dividers, etc.
- How do we know what operation to execute?
 - Our first control bit! `ALUSel(ect)`


But wait! There are many R-Type operations!


0000000	rs2	rs1	000	rd	0110011	ADD
0100000	rs2	rs1	000	rd	0110011	SUB
0000000	rs2	rs1	001	rd	0110011	SLL
0000000	rs2	rs1	010	rd	0110011	SLT
0000000	rs2	rs1	011	rd	0110011	SLTU
0000000	rs2	rs1	100	rd	0110011	XOR
0000000	rs2	rs1	101	rd	0110011	SRL
0100000	rs2	rs1	101	rd	0110011	SRA
0000000	rs2	rs1	110	rd	0110011	OR
0000000	rs2	rs1	111	rd	0110011	AND

ALU == “Arithmetic Logic Unit”

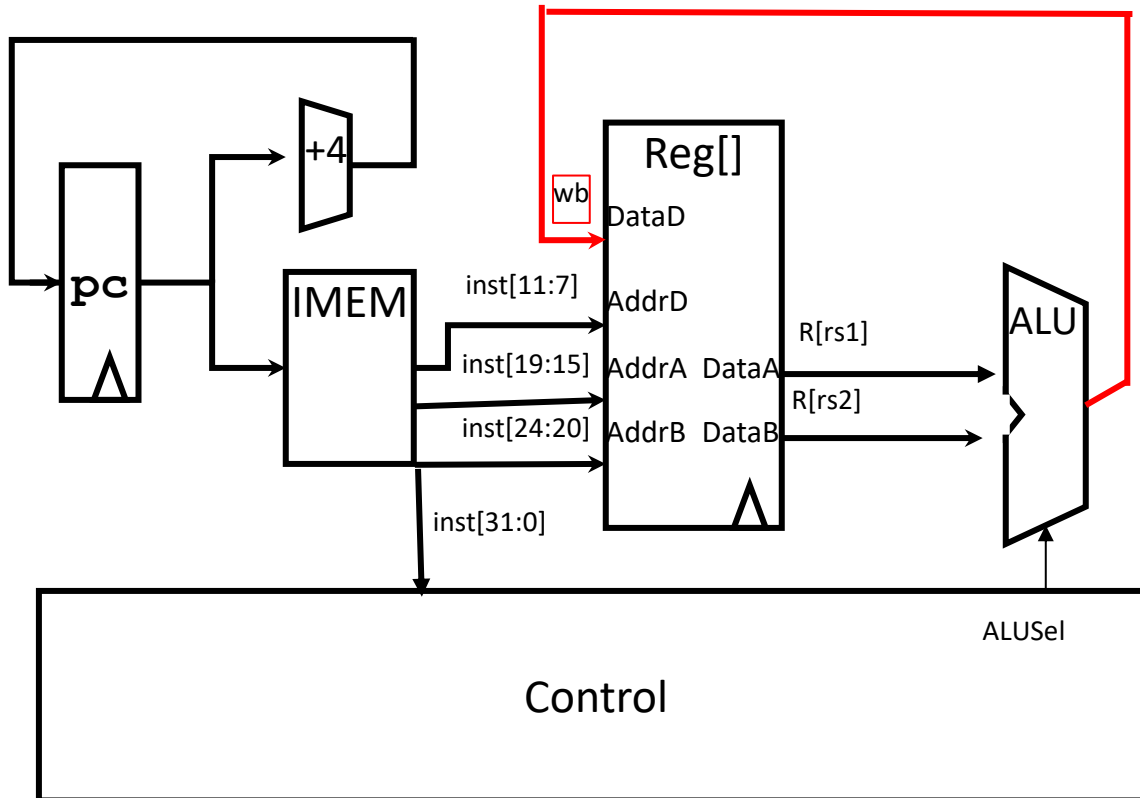
ALUSel is a control bit which encodes the operation we should perform on the given operands

- The value of ALUSel is a mapping from func3 and func7 values to operations

 “if func3 == 000 and func7 == 0000000, perform an add”

 Multiple func3, func7 combinations might lead to the same operation! (ie. add and addi)

Implementing R-Types



(5) Write result to our destination register

- The data we want to write is the result of computing operation on operands, ie. the output from our ALU
- Send it back to the regfile for writing

What changes with an arithmetic 'I-type'?

- A: Get the instruction
- B: Parse instruction fields (rd, rs1, rs2)
- C: Read data based on parsed operands
- D: Perform operation
- E: Write result to our destination register

Think: Do we need to add more hardware, more control, or more of both?

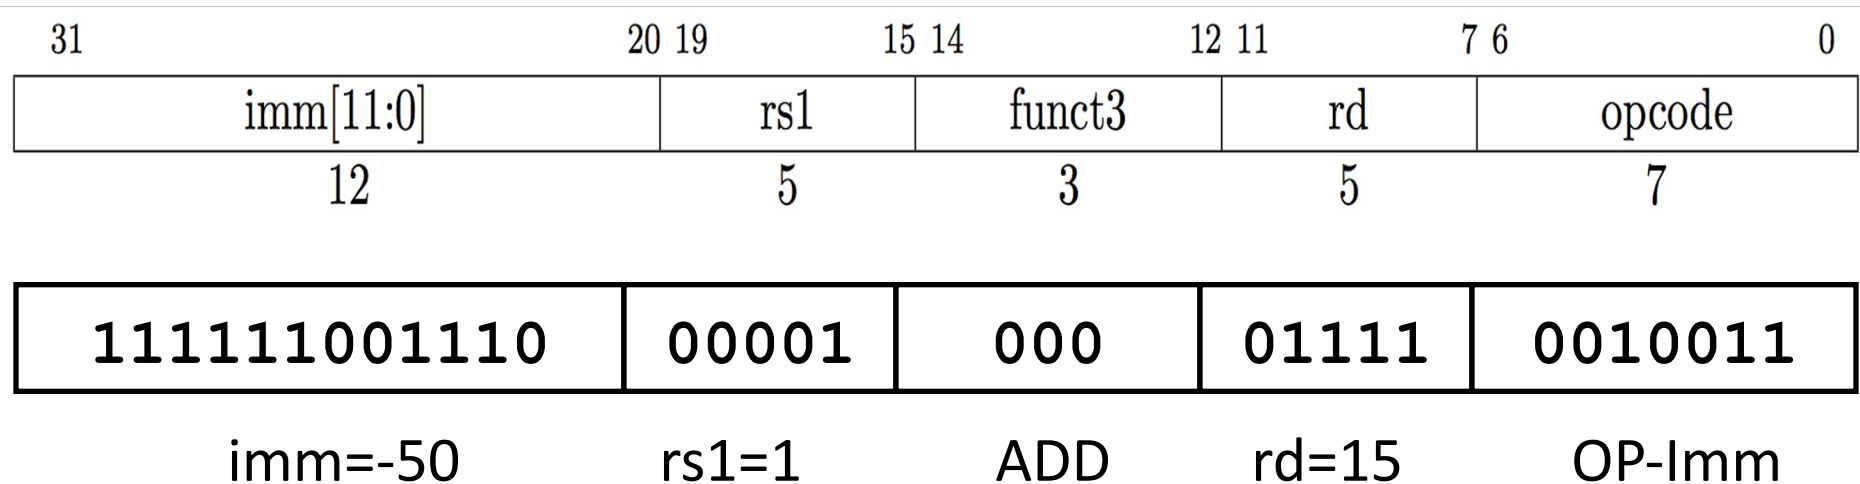
What changes with an arithmetic 'I-type'?

- A: Get the instruction
- B: Parse instruction fields (rd, rs1, rs2, operation...)
 - We also need to parse (and reassemble) our immediate!
- C: Read data based on parsed operands
 - Also an okay answer!
- D: Perform operation
- E: Write result to our destination register

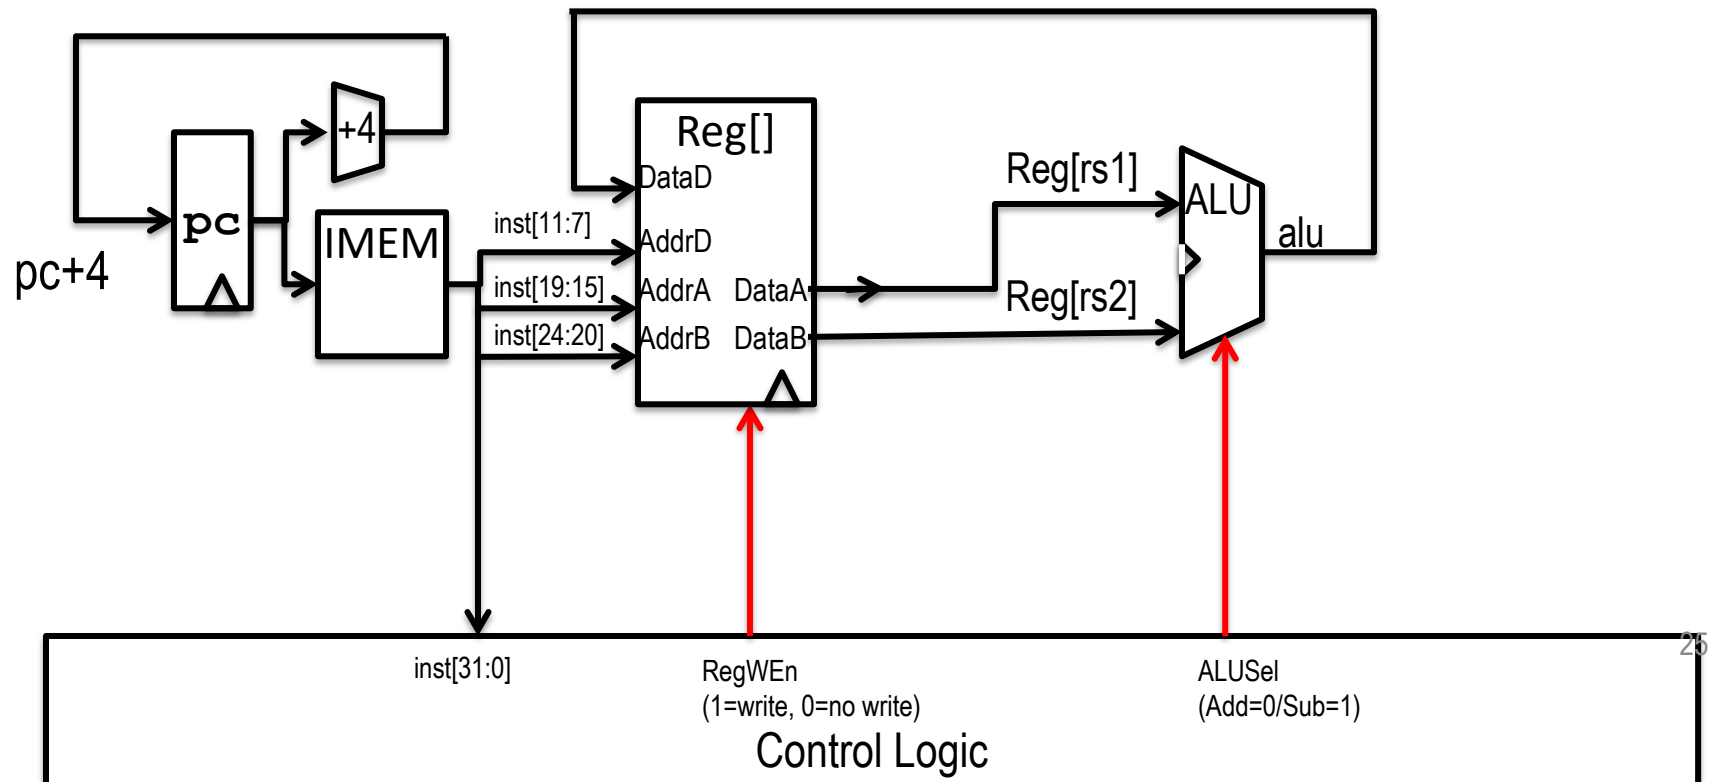
Implementing the `addi` instruction

- RISC-V Assembly Instruction:

