Basic Pipelined Processor

Hung-Wei Tseng



Recap: von Neuman Architecture



509cbd23 (intel) 00c2e800

Processor

Program

nstructions

0f00bb27 509cbd23 00005d24 0000bd24 2ca422a0 130020e4 00003d24 2ca4e2b3

00000008 00c2f000 00000008 00c2f800 00000008 00c30000 0000008

10

00c2e800

Storage

Outline

- Basic Pipeline Processor Design
- Pipeline Hazards
 - Structural Hazards
 - Control Hazards
 - Data Hazards
- Dynamic Branch Predictions

Which version is faster?



- Both version A and B produces the same output. Without compiler optimization, which version of code would have better performance?
 - A. Version A
 - B. Version B
 - C. They are about the same (less than 5% difference)

https://www.pollev.com/hungweitseng close in 1:30

sum+=data[((i << 4) - i) & 131071];</pre>

Performance shift



В

А

С

D

Ε

Which version is faster?



- Both version A and B produces the same output. Without compiler optimization, which version of code would have better performance?
 - A. Version A
 - B. Version B

C. They are about the same (less than 5% difference)



sum+=data[((i << 4) - i) & 131071];</pre>

https://www.pollev.com/hungweitseng close in 1:30

Which swap is faster?



*a ^= *b; *b ^= *a; *a ^= *b; }

- Both version A and B swaps content pointed by a and b correctly. Which version of code would have better performance?
 - A. Version A
 - B. Version B
 - C. They are about the same (sometimes A is faster, sometimes B is)



А

void xorswap(int* a, int* b) {

В

Performance xor

С

D

Ε

Which swap is faster?



 Both version A and B swaps content pointed by a and b correctly. Which version of code would have better performance?

A. Version A

- B. Version B
- C. They are about the same (sometimes A is faster, sometimes B is)



Recap: Why adding a sort makes it faster

• Why the sorting the array speed up the code despite the increased instruction count?

```
if(option)
    std::sort(data, data + arraySize);
for (unsigned i = 0; i < 100000; ++i) {</pre>
    int threshold = std::rand();
    for (unsigned i = 0; i < arraySize; ++i) {</pre>
        if (data[i] >= threshold)
             sum ++;
    }
}
```

Recap: Adding a sort...

```
if(option)
    std::sort(data, data + arraySize);
for (unsigned c = 0; c < arraySize*1000; ++c) {</pre>
        if (data[c%arraySize] >= INT_MAX/2)
             sum ++;
    }
}
```



Basic Processor Design

von Neuman Architecture





509cbd23 (intel)

Processor

Da

Storage

RYZEN

Program

Instructions

0f00bb27 509cbd23 00005d24 0000bd24 2ca422a0 130020e4 00003d24 2ca4e2b3

00c2e800 80000008 00c2f000 ta 00000008 00c2f800 80000008 00c30000 80000008 :10

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



IF

D









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 IF 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)

ID	EX	MEM	WB				
IF	ID	EX	MEM	WB		•	
	IF	ID	EX	MEM	WB		-
		IF	ID	EX	MEM	WB	
			IF	ID	EX	MEM	
				IF	ID	EX	
					IF	ID	
						IF	
			Afte we inst	er thi are c ructi	s poi omp on ea	nt, letinț ach c	



https://www.pollev.com/hungweitseng close in 1:30

Can we get them right?

 Given a simple pipelined RISC-V processor that we discussed so far, how many of the following code snippets can be executed with expected outcome?

I				ll							IV		
а	add	x1,	x2, x3	add	x1,	x2,	x3	add	x1,	x2, x3	add	x1, x2, x3	
b	ld	x4,	0(x1)	ld	χ4,	0(x!	ō)	ld	x4,	0(x5)	ld	x4, 0(x5)	
С	sub	x6,	x7, x8	sub	x6,	x7,	x8	bne	x0,	x7, L	sub	x6, x7, x8	
d	sub	x9,	x10,x11	sub	x9,	x1 ,	x10	sub	x9,	x10,x11	sub	x9,x10,x11	
е	sd	x1,	0(x12)	sd	x11	, 0()	x12)	sd	x1,	0(x12)	sd	x1, 0(x12)	
	A. C)										Pipeline	
	B. 1												
	C. 2	-											
	D. 3	8											
	E. 4	Ļ											

А

С

D

Draw the pipeline diagrams





Both instructions

Can we get them right?

 Given a simple pipelined RISC-V processor that we discussed so far, how many of the following code snippets can be executed with expected outcome?





			V	
3	add	x1,	x2, x3	
	ld	x4,	0(x5)	
	sub	x6,	x7, x8	
-	sub	x9,	x10,x11	
	sd	x1,	0(x12)	

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



Can we get them right?

 Given a simple pipelined RISC-V processor that we discussed so far, how many of the following code snippets can be executed with expected outcome?



