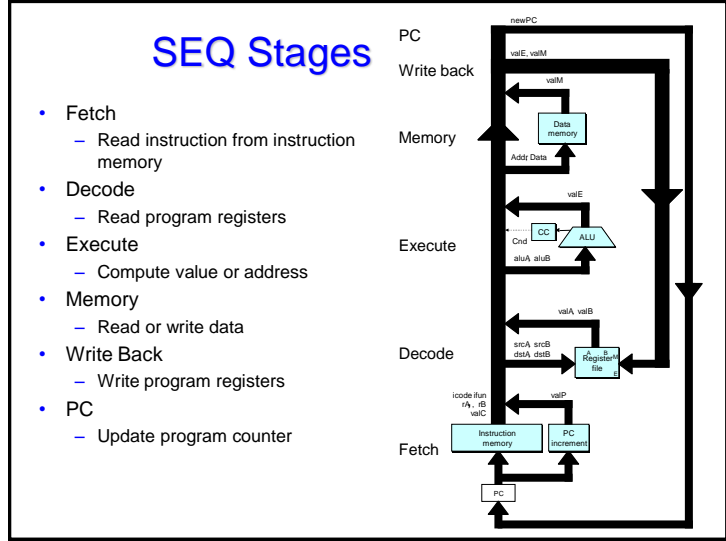
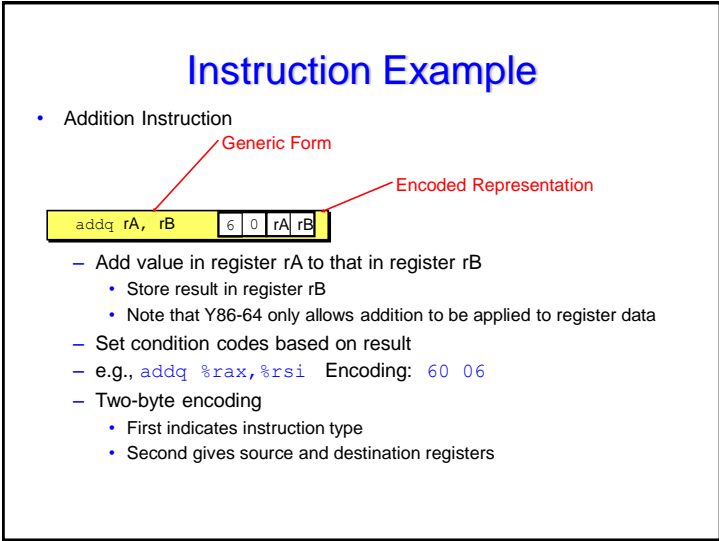


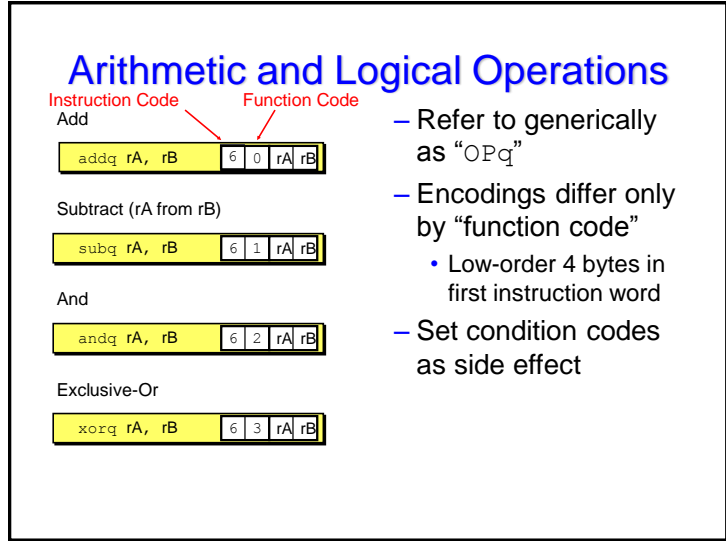
61



62



63



64

Arithmetic and Logical Operations

Instruction Code **Function Code**

Add

`addq rA, rB`

6	0	rA	rB
---	---	----	----

- Refer to generically as "OPq"
- Encodings differ only by "function code"
 - Low-order 4 bits in first instruction word
- Set condition codes as side effect

Subtract (rA from rB)

`subq rA, rB`

6	1	rA	rB
---	---	----	----

And

`andq rA, rB`

6	2	rA	rB
---	---	----	----

Exclusive-Or

`xorq rA, rB`

6	3	rA	rB
---	---	----	----

65

65

Executing an ADD instruction

- How does the processor execute `addq %rax,%rsi`
- The binary encoding is `60 06`

