CS 261 Fall 2017

Mike Lam, Professor



CPU architecture

Topics

- CPU stages and design
- Pipelining
- Y86 semantics

CPU overview

- A CPU consists of
 - Combinational circuits for computation
 - Sequential circuits for (cached) memory
 - Wires/buses for connectivity and intermediate results
 - A clocked register PC for synchronization















③ Beginning of cycle 4







6-stage von Neumann cycle

- 1) Fetch
- 2) Decode
- 3) Execute
- 4) Memory
- 5) Write back
- 6) PC update

Fetch

- Read ten bytes from memory at address PC
- Extract instruction fields
- Compute valP (address of next instruction)
 - oldPC + 1 +
 needsRegIDs +
 8*needsValC



Decode (and Write back)

- Read and write register file
 - Two simultaneous reads and two simultaneous writes
 - Use value 0xF to disable a read or write



Execute

- Perform arithmetic or logic operation
 - Could also be an effective address calculation or stack pointer increment / decrement
- Set condition codes
 - Only if OPq



Memory

- Read or write memory
 - No instruction does both!



PC update

- Set new PC
 - Will depend on whether a branch should be taken

		PC		
		New PC		
1		1	1	1
icode	Cnd	valC	valM	valP

CPU design

- SEQ: sequential Y86 CPU
 - Runs one instruction at a time
 - ssim: simulator
- Components:
 - Clocked register (PC)
 - Hardware units (blue boxes)
 - Combinational/sequential circuits
 - ALU, register file, memory
 - Control logic (grey rectangles)
 - Combinational circuits
 - Details in textbook
 - Wires (white circles)
 - Word (thick lines)
 - Byte (thin lines)
 - Bit (dotted lines)
- Principle: no reading back
 - Stages run simultaneously
 - Effects remain internally consistent



• CPU measurement

- Throughput: instructions executed per second
 - GIPS: billions of ("giga-") instructions per second
 - 1 GIPS \rightarrow each instruction takes 1 nanosecond (a billionth of a second)
- Latency / delay: time required per instruction
 - Picosecond: 10⁻¹² seconds Nanosecond: 10⁻⁹ seconds
 - 1,000 ps = 1 nanosecond
- Relationship: throughput = # instructions / latency
 - Example: 1 / 320ps * (1000ps/ns) = 0.003125 * 1000 ≈ 3.1 GIPS

- Current CPU design is serial
 - One instruction executes at a time
 - Only way to improve is to run faster!
 - Limited by speed of light / electricity
- One approach: make it smaller
 - Shorter circuit = faster circuit
 - Limited by manufacturing technology







- Idea: pipelined design
 - Multiple instructions execute simultaneously ("instruction-level parallelism")
 - Similar to cafeteria line or car wash
 - Split logic into stages and connect stages with clocked registers
 - System design tradeoff: throughput vs. latency



- Idea: pipelined design
 - Multiple instructions execute simultaneously ("instruction-level parallelism")
 - Similar to cafeteria line or car wash
 - Split logic into stages and connect stages with clocked registers
 - System design tradeoff: throughput vs. latency



- Limitation: non-uniform partitioning
 - Logic segments may have significantly different lengths



(a) Hardware: Three-stage pipeline, nonuniform stage delays



(b) Pipeline diagram

- Limitation: dependencies
 - The effect of one instruction depends on the result of another
 - Both data and control dependencies
 - Sometimes referred to as hazards

Data dependency:	Control dependency:
irmovq \$8, %rax	loop:
addq %rax, %rbx	subq %rdx, %rbx
mrmovq 0x300(%rbx), %rdx	jne loop
	irmovq \$10, %rdx

- Limitation: dependencies
 - The effect of one instruction depends on the result of another
 - Both data and control dependencies
 - Sometimes referred to as hazards



- Approaches to avoiding hazards
 - Stalling: "hold back" an instruction temporarily
 - Data forwarding: allow latter stages to feed into earlier stages, bypassing memory or registers
 - Hybrid: stall and forward
 - Branch prediction: guess address of next instruction
 - Halt execution (or throw an exception)
 - For more info, read CS:APP section 4.5

