

CMU 18-447 INTRODUCTION TO COMPUTER ARCHITECTURE, SPRING 2013

HW 2 SOLUTIONS: ISA TRADEOFFS

Instructor: Prof. Onur Mutlu
 TAs: Justin Meza, Yoongu Kim, Jason Lin

1 LC-3b Microcode [40 points]

Instruction	state	IRD	Cond	J	LD.MAR	LD.MDR	LD.IR	LD.BEN	LD.REG	LD.CC	LD.PC	GatePC	GateMDR	GateALU	GateMARMUX	GateSHF	PCMUX	DRMUX	SRIMUX	ADDRIMUX	ADDR2MUX	MARMUX	ALUK	MIO.EN	R.W	DATA.SIZE	LSHF1	
size		1	2	6													2	1	1	1	2	1	2	1	1	1	1	
BR	0	0	2	10010	0	0	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	0	0	X	X	
ADD	1	0	0	10010	0	0	0	0	0	1	1	0	0	0	1	0	0	X	0	1	X	X	X	0	0	X	X	
LDB	2	0	0	11101	1	0	0	0	0	0	0	0	0	0	0	1	0	X	X	1	1	1	1	1	X	0	X	
STB	3	0	0	11000	1	0	0	0	0	0	0	0	0	0	0	1	0	X	X	1	1	1	1	1	X	0	X	
JSR	4	0	3	10100	0	0	0	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	0	X	X	
AND	5	0	0	10010	0	0	0	0	0	1	1	0	0	0	1	0	0	X	0	1	X	X	X	1	0	X	X	
LDW	6	0	0	11001	1	0	0	0	0	0	0	0	0	0	0	1	0	X	X	1	1	1	1	1	X	0	X	
STW	7	0	0	10111	1	0	0	0	0	0	0	0	0	0	0	1	0	X	X	1	1	1	1	1	X	0	X	
RTI	8	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
XOR	9	0	0	10010	0	0	0	0	1	1	0	0	0	1	0	0	0	X	0	1	X	X	X	2	0	X	X	
RES1	10	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
RES2	11	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
JMP	12	0	0	10010	0	0	0	0	0	0	1	0	0	0	0	0	2	X	1	1	0	X	X	0	X	X	X	
JMP	12	0	0	10010	0	0	0	0	0	0	1	0	0	1	0	0	1	X	1	X	X	X	3	0	X	X	X	
SHF	13	0	0	10010	0	0	0	0	1	1	0	0	0	0	0	1	X	0	1	X	X	X	X	0	X	X	X	
LEA	14	0	0	10010	0	0	0	0	1	0	0	0	0	0	1	0	X	X	X	0	2	1	X	0	X	X	1	
TRAP	15	0	0	11100	1	0	0	0	0	0	0	0	0	0	1	0	X	X	X	X	X	0	X	0	X	X	X	
STW	16	0	1	10000	0	0	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	1	1	1	X	
STB	17	0	1	10001	0	0	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	1	1	0	X	
ALL	18	0	0	100001	1	0	0	0	0	0	1	1	0	0	0	0	X	X	X	X	X	X	X	0	X	X	X	
ALL	18	0	0	100001	1	0	0	0	0	0	1	0	0	0	1	0	X	X	X	0	0	1	X	0	X	X	X	
ALL	19	0	0	100001	1	0	0	0	0	0	1	1	0	0	0	0	X	X	X	X	X	X	X	0	X	X	X	
ALL	19	0	0	100001	1	0	0	0	0	0	1	0	0	0	1	0	X	X	X	0	0	1	X	0	X	X	X	
JSR	20	0	0	10010	0	0	0	0	1	0	1	1	0	0	0	0	2	1	1	1	1	0	X	X	0	X	X	
JSR	20	0	0	10010	0	0	0	0	1	0	1	1	0	0	0	0	2	1	1	1	1	1	X	X	0	X	X	
JSR	21	0	0	10010	0	0	0	0	1	0	1	1	0	0	0	0	2	1	X	0	3	X	X	0	X	X	1	
BR	22	0	0	10010	0	0	0	0	0	0	1	0	0	0	0	0	2	X	X	0	2	X	X	0	X	X	1	
STW	23	0	0	10000	0	1	0	0	0	0	0	0	0	1	0	0	X	X	0	X	X	X	3	0	X	X	X	
STW	23	0	0	10000	0	1	0	0	0	0	0	0	0	0	1	0	X	X	0	1	0	1	X	0	X	X	X	
STB	24	0	0	10001	0	1	0	0	0	0	0	0	0	1	0	0	X	X	0	X	X	X	3	0	X	X	X	
STB	24	0	0	10001	0	1	0	0	0	0	0	0	0	0	1	0	X	X	0	1	0	1	X	0	X	X	X	
LDW	25	0	1	11001	0	1	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	1	0	X	X	
FREE	26	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	
LDW	27	0	0	10010	0	0	0	0	1	1	0	0	1	0	0	0	X	0	X	X	X	X	X	0	X	1	X	
TRAP	28	0	1	11100	0	1	0	0	1	0	0	1	0	0	0	0	X	1	X	X	X	X	X	1	0	X	X	
LDB	29	0	1	11101	0	1	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	1	0	X	X	
TRAP	30	0	0	10010	0	0	0	0	0	0	1	0	1	0	0	0	1	X	X	X	X	X	X	0	X	1	X	
LDB	31	0	0	10010	0	0	0	0	1	1	0	0	1	0	0	0	X	0	X	X	X	X	X	0	X	0	X	
ALL	32	1	X	X	0	0	0	1	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	0	X	X	X	
ALL	33	0	1	100001	0	1	0	0	0	0	0	0	0	0	0	0	X	X	X	X	X	X	X	1	0	X	X	
FREE	34	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X	X
ALL	35	0	0	100000	0	0	1	0	0	0	0	0	1	0	0	0	X	X	X	X	X	X	X	0	X	1	X	X

