





https://goo.gl/forms/HKUVwLhVUYzvdat42

- Midterm #2 on 11/29
- 5th and final problem set on 11/22
- 9th and final lab on 12/1

ARM 3-STAGE PIPELINE

- Fetch, Decode, and Execute Stages
- Instructions are decoded in both the Fetch and Decode stages
- Register ports are "Read" in the Decode stage and "Written" at in the end of the Execute stage
- PC+4, register reads for stores (StrData). and BX source (BXreg) are "delayed" for use in later stages



SIMPLE INSTRUCTION FLOW



Consider the following instruction sequence:

Instruction becomes available at the ... end of the Fetch stage sub r0,r1,r2 add r3,r1,#2 Operands at the end of Decode and r1,r1,#1 cmp r0,r4

Destination and PSR are updated at the end of Execute

Time (íin	clock	cvc	les)
		01001		,

	i	i+1	i+2	i+3	i+4	i+5
Fetch	sub r0,r1,r2	add r3,r1,#2	and r1,r1,#1	cmp r0,r4		
Decode		sub r0,r1,r2	2add r3,r1,#2	2and r1,r1,#1	cmp r0,r4	4 O
Execute			sub r0,r1,r2	add r3,r1,#2	and r1,r0,#1	cmp r0,r4

PIPELINE CONTROL HAZARDS



Pipelining HAZARDS are situations where the next instruction cannot execute in the next clock cycle. There are two forms of hazards, CONTROL and STRUCTURAL.

Consider the instruction sequence shown:

• • •		
loop:	add	r0,r0,r0
	cmp	r0,#64
	ble	loop
	and	r1,r0,#7
	sub	r1,r0,r1

	Time (in clock cycles)							
	i	i+1	i+2	i+3	i+4	i+5		
Fetch	add r0,r0,r0	cmp r0,#64	ble loop	and r1,r0,#7	sub r1,r0,r1	add r0,r0,r0		
Decode		add r0,r0,r0	cmp r0,#64	ble loop	and r1,r0,#7	???		
Execute			add r0,r0,r0	cmp r0,#64	ble loop	???		
	T _ When the branch instruction							



I reaches the execute stage the next 2 instructions have already been fetched!

Pipeline





Problem: Two instructions following a branch are fetched before the branch decision is made (to take or not to take)

Solutions:

- I. Program around it. Define the ISA such that the branch does not take effect until after instructions in the "DELAY SLOTS" complete. This is how MIPS pipelines work. It leads to ODD looking code in tight (short) loops. Of course you could always put NOPs in the delay slots.
- 2. Detect the branch decision as early as possible, and ANNUL instructions in the delay slots. This is what ARM does.

EARLY DETECT AND ANNUL





We can detect branch instructions (B, BL, or BX) in the Decode stage. The decision to branch is decided no later than the current instruction in loop: add r0,r0,r0 the Execute stage. Thus, we could make r0,#64 cmp The branch decision in the Decode stage. ble loop We then annul the following instruction by and r1, r0, #7 disabling WERF and PSR updates! Making the sub r1, r0, r1 next instruction a NOP! Time (in clock cycles)

							ŕ
Dinalina		i	i+1	i+2	i+3	i+4	i+5
	Fetch	add r0,r0,r0	cmp r0,#64	ble loop	and r1,r0,#7	add r0,r0,r0	cmp r0,#64
	Decode		add r0,r0,r0	cmp r0,#64	ble loop	and r1,r0,#7	add r0,r0,r0
	Execute			add r0,r0,r0	cmp r0,#64	ble loop	NOP
			•		ţ	If we detect the b stage then the PSA instruction in the	ranch in the decode State of the Execute stage can

pe combined to change the next PC.

THE COST OF TAKEN BRANCHES



When an ARM branch is **taken** the branch instructions are effectively 2 cycles rather than 1 when they aren't. In a MIPS-like instruction set, one can often fill the delay slots with useful instructions, but they are executed whether or not the branch is taken.

The ARM approach is easier to understand, and since it does not "EXPOSE" the pipeline, it also allows for an alternative number of pipeline stages to be implemented in future designs, while conserving code compatibility.

Lastly, using ARM, many *conditional branches can be eliminated using the condition execution*, which pipelines beautifully!



STRUCTURAL PIPELINE HAZARDS

There's another problem with our code fragment!