Y86 pipelining

- It's complicated!
 - Split up the stages and add more clocked registers for intermediate results



Summary

- We've now learned how a CPU is constructed
 - Transistors \rightarrow logic gates \rightarrow circuits \rightarrow CPU
 - Pipelining provides instruction-level parallelism
- This is not a CPU architecture class
 - We won't be closely studying the specifics of SEQ
 - If you're interested, the details are in section 4.3
 - Same for PIPE (the pipelined version), in section 4.5
 - If you're REALLY interested, lobby for CS 456

CS 456: Architecture

- Course objectives:
 - Describe the construction of a pipelined CPU from low-level components
 - Describe hardware techniques for parallelism at various levels
 - Summarize storage and I/O interfacing techniques
 - Apply address decoding and memory hierarchy strategies
 - Evaluate the performance impact of cache designs
 - Implement custom hardware designs in an FPGA
 - Justify the use of hardware-based optimization that fails occasionally
 - Develop a sense for the challenges of hardware debugging

Lessons learned

- Computers are not human; they're complex machines
 - Machines require extremely precise inputs
 - Machine output can be difficult to interpret
- Abstraction helps to manage complexity
 - Use simpler components to build more complex ones
- System design involves tradeoffs
 - Simpler ISA vs. ease of coding
 - Throughput vs. latency
- The details matter (A LOT!)
 - There are many ways to fail
 - Skill and dedication are required to succeed

Y86 semantics

- Semantics: the study of meaning
 - What does an instruction "mean"?
 - For us, this is the effect that it has on the machine
 - We should specify these semantics very formally
 - This will help us think correctly about P4

Stage	HALT	NOP	CMOV	IRMOVQ
Fch	$icode \leftarrow M_1[PC]$	$icode \leftarrow M_1[PC]$	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$	$\texttt{icode:ifun} \ \leftarrow \ \texttt{M}_1[\texttt{PC}]$
			$rA:rB \leftarrow M_1[PC+1]$	$\texttt{rA:rB} \leftarrow \texttt{M}_1[\texttt{PC+1}]$
				$valC \leftarrow M_8[PC+2]$
	$valP \leftarrow PC + 1$	$\texttt{valP} \ \leftarrow \ \texttt{PC} \ \texttt{+} \ \texttt{1}$	$\texttt{valP} \leftarrow \texttt{PC} + \texttt{2}$	$\texttt{valP} \ \leftarrow \ \texttt{PC} \ \texttt{+} \ \texttt{10}$
Dec			$valA \leftarrow R[rA]$	
Exe	cpu.stat = HLT		$valE \leftarrow valA$	$valE \leftarrow valC$
			$\texttt{Cnd} \ \leftarrow \ \texttt{Cond}(\texttt{CC,ifun})$	
Mem				
WB			Cnd ? $R[rB] \leftarrow valE$	$R[rB] \leftarrow valE$
PC	$PC \leftarrow 0$	$\texttt{PC} \leftarrow \texttt{valP}$	$\texttt{PC} \leftarrow \texttt{valP}$	$\texttt{PC} \leftarrow \texttt{valP}$

Aside: syntax notes

- R[RSP] = the value of %rsp
- R[rA] = the value of register with id rA
- M₁[PC] = the value of one byte in memory at address PC
- $M_8[PC+2]$ = the value of eight bytes in memory at address PC+2
- $rA:rB = M_1[PC+1]$ means read the byte at address PC+1
 - Split it into high- and low-order 4-bits for rA and rB
- Cond(CC, ifun) returns 0 or 1 based on CC and ifun
 - Determines whether the given CMOV/JUMP should happen
- Convention: write addresses using hex padded to three chars
- Convention: write integer literals using decimal w/ no padding

