


## Lecture 13

- Homework <sup>2</sup> due today ☺ ✓
  - I hope you had fun with recursion and stack smashing
- Poll:
  - Monday, May 3. Normal lecture or review?
- Today's lecture:
  - Quick review
  - What about load followed by use?
  - What about branches?
  - Crystal ball

1

## Data hazard review


- A data hazard arises if one instruction needs data that isn't ready yet.
  - Below, the AND and OR both need to read register \$2.
  - But \$2 isn't updated by SUB until the fifth clock cycle.
- Dependency arrows that point backwards indicate hazards.



2

## Forwarding

- The desired value (\$1 - \$3) has actually already been computed—it just hasn't been written to the registers yet.
- Forwarding allows other instructions to read ALU results directly from the pipeline registers, without going through the register file.



3

## Example

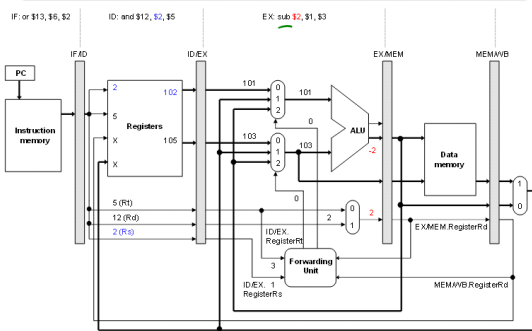
```

sub $2, $1, $3  1
and $12, $2, $5  2
or $13, $6, $2  3
add $14, $2, $2
sw $15, 100($2)
    
```

- Assume again each register initially contains its number plus 100.
  - After the first instruction, \$2 should contain -2 (101 - 103).
  - The other instructions should all use -2 as one of their operands.
- We'll try to keep the example short.
  - Assume no forwarding is needed except for register \$2.
  - We'll skip the first two cycles, since they're the same as before.

