

13. Pipelining: Data Hazards

EECS 370 – Introduction to Computer Organization – Winter 2009

Prof. Todd Austin & Prof. Marios Papaefthymiou

EECS Department
University of Michigan, Ann Arbor, USA


© Austin & Papaefthymiou, 2009

The material in this presentation cannot be
copied in any form without our written permission

Announcements

- ❑ Programming Contest – Entries accepted until Fri 3/6
- ❑ Homework 4 – Due Tue 3/17
- ❑ Exam 1 graded – To be handed back at the end of lecture
- ❑ Exam 1 regrade requests: Hand in a written request with one paragraph per question to be regraded. We reserve the right to regrade the entire exam.
- ❑ Exam 1 regrades are due by Thu 3/12 (one week from today).

Midterm 1 Statistics

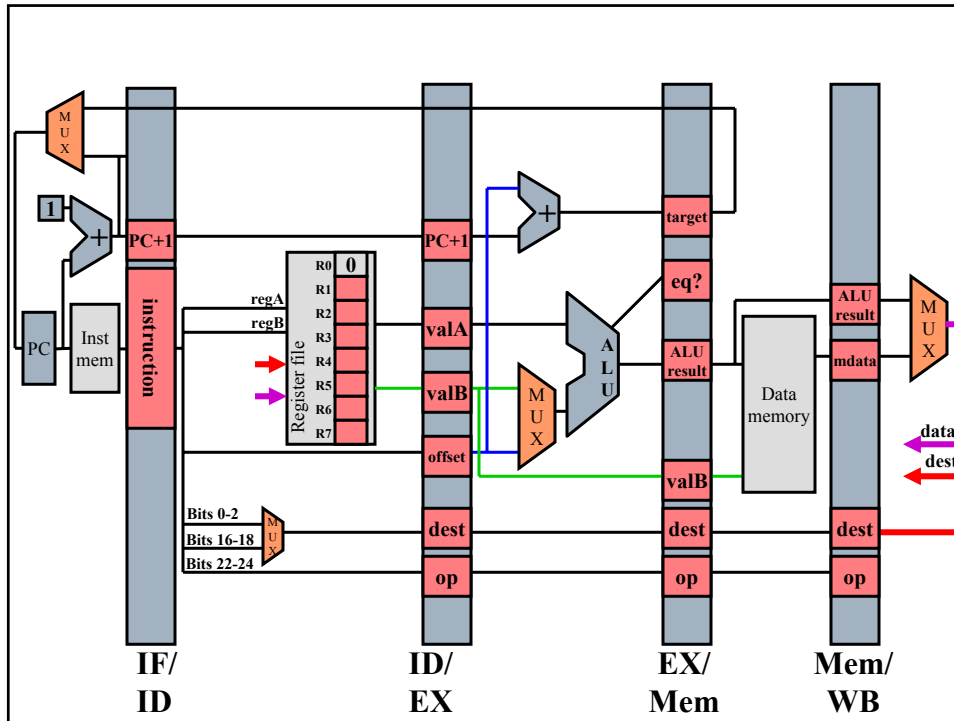


Median = 39

Std. Dev. = 11.5

Pipelining - What can go wrong?

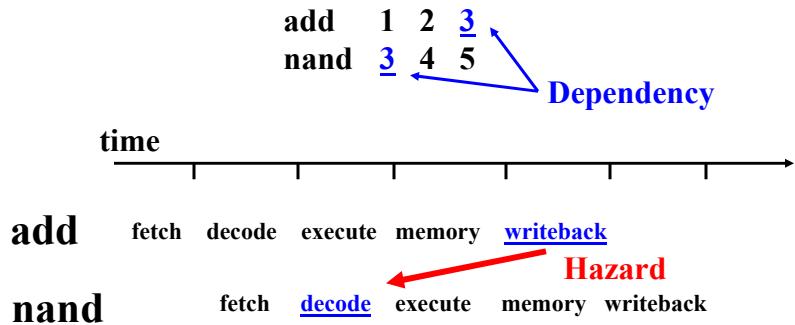
- ❑ **Data hazards:** since register reads occur in stage 2 and register writes occur in stage 5 it is possible to read the wrong value if it is about to be written.
- ❑ **Control hazards:** A branch instruction may change the PC, but not until stage 4. What do we fetch before that?
- ❑ **Exceptions:** How do you handle exceptions in a pipelined processor with 5 instructions in flight?
- ❑ **Today - Data hazards**
 - What are they?
 - How do you detect them?
 - How do you deal with them?



Pipeline function for ADD

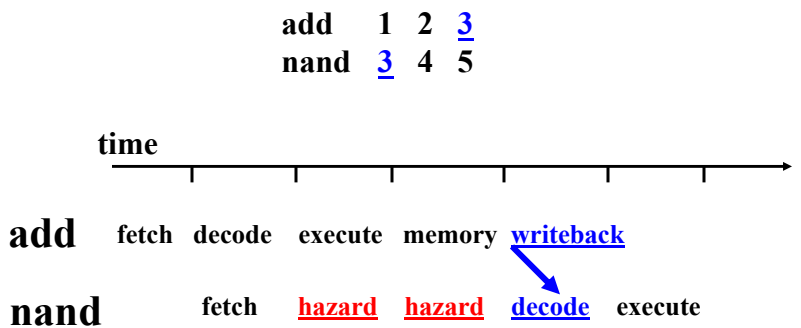
- Fetch: read instruction from memory
- Decode: read source operands from reg
- Execute: calculate sum
- Memory: pass results to next stage
- Writeback: write sum into register file

Data Hazards



If not careful, nand will read the wrong value of R3

Data Hazards



Assume Register File gives the right value of R3 when read/written during same cycle. This is consistent with the book and the MIPS architecture, but not Project 3.

Class Problem

Which data hazards do you see?

add 1 2 3

nand 3 4 5

add 6 3 7

lw 3 6 10

sw 6 2 12

Three approaches to handling data hazards

- Avoid
 - Make sure there are no hazards in the code
- Detect and Stall
 - If hazards exist, stall the processor until they go away.
- Detect and Forward
 - If hazards exist, fix up the pipeline to get the correct value (if possible)

Handling data hazards I: Avoid all hazards

- ❑ Assume the programmer (or the compiler) knows about the processor implementation.
 - Make sure no hazards exist.
 - Put noops between any dependent instructions.

