

CS 473

Final Exam

December 13, 2002

The following exam is open book and open notes. You may feel free to use whatever additional reference material you wish, but **no electronic aids** are allowed. Please note the following instructions. There will be a ten point deduction for failure to comply with them:

- start each problem on a new sheet of paper
- write your social security number, but not your name, on each sheet of paper you turn in
- show your work whenever appropriate. There can be no partial credit unless I see how answers were arrived
- be succinct. You may lose points for facts that, while true, are not relevant to the question at hand

You have until 10:00 AM to finish the exam.

1. (20 points)

(a) Multiply the following two IEEE floating point numbers together (using the floating point multiplication algorithm):

$$3fe00000_{16} * 40400000_{16}.$$

First, interpret the numbers:

0 01111111 110000000000000000000000

Sign: 0

Exponent: 01111111

Significand 1.11

0 10000000 100000000000000000000000

Sign: 0

Exponent: 10000000

Significand: 1.1

Now calculate result:

Sign: $S = S_1 \wedge S_2 = 0$

Exponent: $E = E_1 + E_2 - 01111111 = 10000000$

Significand:

$$\begin{array}{r} 1.11 \\ \times 1.1 \\ \hline 1.11 \\ + .111 \\ \hline 10.101 \end{array}$$

Renormalize:

Exponent = $10000000 + 1 = 10000001$,

Significand = 1.0101 .

Combine:

0 100000001010-0

40a80000

(b) Convert your result into human-readable decimal format.

Sign: 0

Exponent: 10000001

Significand: 1.0101

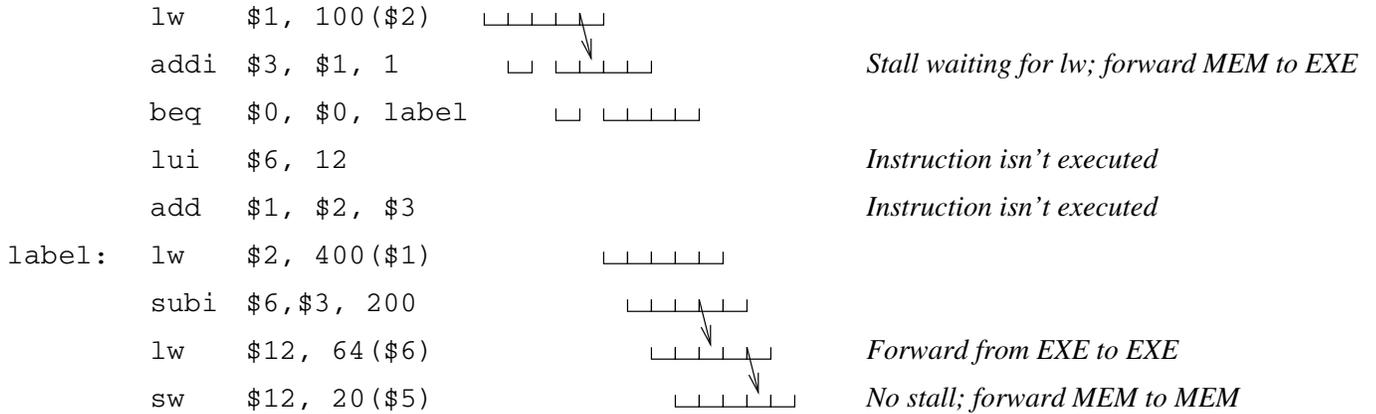
Denormalize: significand becomes $101.01_2 = 5.25_{10}$

Grading:

Problem in five parts: conversion of each operand, actual multiplication, combination of result into IEEE format, conversion of result. Each part 4 points (so 16 points for 1a, 4 for 1b)

Penalty	Error
-2	Renormalization
-2	Hidden bits
-1	Bizarre shifting on multiplication
-1	Forgot to subtract 01111111
-1	Subtracted instead of adding on renormalization

2. (15 points) Draw a Gantt chart showing how the following code will be executed on a MIPS processor, assuming all possible forwarding, stalling on lw when necessary, and a 1-cycle branch delay (no delayed loads or branches). Use arrows to show the forwarding between instructions.



