Basic Pipelined Processor

Hung-Wei Tseng



Outline

- Pipelining
- Pipeline Hazards
- Structural Hazards
- Control Hazards
- Dynamic Branch Predictions

Tasks in RISC-V ISA

- Instruction Fetch (IF) fetch the instruction from memory
- Instruction Decode (ID)
 - Decode the instruction for the desired operation and operands
 - Reading source register values
- Execution (**EX**)
 - ALU instructions: Perform ALU operations
 - Conditional Branch: Determine the branch outcome (taken/not taken)
 - Memory instructions: Determine the effective address for data memory access
- Data Memory Access (MEM) Read/write memory
- Write Back (WB) Present ALU result/read value in the target register
- Update PC
 - If the branch is taken set to the branch target address
 - Otherwise advance to the next instruction current PC + 4 •

Simple implementation w/o branch

- add x1, x2, x3 ID IF EX WB
- ld x4, 0(x5)
- sub x6, x7, x8
- sub x9, x10, x11
- sd x1, 0(x12)









1







- Different parts of the processor works on different instructions simultaneously
- A clock signal controls and synchronize the beginning and the end of each part of the work
- A pipeline register between different parts of the processor to keep intermediate results necessary for the upcoming work



add x1, x2, x3 ld x4, 0(x5) sub x6, x7, x8 sub x9, x10, x11 sd x1, 0(x12) xor x13, x14, x15 and x16, x17, x18 add x19, x20, x21 sub x22, x23, x24 ld x25, 4(x26) sd x27, 0(x28)

IF

ID	EX	MEM	WB				
IF	ID	EX	MEM	WB		I	
	IF	ID	EX	MEM	WB		
		IF	ID	EX	MEM	WB	
			IF	ID	EX	MEM	
				IF	ID	EX	
					IF	ID	
						IF	
			Afte wea	er thi are c	s poi omp	nt, letin	9
			inst	ructi	on e	ach c	



Draw the pipeline diagrams





Both instructions

Pipeline hazards



Three pipeline hazards

- Structural hazards resource conflicts cannot support simultaneous execution of instructions in the pipeline
- Control hazards the PC can be changed by an instruction in the pipeline
- Data hazards an instruction depending on a the result that's not yet generated or propagated when the instruction needs that



Structural Hazards



Dealing with the conflicts between ID/WB

- The same register cannot be read/written at the same cycle
- Solution: insert no-ops (e.g, add x0, x0, x0) between them
- Drawback
 - If the number of pipeline stages changes, the code won't work
 - Slow

add x1, x2, x3 ld x4, 0(x5) sub x6, x7, x8 add x0, x0, x0 sub x9, **x1**, x10







Dealing with the conflicts between ID/WB

- The same register cannot be read/written at the same cycle
- Solution: stall the later instruction, allowing the write to present the change in the register and the later can get the desired value
- Drawback: slow





Dealing with the conflicts between ID/WB

- The same register cannot be read/written at the same cycle
- Better solution: write early, read late
 - Writes occur at the clock edge and complete long enough before the end of the clock cycle.
 - This leaves enough time for outputs to settle for reads
 - The revised register file is the default one from now!







Structural Hazards

- Stall can address the issue but slow
- Improve the pipeline unit design to allow parallel execution

Control Hazards

Dynamic Branch Prediction

Why can't we proceed without stalls/no-ops?

- How many of the following statements are true regarding why we have to stall for each branch in the current pipeline processor
 - The target address when branch is taken is not available for instruction fetch stage of the next cycleYou need a cheatsheet for that — branch target buffer
 - ② The target address when branch is not-taken is not available for instruction fetch
 - stage of the next cycle. You need to predict that history/states The branch outcome cannot be decided until the comparison result of ALU is not out
 - The next instruction needs the branch instruction to write back its result (4)
 - A. 0
 - B. 1
 - - D. 3

E. 4

A basic dynamic branch predictor





2-bit/Bimodal local predictor

- Local predictor every branch instruction has its own state
- 2-bit each state is described using 2 bits
- Change the state based on actual outcome
- If we guess right no penalty
- If we guess wrong flush (clear pipeline registers) for mis-predicted instructions that are currently in IF and ID stages and reset the PC **(**)



	branch PC	target PC	Stat
	0x400048	0x400032	10
Predict Taken	0x400080	0x400068	11
	0x401080	0x401100	00
	0x4000F8	0x400100	01





2-bit local predictor



predict state actual 10 Т Т 11 Т Т 11 Т Т 11 Т Т Т NT 11

90% accuracy! $CPI_{average} = 1 + 20\% \times 10\% \times 2 = 1.04$

Two-level global predictor

Reading: Scott McFarling. Combining Branch Predictors. Technical report WRL-TN-36, 1993.