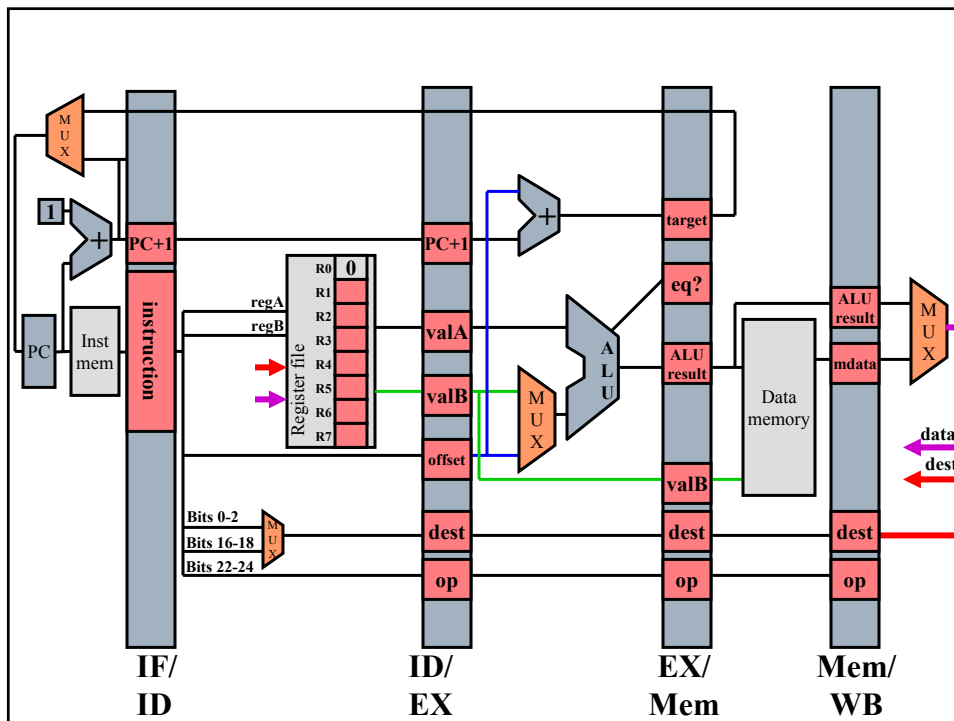
```
add  1  2  3 ← write R3 in cycle 5
noop
noop
nand 3  4  5 ← read R3 in cycle 5
```

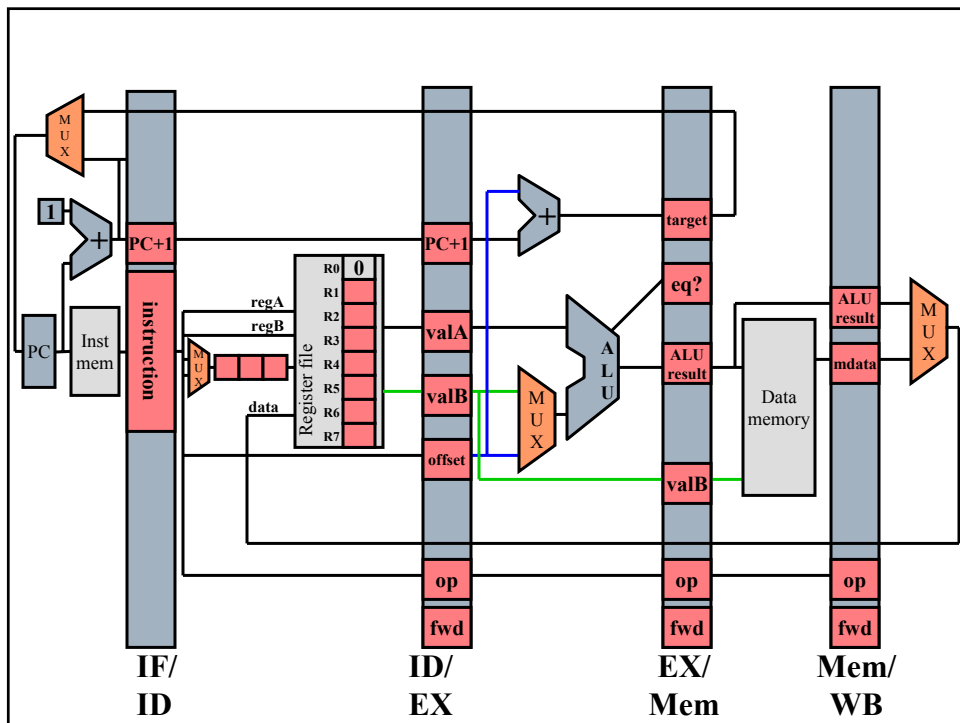
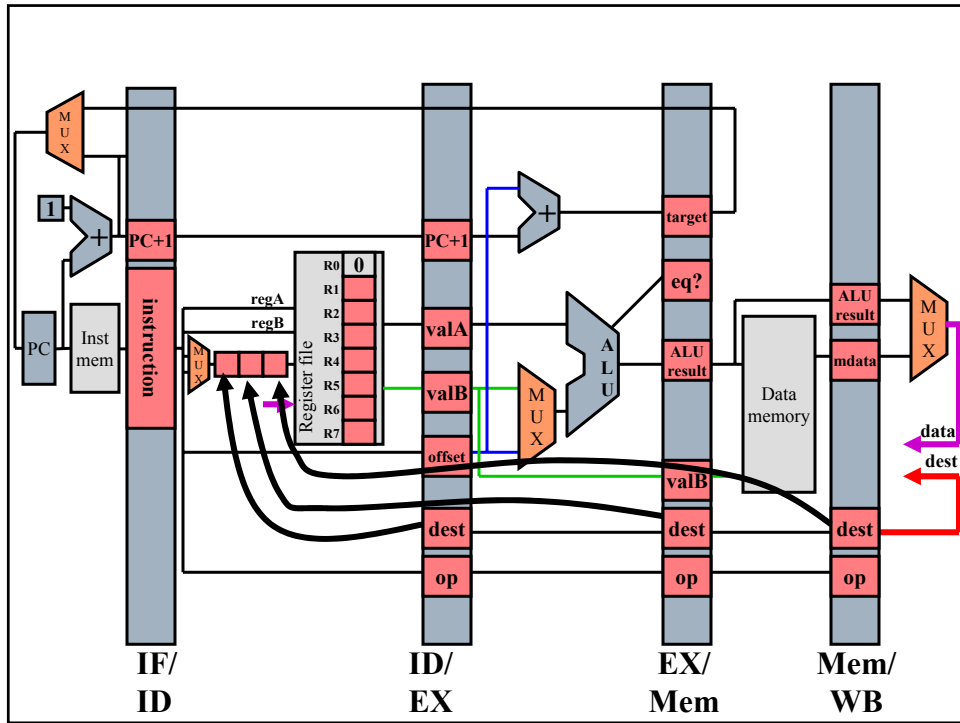
Problems with this solution