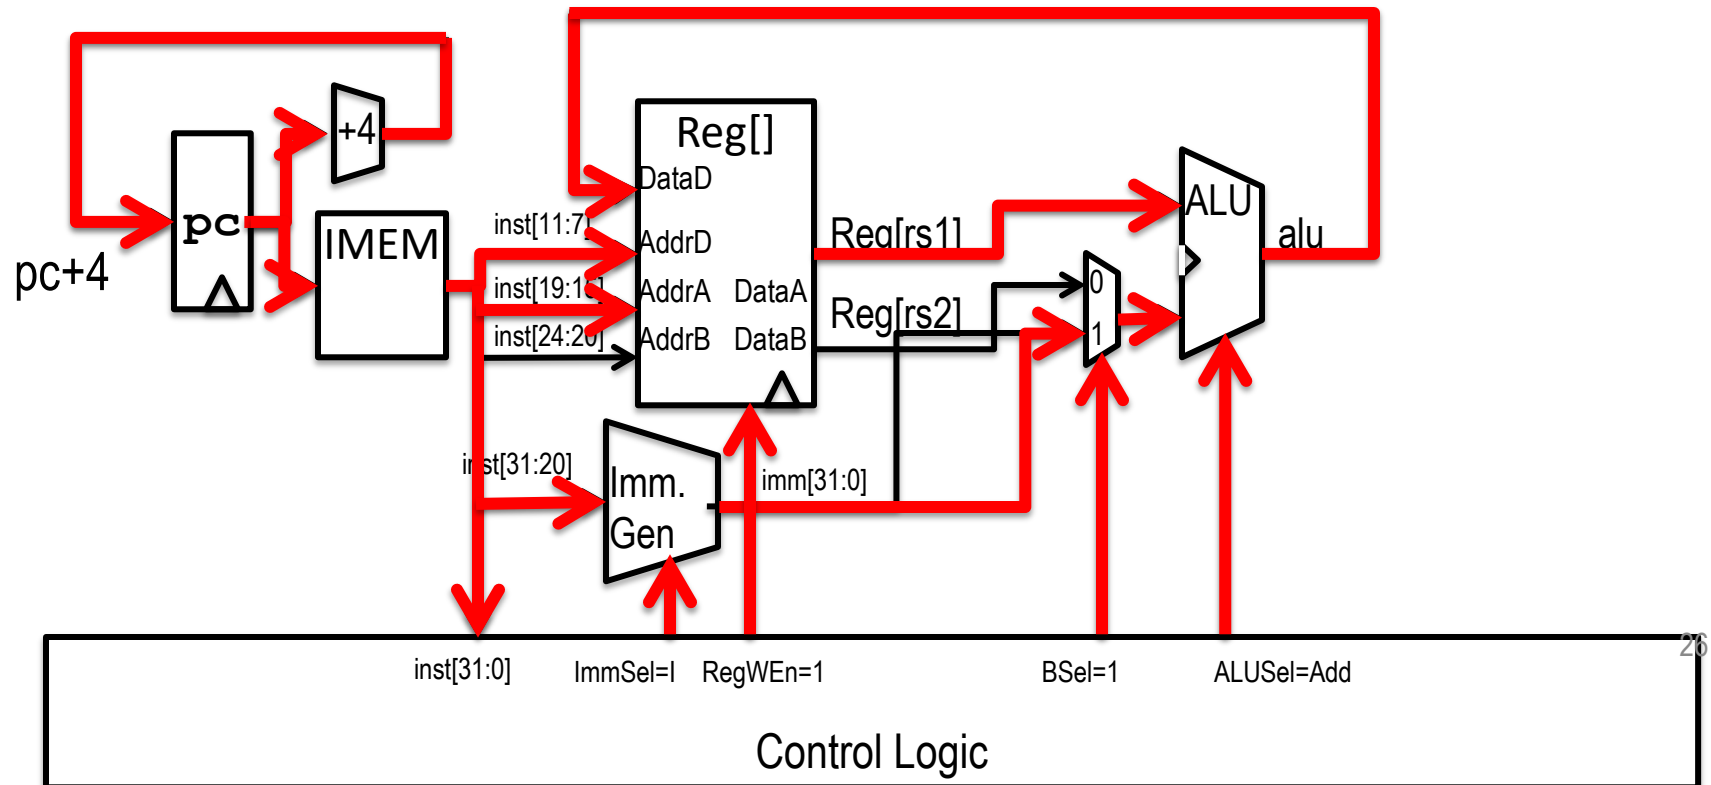
```
addi x15, x1, -50
```




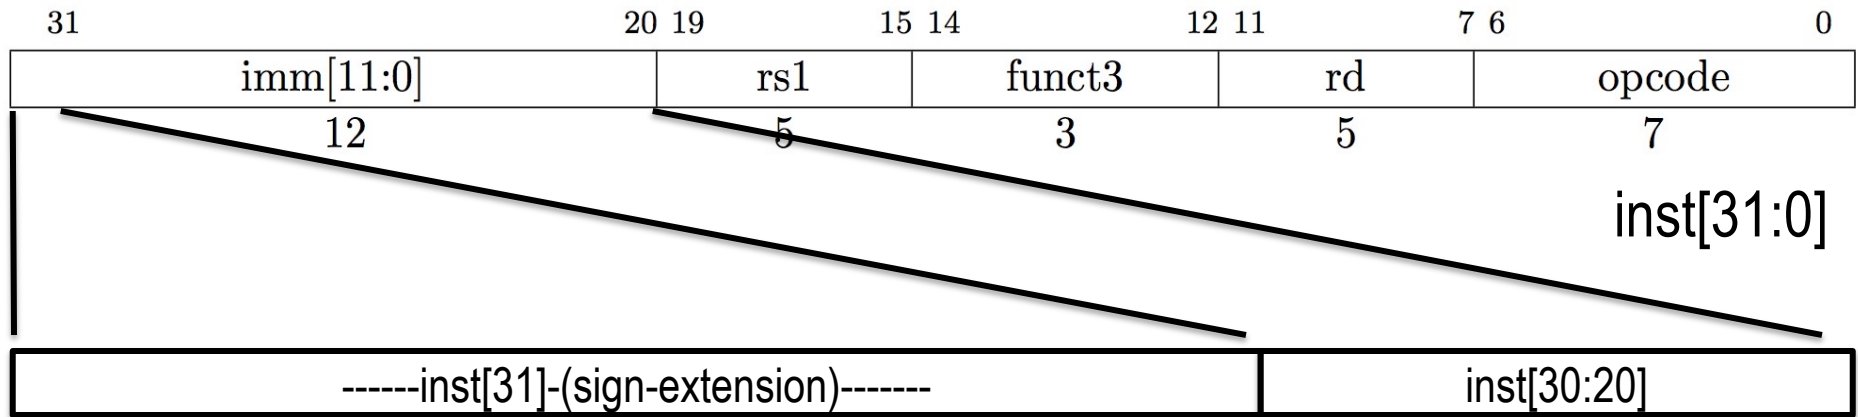
Datapath for add/sub



Adding addi to datapath

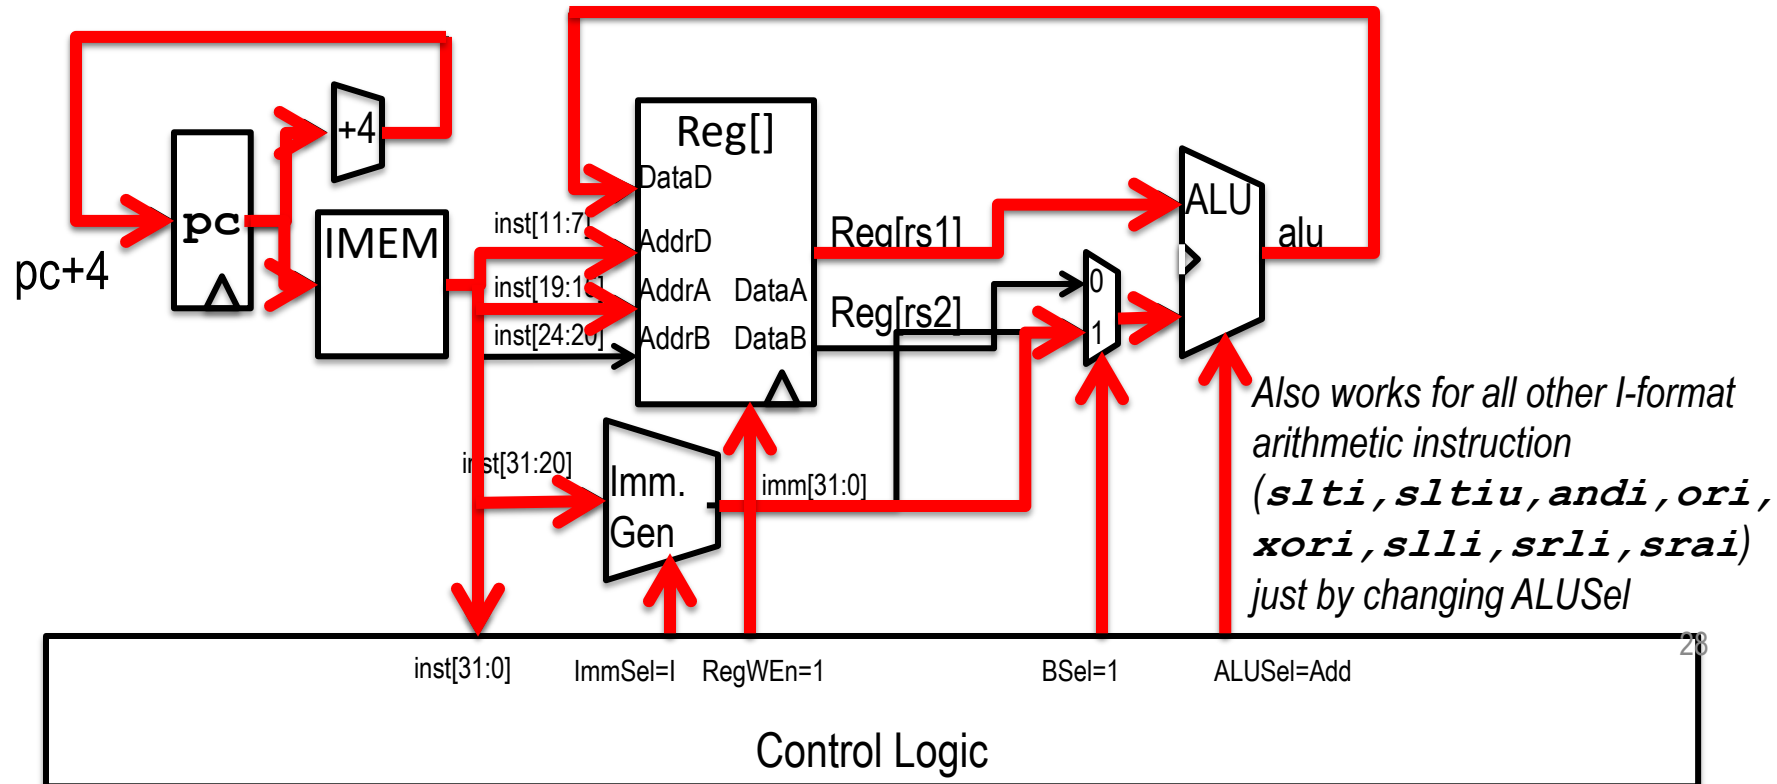


I-Format immediates



- High 12 bits of instruction ($inst[31:20]$) copied to low 12 bits of immediate ($imm[11:0]$)
- Immediate is sign-extended by copying value of $inst[31]$ to fill the upper 20 bits of the immediate value ($imm[31:12]$)

Adding addi to datapath




But wait... Loads are 'I-type's also?

We know we can parse the immediate in the load-word format, but...

