

CSE 30321 – Computer Architecture I – Fall 2009
Lecture 13– In Class Notes and Problems
 October 6, 2009

Question 1:

First, we briefly review the notion of a clock cycle (CC). Generally speaking a CC is the amount of time required for (i) a set of inputs to propagate through some combinational logic and (ii) for the output of that combinational logic to be latched in registers. (The inputs to the combinational logic would usually come from registers too.)

Thus, for the MIPS, *single cycle* datapath that was derived in class last week, the “combinational logic” referred to above would really be the instruction memory, register file, ALU, data memory, and register file. The inputs come from instruction memory and the output might be the main register file (in the case of an ALU instruction) or the PC (in the case of a branch instruction).

Putting the MIPS datapath aside, its generally a good idea to minimize the logic on the critical path between two registers – as this will help shorted the required clock cycle time, will result in an increase in clock rate, which generally means better “performance”.

To be more quantitative, hypothetically, let’s say that there are two gate mappings that can implement the logic function that we need to evaluate. The inputs to these gates would come from 1 set of registers, and the outputs from these gates would be stored in another set of registers. The composition of each is specified in the table below:

	AND gates	NOR gates	XOR gates
Design 1	7	17	13
Design 2	24	4	7

Furthermore, the delay associated with each gate is also listed below:

AND gate	NOR gate	XOR gate
2 picoseconds*	1 picosecond	3 picoseconds

Which design will lead to the shorted CC time?

Design 1: = (7 ANDs * 2 ps/AND) + (17 NORs * 1 ps/NOR) + (13 XORs * 3 ps/XOR)
 = 14 ps + 17 ps + 39 ps = **70 ps**

Design 2: = (24 ANDs * 2 ps/AND) + (4 NORs * 1 ps/NOR) + (7 XORs * 3 ps/XOR)
 = 48 ps + 4 ps + 21 ps = **73 ps**

What is the clock rate for that design (in GHz)?

- Clock rate = 1 / CC
- Clock rate = 1 / 70ps
- Clock rate = 1 / (70*10⁻¹²)
- Clock rate = 1.428 x 10¹¹ Hz
- Clock rate = 1.428 x 10¹¹ Hz * (1 GHz / 10⁹ Hz)
- Clock rate = 14.3 GHz


Take Away: The longest critical path determines the clock cycle time

Question 2:


Without looking back in your notes, write out the register transfer language (RTL) for a generic MIPS R-type, I-type, and J-type instruction. (Don't be overly concerned with getting all of the sign extensions, etc. exactly right.)

For your reference, the instruction encodings for the R, I, and J instruction types are shown below.


R-Type



I-Type



J-Type



Answers:

	Your Solution	My Solution
R-Type		<ul style="list-style-type: none"> - get \leftarrow Mem(PC) - $\text{Reg(Rd)} \leftarrow \text{Reg(Rs)} \text{ op } \text{Reg(Rt)}$ - $\text{PC} \leftarrow \text{PC} + 4$
I-Type (lw/sw)		<ul style="list-style-type: none"> - get Mem(PC) - $\text{Addr} \leftarrow \text{Reg(Rs)} + \langle 0000_h \rangle \langle \text{imm}(15:0) \rangle$ - if lw <ul style="list-style-type: none"> - $\text{Reg(Rt)} \leftarrow \text{Mem(Addr)}$ - if sw <ul style="list-style-type: none"> - $\text{Mem(addr)} \leftarrow \text{Reg(Rs)}$ - $\text{PC} \leftarrow \text{PC} + 4$ <p style="color: red; margin-top: 10px;"> (?) – lw or sw = longer latency? (?) – # of register/memory transfers in each? </p>