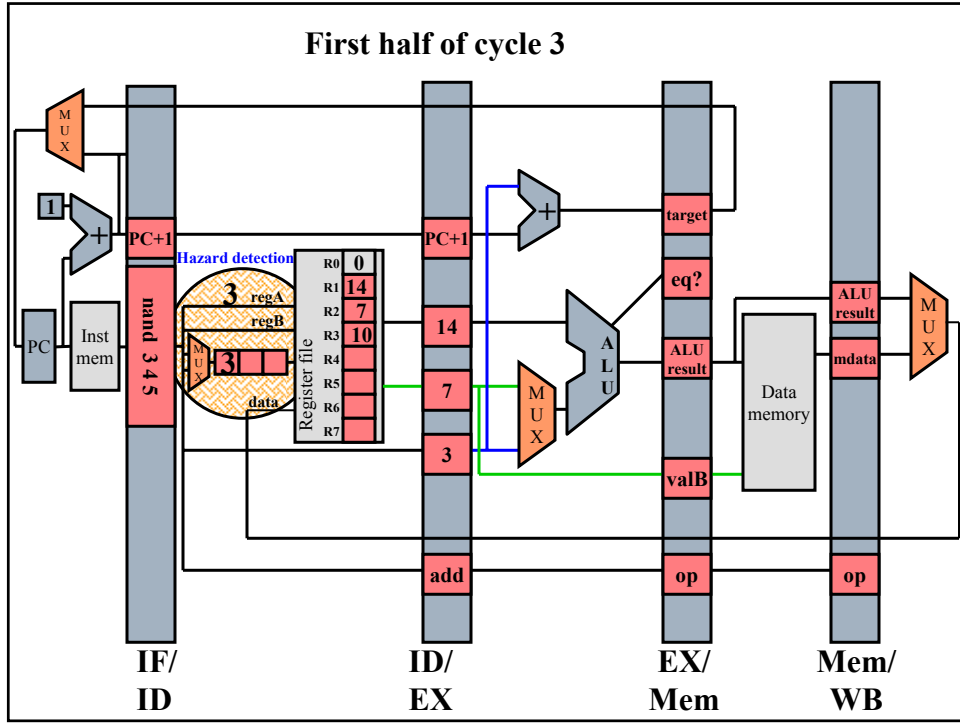
- ❑ Old programs (legacy code) may not run correctly on new implementations
 - Longer pipelines need more noops
- ❑ Programs get larger as noops are included
 - Especially a problem for machines that try to execute more than one instruction every cycle
 - Intel EPIC: Often 25% - 40% of instructions are noops
- ❑ Program execution is slower
 - CPI is 1, but some instructions are noops


Handling data hazards II: Detect and stall until ready

- ❑ Detect:
 - Compare regA with previous DestRegs
 - 3 bit operand fields
 - Compare regB with previous DestRegs
 - 3 bit operand fields
- ❑ Stall:
 - Keep current instructions in fetch and decode
 - Pass a noop to execute



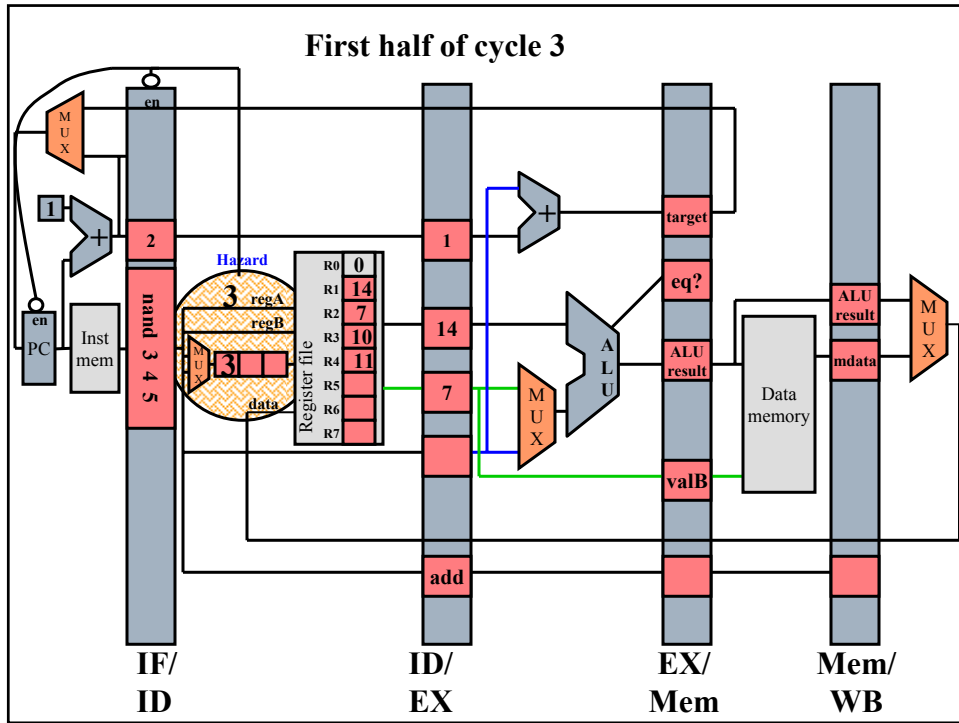






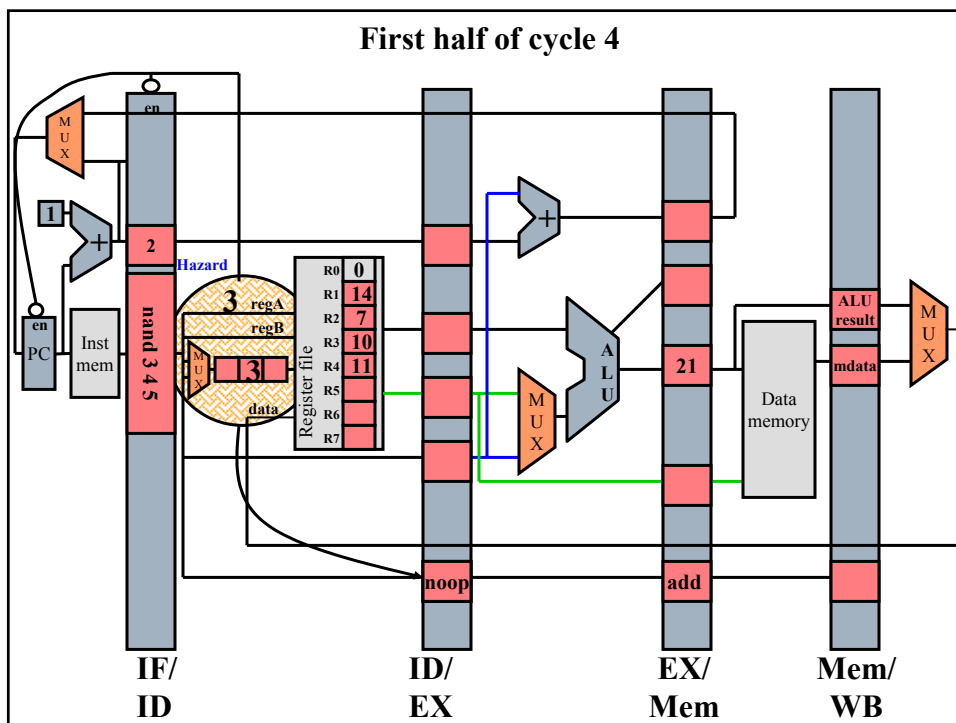
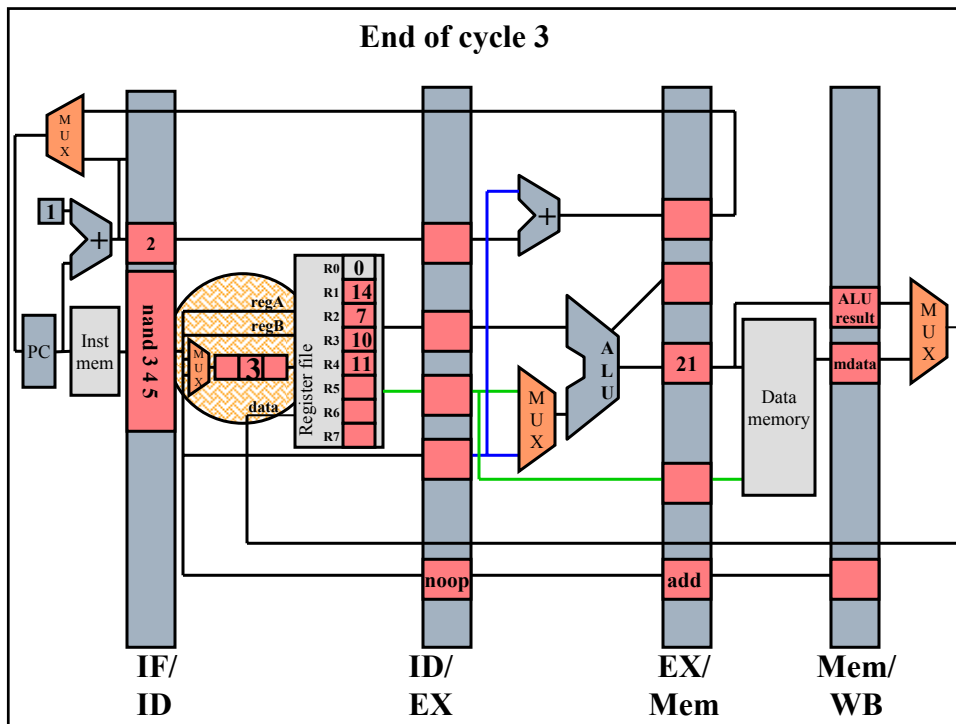
Handling data hazards II: Detect and stall until ready