Instruction Code **Function Code**

Add

`addq rA, rB`

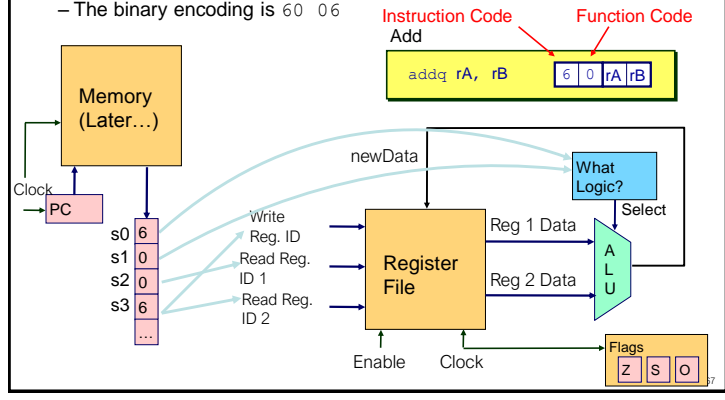
6	0	rA	rB
---	---	----	----

66

66

Executing an ADD instruction

- How does the processor execute `addq %rax,%rsi`
- The binary encoding is `60 06`



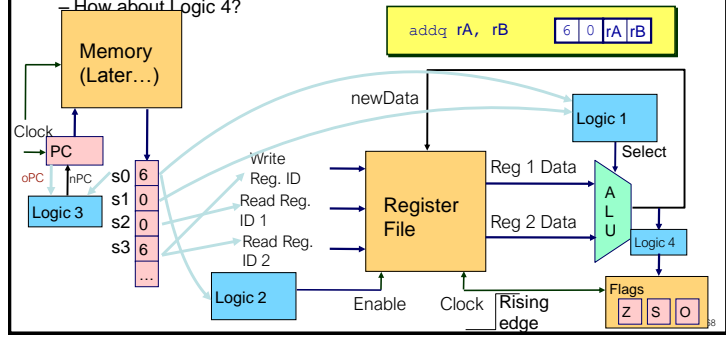
67

67

Executing an ADD instruction

- Logic 1: if (s0 == 6) select = s1;
- Logic 2: if (s0 == 6) Enable = 1; else Enable = 0;
- Logic 3: if (s0 == 6) nPC = oPC + 2;
- How about Logic 4?

How do these logics get implemented?

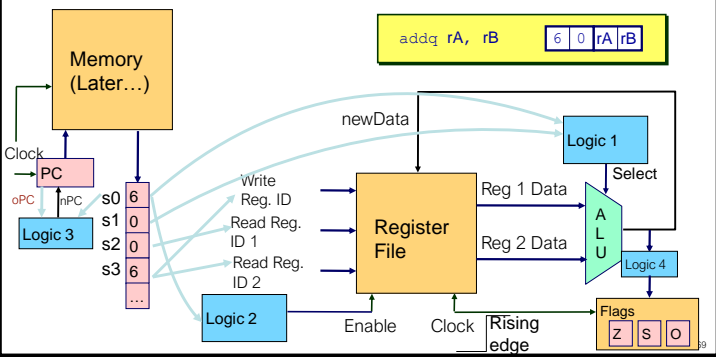


68

68

Executing an ADD instruction

- When the rising edge of the clock arrives, the RF/PC/Flags will be written.
- So the following has to be ready: newData, nPC, which means Logic1, Logic2, Logic3, and Logic4 has to finish.



69

Jump Instructions

Jump (Conditionally)



- Refer to generically as “jxx”
- Encodings differ only by “function code” fn
- Based on values of condition codes
- Same as x86-64 counterparts
- Encode full destination address
 - Unlike PC-relative addressing seen in x86-64