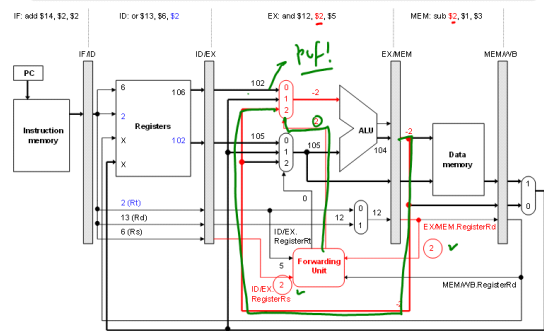
4

## Clock cycle 3

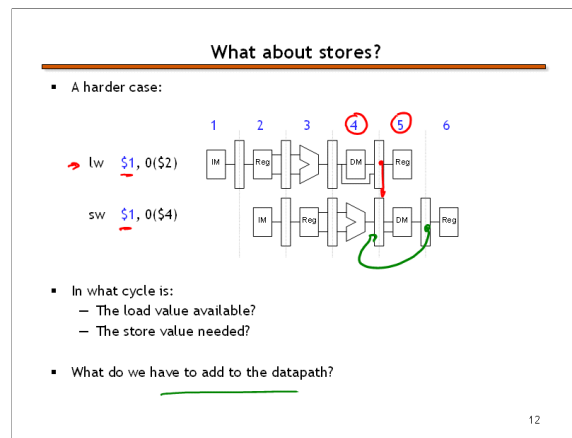
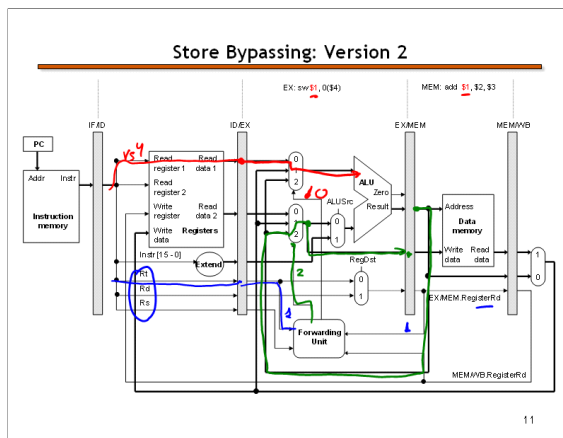
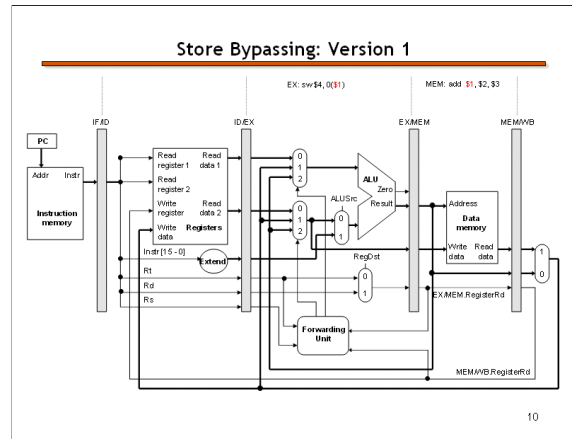
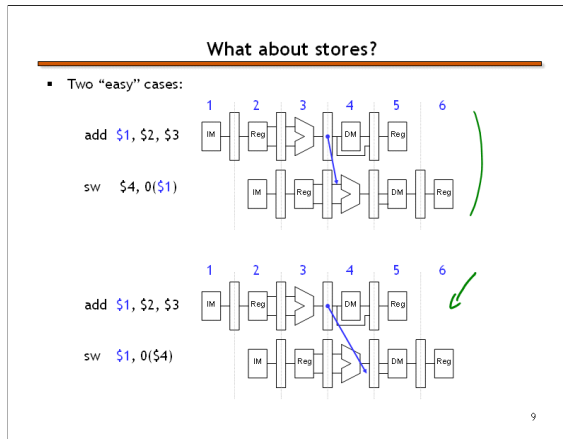
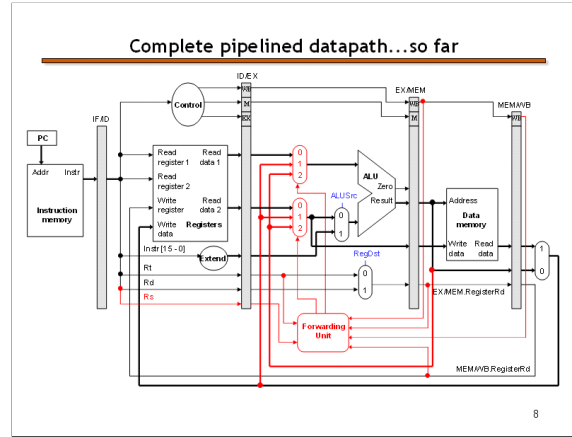
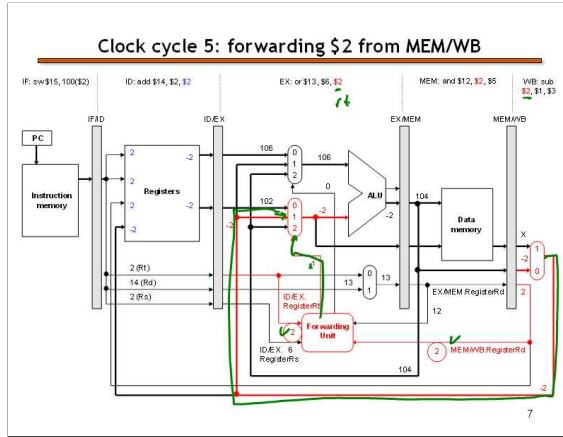