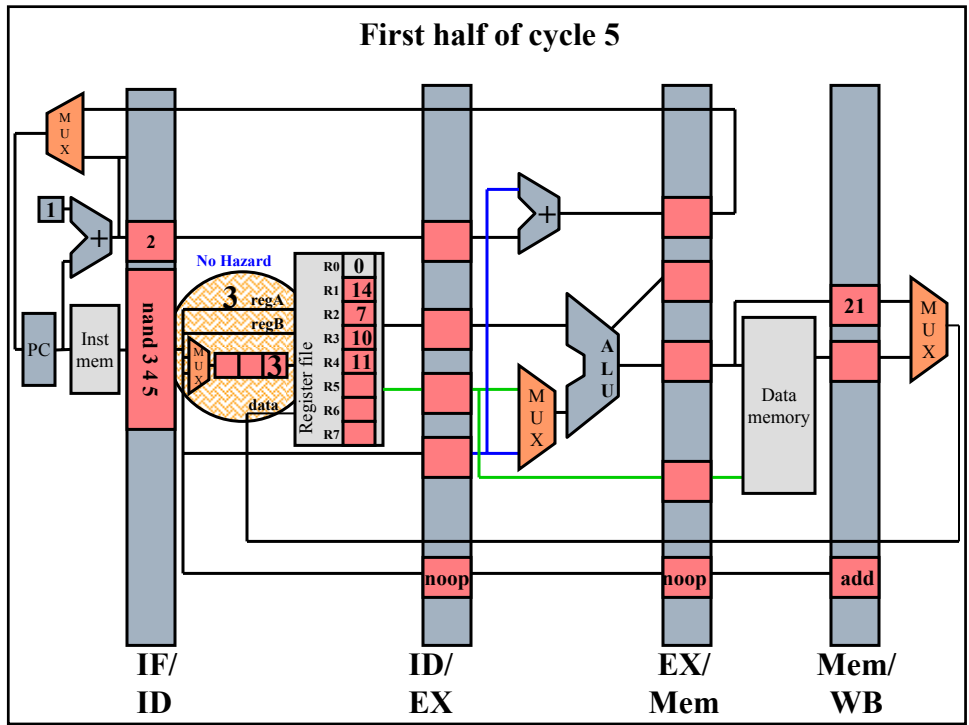
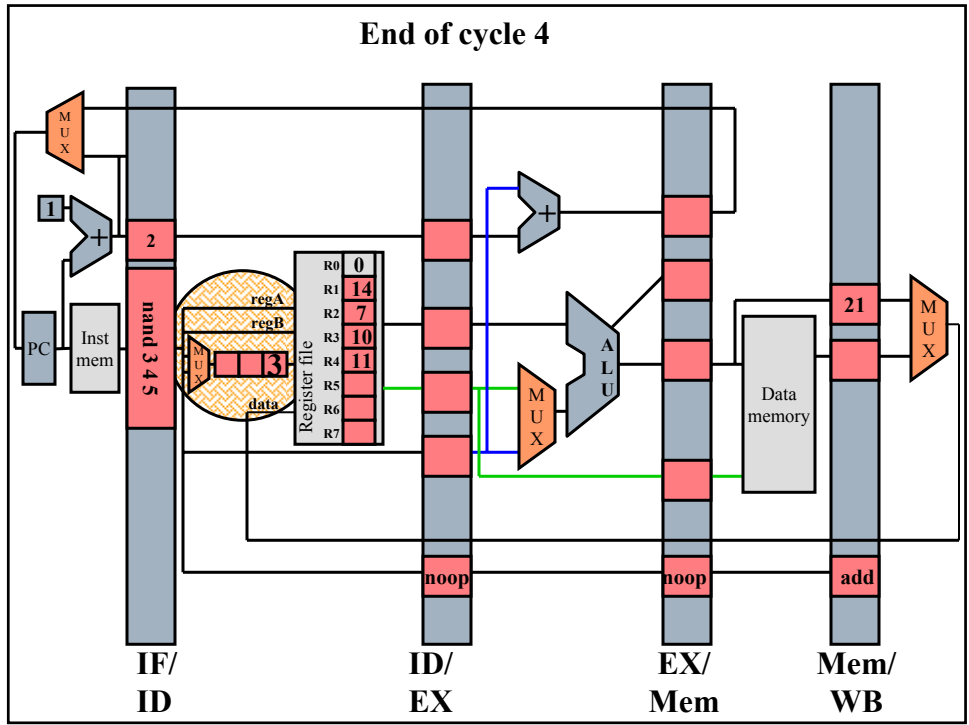
- ❑ Detect:
 - Compare regA with previous DestReg
 - 3 bit operand fields
 - Compare regB with previous DestReg
 - 3 bit operand fields
- ❑ Stall:
 - Keep current instructions in fetch and decode**
 - Pass a noop to execute