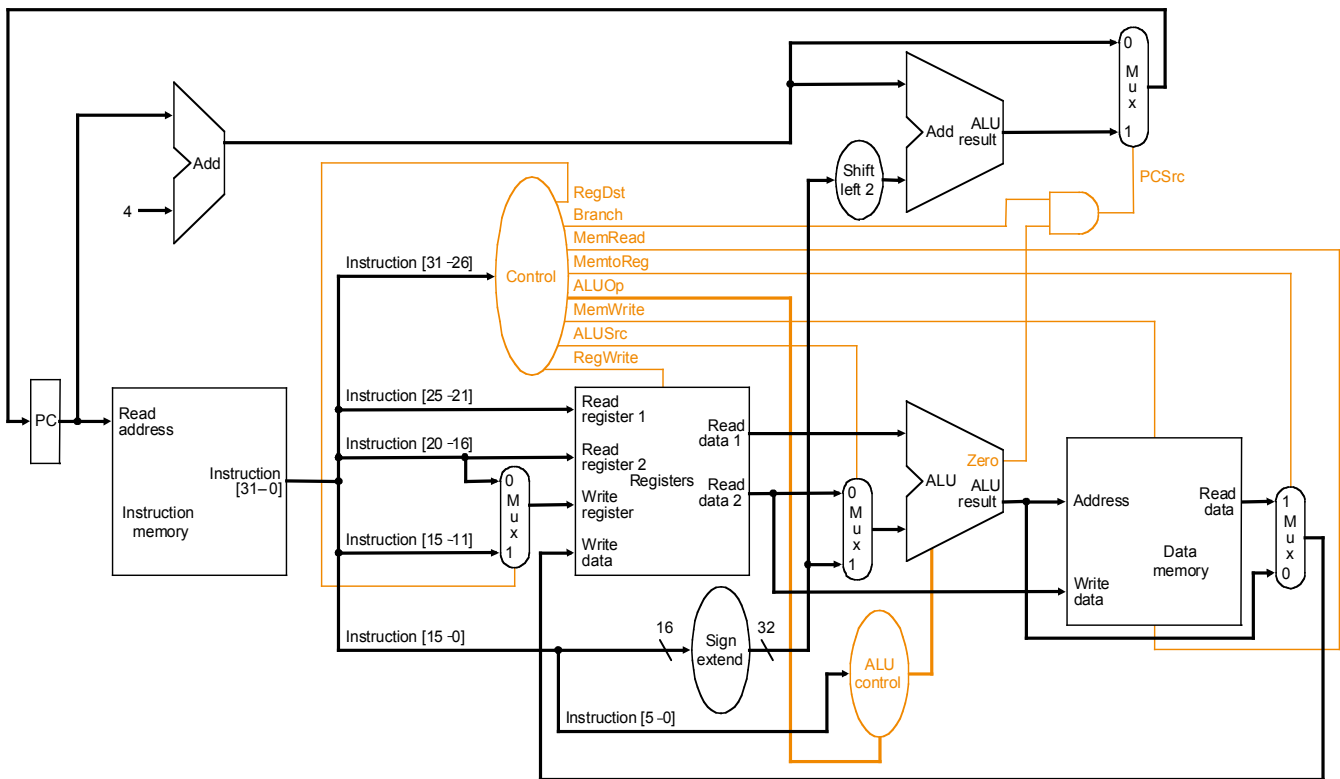
I-Type (beq)		<ul style="list-style-type: none"> - get \leftarrow Mem(PC) - Cond \leftarrow Reg(Rs) – Reg(Rt) - if cond met <ul style="list-style-type: none"> - PC \leftarrow new value - PC \leftarrow PC + 4 + signext(imm(15:0))x4 - if cond <i>not</i> met <ul style="list-style-type: none"> - PC \leftarrow PC + 4
J-Type		<ul style="list-style-type: none"> - get \leftarrow Mem(PC) - PC \leftarrow <PC extension> <IR(25:0)>

Take Away: There's a lot of stuff common between instructions! For example:

- For R-Type:
 - o Everything the same except what the ALU does
- Looking at R-type vs. I-type
 - o Still, lots of stuff the same – register reads, ALU operation, register write back

Question 3:

In class last week, you derived the single cycle datapath for the MIPS ISA. A block diagram is shown below:



Assume that the following latencies would be associated with the datapath above:

Operation	Time
Get instruction encoding – Mem(PC)	2 ns
Get data from register file	2 ns
Perform operation on data in registers	3 ns
Increment PC by 4	1 ns
Change PC depending on conditional instruction	1 ns
Access data memory	2 ns
Write data back to the register file (from the ALU or from memory)	2 ns

Questions:

- (a) What type/class of instruction in the MIPS ISA will set CC time (and hence clock rate)? (e.g. ALU type, conditional branch, load, store, etc.)
- (b) Given the latencies in the above table, what would the clock rate be for our single cycle datapath?

Note: for both (a) and (b), pay particularly close attention to the diagram shown above and try to *define the critical path*.

Answers:

First, the load instruction sets the critical path. For the load, we must:

- Get the instruction encoding from memory
- Read data from the register file
- Perform an ALU operation
- Access data memory
- Write data back to the register file

Note that we can ignore the overhead associated with $PC \leftarrow PC + 4$ b/c this would happen in parallel with operations

Second, the clock rate would just be the inverse of the some of the delays:

- (2 ns + 2 ns + 3 ns + 2 ns + 2 ns)
- 11 ns
- $1 / 11 \text{ ns} = 91 \text{ MHz}$

Take Away:

- In a single cycle implementation, the instruction with the longest critical path sets the clock rate for EVERY instruction.
- Think about Amdahl's Law – this is good if we can improve something that we use often...
 - o ...but here, we're sort of making a lot of things worse!
 - o (How worse? We'll see next.)

Question 4:

Let's look at the impact of the load instruction's domination of the CC time a bit more:

Recall, that a load (lw) needs to do the following:

- a) Get Memory(PC)
- b) Read data from the register file and "sign extend" an immediate value
- c) Calculate an address
- d) Get data from Memory(Address)
- e) Write data back to the register file

With this design, let's assume that each step takes the amount of time listed in the table

(a)	(b)	(c)	(d)	(e)
4 ns	2 ns	2 ns	4 ns	2 ns

(for simpler math)

Part 1:

With these numbers, what is the CC time?