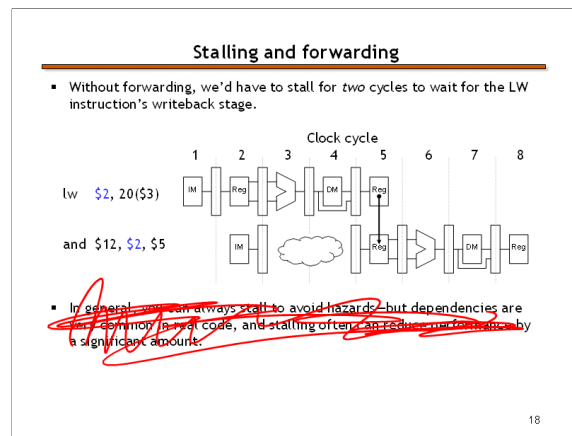
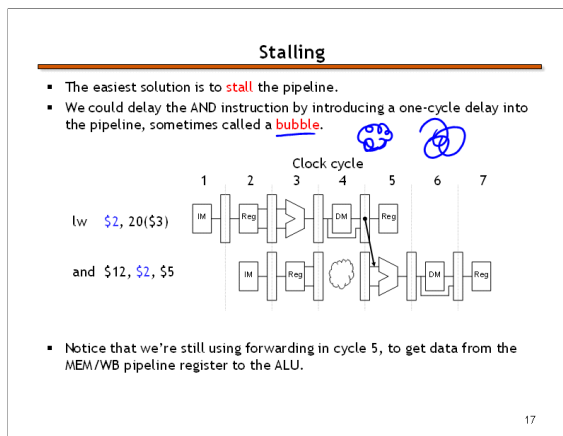
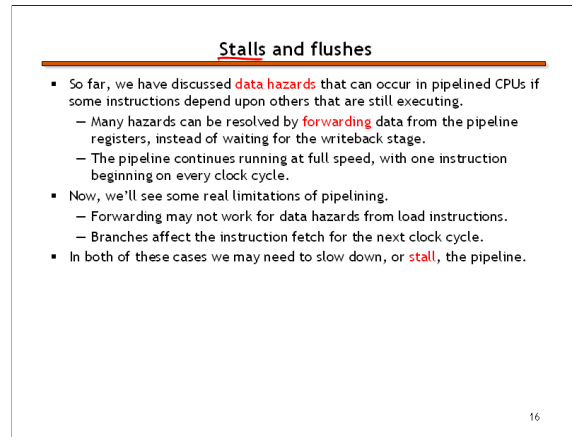
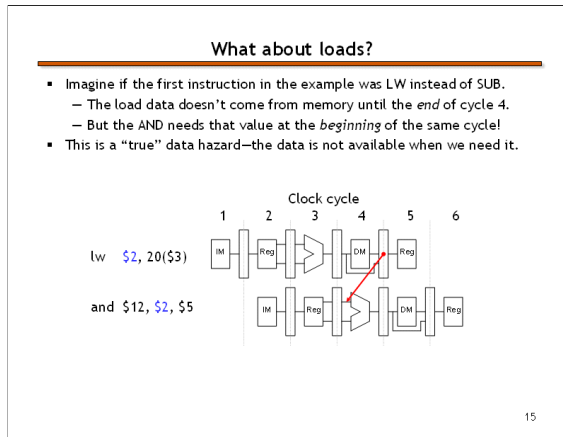
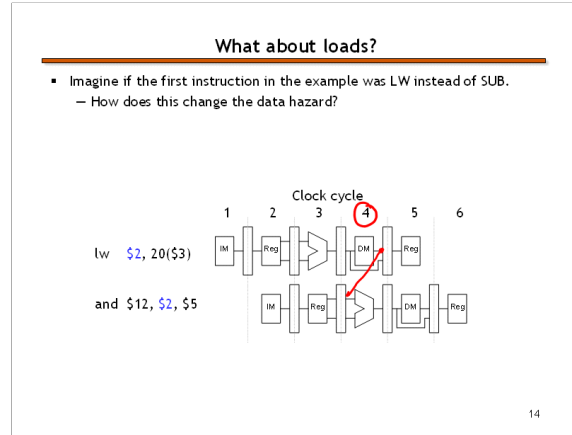
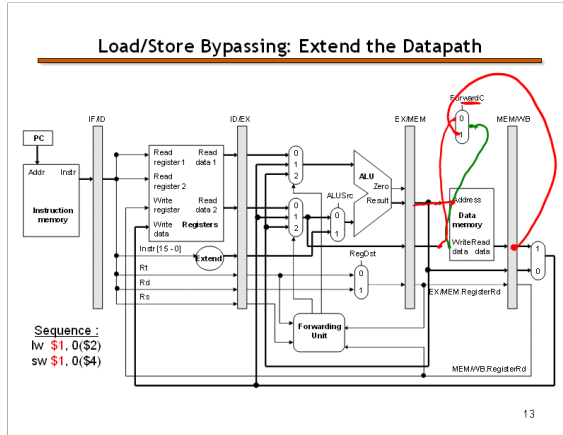
5

## Clock cycle 4: forwarding \$2 from EX/MEM




6





### Stalling delays the entire pipeline


- If we delay the second instruction, we'll have to delay the third one too.
  - Why?



19

### Stalling delays the entire pipeline


- If we delay the second instruction, we'll have to delay the third one too.
  - It prevents problems such as two instructions trying to write to the same register in the same cycle.
  - Also allows forwarding between AND and OR.



20

### Implementing stalls

- One way to implement a stall is to force the two instructions after LW to pause and remain in their ID and IF stages for one extra cycle.




- This is easily accomplished.
  - Don't update the PC, so the current IF stage is repeated.
  - Don't update the IF/ID register, so the ID stage is also repeated.

21

### What about EXE, MEM, WB


- But what about the ALU during cycle 4, the data memory in cycle 5, and the register file write in cycle 6?



- Those units aren't used in those cycles because of the stall, so we can set the EX, MEM and WB control signals to all 0s.