70

Jump Instructions

Jump Unconditionally



Jump When Less or Equal



Jump When Less



Jump When Equal



Jump When Not Equal



Jump When Greater or Equal



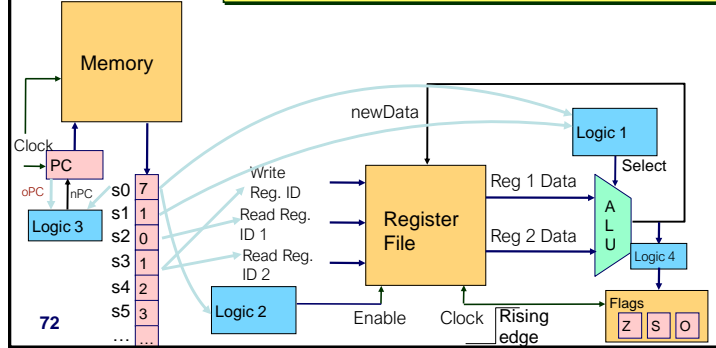
Jump When Greater



71

Executing a JLE instruction

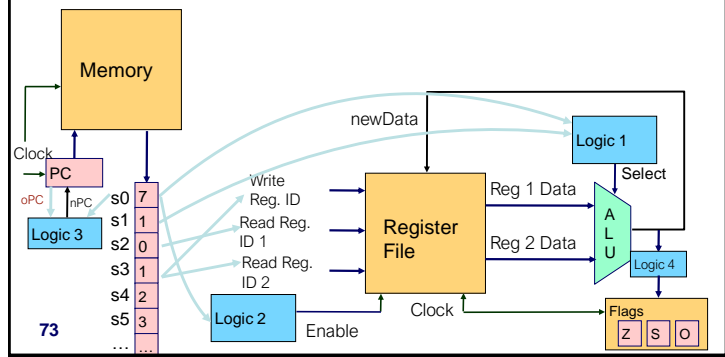
- Let's say the binary encoding for jle .L0 is 71 0123000000000000
- What are the logics now?



72

Executing a JLE instruction

- Logic 1: if (s0 == 6) select = s1;
- Logic 2: if (s0 == 6) Enable = 1; else Enable = 0;

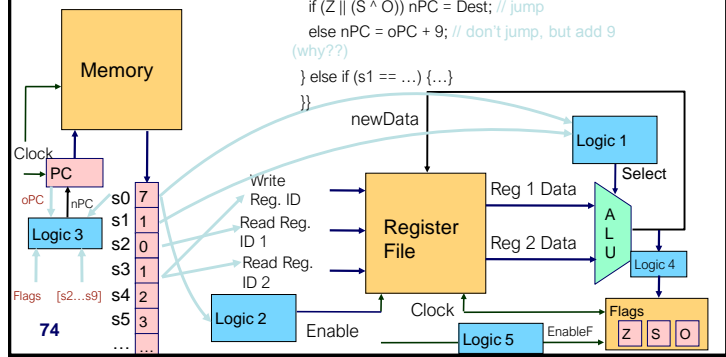


73

Executing a JLE instruction

```

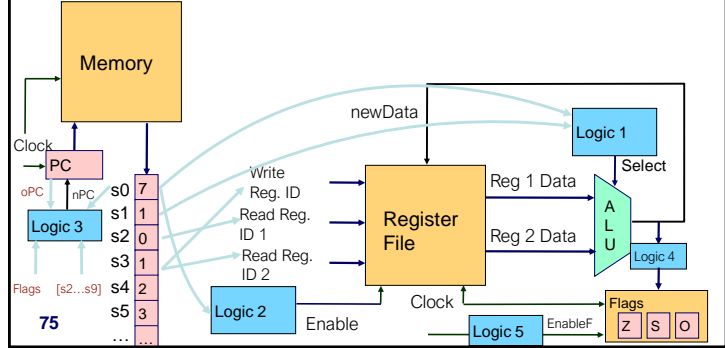
- Logic 3??
if (s0 == 6) nPC = oPC + 2;
else if (s0 == 7) {
  if (s1 == 1) { // JLE
    if (Z || (S ^ O)) nPC = Dest; // jump
    else nPC = oPC + 9; // don't jump, but add 9
    (why??)
  } else if (s1 == ...) {...}
}
    
```



74

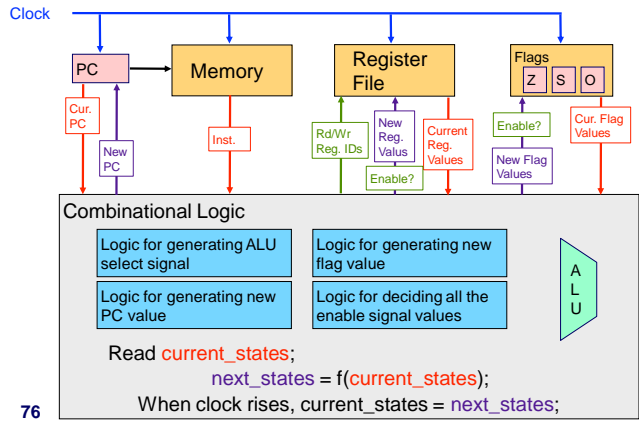
Executing a JLE instruction

- Logic 4? Does JLE write flags?
- Need another piece of logic.
- Logic 5: if (s0 == 7) EnableF = 0; else if (s0 == 6) EnableF = 1;



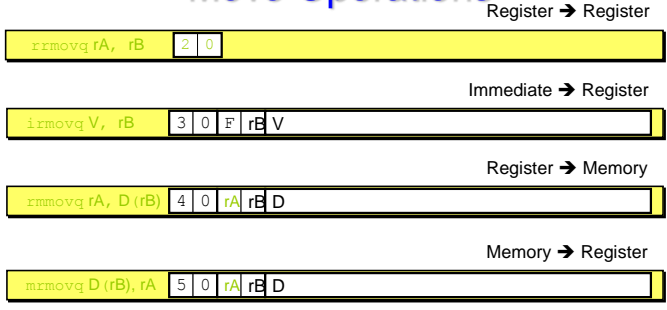
75

Microarchitecture (So far)