			V	
3	add	x1,	x2, x3	
	ld	x4,	0(x5)	
	sub	x6,	x7, x8	
-	sub	x9,	x10,x11	
	sd	x1,	0(x12)	

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 sd x11, 0(x12)

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!

https://www.pollev.com/hungweitseng close in 1:30 Why "split L1" cache

- Modern processors typically adopts a split L1 cache design that separates the instruction access from data accesses. What pair of instructions will be problematic if we don't use a split cache?
 - a: ld x1, 0(x2) b: add x3, x4, x5 c: sub x6, x7, x8 d: sub x9,x10,x11 e: sd x1, 0(x12) A. a&bB. a&c C. b&e D. c&e E. None

L2\$

L3\$

Why "split L1" cache

• Modern processors typically adopts a split L1 cache design that separates the instruction access from data accesses. What pair of instructions will be problematic if we don't use a split cache?

Why always 5 stages

- What pair of instructions will be problematic if we allow ALU instructions to skip the "MEM" stage?
 - a: ld x1, 0(x2)
 - b: add x3, x4, x5
 - c: sub x6, x7, x8
 - d: sub x9,x10,x11
 - e: sd x1, 0(x12)
 - A. a&b
 - B. a&c
 - C. b&e
 - D. c&e
 - E. None

https://www.pollev.com/hungweitseng close in 1:30

StructuralHazard

А

В

Such the presentation to see live-content. For screen share software, share the entire screen. Get help at policy.com/spc

С

D

Ε

Why always 5 stages

 What pair of instructions will be problematic if we allow ALU instructions to skip the "MEM" stage?

Structural Hazards

- Stall can address the issue but slow
- Improve the pipeline unit design to allow parallel execution
 - Write-first, read later register files
 - Split L1-Cache
 - All instructions need to go through 5 stages

Control Hazards

https://www.pollev.com/hungweitseng close in 1:30 The impact of control hazards

- Assuming that we have an application with 20% of branch instructions and the instruction stream incurs no data hazards. When there is a branch, we disable the instruction fetch and insert no-ops until we can determine the PC (happens in the EXE stage). What's the average CPI if we execute this program on the 5-stage **RISC-V** pipeline?
 - A. 1
 - B. 1.2
 - C. 1.4
 - D. 1.6
 - E. 1.8

ControlHazard

А

С

The impact of control hazards

 Assuming that we have an application with 20% of branch instructions and the instruction stream incurs no data hazards. When there is a branch, we disable the instruction fetch and insert no-ops until we can determine the PC (happens in the EXE stage). What's the average CPI if we execute this program on the 5-stage RISC-V nineline?

	add x1,	x2, x3	IF	ID	EX	MEM	WB						
A. 1	ld x4,	0(x5)		IF	ID	EX	MEM	WB					
R 12	bne x0,	x7, L			IF	ID	EX	MEM	WB				
D. 1.2	add x0,	x0, x0				IF	ID	EX	MEM	WB			
C. 1.4	add x0,	x0, x0					IF	ID	EX	MEM	WB		
D. 1.6	sub x9,	x10,x11						IF	ID	EX	MEM	WB	
	sd x1,	0(x12)							IF	ID	EX	MEM	WB
E. 1.8	1 + 20%	$\times 2 = 1.4$	ŀ										
				54									

https://www.pollev.com/hungweitseng close in 1:30 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 (1)the next cycle
 - The target address when branch is not-taken is not available for instruction fetch 2 stage of the next cycle
 - The branch outcome cannot be decided until the comparison result of ALU is not out (3)
 - The next instruction needs the branch instruction to write back its result (4)

Stall

Sort the presentation to see live content. For screen share software, share the entire screen. Get help at polytocompapy

С

D

Ε

В

А

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 cycle
 - ② The target address when branch is not-taken is not available for instruction fetch stage of the next cycle
 - The branch outcome cannot be decided until the comparison result of ALU is not out
 - ④ The next instruction needs the branch instruction to write back its result
 - A. 0
 - B. 1
 - C. 2
 - D. 3

n result of ALU is not out ck its result

Quizzes **Midterm** Grades

Piazza

Collaborations

- Average is 72 (last year was 66, online, open-book)
- Your overall grade decides your final letter grade, not just the midterm
- We still have 45% of the grades to be offered
- Pick up outside of Hung-Wei's office @ WCH 406

Current "Total" in eLearn and "Projected" Letter Grades

Partic	ipation
Particip	pation

Participation	-	100
Project	N/A	0.00
Reading Quizzes	0%	0.00 200.0
Assignments	0%	0.00 100.0
Participation	N/A	0.00
Imported Assignments	0%	0.00 100.0
Midterm	N/A	0.00
Final	N/A	0.00
Total	0%	

Announcements

- Reading quiz due next Monday
- Project is released
 - Please check website to the link of GitHub repo
 - You may discuss, but each needs an individual/distinguishable version of code
 - You need to write a brief report
 - Due 11/29 no extension