Example: IRMOVQ

0x016: 30f480000000000000000000 | irmovq \$128,%rsp

Stage	IRMOVQ	
Fch	$icode:ifun \leftarrow M_1[PC]$	icode: if un \leftarrow M ₁ [0x016] = 3:0
	$rA:rB \leftarrow M_1[PC+1]$	$rA:rB \leftarrow M_1[0x017] = f:4$
	valC \leftarrow M ₈ [PC+2]	valC \leftarrow M ₈ [0x018] = 128
	valP \leftarrow PC + 10	valP \leftarrow 0x016 + 10 = 0x020
Dec		
Exe	$\texttt{valE} \leftarrow \texttt{valC}$	valE \leftarrow 128
Mom		
WB	$R[rB] \leftarrow valE$	$R[\%rsp] \leftarrow valE = 128$
PC -	$PC \leftarrow valP$	$PC \leftarrow valP = 0x020$

This instruction sets %rsp to 128 and increments the PC by 10

Example: POPQ

0x02c: b00f		popq %rax	
	R[%rsp] = 120	$M_8[120] = 9$	
Stage Fch	$\begin{array}{l} \texttt{POPQ} \\ \texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}] \\ \texttt{rA:rB} \leftarrow \texttt{M}_1[\texttt{PC+1}] \end{array}$	icode:ifun $\leftarrow M_1[0x02c] = b:0$ rA:rB $\leftarrow M_1[0x02d] = 0:f$ valP $\leftarrow 0x02c + 2 = 0x02e$	
Dec	$\begin{array}{rl} \texttt{valP} &\leftarrow \texttt{PC} + 2 \\ \texttt{valA} &\leftarrow \texttt{R[RSP]} \\ \texttt{valB} &\leftarrow \texttt{R[RSP]} \end{array}$	valA \leftarrow R[%rsp] = 120 valB \leftarrow R[%rsp] = 120 valE \leftarrow 120 + 8 = 128	
Exe Mem WB	$valE \leftarrow valB + 8$ $valM \leftarrow M_8[valA]$ $R[RSP] \leftarrow valE$	valM \leftarrow M ₈ [120] = 9 R[%rsp] \leftarrow 128 R[%rax] \leftarrow 9	
PC	$R[rA] \leftarrow valM$ PC $\leftarrow valP$	$PC \leftarrow 0x02e$	

This instruction sets %rax to 9, sets %rsp to 128, and increments the PC by 2

Example: CALL

0x037: 804100000000000000000 | call proc

R[%rsp] = 128

Stage	CALL		
Fch	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$	icode:ifun \leftarrow M ₁ [0x037]=8:0	
	valC \leftarrow M ₈ [PC+1] valP \leftarrow PC + 9	valC \leftarrow M ₈ [0x038] = 0x041 valP \leftarrow 0x037 + 9 = 0x040	
Dec		valB \leftarrow R[%rsp] = 128	
Exe	vale \leftarrow vale - 8	valE ← 128 - 8 = 120	
Mem	$M_8[valE] \leftarrow valP$	$M_8[120] \leftarrow 0x040$	
WB	$R[RSP] \leftarrow valE$	R[%rsp] ← 120	
P C	PC ← valC	$PC \leftarrow 0x041$	

This instruction sets %rsp to 120, stores the return address 0x040 at [%rsp], and sets the PC to 0x041