76

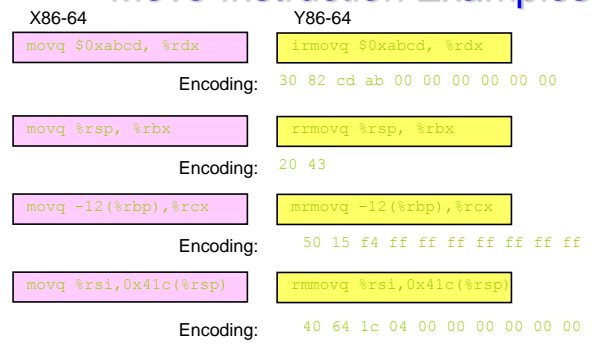
Move Operations



- Like the x86-64 `movq` instruction
- Simpler format for memory addresses
- Give different names to keep them distinct

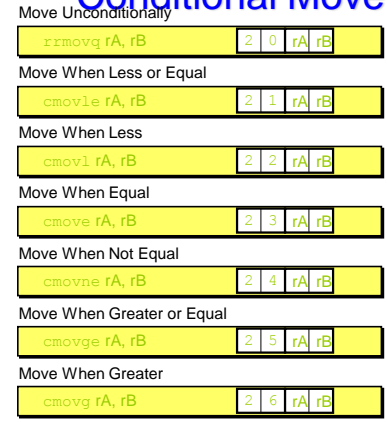
77

Move Instruction Examples



78

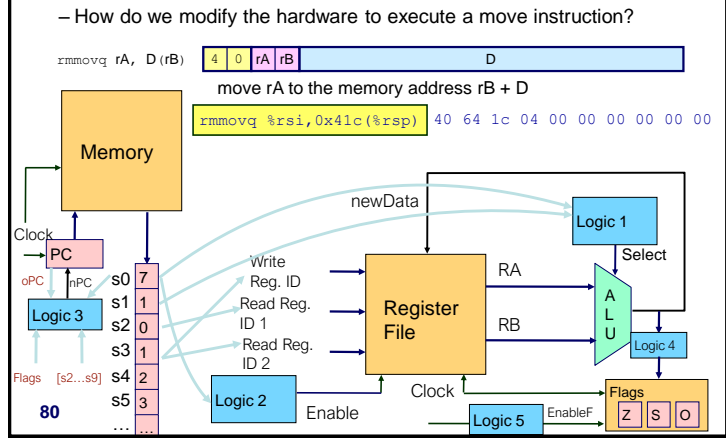
Conditional Move Instructions



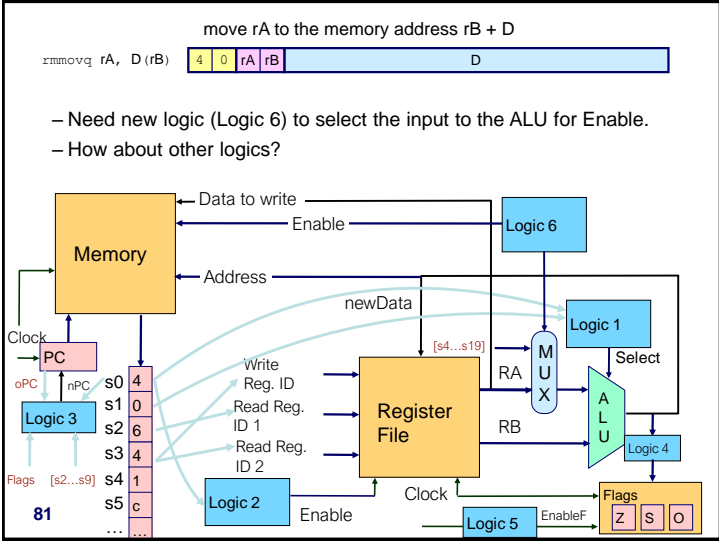
- Refer to generically as “`cmovXX`”
- Encodings differ only by “function code”
- Based on values of condition codes
- Variants of `rmmovq` instruction
- (Conditionally) copy value from source to destination register