Grading:

7 instructions executed: 2 points per instruction. 1 point freebie

Penalty	Error
-2	Executed skipped instructions
-1	Missed/extra stall
-1	Missed/extra forward
0	4-cycle instruction
-2	Issue 4 instructions per cycle

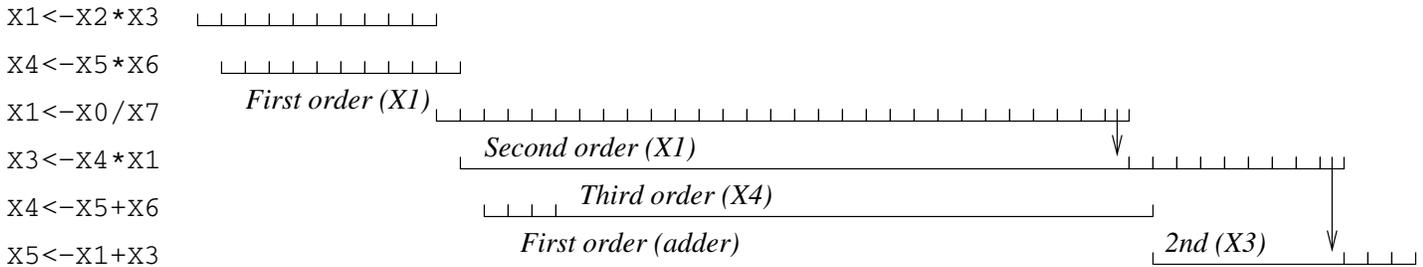
Common errors:

Forgot that stalls "ripple" - instruction following a stalled instruction has to delay a cycle due to structural hazard (second and third instruction)

Forgot that you can forward from a lw to a sw without a stall (last two instructions)

Didn't take branch (instructions 4 and 5)

3. (10 points) Draw a Gantt chart showing how the following code will be executed on a CDC, using the assumptions from the notes.



Grading:

6 instructions; 2 points per instruction. 3 points freebie

<i>Penalty</i>	<i>Error</i>
-1	4-cycle multiply
-1	Missed conflict
-1	10-cycle add
0	Didn't show forwarding (I forgot to ask for it)
-1	Multiple issue (not go) per cycle

4. (35 points) Suppose a computer system, using a 32 bit physical address, has a 16KB, 4-way set-associative cache with a 16 byte block size.

(a) How would a physical address be broken down for cache lookups?

The 16 byte block size implies a 4 bit offset field.

There are 16KB/16B = 1024 blocks; since it's 4-way set associative there are 1024/4 = 256 sets so the index field is 8 bits.

The tag is 32 - (8 + 4) = 20 bits.

Now suppose this same computer system also has a 32 bit virtual address, a virtual memory system just like Intel's, and a 64-entry, 4-way set-associative TLB.

(b) How would a virtual address be broken down for TLB lookups?

Since the VM system is borrowed from Intel it uses a 12 bit offset. The TLB has 64/4 = 16 sets, so there is a 4 bit set number.

The TLB tag is 32 - (12 + 4) = 16 bits.

(c) Would it be possible to perform TLB and cache lookups simultaneously in this computer?

Yes. The cache offset+index fields are no wider than the VM offset field.

Here are some relevant contents of the computer's TLB, cache, and the PDDBR:

TLB

Index	Tag	Valid	Contents
7	22fa	1	000c2025
7	5767	0	5ca8b39e
7	696c	0	01a2fe7f
7	696c	1	0009a043
b	22fa	1	000d7003
b	5767	0	000c2025
b	5767	1	5ca8b39e
b	696c	0	0039707f
c	22fa	1	0039707f
c	22fa	1	000d7003
c	696c	0	000c2025
c	696c	1	00149025

Cache

Index	Tag	Valid	Contents (4 bytes at given offset)			
			c	8	4	0
11	00397	0	b55d3fd1	2cb265fc	729df029	79f0c9a5
11	0acab	0	a1112079	7e0f2618	d2225f7e	3cad0d5
11	28fd3	0	6a7fd81e	e2fea616	eacf5cf6	71201139
11	77223	1	4eae81d	06967fac	fd96b0a3	7ee0fbc
47	00149	0	89ccea47	96d701e8	84ce6123	5839bff8
47	2bc3b	1	01457a16	046af216	1658f4ad	05691ab3
47	6d224	0	00f1a2cd	3cfb6a72	af09eedd	bd6f7564
47	7b483	0	8fe23130	fde64f24	2c6e3d14	5a92f7b0
c1	000d7	1	4ca424f7	a30ce94f	76f7b80c	ebc1bee0
c1	6c45e	0	0a87c67c	4adeb8a5	42b931ca	5acc9105
c1	759e2	0	c43123a1	0e3e097d	40a6a704	12d0d9bd
c1	78be5	1	34bbc2cf	ff113e65	ebc1bee0	e33fea7a

Memory

Address	Data
0004c9dc	0039707f
0009aeac	000d7003
000c2b30	00149025
000d7c14	76f7b80c
0012a22c	0009a043
0012a574	0004c067
0012a694	000c2025
0014947c	5ca8b39e
0039711c	01a2fe7f