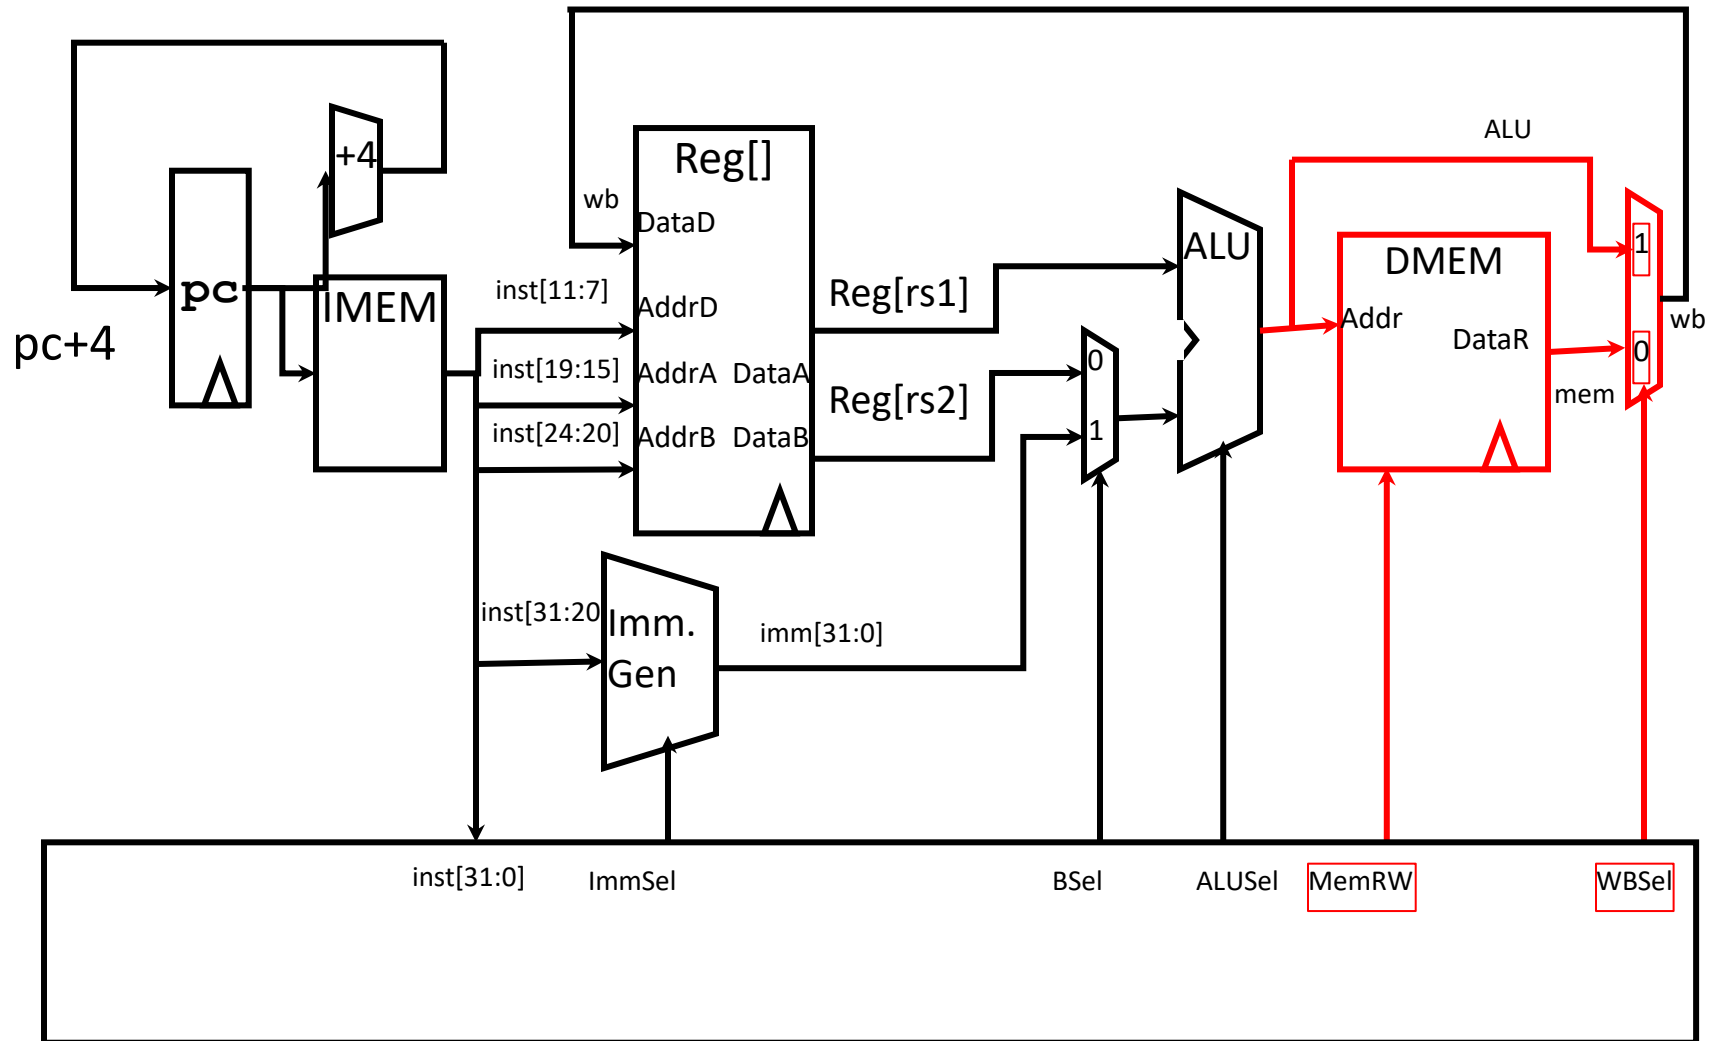
 What do we do with the immediate?

 What operation should we perform?

 Maybe a better question: How do we know what the instruction does?


Any ideas?

Adding `lw` to datapath










Storage Element: Idealized Memory

- Memory (idealized)
 - One input port: Data In
 - One output port: **Data Out**
- Memory access:
 - Read: Write Enable = 0, data at Address is placed on **Data Out**
 - Write: Write Enable = 1, Data In written to Address
- Clock input (CLK)
 - CLK input is a factor ONLY during write operation
 - During read, behaves as a combinational logic block: Address valid → **Data Out** valid after “access time”



A few notes on our new datapath...

-  We have a lot of different components!
 -  IMEM, Register file, ALU, DMEM
-  Does every instruction need every component?
 -  No! We got through all of the R-Types (and some of the I-Types) without DMEM
-  Does any instruction need every component?
 -  Yep! Loads!
 -  This is the instruction which exercises our “critical path”

All RV32 Load Instructions

imm[11:0]	rs1	000	rd	0000011	LB
imm[11:0]	rs1	001	rd	0000011	LH
imm[11:0]	rs1	010	rd	0000011	LW
imm[11:0]	rs1	100	rd	0000011	LBU
imm[11:0]	rs1	101	rd	0000011	LHU





↑ funct3 field encodes size and signedness of load data


- Supporting the narrower loads requires additional circuits to extract the correct byte/halfword from the value loaded from memory, and sign- or zero-extend the result to 32 bits before writing back to register file.
- We'll assume these are implemented in the DMEM module (not our datapath) and won't add them to our schematic

Agenda

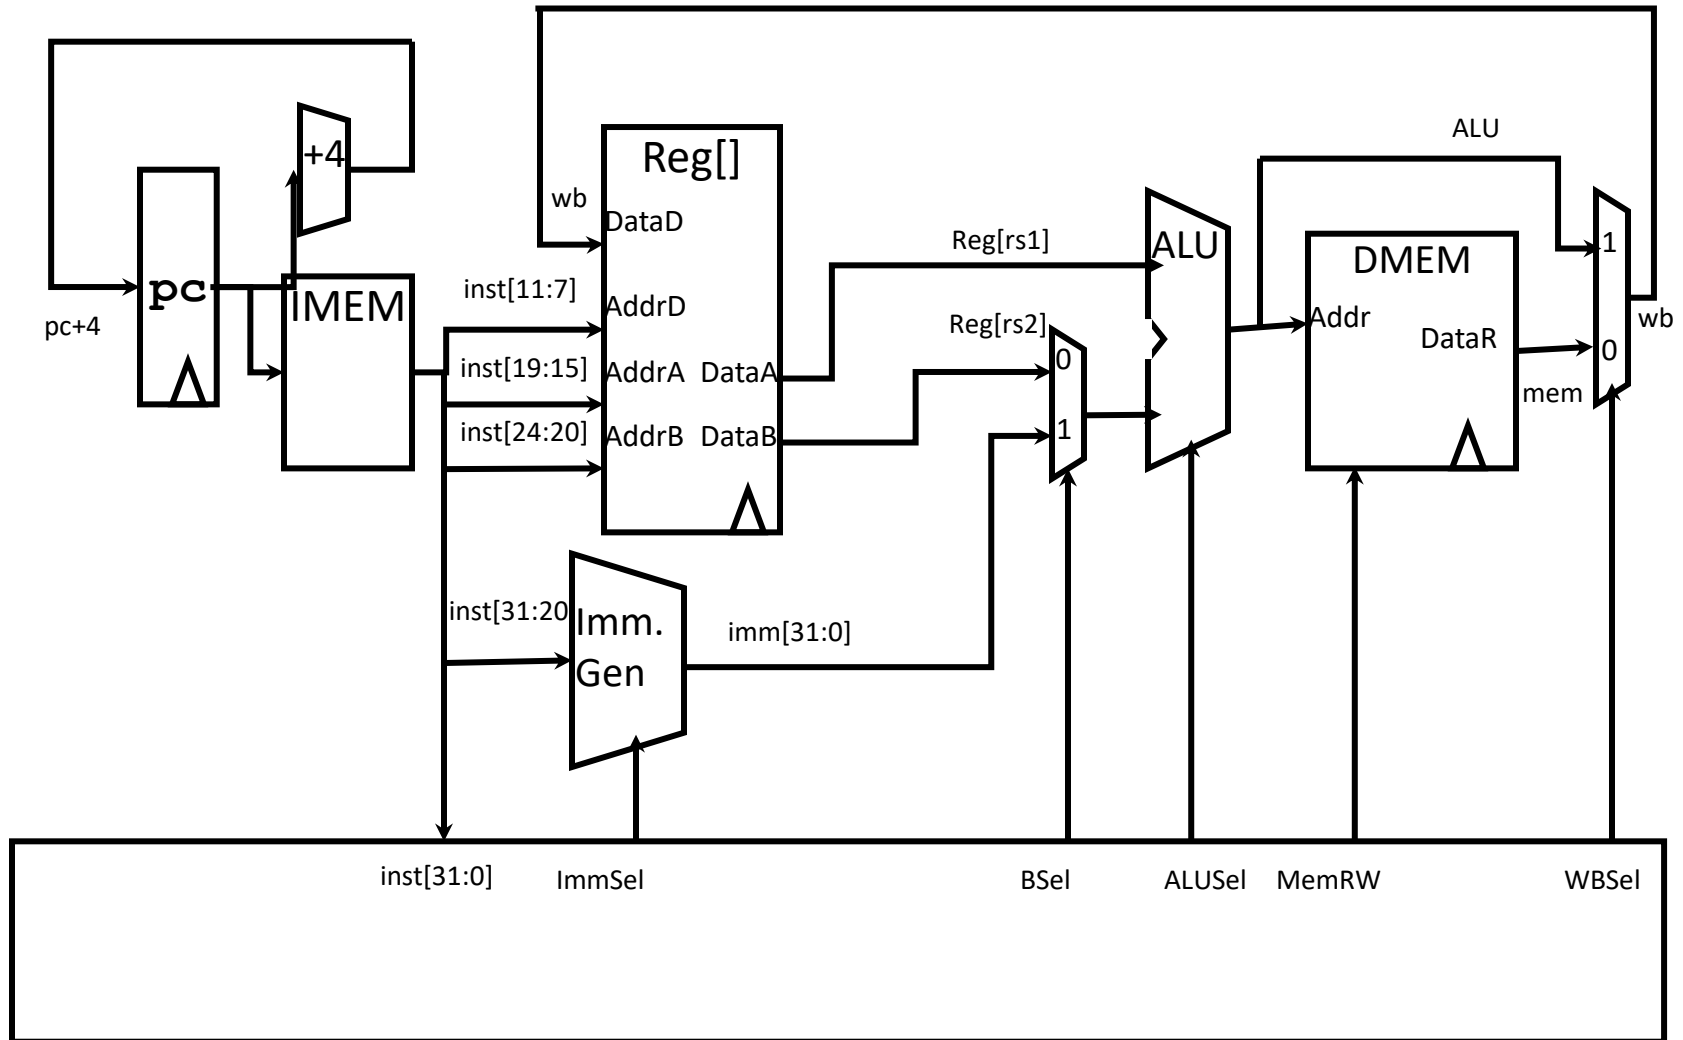
- Building from what we know
- Our CPU
- Processor Design Principles

This lecture is long... what have we done again?