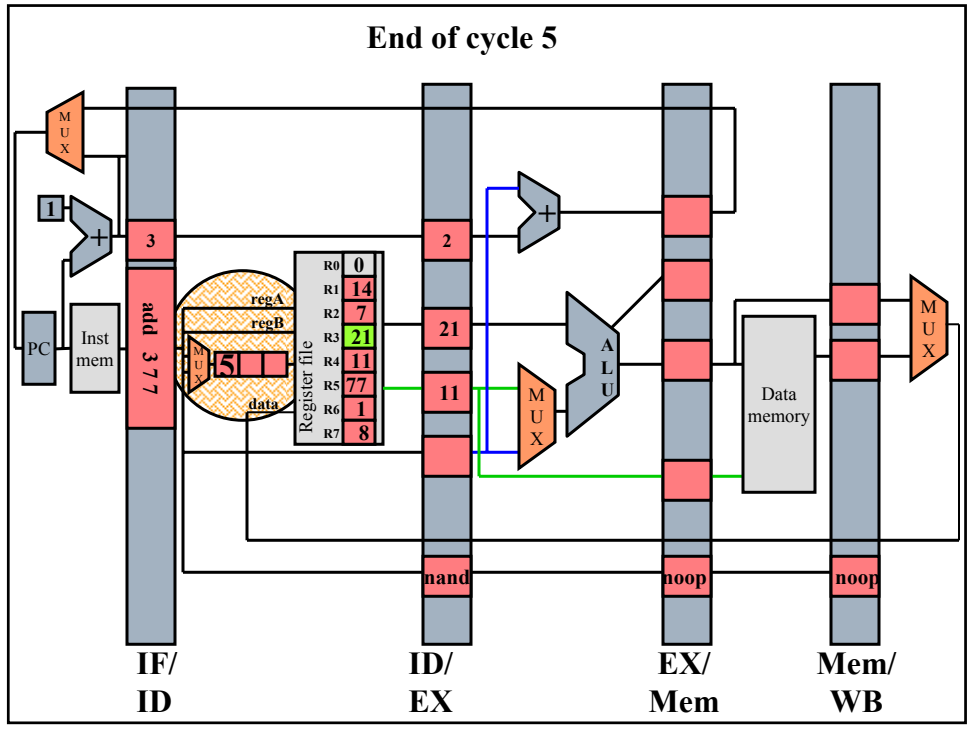


Handling data hazards II: Detect and stall until ready

- ❑ Detect:
 - Compare regA with previous DestReg
 - 3 bit operand fields
 - Compare regB with previous DestReg
 - 3 bit operand fields
- ❑ Stall:
 - Keep current instructions in fetch and decode
 - **Pass a noop to execute**







Time Graph

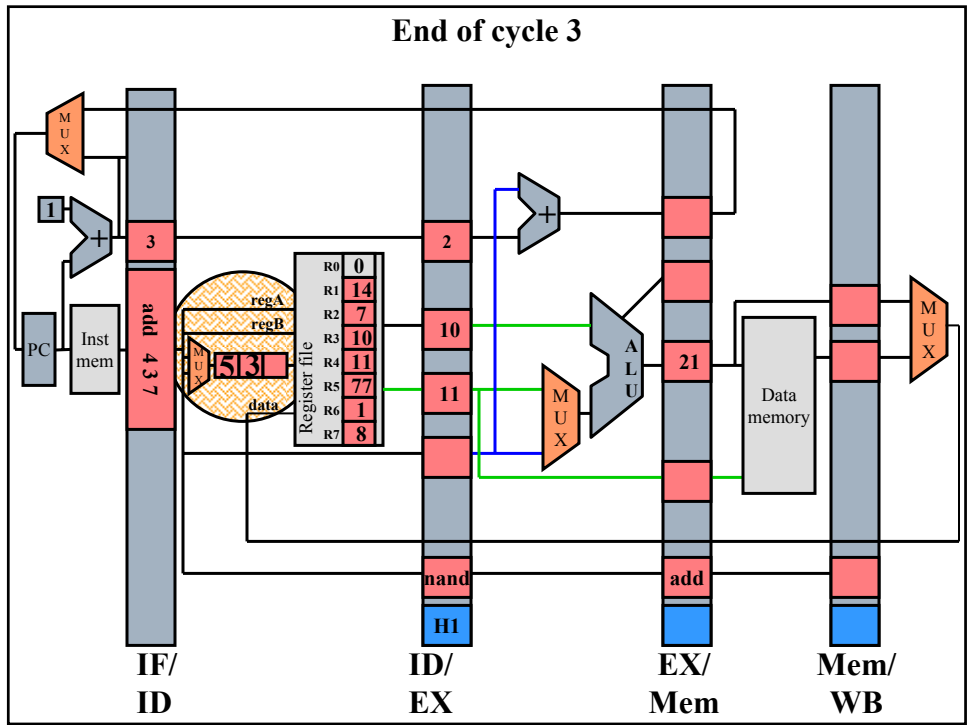

Time:	1	2	3	4	5	6	7	8	9	10	11	12	13
add 1 2 3	IF	ID	EX	ME	WB								
nand 3 4 5		IF	no op	no op	ID	EX	ME	WB					
add 6 3 7					IF	ID	EX	ME	WB				
lw 3 6 10						IF	ID	EX	ME	WB			
sw 6 2 12							IF	no op	no op	ID	EX	ME	WB

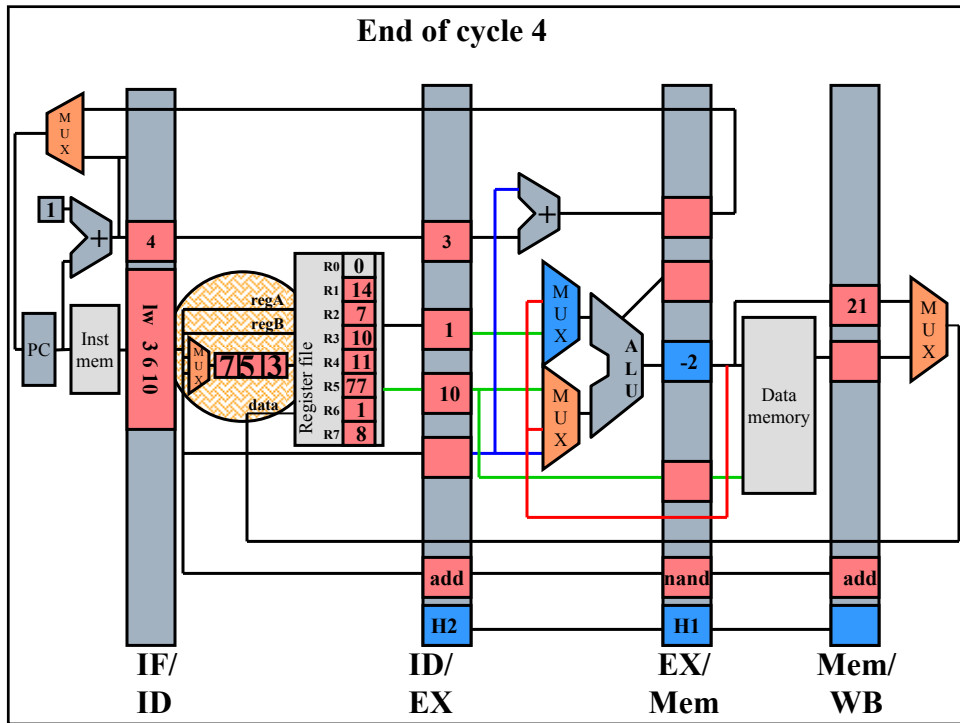
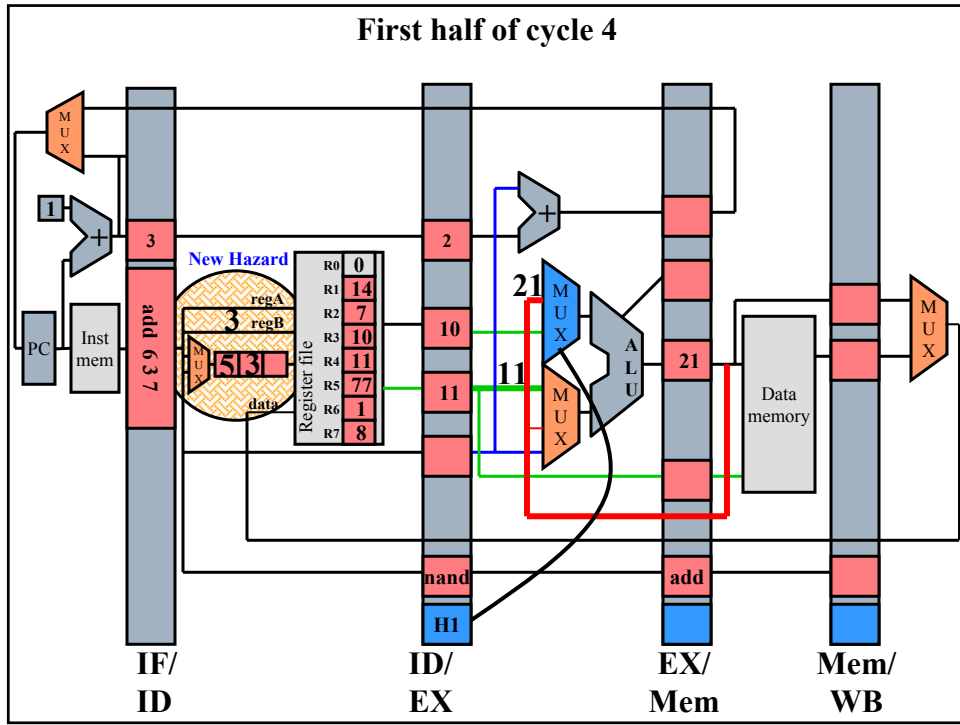
Problems with detect and stall