22

### Stall = Nop conversion



- The effect of a load stall is to insert an empty or nop instruction into the pipeline


23

### Detecting stalls

- Detecting stall is much like detecting data hazards.

- Recall the format of hazard detection equations:


if (EX/MEM.RegWrite = 1  
 and EX/MEM.RegisterRd = ID/EX.RegisterRs)  
 then Bypass Rs from EX/MEM stage latch



24

### Detecting Stalls, cont.

- When should stalls be detected?



- What is the stall condition?

if (  $ID/EX.MemRead = 1$  **OR** ( $ID/EX.Rt == IF/ID.Rs$  **||**  $ID/EX.Rt == IF/ID.Rt$ ) )  
 then stall

### Detecting stalls

- We can detect a load hazard between the current instruction in its ID stage and the previous instruction in the EX stage just like we detected data hazards.
- A hazard occurs if the previous instruction was LW...

$ID/EX.MemRead = 1$

...and the LW destination is one of the current source registers.

$ID/EX.RegisterRt = IF/ID.RegisterRs$

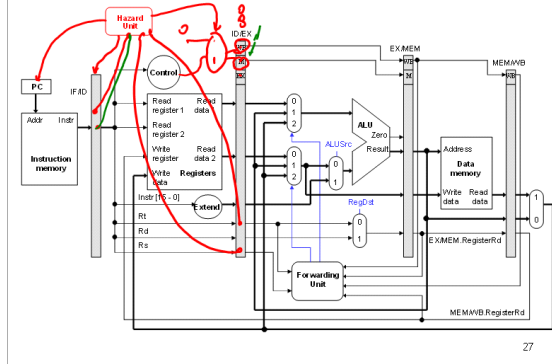
or

$ID/EX.RegisterRt = IF/ID.RegisterRt$

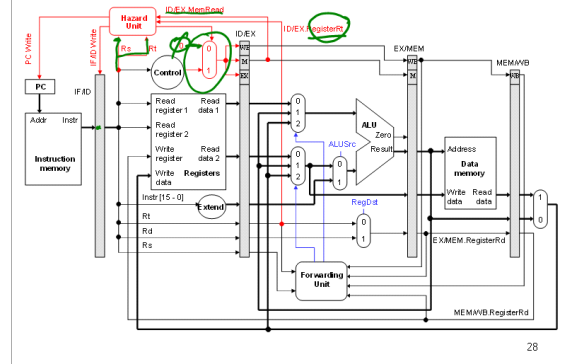
- The complete test for stalling is the conjunction of these two conditions.

if ( $ID/EX.MemRead = 1$  and  
 ( $ID/EX.RegisterRt = IF/ID.RegisterRs$  or  
 $ID/EX.RegisterRt = IF/ID.RegisterRt$ ))  
 then stall

### Adding hazard detection to the CPU



### Adding hazard detection to the CPU




### The hazard detection unit

- The hazard detection unit's inputs are as follows.
  - $IF/ID.RegisterRs$  and  $IF/ID.RegisterRt$ , the source registers for the current instruction.
  - $ID/EX.MemRead$  and  $ID/EX.RegisterRt$ , to determine if the previous instruction is LW and, if so, which register it will write to.
- By inspecting these values, the detection unit generates three outputs.
  - Two new control signals  $PCWrite$  and  $IF/ID Write$ , which determine whether the pipeline stalls or continues.
  - A  $mux select$  for a new multiplexer, which forces control signals for the current EX and future MEM/WB stages to 0 in case of a stall.

### Generalizing Forwarding/Stalling

- What if data memory access was so slow, we wanted to pipeline it over 2 cycles?

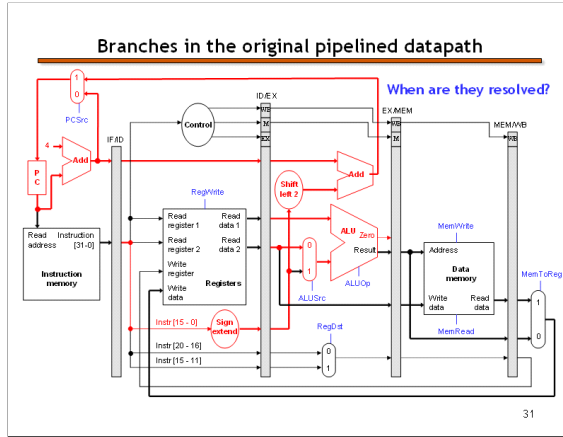


- How many bypass inputs would the muxes in EXE have?
- Which instructions in the following require stalling and/or bypassing?

lw r13, 0(r11)  
 add r7, r8, r9  
 add r15, r7, r13

IF	ID	EX	MEM1	MEM2	WB

IF ID EX MEM1 MEM2 WB



### Branches

- Most of the work for a branch computation is done in the EX stage.
  - The branch target address is computed.
  - The source registers are compared by the ALU, and the Zero flag is set or cleared accordingly.
- Thus, the branch decision cannot be made until the end of the EX stage.
  - But we need to know which instruction to fetch next, in order to keep the pipeline running!
  - This leads to what's called a **control hazard**.

