


University of California, Berkeley – College of Engineering

Department of Electrical Engineering and Computer Sciences

Spring 2007

Instructor: Dan Garcia

2007-05-12



After the exam, indicate on the line above where you fall in the emotion spectrum between “sad” & “smiley”...

Last Name													
First Name													
Student ID Number													
Login	cs61c-												
Login First Letter (please circle)	a	b	c	d	e	f	g	h	i	j			
Login Second Letter (please circle)	a	b	c	d	e	f	g	h	i	j	k	l	m
	n	o	p	q	r	s	t	u	v	w	x	y	z
Your LAB TA's name (please circle)	Aaron	Alex	Brian	David	Matt	Michael	Valerie						
Name of the person to your Left													
Name of the person to your Right													
All the work is my own. I have no prior knowledge of the exam contents nor will I share the contents with others in CS61C who have not taken it yet. (please sign)													

Instructions (Read Me!)

- This booklet contains 9 numbered pages including the cover page. Put all answers on these pages (feel free to use the back of any page for scratch work); don't hand in any stray pieces of paper.
- Please **turn off** all pagers, cell phones & beepers. Remove all hats & headphones. Place your backpacks, laptops and jackets at the front. Sit in *every other* seat. Nothing may be placed in the “no fly zone” spare seat/desk between students. The exam is closed book, no computers, PDAs or calculators.
- Fill in the front of this page and put your name & login on every sheet of paper.
- You may use 2 pages (US Letter, front and back) of notes, plus the green COD 3/e reference sheet.
- There may be partial credit for incomplete answers; write as much of the solution as you can. We will deduct points if your solution is far more complicated than necessary. When we provide a blank, please fit your answer within the space provided. “IEC format” refers to the mebi, tebi, etc prefixes.
- **You must complete ALL THE QUESTIONS, regardless of your score on the midterm.** Clobbering only works from the Final to the Midterm, not vice versa. You have 3 hours...relax.

Problem	M1	M2	M3	Ms
Minutes	20	20	20	60
Points	10	10	10	30
Score				

F1	F2	F3	F4	F5	Fs
24	24	24	24	24	120
18	18	18	18	18	90

Total
180
120

Midterm Revisited

M1) Do you remember when... we used to sing, sha la la la...

a) Some say bubble-gum pop music is garbage, so let's start with a question about memory management and running programs! BRIEFLY tell us why the first thing is better than the second thing:

<u>Mark & sweep</u> is better than <u>reference counting</u> because it...	
<u>Reference counting</u> is better than <u>copying</u> because it...	
<u>Copying</u> is better than <u>mark & sweep</u> because it...	
The <u>buddy scheme</u> is better than the <u>slab allocator</u> because it...	
The <u>slab allocator</u> is better than the <u>plain K&R free list</u> because it...	
<u>Compiling</u> is better than <u>interpreting</u> because...	
<u>Interpreting</u> is better than <u>compiling</u> because...	
<u>Dynamic linking</u> is better than <u>Static linking</u> because...	
<u>Static linking</u> is better than <u>Dynamic linking</u> because...	


b) Decode the binary numbers into MIPS instructions *with proper register names* (\$s0, \$t0, etc.). If there are any memory addresses, represent them in hex.

Address	32-bit Binary Instruction	Type (R, I, J)	MIPS Instruction w/args
0xAFFFFFFF8	0000 0001 0000 1000 0100 0000 0010 0110		
0xAFFFFFFFC	0001 0100 0000 1000 1111 1111 1111 1110		
0xB0000000	0000 1000 0000 0000 0000 0000 0000 0001		
0xB0000004	...whatever...	...whatever...	ori \$v0, \$0, 0x61C
0xB0000008	...whatever...	...whatever...	jr \$ra


c) You can replace the first instruction with a *new* instruction and save 2 clock cycles on a single-cycle non-delayed branch MIPS machine. What is it (in MIPS)? *Careful!* _____

M2) Michael Moore should do his next film on *this* encoding!

Normal floating-point numbers are “bunched up” around zero so that they can represent very small numbers with high accuracy, but they increasingly spread out as the magnitude of the number increases:



However, now instead of having accuracy around zero *only*, we would like accuracy around *every* integer. We'll call them “floating point integers,” or `flints`.



We will have *four* fields in our **16-bit** `flint` representation: a `mini_int` part (a signed integer representation), and a `mini_float` part (complete with the usual sign bit, exponent field, and mantissa field):

Integer part	Float sign	Exponent	Mantissa
8 bits (signed representation) <i>“mini_int part”</i>	1 bit	4 bits	3 bits
	<i>“mini_float part”</i>		

$$\text{Value} = \text{value_of}(\text{mini_int part}) + \text{value_of}(\text{mini_float part})$$

We decode the raw bit representation into a number using the formula above. The `mini_float` pattern will *not* have $\pm \infty$ nor NaNs, but it *will* have denorms. We wish to adjust the exponent field so $\text{abs}(\text{mini_float}) < \frac{1}{2}$. This allows us to “copy & paste” this `mini_float` distribution around every integer (via the `mini_int offset`) to get the overall `flint` representation number distribution sketched above. Cool, eh?