PDBR = 0012a000

- (d) For each of the following attempted memory operations, what happens? Be sure to say whether TLB is a hit or a miss (and why, if a miss) cache is a hit or a miss (and why, if a miss), whether there are any protection violations, and whether there is a page miss. If the result is that data is returned to the program, give the value returned (assume a 4-byte read). If you have a page fault or a protection violation, there's no need to look in the cache. *Remember, the cache uses physical addresses.*

	Address	Mode	Operation
i.	696cc47c	User	Write
ii.	22fabcd14	Kernel	Read
iii.	5767711c	Kernel	Write

- i. First, see if we can find the translation in the TLB. For TLB lookup we divide the address as $696c \ c \ 47c$. Looking in TLB set c , we find a tag of $696c$ on a valid translation, so our page table entry is 00149025 (the fact that there is also an invalid entry with the same index and tag is no problem; only valid entries matter. The fact that there are two valid entries in the set with tag $22fa$ is an error on my part, but since page $22fac$ doesn't turn up in the problem it's not a serious one). The PTE tells us that the page is in memory (bit 0) and a user can access the page (bit 2), but can't write (bit 1). So it's a protection fault.
- ii. Again, start with the TLB. Looking in set b , tag $22fa$, we have another TLB hit; this time the PTE is $000d7003$. This is also a page hit; we're in kernel mode so the U bit being 0 doesn't matter and we're reading so the W bit being 0 doesn't matter. Our physical address is $000d7c14$. Looking in cache set $c1$ we find tag $000d7$ is valid, so we read the value $76f7b80c$ from offset 4.
- iii. TLB set 7 tag 5767 appears, but is not valid. So we need to go to the tables in memory. The directory entry is at $0012a574$, and contains $0004c067$. We're valid and have no protection violations, so we go on to the page table entry at $0004c9dc$, which contains $0039707f$. Once again no protection violations, so the address we're writing to is $0039711c$. Looking in cache again, set 11 tag 00397 is invalid. So data is read from the 16 bytes at 00397110 into any of the three invalid blocks of cache set 11, and the new value is written into offset c of the block (oops – I didn't give you any of those values...). If it's a write-through cache (I didn't specify), it'll also write the value to memory.

Grading:

Parts 4a-4b: 5 points each, 4c 2 points, 4(d)i 5, 4(d)ii 8 points, 4(d)iii 10 points

Penalty	Error
-2	field width wrong
-2	4c forgot "implied why"
-2	4d didn't look in TLB
-2	4(d)ii, 4(d)iii Didn't look in cache
-1	Page fault vs. protection violation
-1	4(d)i TLB miss
-1	4(d)iii returned value on write

5. (20 points) A particular computer system is capable of executing 2 billion instructions per second. It has a disk drive that has an average seek time of 10 msec (a bit slow by today's standards, but easy to work with), and turns at 6000 RPM (also a bit slow, but also easy to work with). The disk is capable of transferring 20 million bytes per second. This computer system is executing a program that requires it to repeatedly execute transactions requiring the following sequence of operations:

- Spend 1 million instructions calculating its next sequence of reads and writes.
- Read 5 blocks from random locations on the disk. Each block is 4000 bytes.
- Spend 10 millions instructions processing the data.

- Write one (4000 byte) block to a random location on disk.

(a) How many of these transactions can this system execute per second?

The times for the transaction steps are:

Processing: $(10 + 1) \times 10^6 / 2 \times 10^9 = 5.5 \times 10^{-3}$

IO: $(5 + 1)(10 \times 10^{-3} + 5 \times 10^{-3} + 4 \times 10^3 / 20 \times 10^6) = 6(10 + 5 + .2) \times 10^{-3} = 91.2 \times 10^{-3}$

So, each transaction takes 96.7×10^{-3} seconds; the system can process slightly more than ten transactions per second.

(b) You are given the option of doubling the performance of one of the following aspects of the system: the processor speed, the disk RPM, or the disk transfer rate. Which one should you pick? Why, in the practical world, is that the hardest one of the three to improve?

The CPU costs 11 msec, the disk RPM costs 30 msec, and the transfer rate costs 3 msec. The one to double is the disk RPM, since that'll make the biggest difference to the time to perform a transaction. It's the hardest one of the three to increase because it's a mechanical system.

Grading: 5a 10 points, 5b 5 points.

Penalty	Error
-2	5b Reason other than physical limits
-5	5a Forgot access time
-3	5b No reason
-5	5a Partial setup with most numbers
-4	5b Transfer rate
-3	5a No disk seek on read (???)
-2	Didn't use half of rotation