79

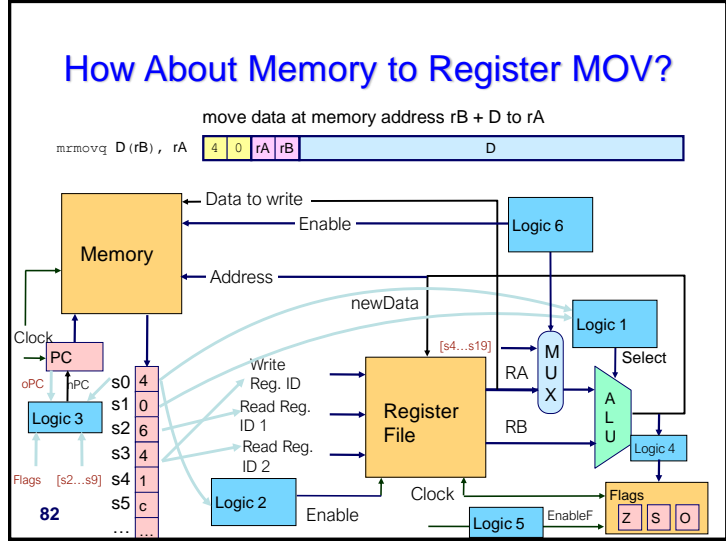
Executing a MOV instruction



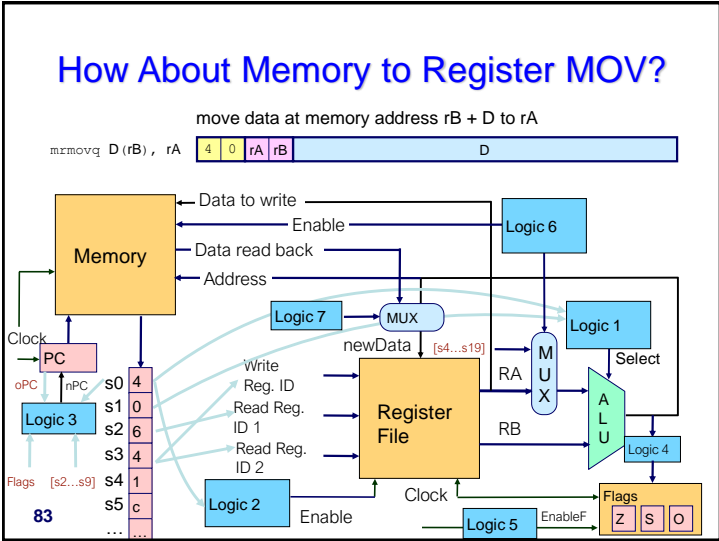
80



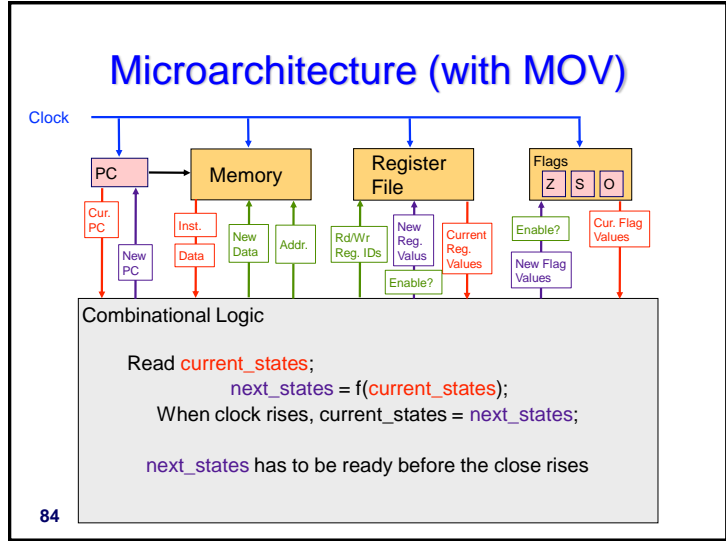
81



82



83



84


Microarchitecture Overview

Think of it as a state machine

Every cycle, one instruction gets executed. At the end of the cycle, architecture states get modified.


States (All updated as clock rises)

- PC register
- Cond. Code register
- Data memory
- Register file




85

85




Cycle 1:	0x000: irmovq \$0x100,%rbx # %rbx <-- 0x100
Cycle 2:	0x00a: irmovq \$0x200,%rdx # %rdx <-- 0x200
Cycle 3:	0x014: addq %rdx,%rbx # %rbx <-- 0x300 CC <-- 000
Cycle 4:	0x016: je dest # Not taken
Cycle 5:	0x01f: rmmovq %rbx,0(%rdx) # M[0x200] <-- 0x300




- state set according to second irmovq instruction
- combinational logic starting to react to state changes

86




Cycle 1:	0x000: irmovq \$0x100,%rbx # %rbx <-- 0x100
Cycle 2:	0x00a: irmovq \$0x200,%rdx # %rdx <-- 0x200
Cycle 3:	0x014: addq %rdx,%rbx # %rbx <-- 0x300 CC <-- 000
Cycle 4:	0x016: je dest # Not taken
Cycle 5:	0x01f: rmmovq %rbx,0(%rdx) # M[0x200] <-- 0x300



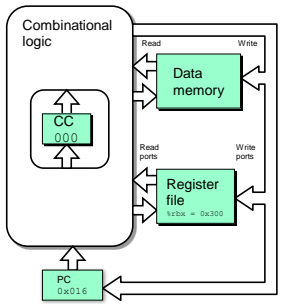
- state set according to second irmovq instruction
- combinational logic generates results for addq instruction