2 Data Reuse in Different ISAs [0 points]

(a) Assembly code sequences for code segments A and B for machines X and Y

Code Segment A

Machine X

```
ADD A, A, 1
ADD B, B, 2
ADD C, C, 3
ADD D, D, 4
```

Machine Y

```
LD R1, A
ADD R1, R1, 1
ST R1, A
LD R2, B
ADD R2, R2, 2
ST R2, B
LD R3, C
ADD R3, R3, 3
ST R3, C
LD R4, D
ADD R4, R4, 4
ST R4, D
```

Code Segment B

Machine X

```
ADD B, B, A
ADD C, B, A
SUB A, B, C
SUB D, A, C
```

Machine Y

```
LD R1, A
LD R2, B
ADD R3, R1, R2
ADD R4, R3, R1
SUB R5, R3, R4
SUB R6, R5, R4
ST R3, B
ST R4, C
ST R5, A
ST R6, D
```

(b) Total number of bytes transferred

	Code Segment A	Code Segment B
Machine X	60	76
Machine Y	80	64

- (c) (i) No. For code segment A, machine X has lower number of total bytes transferred. For code segment B, machine Y has lower number of total bytes transferred.

(ii) Code segment A has no data reuse. It just adds an immediate value to A, B and C each, and does not reuse the values of A, B or C. A memory-memory machine needs just three instructions to express code segment A. Each instruction fetches a memory location, adds an immediate value to it and stores it back. On the other hand, a load-store machine incurs several instructions to load the values of A, B and C into registers and store them back. This is unnecessary and merely increases the number of instructions (and hence number of instruction bytes transferred) for code segment A, as there is no reuse of data.

Code segment B reuses the values of A and B. In this case, a load-store machine loads A and B into registers first and reuses the values stored in registers. On the other hand, a memory-memory machine does repeated memory accesses to memory to fetch A and B every time they are used. Therefore, a load-store machine reduces the number of data bytes transferred significantly, for code segment B, resulting in a smaller number of total bytes transferred.

3 Addressing Modes [12 points]

- (a) Auto increment
- (b) Scale indexed
- (c) Register indirect
- (d) Memory indirect

4 Addressability [0 points]

- (a) (i) 29 bits
(ii) 26 bits
(iii) 23 bits
(iv) 21 bits
- (b) # load the two words
- ```
lw $8, 0($5)
lw $9, 0($7)
```

```
multiply the offset by 4 to get the shift amount
sll $4, $4, 2
sll $6, $6, 2
```

```
shift the loaded word to obtain the nibble at the desired offset
srlv $8, $8, $4
srlv $9, $9, $6
```

```
zero out the rest of the word to keep the nibble
andi $8, $8, 0xF
andi $9, $9, 0xF
```

```
add $2, $8, $9
```


### 5 Microarchitecture vs. ISA [15 points]

- (a) The ISA level is the interface a machine exposes to the software. The microarchitecture is the actual underlying implementation of the machine. Therefore, the microarchitecture and changes to the microarchitecture are transparent to the compiler/programmer (except in terms of performance), while changes to the ISA affect the compiler/programmer.

The compiler does *not* need to know about the microarchitecture of the machine in order to compile the program correctly.