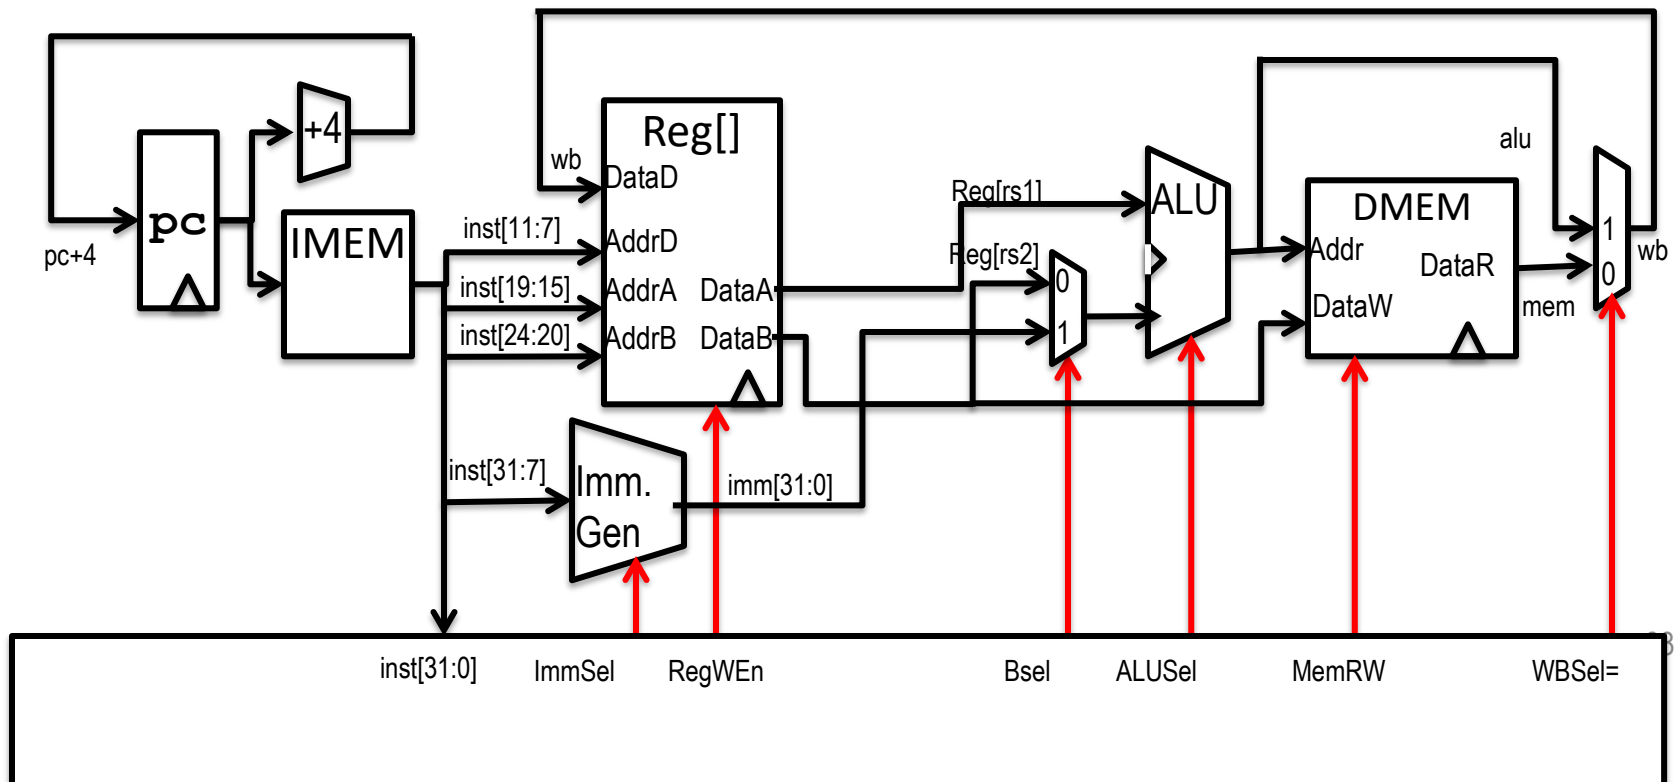
-  Added 'R-type' and 'I-type' instructions to our data path!
-  We still need to figure out how to do S, SB, U, and UJ types
 -  Don't worry, we actually have all the big pieces we need!
-  We talked about control bits and how they instruct our hardware to deal with many different kinds of instructions

-  Let's keep going!