%rbx
<--
0x300

87

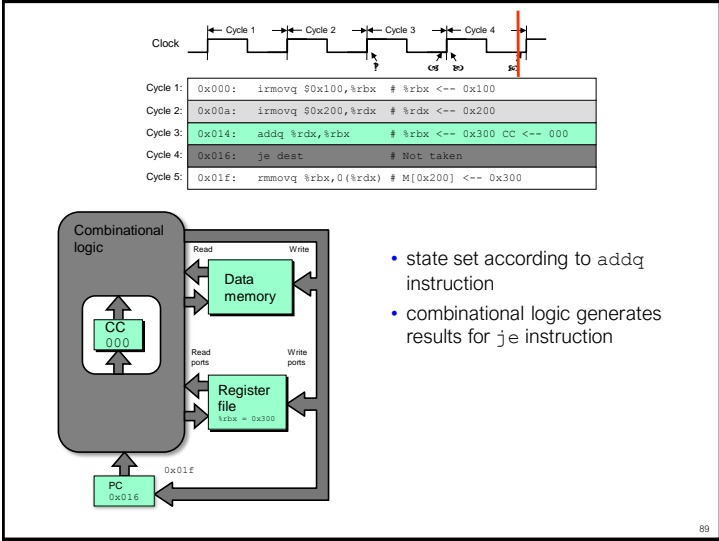


Cycle 1:	0x000: irmovq \$0x100,%rbx # %rbx <-- 0x100
Cycle 2:	0x00a: irmovq \$0x200,%rdx # %rdx <-- 0x200
Cycle 3:	0x014: addq %rdx,%rbx # %rbx <-- 0x300 CC <-- 000
Cycle 4:	0x016: je dest # Not taken
Cycle 5:	0x01f: rmmovq %rbx,0(%rdx) # M[0x200] <-- 0x300



- state set according to addq instruction
- combinational logic starting to react to state changes

88



89

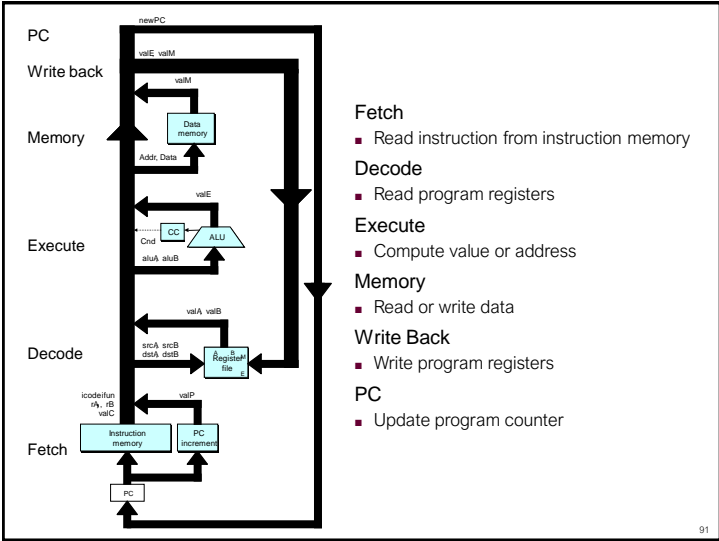
Another Way to Look At the Microarchitecture

Principles:

- Execute each instruction one at a time, one after another
- Express every instruction as series of **simple steps**
- Dedicated hardware structure for completing each step
- Follow same general flow for each instruction type

Fetch: Read instruction from instruction memory
 Decode: Read program registers
 Execute: Compute value or address
 Memory: Read or write data
 Write Back: Write program registers
 PC: Update program counter

90



91

Stage Computation: Arith/Log. Ops

$OPq \ rA, \ rB$ $\epsilon \ fn \ rA \ rB$

Fetch	$icode:ifun \leftarrow M_1[PC]$ $rA:rB \leftarrow M_1[PC+1]$ $valP \leftarrow PC+2$	Read instruction byte Read register byte Compute next PC
Decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$	Read operand A Read operand B
Execute	$valE \leftarrow valB \text{ OP } valA$ Set CC	Perform ALU operation Set condition code register
Memory		
Write back	$R[rB] \leftarrow valE$	Write back result
PC update	$PC \leftarrow valP$	Update PC

92

Stage Computation: `rmmovq`



	<code>rmmovq rA, D(rB)</code>	
Fetch	$icode:ifun \leftarrow M_1[PC]$ $rA:rB \leftarrow M_1[PC+1]$ $valC \leftarrow M_8[PC+2]$ $valP \leftarrow PC+10$	Read instruction byte Read register byte Read displacement D Compute next PC
Decode	$valA \leftarrow R[rA]$ $valB \leftarrow R[rB]$	Read operand A Read operand B
Execute	$valE \leftarrow valB + valC$	Compute effective address
Memory	$M_8[valE] \leftarrow valA$	Write value to memory
Write back		
PC update	$PC \leftarrow valP$	Update PC