- (b) (i) ISA
- (ii) Microarchitecture
- (iii) ISA
- (iv) ISA
- (v) Microarchitecture
- (vi) ISA
- (vii) Microarchitecture
- (viii) Microarchitecture

### 6 Single-Cycle Processor Datapath [30 points]



### 7 Pipelining [30 points]

(a) A non-pipelined machine

$$9 + 6 + 6 + 9 + 6 + 9 = 45 \text{ cycles}$$

(b) A pipelined machine with scoreboarding and five adders and five multipliers without data forwarding

| Cycles          | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 | 15 | 16 | 17 | 18 | 19 | 20 | 21 | 22 | 23 | 24 | 25 | 26 | 27 | 28 |
|-----------------|---|---|---|---|---|---|---|---|---|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|----|
| MUL R3, R1, R2  | F | D | E | E | E | E | E | W | W |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |
| ADD R5, R4, R3  | F | - | - | - | - | - | - | - | D | E  | E  | W  | W  |    |    |    |    |    |    |    |    |    |    |    |    |    |    |    |
| ADD R6, R4, R1  |   |   |   |   |   |   |   |   | F | D  | E  | E  | W  | W  |    |    |    |    |    |    |    |    |    |    |    |    |    |    |
| MUL R7, R8, R9  |   |   |   |   |   |   |   |   |   | F  | D  | E  | E  | E  | E  | E  | W  | W  |    |    |    |    |    |    |    |    |    |    |
| ADD R4, R3, R7  |   |   |   |   |   |   |   |   |   | F  | -  | -  | -  | -  | -  | -  | -  | -  | D  | E  | E  | W  | W  |    |    |    |    |    |
| MUL R10, R5, R6 |   |   |   |   |   |   |   |   |   |    |    |    |    |    |    |    |    |    |    | F  | D  | E  | E  | E  | E  | W  | W  |    |

28 cycles (or 26 cycles with internal register file data forwarding)

(c) A pipelined machine with scoreboarding and five adders and five multipliers with data forwarding.

|                 |                                                          |
|-----------------|----------------------------------------------------------|
| Cycles          | 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 |
| MUL R3, R1, R2  | F D E E E E E W W                                        |
| ADD R5, R4, R3  | F - - - - D E E W W                                      |
| ADD R6, R4, R1  | F D E E W W                                              |
| MUL R7, R8, R9  | F D E E E E E W W                                        |
| ADD R4, R3, R7  | F - - - - D E E W W                                      |
| MUL R10, R5, R6 | F D E E E E E W W                                        |

**22 cycles**

(d) A pipelined machine with scoreboarding and one adder and one multiplier without data forwarding

|                 |                                                                               |
|-----------------|-------------------------------------------------------------------------------|
| Cycles          | 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 24 25 26 27 28 29 |
| MUL R3, R1, R2  | F D E E E E E W W                                                             |
| ADD R5, R4, R3  | F - - - - - - D E E W W                                                       |
| ADD R6, R4, R1  | F - D E E W W                                                                 |
| MUL R7, R8, R9  | F - D E E E E E W W                                                           |
| ADD R4, R3, R7  | F D E E W W                                                                   |
| MUL R10, R5, R6 | F D E E E E E W W                                                             |

**29 cycles** (or 27 cycles with internal register file data forwarding)

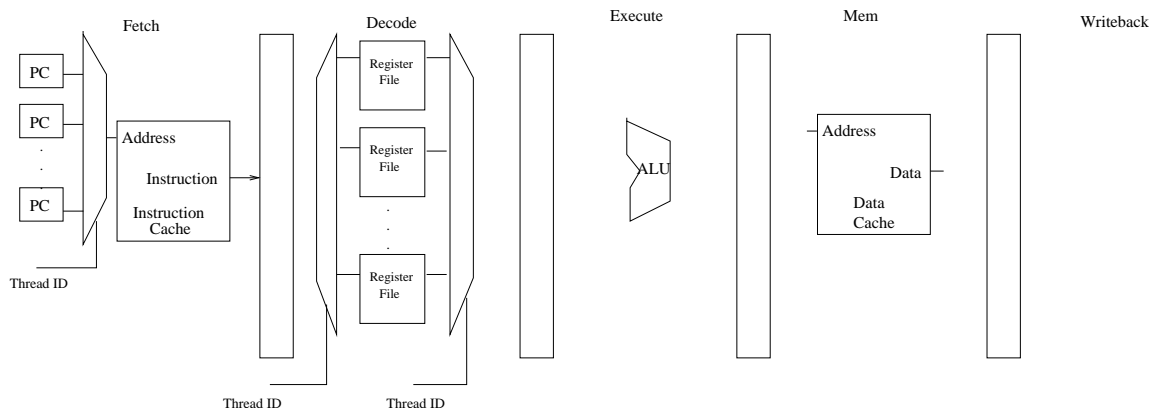
(e) A pipelined machine with scoreboarding and one adder and one multiplier with data forwarding

|                 |                                                             |
|-----------------|-------------------------------------------------------------|
| Cycles          | 1 2 3 4 5 6 7 8 9 10 11 12 13 14 15 16 17 18 19 20 21 22 23 |
| MUL R3, R1, R2  | F D E E E E E W W                                           |
| ADD R5, R4, R3  | F - - - - D E E W W                                         |
| ADD R6, R4, R1  | F - D E E W W                                               |
| MUL R7, R8, R9  | F D E E E E E W W                                           |
| ADD R4, R3, R7  | F - - - - D E E W W                                         |
| MUL R10, R5, R6 | F D E E E E E W W                                           |