A: 14 ns

- However, store (sw) instruction only needs to do steps A, B, C, and D
 - o Therefore a store only really needs 12 ns
- Similarly, add instruction only needs to do steps A, B, C, and E
 - o Therefore could do an add in 10 ns
- *Still, with single CC design, both of the above instructions take 14 ns.*

Take away: Our CC time is limited by the longest instruction.

Part 2:

Let's quantify performance hit that we're taking – and assume that we could somehow execute each instruction in the amount of time that it actually takes:

Assume the following:

ALU	lw	sw	Branch / jump
45 %	20 %	15 %	20 %

Per the previous discussion, we can assume that ALU instructions take 10 ns, lw's take 14 ns, and sw's take 12 ns. We can also estimate how long it takes to do a branch/jump:

- Assume that steps A, B, C, and E are needed
 - o (A) Fetch, (B) Read data to compare, (C) Subtract to do comparison, (D) Update PC (like a register write back)

$$\begin{aligned}
 \text{CPU time} &= (\text{instruction / program}) \times (\text{seconds / instruction}) \\
 \text{CPU time (single CC)} &= i \times 14 \text{ ns} = 14(i) \\
 \text{CPU time (variable CC)} &= i \times [(.45 \times 10) + (.2 \times 14) + (.15 \times 12) + (.2 \times 10)] = 11.1(i) \\
 \text{Potential performance gain} &= 14(i) / 11.1(i) = 27\%
 \end{aligned}$$

Take away: By using a single, long CC, we're losing performance.

Part 3:

Unfortunately, such a “variable CC” is not practical to implement... however, there is a way to get some of the above performance back.

- To make this solution work, want to balance the amount of work done per step. **Why?**
 - o Because if every step has to take the same amount of time, if we have 2, 2, 2, 2, and 4 ns, we’re still at 4!

As an example, let’s see what happens if we can make each step take 3 ns:

$$\begin{aligned}\text{CPU time (3 ns, multi-cycle)} &= (i) \times (\text{CC} / \text{instruction}) \times (s / \text{CC}) \\ &= (i) \times [(.45 \times 4) + (.2 \times 5) + (.15 \times 4) + (.15 \times 4)] \times 3 \\ &= 11.8(i)\end{aligned}$$

11.8(i) is not as good as 11(i), but its better than 14! → 19% vs. 27%

Now, we have a shorter clock rate and each instruction takes an integer number of CCs.

Take away: There is a “catch” to this approach. We have to store the intermediate results. Remember, a CC is defined as the time some logic was evaluated and the result was latched.... and that inputs often come from another, previous latch.

Realistically, the latching time of the intermediate registers will add a nominal overhead to each stage.

Part 4:

Let’s see how the overhead might affect performance. We’ll assume its 0.1 ns.

$$\begin{aligned}\text{CPU time (3 ns, overhead)} &= (i) \times (\text{CC} / \text{instruction}) \times (s / \text{CC}) \\ &= (i) \times [(.45 \times 4) + (.2 \times 5) + (.15 \times 4) + (.15 \times 4)] \times (3+0.1) \\ &= (i) \times [1.8 + 1 + 0.6 + 0.6] \times 3.1 \\ &= (i) \times [4] \times 3.1 \\ &= 12.4(i)\end{aligned}$$

Take away:

- Again, we lose a little performance, but we’re still better than the single cycle method
 - o $14(i) / 12.4(i) = 13\%$
- Every approach is better than the single cycle approach:
 - o $11.1(i) < 11.8(i) < 13(i) < 14(i)$
- The 11.8(i) and 13(i) approach are feasible

Question 5:

If we take the above approach, how should we break up the instructions?

Think about each step ... what needs to be done to ensure R, I, and J type instructions all work?

- Hint: Think about each instruction and think about what can be done in parallel.

(1) IF	(2) ID	(3) EX	(4) M / R-Type	(5) WB
(every) Fetch instruction at address in PC	(every) Decode instruction type (i.e. send opcode to control logic)	(lw/sw) Compute memory address	(lw) Read memory, store result in MDR (memory-data-reg)	(lw) Write MDR to register file.
(every) Store result in instruction register (IR)	(every) Store results read from the register file	(R-type) Perform correct logical operation	(R-type) Write ALUout to reference	
(every) increment PC	(jump) Compute new PC address****	(Branch) Update PC post comparison	(sw) Write to memory	
		(All) Store intermediate result in ALU (Note) All use ALU		

Let's talk about **** more...

- What if we have an R-type instruction: | opcode | rs | rt | rd | shamt | function code | | ?
 - o Would send rs, and rt encodings to RF as indices
 - o Data would be stored in intermediate registers A and B
- What if we have a J-type instruction: | opcode | offset | | ?
 - o Answer:
 - Do the same as above anyway!
 - There's no harm and no control signals based on opcode available to end of CC anyhow...
 - o Similarly, we can calculate a new PC address – i.e. for a jump or branch – even if we don't know if we have a jump or a branch
 - Again, it does no harm + it subscribes to ASAP philosophy
 - We can do this now, and save a CC later