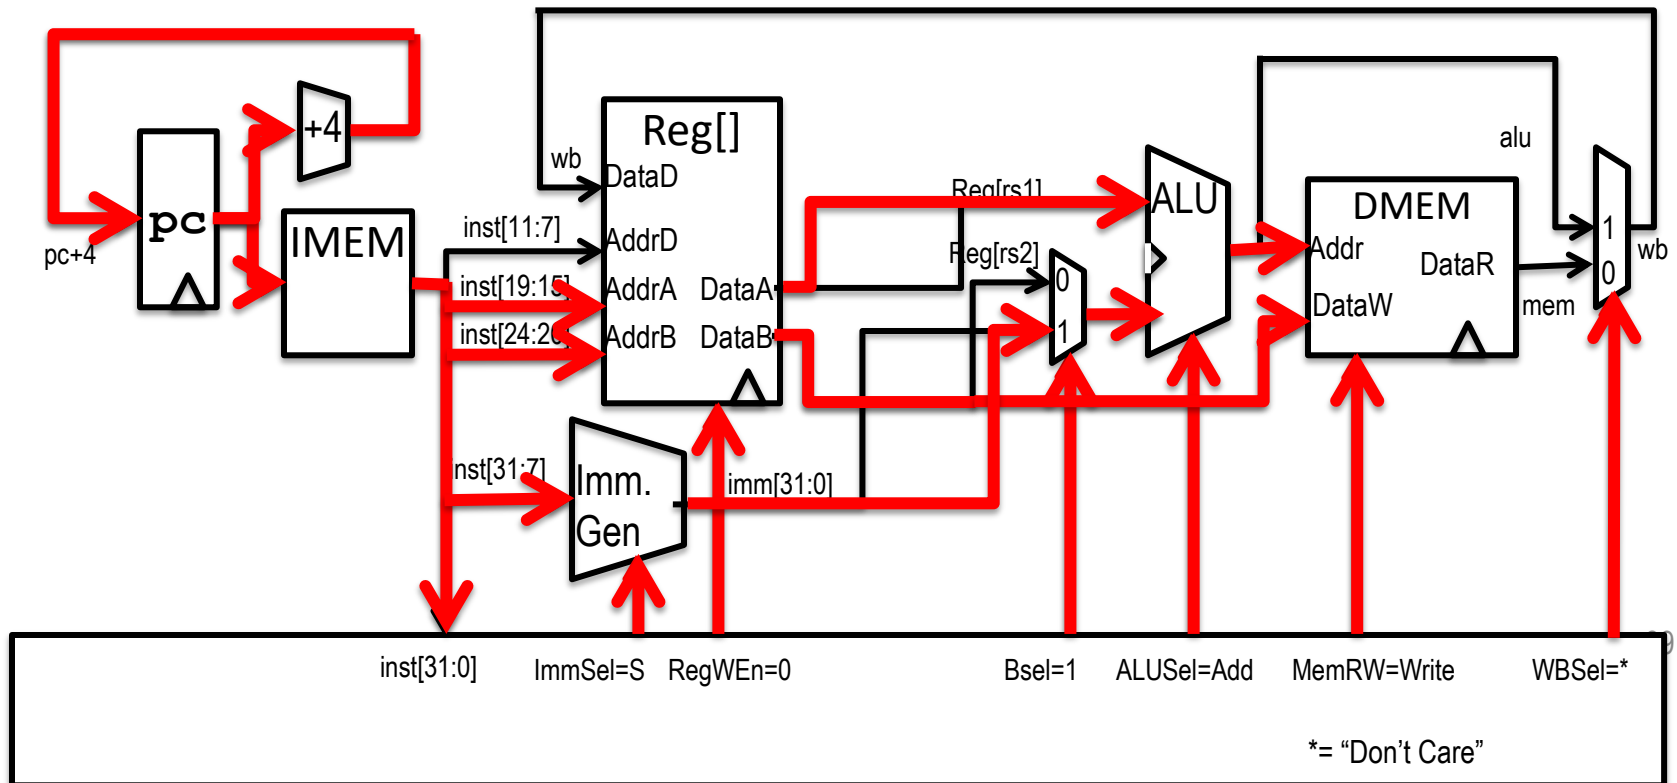
Current Datapath




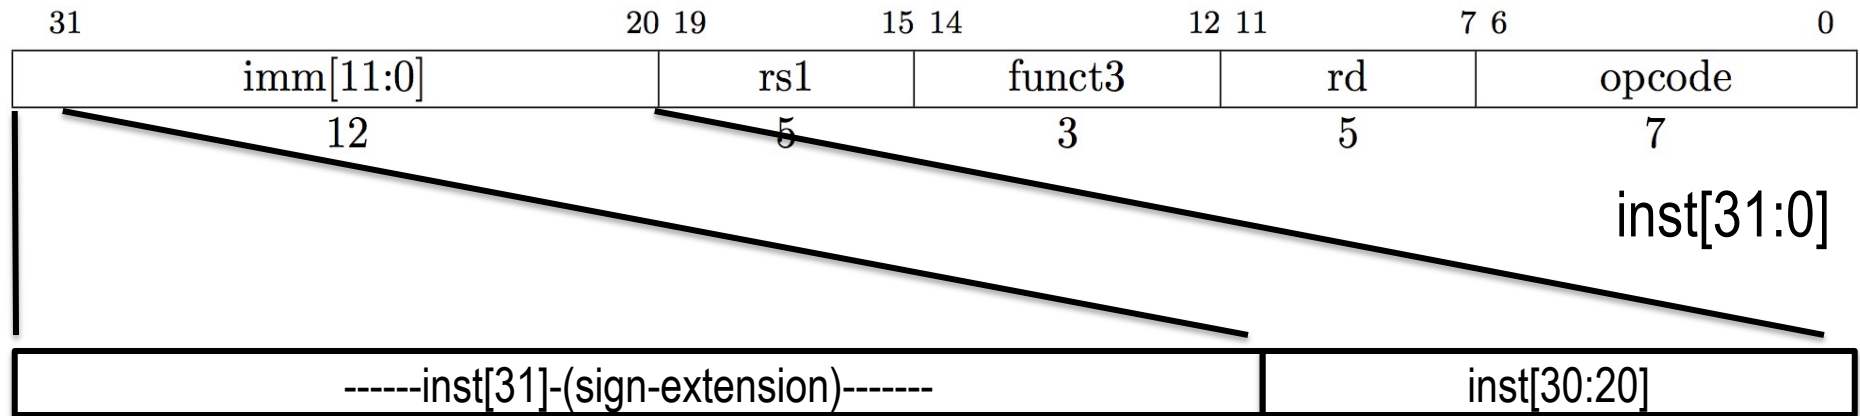
Adding sw to datapath



Adding sw to datapath

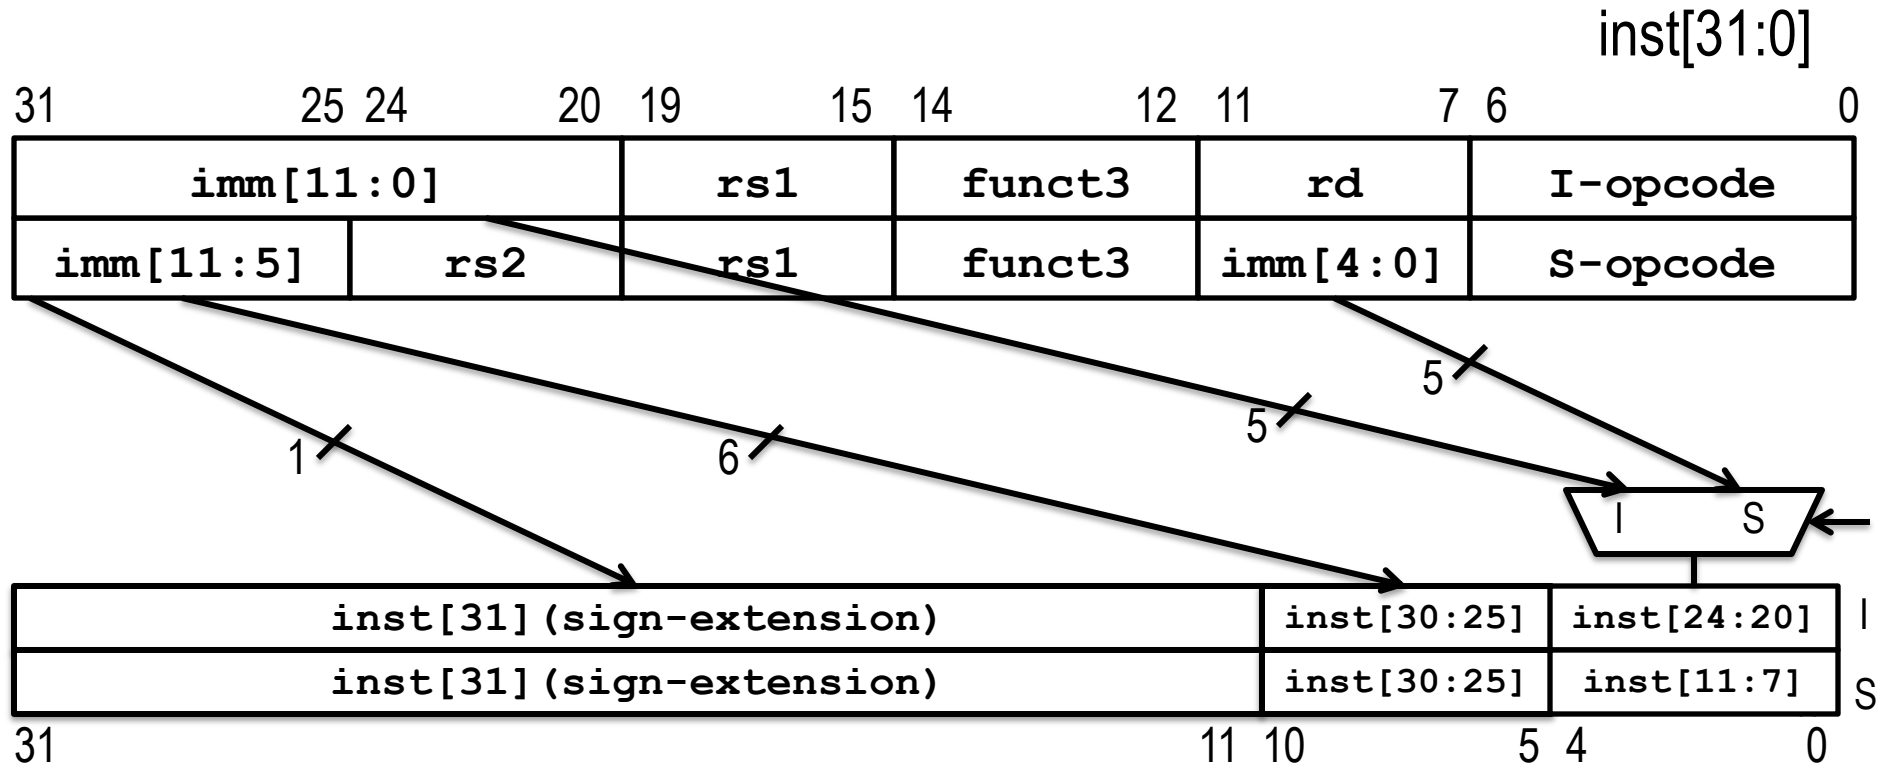


I-Format immediates



- High 12 bits of instruction (**inst[31:20]**) copied to low 12 bits of immediate (**imm[11:0]**)
- Immediate is sign-extended by copying value of **inst[31]** to fill the upper 20 bits of the immediate value (**imm[31:12]**)


I & S Immediate Generator



- Just need a 5-bit mux to select between two positions where low five bits of immediate can reside in instruction
- Other bits in immediate are wired to fixed positions in instruction

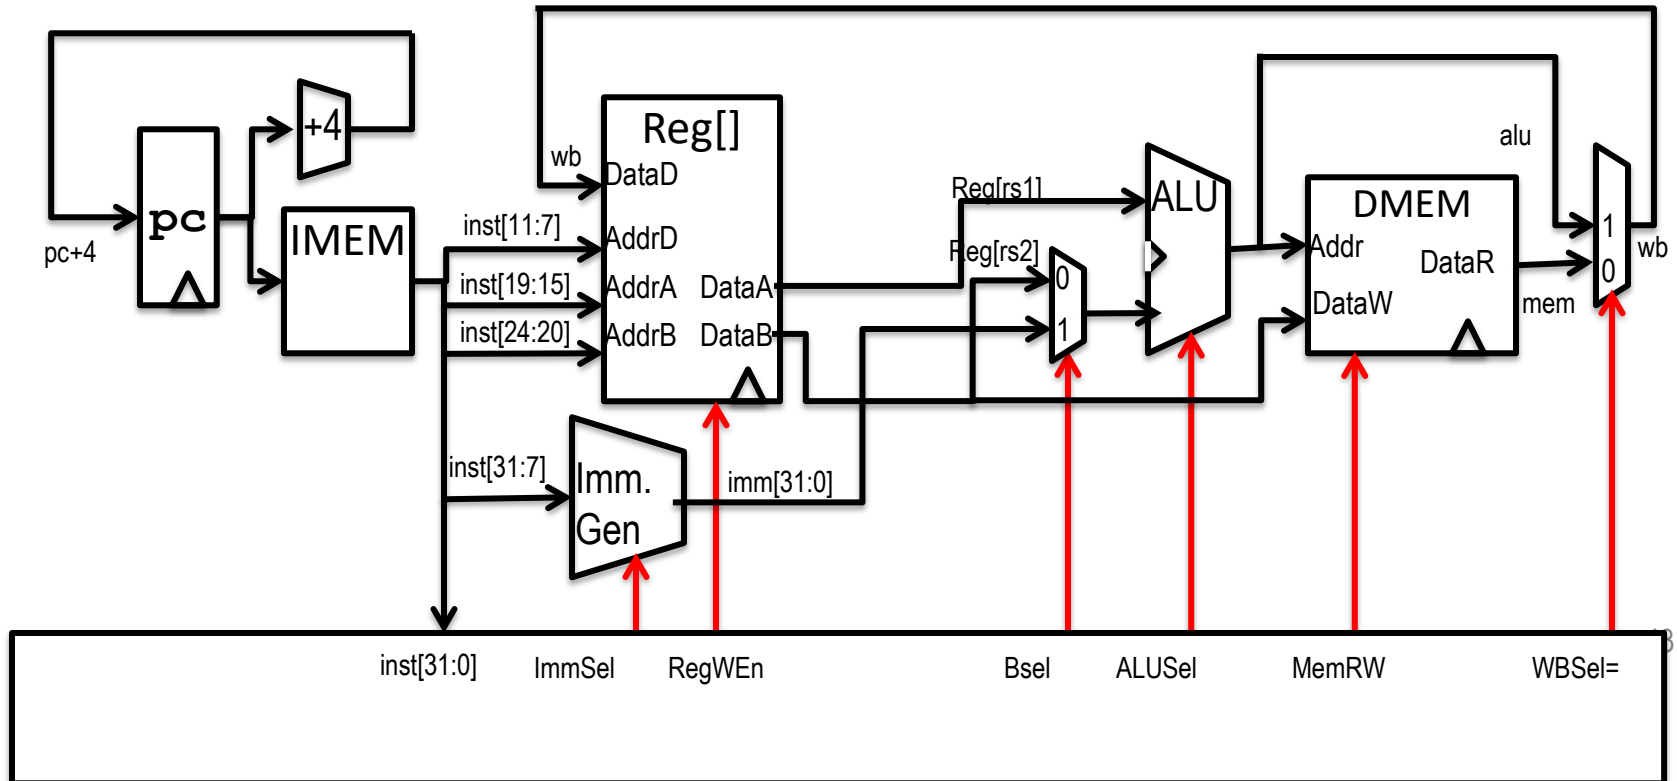
imm[31:0]

Implementing Branches

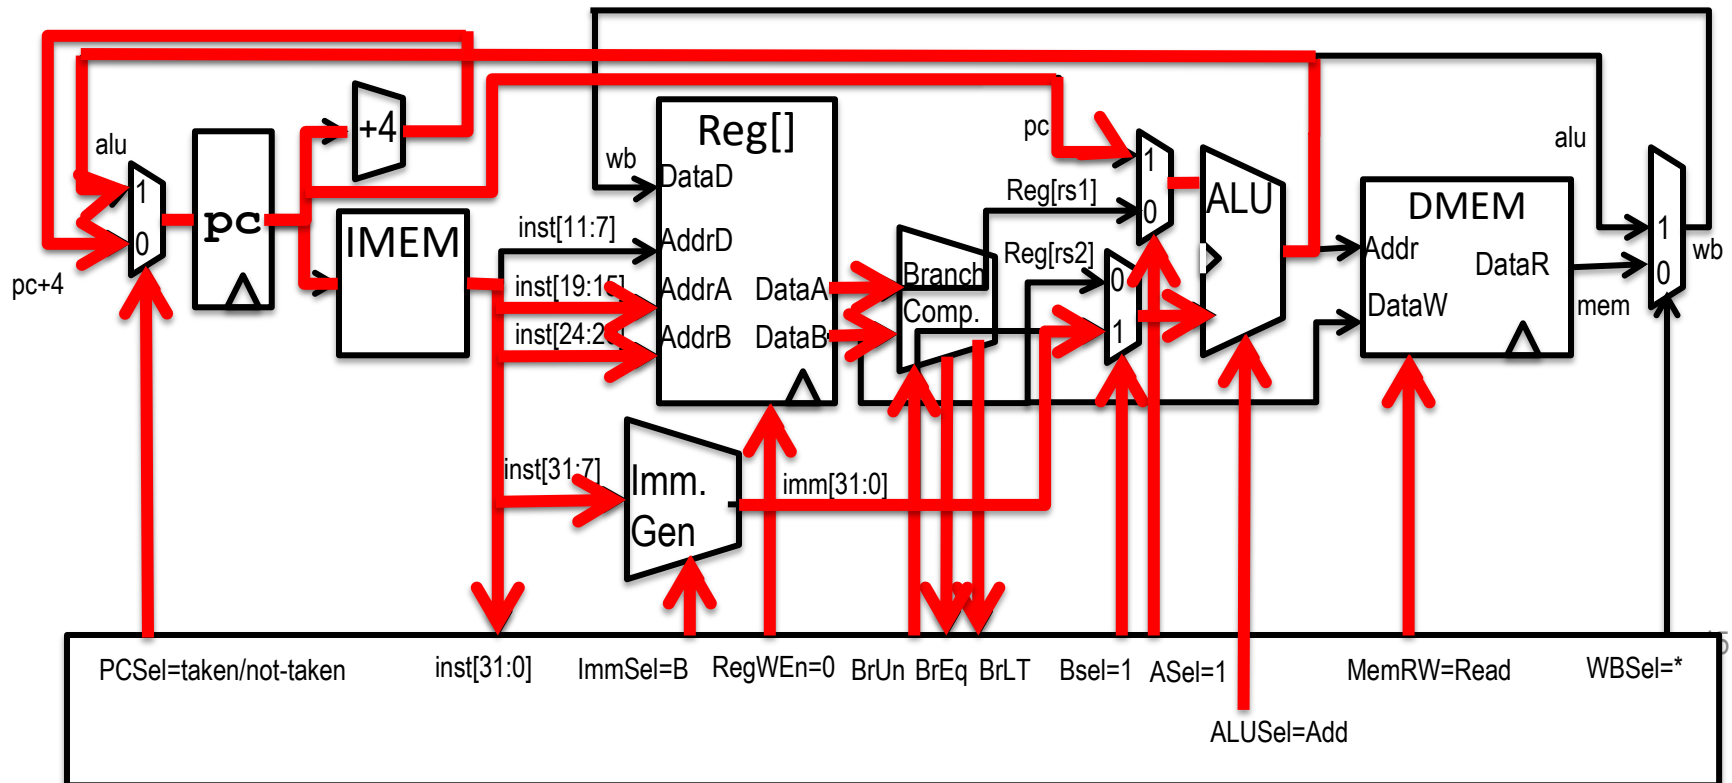


- ❖ B-format is mostly same as S-Format, with two register sources (rs1/rs2) and a 12-bit immediate
- ❖ But now immediate represents values -4096 to +4094 in 2-byte increments
- ❖ The 12 immediate bits encode *even* 13-bit signed byte offsets (lowest bit of offset is always zero, so no need to store it)