32

### Stalling is one solution

- Again, stalling is always one possible solution.

- Here we just stall until cycle 4, after we do make the branch decision.

33

### Branch prediction

- Another approach is to guess whether or not the branch is taken.
  - In terms of hardware, it's easier to assume the branch is not taken.
  - This way we just increment the PC and continue execution, as for normal instructions.
- If we're correct, then there is no problem and the pipeline keeps going at full speed.

34

### Branch misprediction

- If our guess is wrong, then we would have already started executing two instructions incorrectly. We'll have to discard, or **flush**, those instructions and begin executing the right ones from the branch target address, Label.

35


### Performance gains and losses

- Overall, branch prediction is worth it.
  - Mispredicting a branch means that two clock cycles are wasted.
  - But if our predictions are even just occasionally correct, then this is preferable to stalling and wasting two cycles for every branch.
- All modern CPUs use branch prediction.
  - Accurate predictions are important for optimal performance.
  - Most CPUs predict branches dynamically—statistics are kept at run-time to determine the likelihood of a branch being taken.
- The pipeline structure also has a big impact on branch prediction.
  - A longer pipeline may require more instructions to be flushed for a misprediction, resulting in more wasted time and lower performance.
  - We must also be careful that instructions do not modify registers or memory before they get flushed.

36

## Implementing branches

- We can actually decide the branch a little earlier, in ID instead of EX.
  - Our sample instruction set has only a BEQ.
  - We can add a small comparison circuit to the ID stage, after the source registers are read.
- Then we would only need to flush one instruction on a misprediction.



37

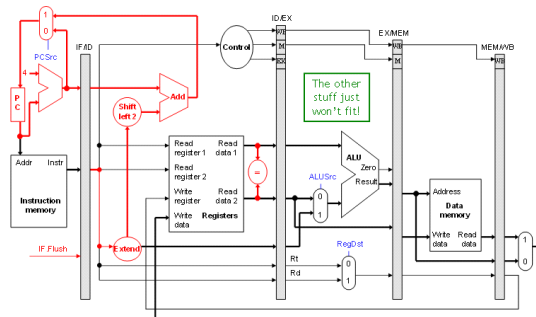
## Implementing flushes

- We must flush one instruction (in its IF stage) if the previous instruction is BEQ and its two source registers are equal.
- We can flush an instruction from the IF stage by replacing it in the IF/ID pipeline register with a harmless nop instruction.
  - MIPS uses `sll $0, $0, 0` as the nop instruction.
  - This happens to have a binary encoding of all 0s: 0000 .... 0000.
- Flushing introduces a bubble into the pipeline, which represents the one-cycle delay in taking the branch.
- The `IF.Flush` control signal shown on the next page implements this idea, but no details are shown in the diagram.



38

## Branching without forwarding and load stalls



39

## Timing

- If no prediction:
  - IF ID EX MEM WB
  - IF IF ID EX MEM WB ... lost 1 cycle
- If prediction:
  - If Correct
    - IF ID EX MEM WB
    - IF ID EX MEM WB -- no cycle lost
  - If Misprediction:
    - IF ID EX MEM WB
    - IF0 IF1 ID EX MEM WB ... 1 cycle lost

40

## Summary

- Three kinds of hazards conspire to make pipelining difficult.
- Structural hazards** result from not having enough hardware available to execute multiple instructions simultaneously.
  - These are avoided by adding more functional units (e.g., more adders or memories) or by redesigning the pipeline stages.
- Data hazards** can occur when instructions need to access registers that haven't been updated yet.
  - Hazards from R-type instructions can be avoided with forwarding.
  - Loads can result in a "true" hazard, which must stall the pipeline.
- Control hazards** arise when the CPU cannot determine which instruction to fetch next.
  - We can minimize delays by doing branch tests earlier in the pipeline.
  - We can also take a chance and predict the branch direction, to make the most of a bad situation.

41