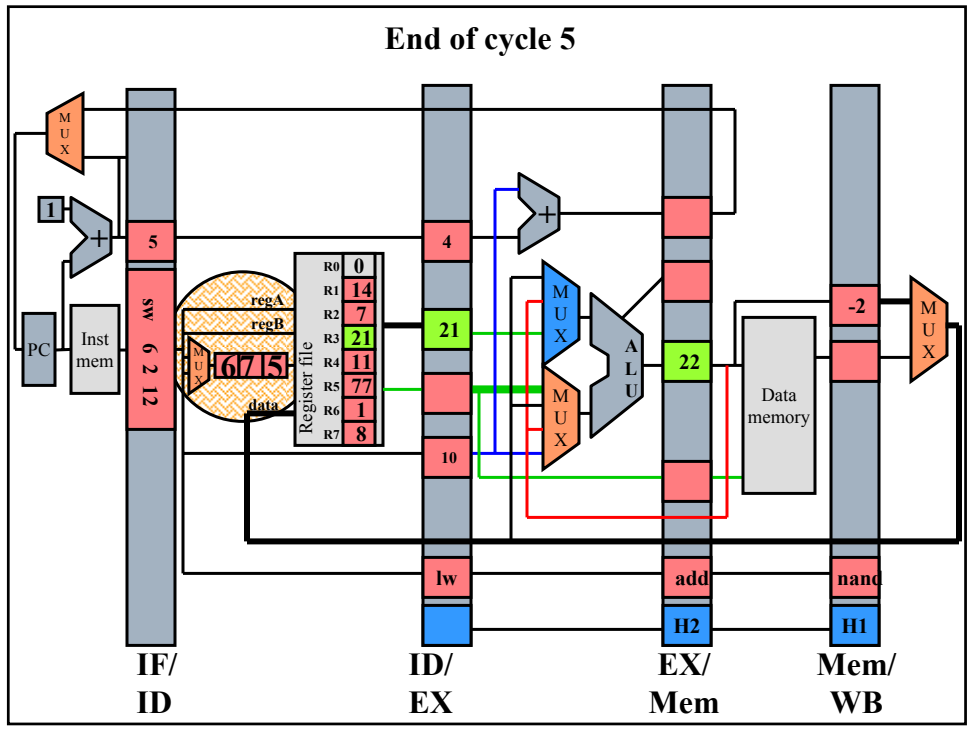
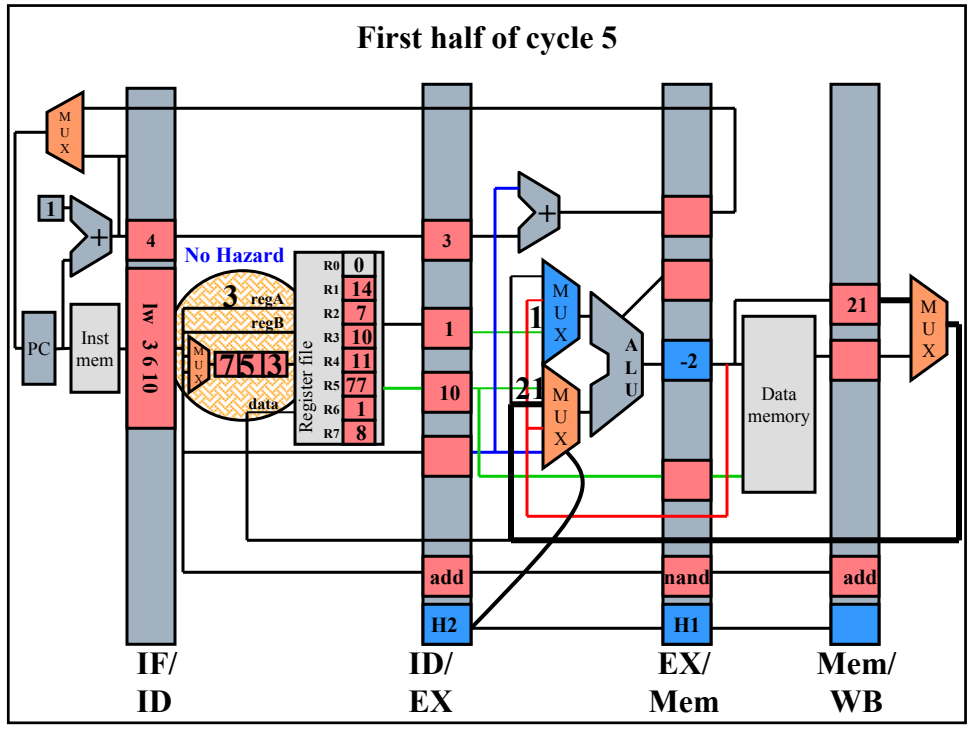
- ❑ CPI increases every time a hazard is detected!
- ❑ Is that necessary? Not always!
 - Re-route the result of the add to the nand
 - nand no longer needs to read R3 from reg file
 - It can get the data later (when it is ready)
 - This lets us complete the decode this cycle
 - But we need more control to remember that the data that we aren't getting from the reg file at this time will be found elsewhere in the pipeline at a later cycle.