Adding sw to datapath



Adding branches to datapath



Branch Comparator



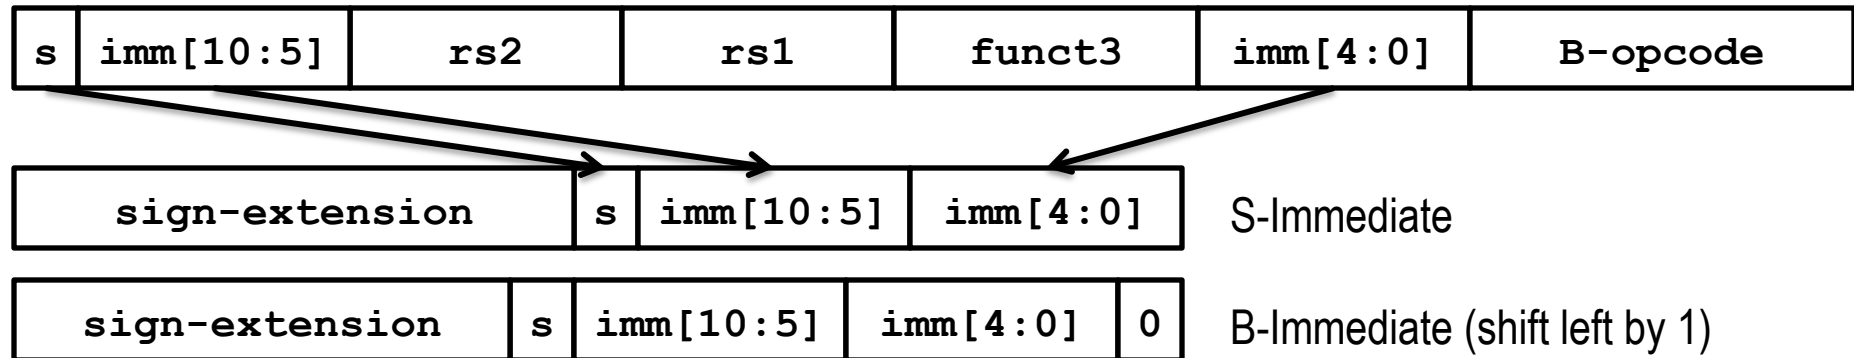
- ❖ $\text{BrEq} = 1$, if $A=B$
- ❖ $\text{BrLT} = 1$, if $A < B$
- ❖ $\text{BrUn} = 1$ selects unsigned comparison for BrLT , 0 =signed
- ❖ BGE branch: $A \geq B$, if $\neg(A < B)$

Break!



Multiply Branch Immediates by Shift?





- ❖ 12-bit immediate encodes PC-relative offset of -4096 to +4094 bytes in multiples of 2 bytes
- ❖ Standard approach: treat immediate as in range -2048..+2047, then shift left by 1 bit to multiply by 2 for branches






Each instruction immediate bit can appear in one of two places in output immediate value – so need one 2-way mux per bit

RISC-V Immediate Encoding

Instruction Encodings, inst[31:0]



32-bit immediates produced, imm[31:0]




← Upper bits sign-extended from inst[31] always

Only bit 7 of instruction changes role in immediate between S and B


RISC-V Branch Immediates

- ❖ 12-bit immediate encodes PC-relative offset of -4096 to +4094 bytes in multiples of 2 bytes
- ❖ RISC-V approach: keep 11 immediate bits in fixed position in output value, and rotate LSB of S-format to be bit 12 of B-format



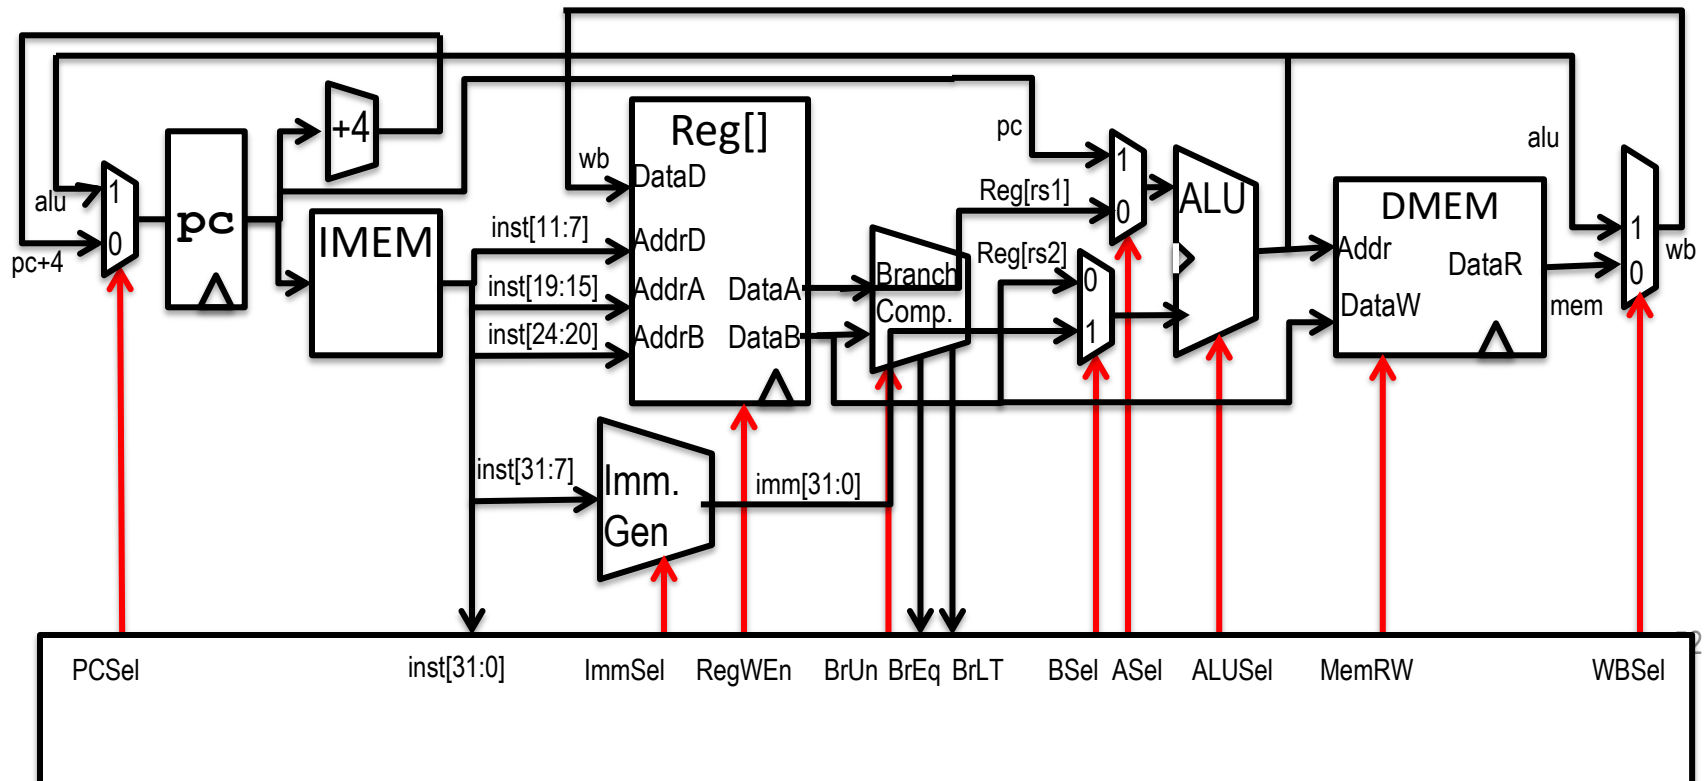
Only one bit changes position between S and B, so only need a single-bit 2-way mux

Implementing JALR Instruction (I-Format)

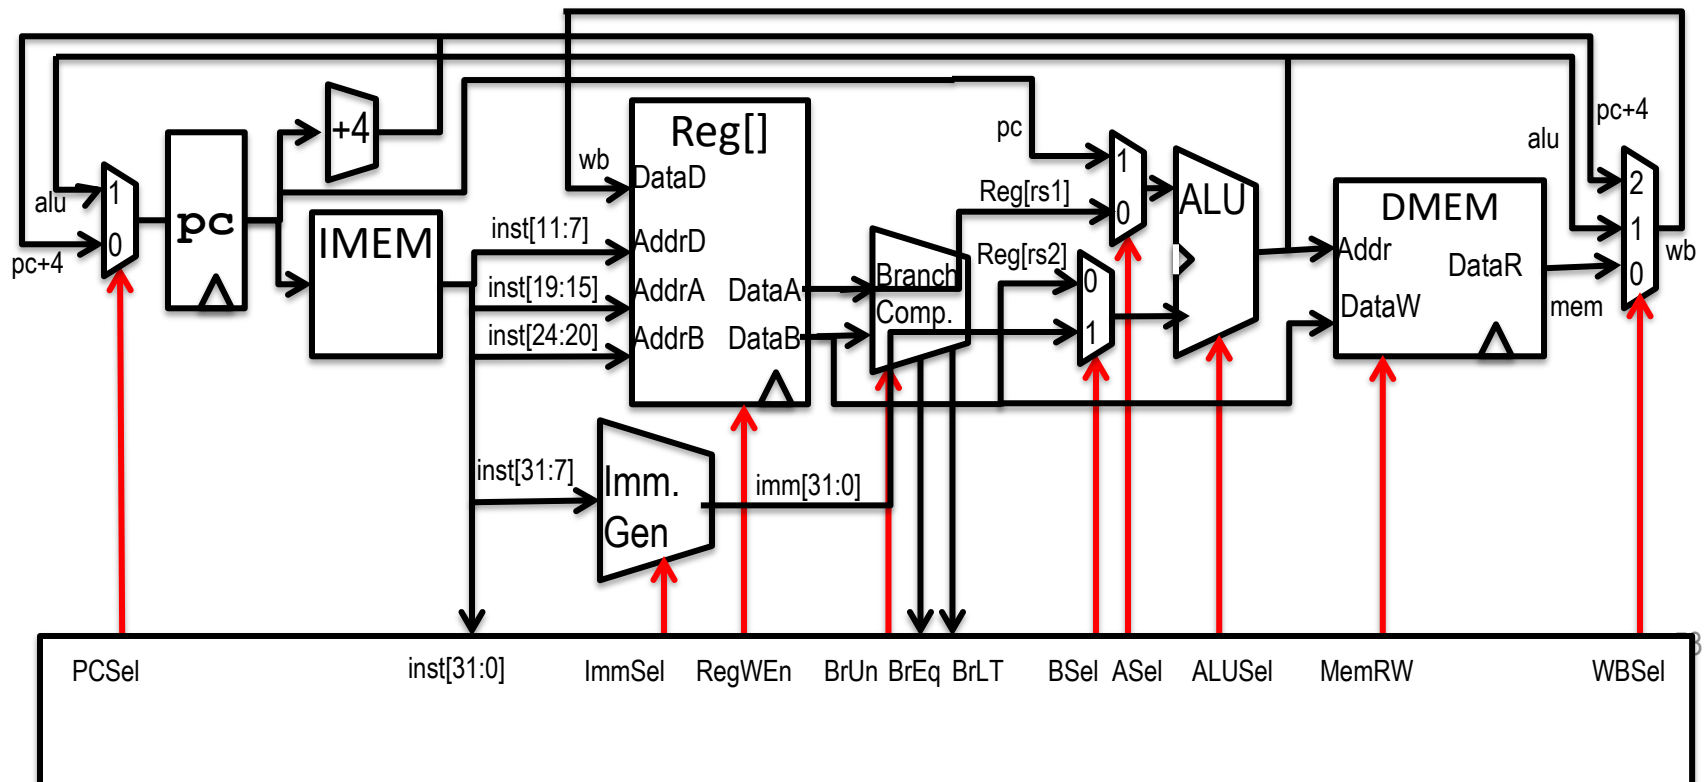


- ❖ JALR rd, rs, immediate
 - Writes PC+4 to Reg[rd] (return address)
 - Sets PC = Reg[rs1] + immediate
 - Uses same immediates as arithmetic and loads
 - *no* multiplication by 2 bytes