93

Stage Computation: Jumps

	<code>jXX Dest</code>	
Fetch	$icode:ifun \leftarrow M_1[PC]$ $valC \leftarrow M_8[PC+1]$ $valP \leftarrow PC+9$	Read instruction byte Read destination address Fall through address
Decode		
Execute	$Cnd \leftarrow Cond(CC,ifun)$	Take branch?
Memory		
Write back		
PC update	$PC \leftarrow Cnd ? valC : valP$	Update PC

- Compute both addresses
- Choose based on setting of condition codes and branch condition

94

Processor Microarchitecture

- Sequential, single-cycle microarchitecture implementation
 - Basic idea
 - Hardware implementation
- Pipelined microarchitecture implementation
 - Basic Principles
 - Difficulties: Control Dependency
 - Difficulties: Data Dependency

95

Possible Implementation Strategies

- Single-cycle control
- Multi-cycle control
- Pipelined (in-order execution)
- Superscalar and out-of-order execution


10/13/2020

96

96

SEQ Hardware Structure

- State
 - Program counter register (PC)
 - Condition code register (CC)
 - Register File
 - Memories
 - Access same memory space
 - Data: for reading/writing program data
 - Instruction: for reading instructions
- Instruction Flow
 - Read instruction at address specified by PC
 - Process through stages
 - Update program counter



10/13/2020

97

Single-Cycle Control

- Cannot use any hardware unit twice in one cycle
 - Multiple datapath elements (e.g., ALU and PC incrementer) if needed for multiple purposes
 - Multiple reads or writes to memory or register file require multiplexing
- Clock cycle must accommodate instruction with the longest latency

10/13/2020

98

98

Multi-cycle Control

- Split single instruction into multiple pieces
- Execute individual pieces using hardwired finite state machine (FSM) or microprogramming
- Can finish simple instructions in fewer cycles
- Can run clock faster
- Still execute a single instruction at a time

10/13/2020

99

99

Pipelined Datapath

- Multiple instructions overlapped in execution
- Improve throughput, not individual instruction execution time
- Exploit parallelism among instructions in a sequential stream
- Balance length of each stage. Ideally -
 - $\text{Time between instrs}_{\text{pipelined}} = \text{Time between instrs}_{\text{non-pipelined}} / \text{Number of pipe stages}$

10/13/2020

100

100

Real-World Pipelines: Car Washes




Idea

- Divide process into independent stages
- Move objects through stages in sequence
- At any given times, multiple objects being processed

101

101

Computational Example




System

- Computation requires total of 300 picoseconds
- Additional 20 picoseconds to save result in register
- Must have clock cycle time of at least 320 ps

102

102

3-Stage Pipelined Version



System


- Divide combinational logic into 3 blocks of 100 ps each
- Can begin new operation as soon as previous one passes through stage A.
 - Begin new operation every 120 ps
- Overall latency increases
 - 360 ps from start to finish

103


103

Pipeline Diagrams

Unpipelined



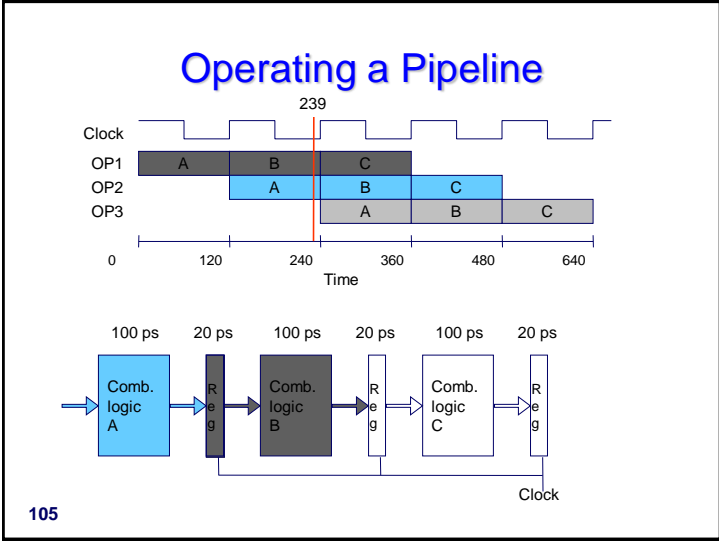
- Cannot start new operation until previous one completes