**23 cycles**

## 8 Fine Grain Multi-threading [40 points]

(a) The figure shows the solution



(b)

3

Why?

Since branches are resolved in the Execute stage, it is necessary that the Fetch stage does not fetch for a thread until the thread's previous instruction has passed Execute. Hence three threads are needed to cover Fetch, Decode, Execute.

(c)

4

Why?

The designer must ensure that when an instruction is in Writeback, the next instruction in the same thread has not reached Decode yet. Hence, at least 4 threads are needed.

(d) Is the number of threads required to eliminate branch-related stalls in Machine II the same as in Machine I?

YES  NO (Circle one)

If yes, why?

Branches are resolved at the third pipeline stage in both machines, and distance from fetch to branch resolution determines the minimum number of threads to avoid branch stalls.

(e)

3 (if no flow dependence stalls occur)

(f) Does Machine II require the same minimum number of threads as Machine I to avoid the need for flow-dependence stalls?

YES  NO (Circle one)

how many threads are required?

12 (the Decode, Execute 1 – 8, Memory, and Writeback stages must all have instructions from independent threads.)

(g)

12

(h)

The additional FGMT-related logic (MUXes and thread selection logic) could increase the critical path length, which will reduce maximum frequency and thus performance.

## 9 Branch Prediction and Dual Path Execution [35 points]

(a) 5 instructions.

(b) *Note that if you assumed the wrong number of instructions in Part (a), you will only be marked wrong for this in Part (a). You can still get full credit on this and the following parts.*

$$\begin{aligned}\text{Correct path instructions} &= N \\ \text{Incorrect path instructions} &= N(0.2)(1 - A)5 = N(1 - A) \\ \text{Fetched instructions} &= \text{Correct path instructions} + \text{Incorrect path instructions} \\ &= N + N(1 - A) \\ &= \boxed{N(2 - A)}\end{aligned}$$

$$\begin{aligned}\text{Correct path instructions} &= N \\ \text{Incorrect path instructions} &= N(0.2)5 \\ \text{(c) Fetched instructions} &= \text{Correct path instructions} + \text{Incorrect path instructions} \\ &= N + N(1 - 0.8)5 \\ &= \boxed{2N}\end{aligned}$$

*This solution assumes you have enough hardware in the frontend of the machine to fetch concurrently from both paths. If you assumed that both paths are fetched from on alternate cycles, that high-level approach is also OK, although note that you would need additional branch taken and not taken information to solve it completely.*

$$\begin{aligned}
& \text{Correct path instructions} = N \\
& \text{Incorrect path instructions due to...} \\
& \quad \text{lack of confidence} \\
& \quad = N(0.2)(1 - C)5 = N(1 - C) \\
& \quad \text{incorrect high confidence estimate} \\
& \quad = N(0.2)CM5 = NCM \\
(d) \quad \text{Fetched instructions} &= \text{Correct path instructions} \\
& \quad + \text{Incorrect path instructions due to} \\
& \quad \quad \text{lack of confidence} \\
& \quad + \text{Incorrect path instructions due to} \\
& \quad \quad \text{incorrect high confidence estimate} \\
& = N + N(1 - C) + NCM \\
& = \boxed{N[2 + C(M - 1)]}
\end{aligned}$$

*Like above, if you assumed a different execution model for Part (c), you will not be penalized for using it in this part.*

## 10 Mysterious Instruction [40 points]

The instruction finds the position of the most significant set bit in SR1 and places this in DR.

Mathematically, this can be expressed as  $DR = \lfloor \log_2 SR1 \rfloor$

Note that this instruction is semantically the same as FINDFIRST in the VAX ISA.

Specifically, here is what happens in each state:

```

State 10: DR = 15
State 40: TEMP = SR1
 if (SR1 is negative)
 Go to State 18/19 (Fetch)
 else
 Go to State 50
State 50: DR = DR - 1
State 51: Left Shift TEMP
 if (TEMP is negative)
 Go to State 18/19 (Fetch)
 else
 Go to State 50

```