2-bit local predictor

• What's the overall branch prediction (include both branches) accuracy for this nested for loop?

```
i = 0;
                                                    i b
do {
   if( i % 2 != 0) // Branch X, taken if i % 2 == 0
      a[i] *= 2;
                         This pattern
   a[i] += i;
} while ( ++i < 100)// Branch</pre>
(assume all states state peats all the tin
 A. ~25%
                                                    3
  B. ~33%
                                                    4
  C. ~50%
                                                    4
 D. ~67%
                                                    5
                             For branch Y, almost 100%,
                                                    5
                             For branch X, only 50%
  E. ~75%
                                                    6
```

6

ranch?	state	prediction	actual
Х	00	NT	Т
Y	00	NT	Т
Х	01	NT	NT
Y	01	NT	Т
Х	00	NT	Т
ne	10	Т	Т
X	01	NT	NT
Y	11	Т	Т
Х	00	NT	Т
Y	11	Т	Т
Х	01	NT	NT
Y	11	Т	Т
Х	00	NT	Т
Y	11	Т	Т







Performance of GH predictor



Near perfect after this

i	branch?	GHR	state	prediction	actual
0	Х	000	00	NT	Т
0	Y	001	00	NT	Т
1	Х	011	00	NT	NT
1	Y	110	00	NT	Т
2	Х	101	00	NT	Т
2	Y	011	00	NT	Т
3	Х	111	00	NT	NT
3	Y	110	01	NT	Т
4	Х	101	01	NT	Т
4	Y	011	01	NT	Т
5	Х	111	00	NT	NT
5	Y	110	10	Т	Т
6	Х	101	10	Т	Т
6	Y	011	10	Т	Т
7	Х	111	00	NT	NT
7	Y	110	11	Т	Т
8	Х	101	11	Т	Т
8	Y	011	11	Т	Т
9	Х	111	00	NT	NT
9	Y	110	11	Т	т
10	Х	101	11	Т	Т
10	Y	011	11	Т	Т

Hybrid predictors

gshare predictor







gshare predictor

 Allowing the predictor to identify both branch address but also use global history for more accurate prediction



00

Local History Predictor branch PC local history

x400048	1000
x400080	0110
x401080	1010
x4000F8	0110

Predict Taken

Tournament Predictor

- The state predicts "which predictor is better"
 - Local history
 - Global history
- The predicted predictor makes the prediction



Branch predictor in processors

- The Intel Pentium MMX, Pentium II, and Pentium III have local branch predictors with a local 4-bit history and a local pattern history table with 16 entries for each conditional jump.
- Global branch prediction is used in Intel Pentium M, Core, Core 2, and Silvermont-based Atom processors.
- Tournament predictor is used in DEC Alpha, AMD Athlon processors
- The AMD Ryzen multi-core processor's Infinity Fabric and the Samsung Exynos processor include a perceptron based neural branch predictor.



Announcement

- Project is up check the website
- Assignment #3 due next Monday
- Midterm
 - Release Tuesday 0:00am, turn in before next Friday 11:59pm
 - You can only open it once and you have to finish a total of 30 questions within 80 minutes.
 - You may open book, but you have to bare the risks of not being able to finish them
- Attendance
 - The attendance throughout the quarter count as one assignment
 - You only need to answer 50% of the Zoom polls to receive full credits
 - Please don't email me for absence we count only 50% to give you flexibility
 - If you just login but never answer questions, you won't receive any.
- Reading Quizzes 2 attempts, average
- Office Hours on Zoom (the office hour link, not the lecture one)
 - Hung-Wei/Prof. Usagi: M 8p-10p (make up for the last week), W 2p-3p
 - Quan Fan: F 1p-3p

hin 80 minutes. h them Computer Science & Engineering