Y86 semantics

Stage	HALT	NOP	CMOV	IRMOVQ
Fch	icode $\leftarrow M_1[PC]$	icode $\leftarrow M_1[PC]$	$icode:ifun \leftarrow M_1[PC]$	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$
			$rA:rB \leftarrow M_1[PC+1]$	$rA:rB \leftarrow M_1[PC+1]$
				$valC \leftarrow M_8[PC+2]$
	$valP \leftarrow PC + 1$	$valP \leftarrow PC + 1$	$valP \leftarrow PC + 2$	$valP \leftarrow PC + 10$
Dec			$valA \leftarrow R[rA]$	
Exe	cpu.stat = HLT		$valE \leftarrow valA$	$valE \leftarrow valC$
			$Cnd \leftarrow Cond(CC, ifun)$	
Mem				
WB			Cnd ? $R[rB] \leftarrow valE$	$R[rB] \leftarrow valE$
PC	$PC \leftarrow 0$	$PC \leftarrow valP$	$PC \leftarrow valP$	$PC \leftarrow valP$
Stage	RMMOVQ	MRMOVQ	OPq	JUMP
Fch	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$	$icode:ifun \leftarrow M_1[PC]$	$icode:ifun \leftarrow M_1[PC]$	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$
	$rA:rB \leftarrow M_1[PC+1]$	$rA:rB \leftarrow M_1[PC+1]$	$rA:rB \leftarrow M_1[PC+1]$	
	$valC \leftarrow M_8[PC+2]$	$valC \leftarrow M_8[PC+2]$		$valC \leftarrow M_8[PC+1]$
	$valP \leftarrow PC + 10$	$valP \leftarrow PC + 10$	$valP \leftarrow PC + 2$	$valP \leftarrow PC + 9$
Dec	$valA \leftarrow R[rA]$		$valA \leftarrow R[rA]$	
	$\texttt{valB} \leftarrow \texttt{R[rB]}$	$valB \leftarrow R[rB]$	$valB \leftarrow R[rB]$	
Exe	$valE \leftarrow valB + valC$	$valE \leftarrow valB + valC$	$valE \leftarrow valB OP valA$	$Cnd \leftarrow Cond(CC, ifun)$
Mem	$M_8[valE] \leftarrow valA$	$valM \leftarrow M_8[valE]$		
WB		$R[rA] \leftarrow valM$	$R[rB] \leftarrow valE$	
PC	$PC \leftarrow valP$	$PC \leftarrow valP$	$PC \leftarrow valP$	PC ← Cnd?valC:valP
Stage	CALL	RET	PUSHQ	POPQ
Fch	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$	$icode:ifun \leftarrow M_1[PC]$	$\texttt{icode:ifun} \leftarrow \texttt{M}_1[\texttt{PC}]$
			$rA:rB \leftarrow M_1[PC+1]$	$rA:rB \leftarrow M_1[PC+1]$
	$valC \leftarrow M_8[PC+1]$			
	$valP \leftarrow PC + 9$	$valP \leftarrow PC + 1$	$valP \leftarrow PC + 2$	$valP \leftarrow PC + 2$
Dec		$valA \leftarrow R[RSP]$	$valA \leftarrow R[rA]$	$valA \leftarrow R[RSP]$
	$\texttt{valB} \leftarrow \texttt{R[RSP]}$	$\texttt{valB} \leftarrow \texttt{R[RSP]}$	$valB \leftarrow R[RSP]$	$\texttt{valB} \leftarrow \texttt{R[RSP]}$
Exe	valE \leftarrow valB - 8	$valE \leftarrow valB + 8$	$valE \leftarrow valB - 8$	$valE \leftarrow valB + 8$
Mem	$M_8[valE] \leftarrow valP$	$valM \leftarrow M_8[valA]$	$M_8[valE] \leftarrow valA$	$valM \leftarrow M_8[valA]$
WB	$R[RSP] \leftarrow valE$	$R[RSP] \leftarrow valE$	$R[RSP] \leftarrow valE$	$R[RSP] \leftarrow valE$
				$R[rA] \leftarrow valM$
PC	$PC \leftarrow valC$	$PC \leftarrow valM$	$PC \leftarrow valP$	$PC \leftarrow valP$

Y86 CPU (P4)

von Neumann architecture

- 1) Fetch $\leftarrow P3!$
 - Splits instruction at PC into pieces
 - Save info in y86_inst_t struct
- 2) Decode (register file)
 - Reads registers
 - P4: Sets valA
- 3) Execute (ALU)
 - Arithmetic/logic operation, effective address calculation, or stack pointer increment/decrement
 - P4: Sets valE and Cnd
- 4) Memory (RAM)
 - Reads/writes memory
- 5) Write back (register file)
 - Sets registers
- 6) PC update
 - Sets new PC