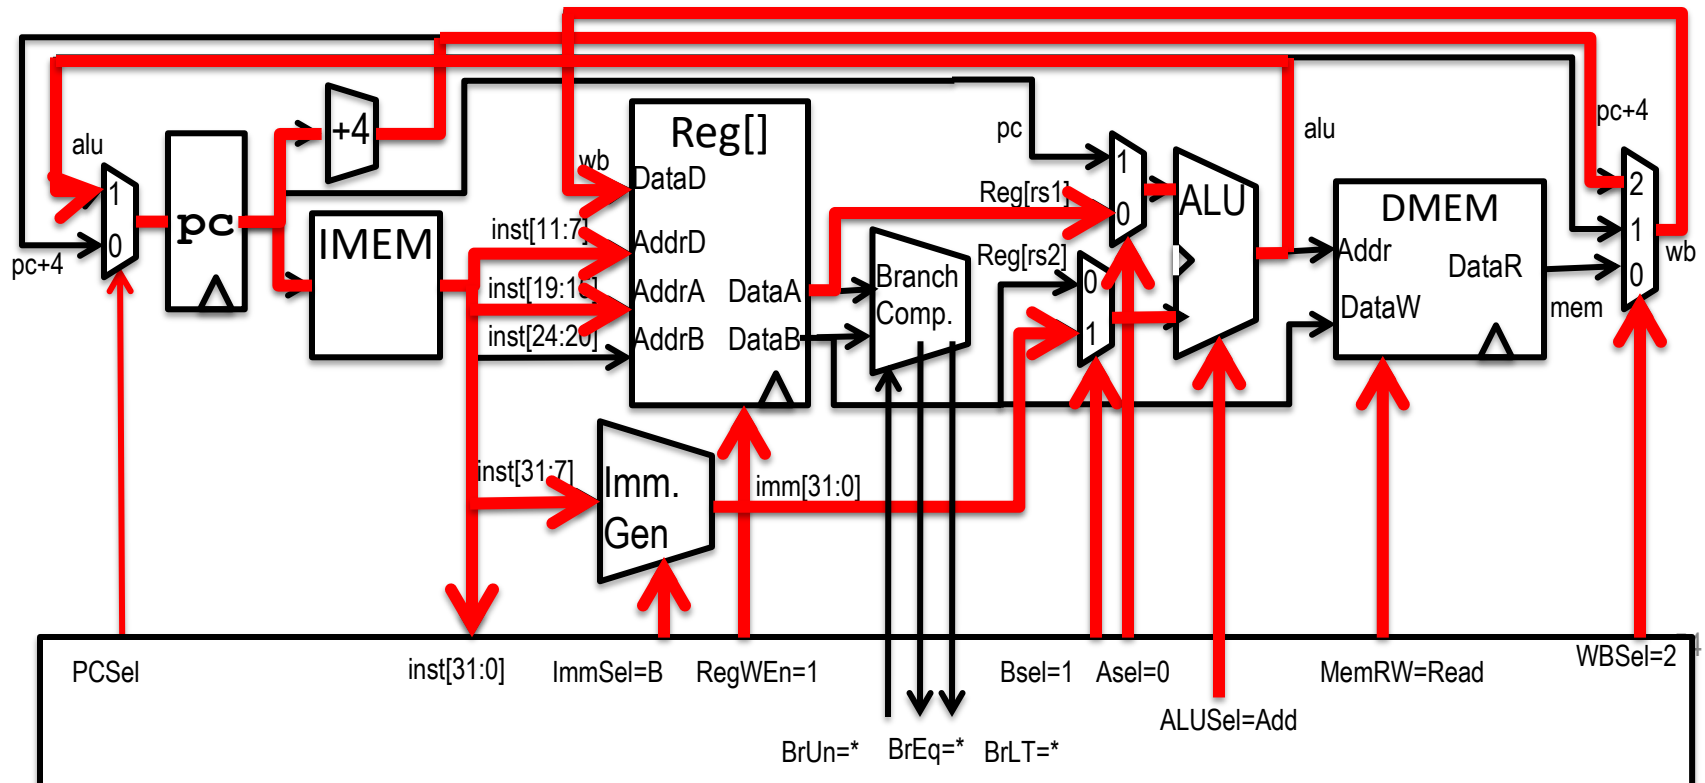
Adding branches to datapath




Adding jalr to datapath



Adding jalr to datapath

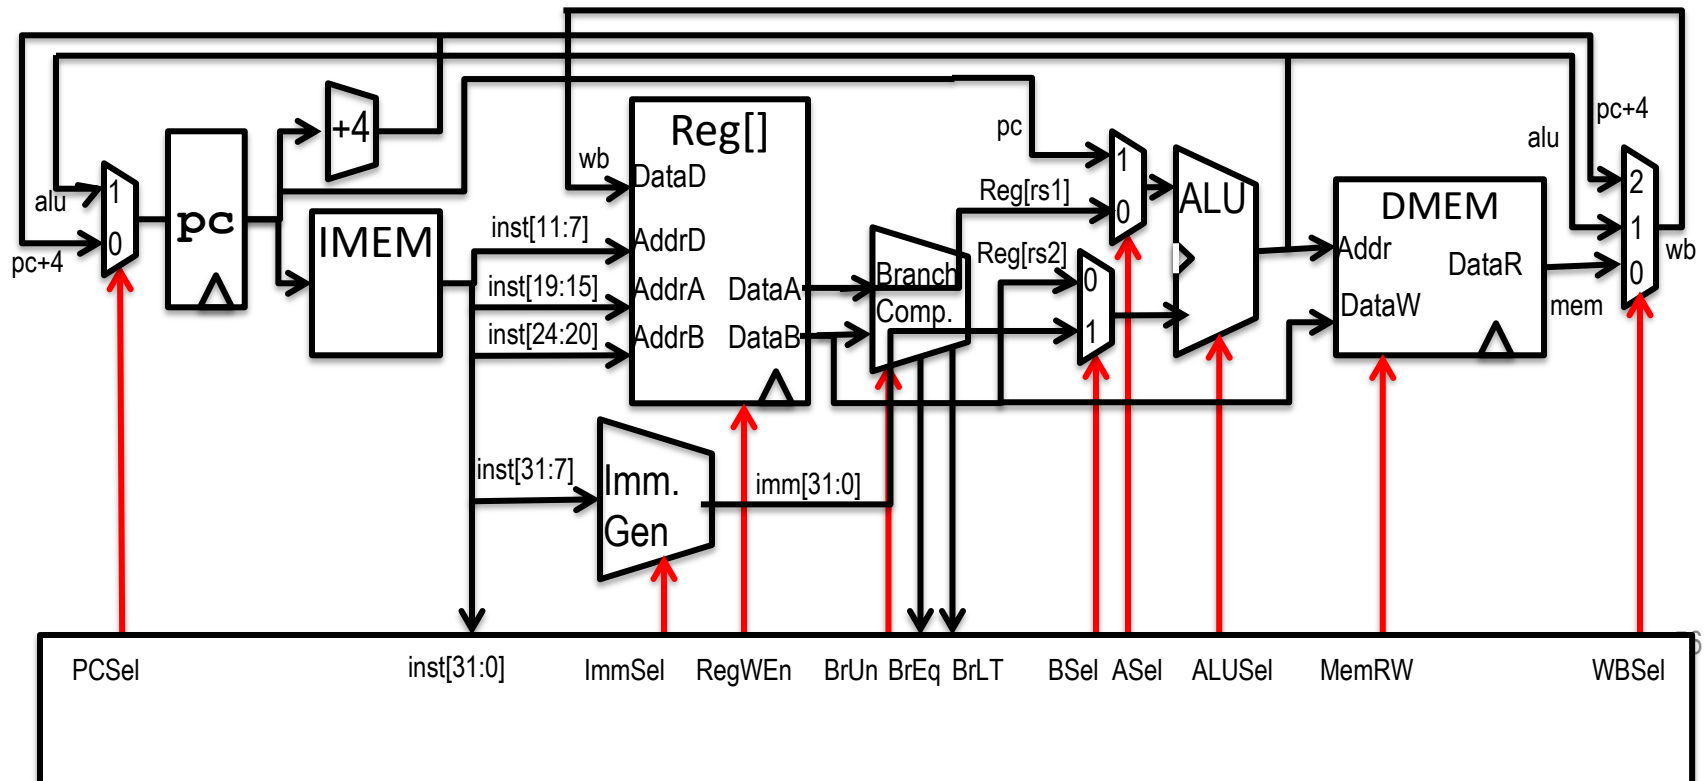


Implementing jal Instruction

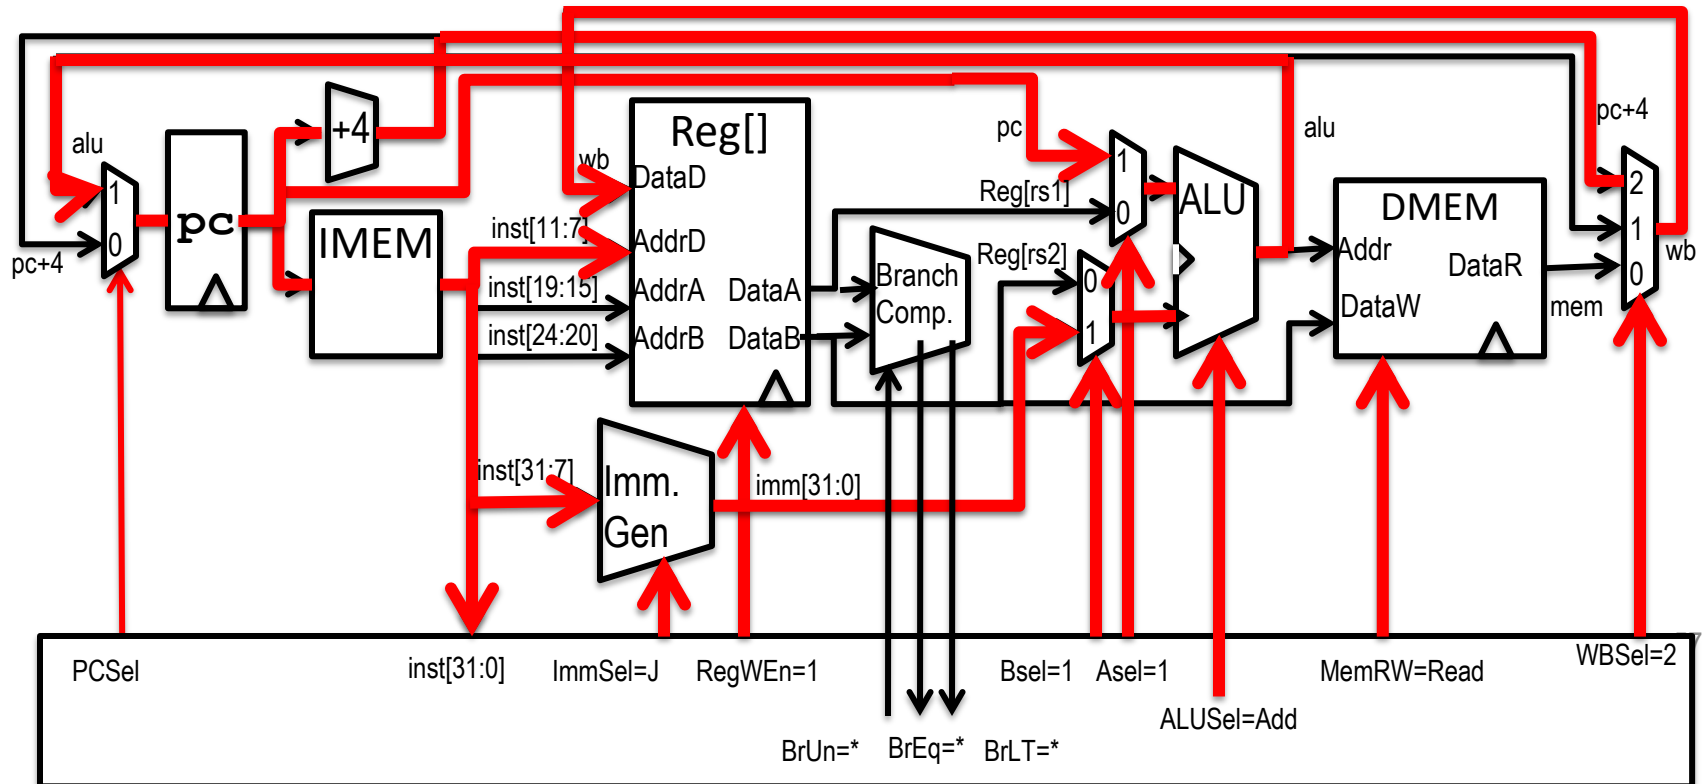


- ❖ JAL saves PC+4 in Reg[rd] (the return address)
- ❖ Set PC = PC + offset (PC-relative jump)
- ❖ Target somewhere within $\pm 2^{19}$ locations, 2 bytes apart
 - $\pm 2^{18}$ 32-bit instructions
- ❖ Immediate encoding optimized similarly to branch instruction to reduce hardware cost

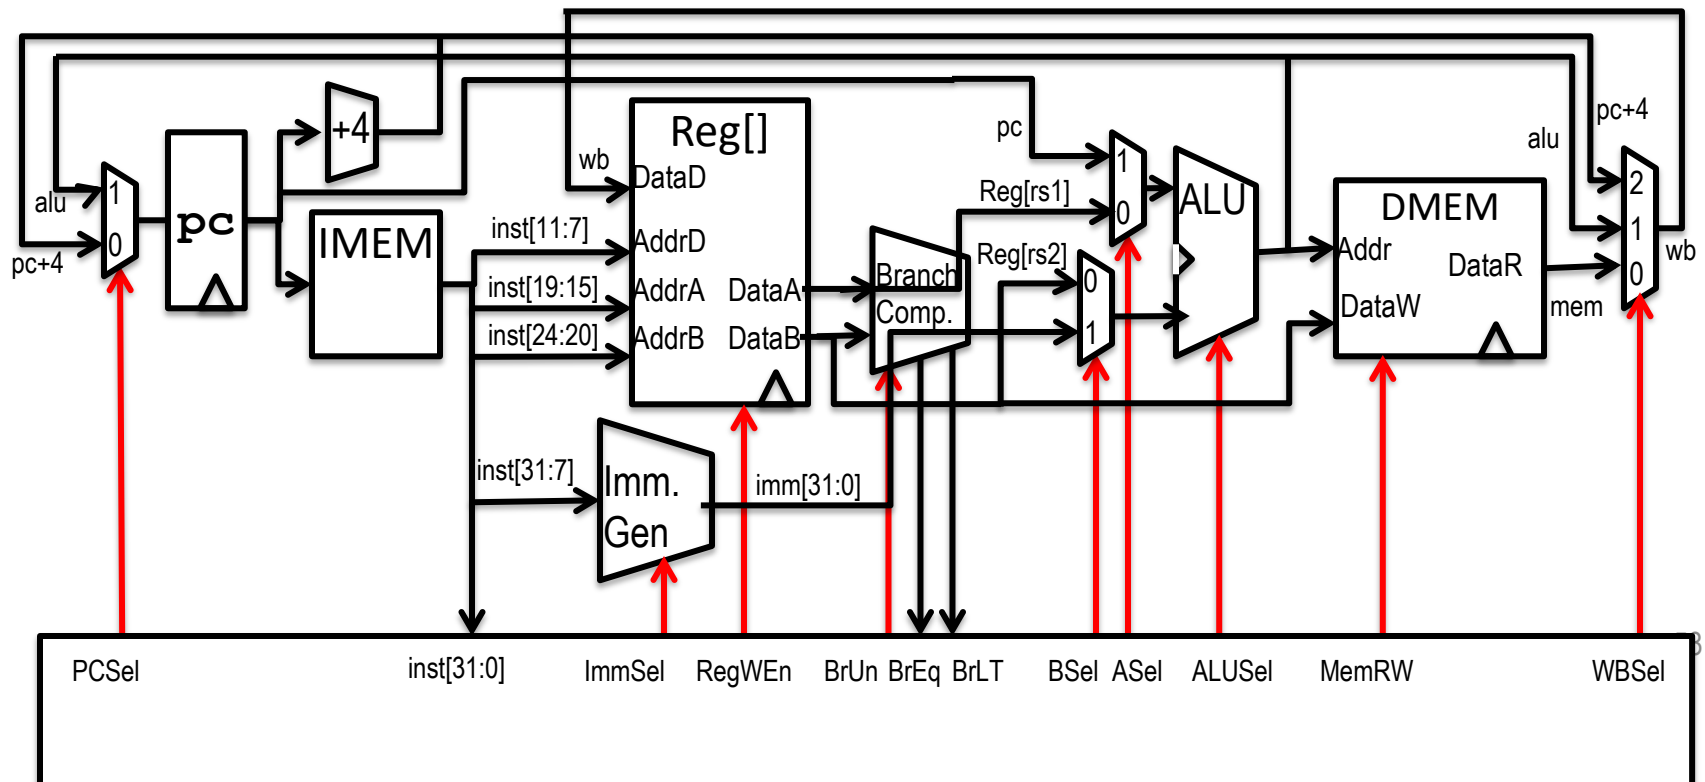
Adding jal to datapath




Adding jal to datapath



Single-Cycle RISC-V RV32I Datapath




“Upper Immediate” instructions

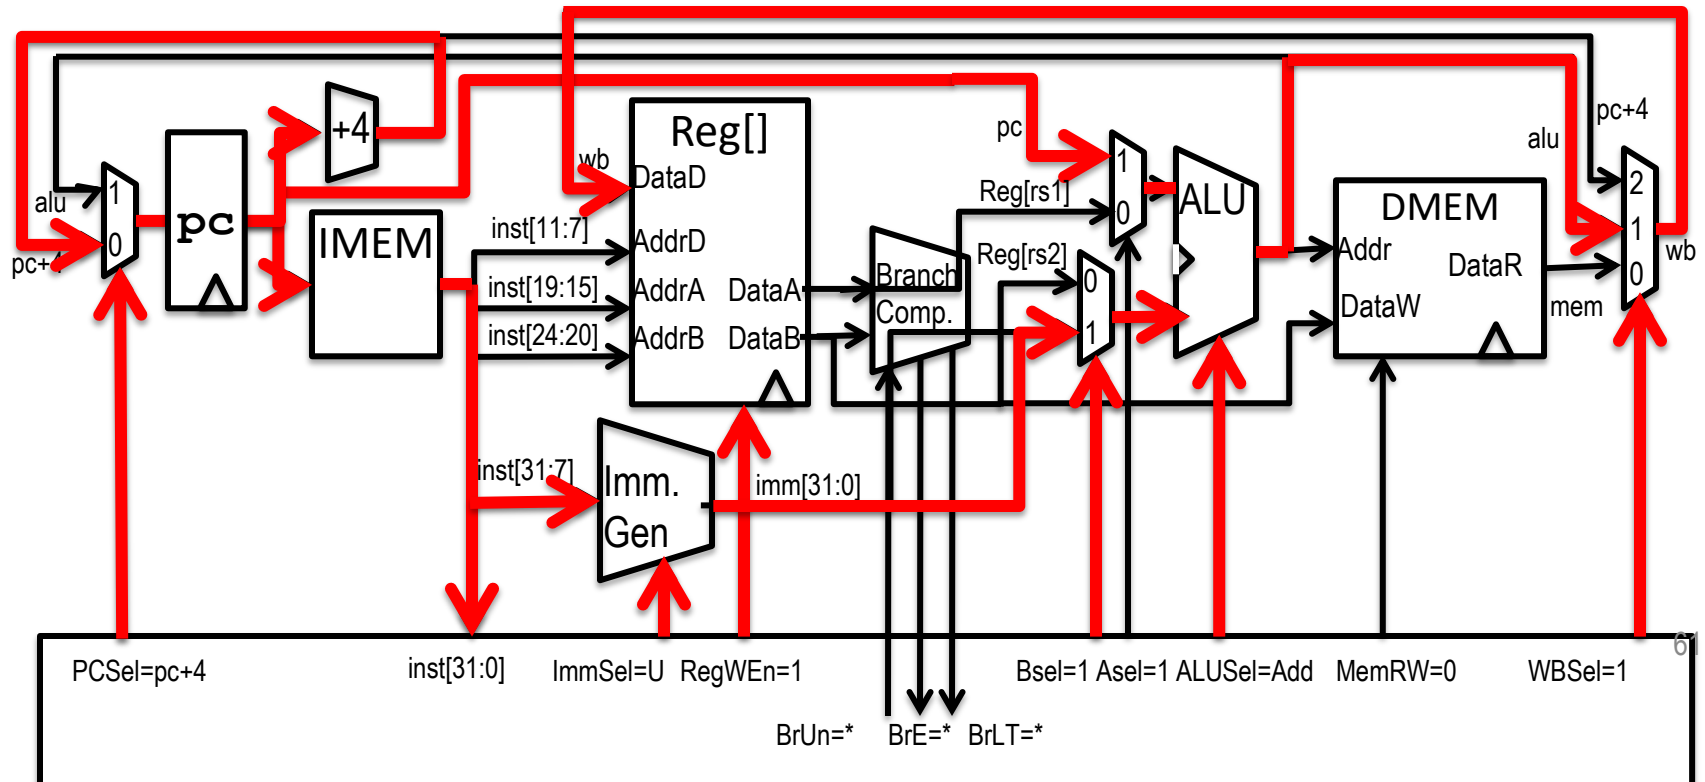


- ❖ Has 20-bit immediate in upper 20 bits of 32-bit instruction word
- ❖ One destination register, rd
- ❖ Used for two instructions
 - LUI – Load Upper Immediate (add to zero)
 - AUIPC – Add Upper Immediate to PC

Implementing lui




Implementing `auipc`



Design Principles

- 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



Design Principles

- Determining control signals
 - Any time a datapath element has an input that changes behavior, it requires a control signal (e.g. ALU operation, read/write)
 - Any time you need to pass a different input based on the instruction, add a **MUX** with a control signal as the selector (e.g. next PC, ALU input, register to write to)
- Your control signals will change based on your exact datapath
- Your datapath will change based on your ISA

Summary !

- Universal datapath
 - Capable of executing all RISC-V instructions in one cycle each
 - Not all units (hardware) used by all instructions
- 5 Phases of execution
 - IF (Instruction Fetch), ID (Instruction Decode), EX (Execute), MEM (Memory), WB (Write Back)
 - Not all instructions are active in all phases (except for loads!)
- Controller specifies how to execute instructions
 - Worth thinking about: what new instructions can be added with just most control?