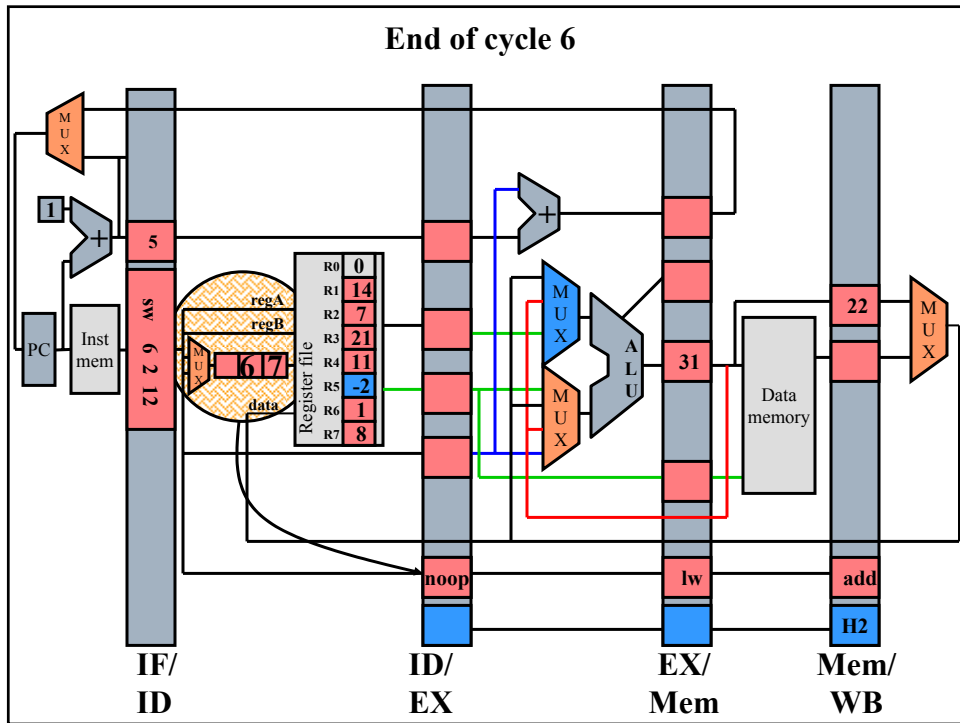
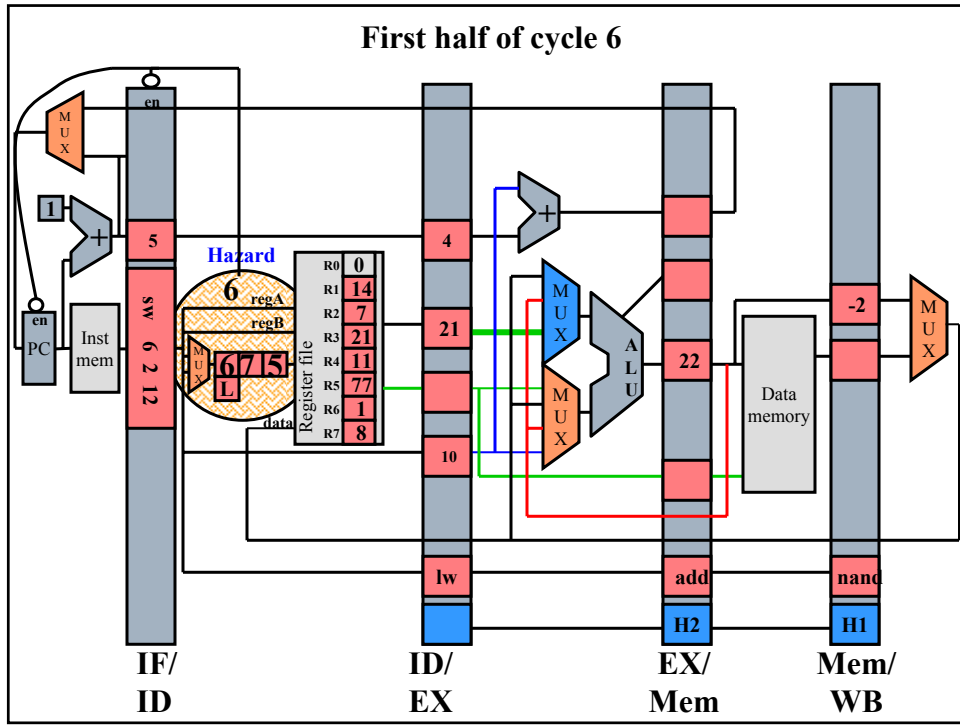
Handling data hazards III: Detect and forward

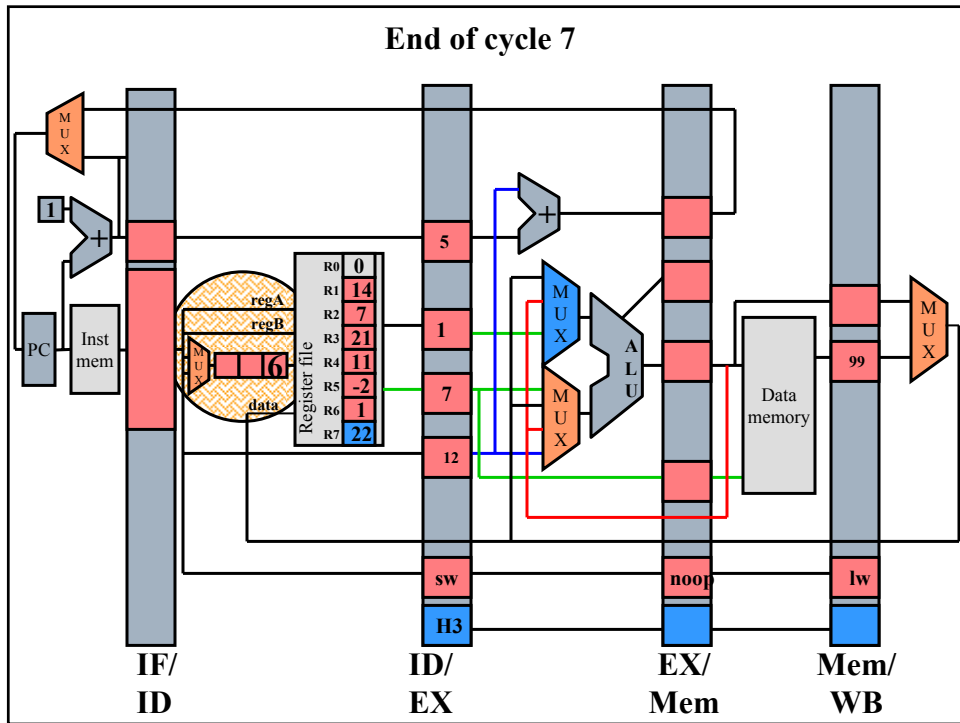
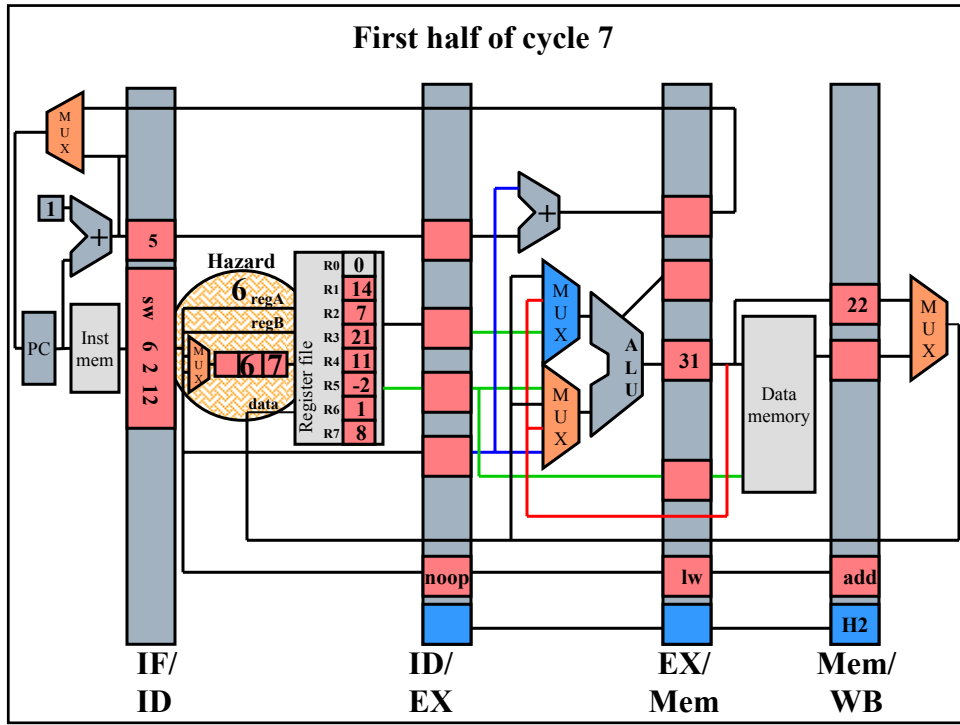
- ❑ Detect: same as detect and stall
 - Except that all 4 hazards are treated differently
 - i.e., you can't logical-OR the 4 hazard signals
- ❑ Forward:
 - New **bypass datapaths** route computed data to where it is needed
 - New MUX and control to pick the right data
- ❑ **Beware:** Stalling may still be required even in the presence of forwarding