- Up to 3 operations in process simultaneously

104


104



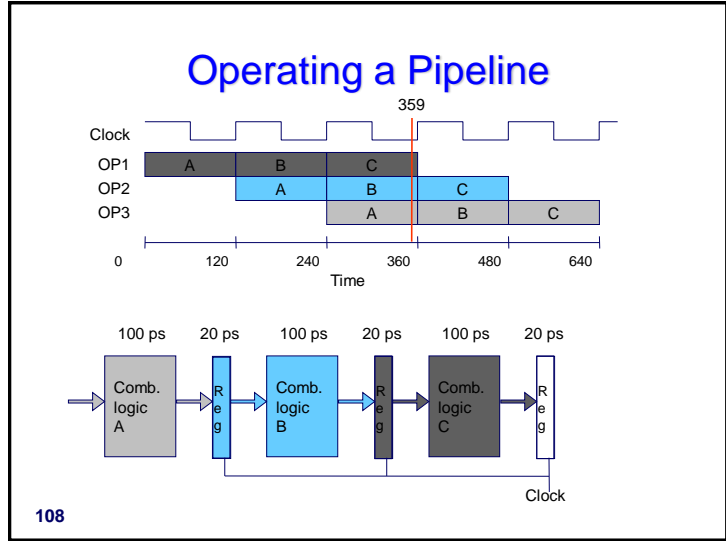
105



106



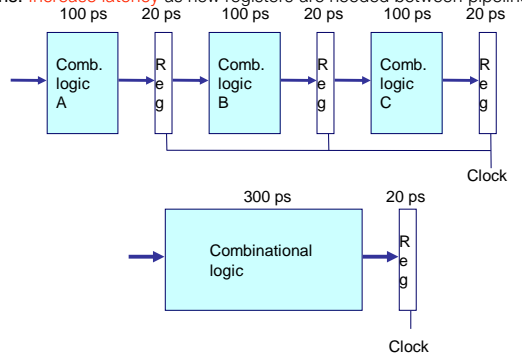
107



108

Pipeline Trade-offs

- Pros: **Increase throughput**. Can process more instructions in a given time span.
- Cons: **Increase latency** as new registers are needed between pipeline stages.

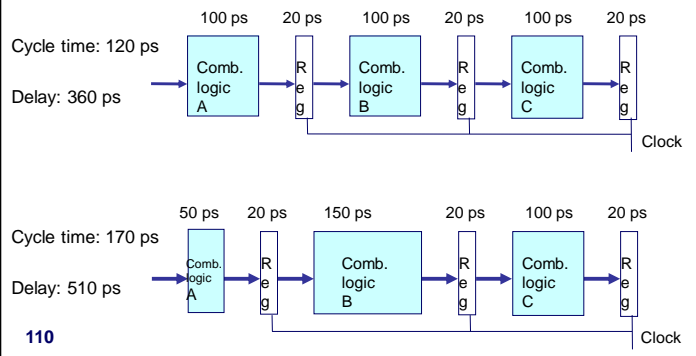


109

109

Unbalanced Pipeline

- A pipeline's delay is limited by the slowest stage. This limits the cycle time and the throughput





110

110

Unbalanced Pipeline

- A pipeline's delay is limited by the slowest stage. This limits the cycle time and the throughput




111

111

Mitigating Unbalanced Pipeline

- Solution 1: Further pipeline the slow stages
 - Not always possible. What to do if we can't further pipeline a stage?
- Solution 2: Use multiple copies of the slow component




What logic do you need there?
Hint: it needs to control the clock signals of the two registers and the select signal of the MUX.

112

112

Mitigating Unbalanced Pipeline

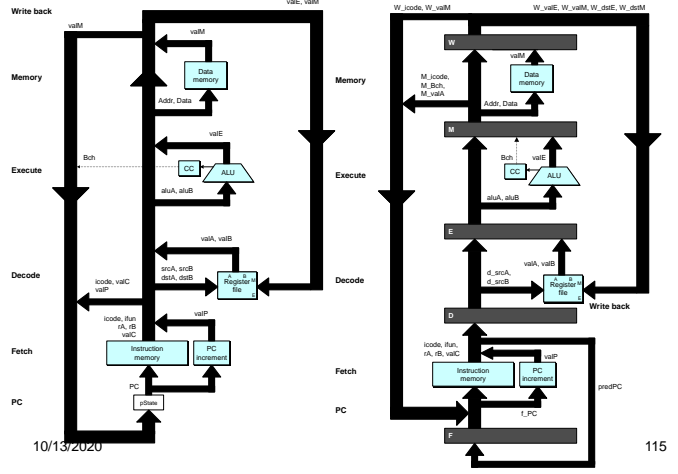
- Data sent to copy 1 in odd cycles and to copy 2 in even cycles.
- This is called 2-way interleaving. Effectively the same as pipelining Comb. logic B into two sub-stages.
- The cycle time is reduced to 70 ps (as opposed to 120 ps) at the cost of extra hardware.



113

113

Adding Pipeline Registers



10/13/2020

115

115