The destination register of instructions are written at the end of the Execute stage. However the following instruction might use this result as a source operand.

loop:	add	r0,r0,r0
•	cmp	r0,#64
	ble	loop
	and	r1,r0,#7
	sub	r1,r0,r1

Time (in clock cycles) i+5 i i+1 i+2 i+3 i+4 **Fetch** add r0,r0,r0 cmp r0,#64 ble loop and r1.r0.#7 add r0.r0.r0 cmp r0,#64 Pipeline cmpr0,#64 Decode add r0.r0.r0 and r1.r0.#7 add r0.r0.r0 ble loop add r0,r0,r0 cmp r0,#64 ble loop NOP Execute The "CMP" instruction needs to access the contents of RO before it is actually written are the end of i+2 1/20/2017

DATA HAZARDS



Problem: When a register source is needed from a later stage of the pipeline before it is written.

Solutions:

- 1. Program around it. One could document the weird semantics-- "You can't reference the destination register of an instruction in the immediately following instruction." Would make make assembly language even harder to understand. Would expose the pipeline, once again making future improvements difficult to implement while maintaining code compatibility.
- 2. Hardware bypass multiplexers.

SOURCE BYPASSING



The idea here is to load the value that will be saved in the destination register into the pipeline registers hold the ALU operands.

We also need bypass MUXes on the StrReg and BXreg pipeline registers.



		i	i+1	i+2	
Dinolino	Fetch	add r0,r0,r0	cmp r0,#64	ble loop	
	Decode		add r0,r0,r0	cmp r0,#64	
	Execute			add r0,r0,r0	

The new value for RO will be computed just prior to the rising clock edge between i+2 and i+3, we can take the output of



LOAD/STORE STALLS



Load and Store memory accesses are the actual bottleneck of the ARM pipeline. Also, recall that instructions and load/stores actually come from the same memory. Thus, we need to stall instruction fetching to allow loop: r0,[r1,#4] ldr for loads and stores. add r0,r0,#4

r0,[r1,#4] str r0,r0,#4 sub and r2, r2, #f

			Time (in clock cycles)								
		i	i+1	i+2	i+3	i+4	i+5	i+6			
Pipeline	Fetch	ldr r0,[r1,#4]	add r0,r0,#4	str r0,[r1,#4]	Mem read	sub r0,r0,#4	and r2,r2,#f	Mem store			
	Decode		ldr r0,[r1,#4]	add r0,r0,#4		str r0,[r1,#4]	sub r0,r0,#4				
	Execute			ldr r0,[r1,#4]		add r0,r0,#4	str r0,[r1,#4]				



LOAD/STORE STALL IMPLEMENTATION



Disable loading of pipeline registers for one clock when a load or store instruction reaches the execute stage.

- 1. Adding enable lines to the PC and pipeline registers on the control path
- 2. A simple 2-state state machine to stall the pipeline for 1 state to allow for the load/store memory cycle.



WHERE DOES THIS LEAVE US

Overall we can now nearly triple the clock rate. Instructions have a throughput of one-per-clock with the following caveats:

- 1. Taken branches take 2 cycles.
- 2. Loads and store take 2 cycles.



You can pipeline an ARM CPU even more. There exist ARM implementations with 7, 8, and 9 pipeline stages. But the overhead of bypass paths and stall cases increase.



REALITY VS SPECMANSHIP



Assuming approximately 10% of instructions executed are branches, and of those 80% of the time they are taken, and 12.5% of instruction executed are loads or stores, what sort of real speed up do we expect?

$$Perf_{before} = (100) * 1 = 100 \text{ Clocks} * 10 * 10^{-9} \text{ sec/clock} = 1000 * 10^{-9} \text{ secs}$$

$$Perf_{after} = (10)(0.8) * 2 + 12.5 * 2 + 79.5 * 1 = 120.5 \text{ Clocks}$$

$$10.5 * 3333 * 10^{-9} \text{ sec} (clock = 40)(c(c * 10^{-9} \text{ secs}))$$

Speedup =
$$\frac{\text{Perf}_{before}}{\text{Perf}_{after}} = 1000/401.666 = 2.490 X$$

11/20/2017

15

It appears memory access time is our real bottleneck. What tricks can be applied to improving CPU performance in this case?

- Interleaving
- Block-transfers
- Caching

NEXT TIME