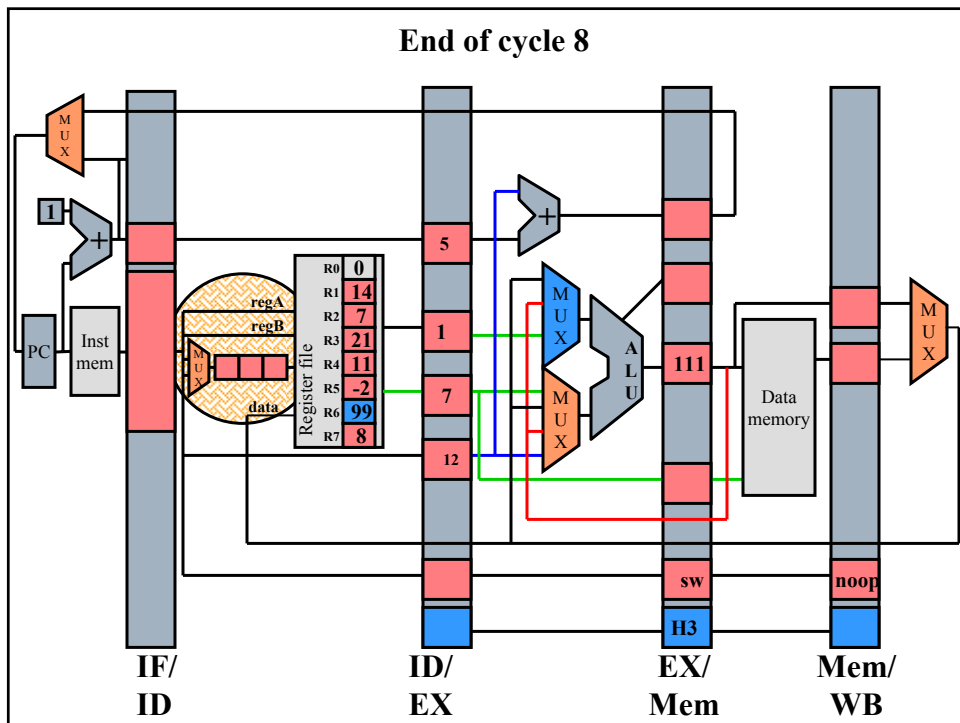
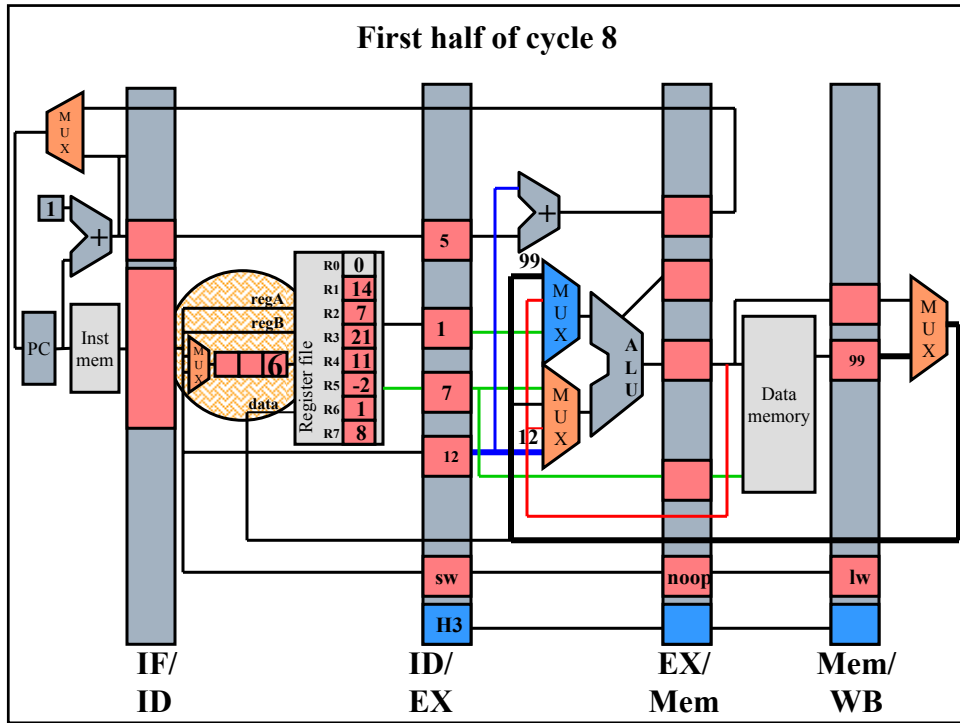












Time Graph

Time:	1	2	3	4	5	6	7	8	9	10	11	12	13
add 1 2 3	IF	ID	EX	ME	WB								
nand 3 4 5		IF	ID	EX	ME	WB							
add 6 3 7			IF	ID	EX	ME	WB						
lw 3 6 10				IF	ID	EX	ME	WB					
sw 6 2 12					IF	no op	ID	EX	ME	WB			

Next time

- ❑ Control hazards
 - How can the pipeline handle branch and jump instructions?