A. What should the value of `bias` be? (write your answer in decimal)

B. Given the field widths in the diagram and the `bias` value from (A), complete the table. Expressions are OK, but the rightmost column answers **MUST be in base 10!**

Commentary on number	Bit Encoding	Represented Number (in base 10)
Biggest representable number smaller than 10	0x	
Nothing special, just a nice number	0x	-15.25
Most-negative number (closest to $-\infty$)	0x	

M3) Because I MIPS you, baby, and I don't want to C a thing...

What follows is a self-modifying MAL MIPS function. Read it carefully, and answer the questions below.

- A) What is the equivalent C code for `mystery`?
 Assume for this part that we only call `mystery` once.

```
// Precondition: 0 < a1 < 215
_____ mystery (_____ a0, int a1) {

}

```

```
mystery:
    la    $t0, loop
    addu  $v0, $0, $0

loop:  beq   $a1, $0, done
       lw    $t1, 0($a0)
       addu  $v0, $v0, $t1
       lw    $t1, 4($t0)
       addiu $t1, $t1, 4
       sw    $t1, 4($t0)
       addiu $a1, $a1, -1
       j     loop

done:  nop
       nop
       nop
       jr   $ra

```

- B) When we call `mystery` the second time with the *same arguments we used the first time*, do we get the same answer? If not, what do we get? (assume it's not an error)

- C) Replace as few `nops` as possible by the `done:` label so that `mystery` behaves as it did in part A every time it's called, not just the first time.

```
done: _____
       _____
       _____
       jr  $ra

```

Name: _____ Login: cs61c-_____

Post-Midterm Questions

F1) Code bugs are terrified of *RA-AI-AID*!!! (Poof!)

1) If you put 5 drives with a mean time to failure (MTTF) of 10 years in a RAID 0 array, what will be the MTTF of the array?

2) In an interrupt service routine, should the ready bit of a device be checked before accessing it?

3) A CPU running some program was found to have a CPI is 2.5 and, on average, 5 ns of CPU time were used *per instruction*. What is the clock speed in Gigahertz?

4) What are the two largest challenges that prevent parallel programs from achieving perfect speedup equal to the number of processors?

a)

b)

5) What does Professor Patterson (and team) want to put into the hands of every systems researcher very soon?

6) Rather than endlessly trying to create bug-free programs and somehow find users & operators that never make mistakes (or even trying to eliminate the human element altogether), where does Professor Patterson believe we should focus our efforts?

7) What does the receiver do with a packet if its checksum indicates it was corrupted in transit?


F2) Tune in to 101 on your FSM dial...

We are designing a *palindrome*-finder circuit with a 1-bit input $I(t)$ and a 1-bit output $O(t)$, that will produce, at time t , whether the sequence $\{I(t-2), I(t-1), I(t)\}$ is the same backwards and forwards (e.g, 101). We'll assume $I(t)$ has been 1 for all negative time (i.e., before the finder circuit starts). As an example,

the input: $I: 1\ 1\ 0\ 0\ 1\ 1\ 1\ 0\ 0\ 0\ 1\ 0\ 1\ 1\ 0\ 0$
 will produce the output: $O: 1\ 1\ 0\ 0\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 1\ 1\ 0\ 0\ 0$

P1	P0	I	O	N1	N0
0	0	0			
0	0	1			
0	1	0			
0	1	1			
1	0	0			
1	0	1			
1	1	0			
1	1	1			

- a) Complete the FSM diagram below. Our states have been labeled s_{xy} indicating that the previous 2 bits, $\{I(t-2), I(t-1)\}$ would be $\{x, y\}$. Fill in the truth table on the right. The previous state is encoded in $(P1, P0)$, the next state is encoded in $(N1, N0)$, and the output is encoded as o . Make sure to indicate the value of the *output* on your state transitions AND to indicate the starting state with an "incoming arrow".



- b) Provide a *fully reduced* (i.e., fewest gates to implement...you can use any n -input gates) Boolean expression for the Output o as a function of $P1$, $P0$ and I . If there is a name for the circuit, write it in the box above. E.g., "The always-1", "3-input NAND", etc. A 2-input XOR has the symbol of \oplus .

Name: _____

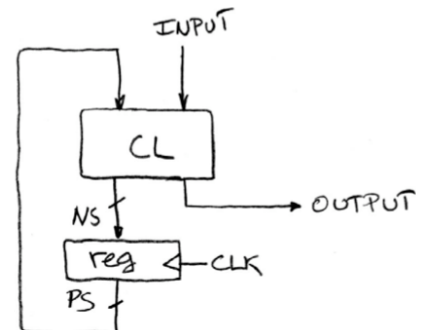
$o =$ _____

- c) How many different answers could I have put in the box for "b" above? Said another way, how many different circuits can a 3-LUT imitate?

- d) We're always concerned about testing. What is the shortest length of an $I(t)$ stream that can guarantee you've tested this particular circuit exhaustively?

- e) Finally, we wish to build our circuit as we normally do for SDS systems (shown below). Given the four standard spec times from the chip manufacturer (τ_{setup} , τ_{hold} , $\tau_{clk-to-q}$, and τ_{CL}), what is the smallest clock period τ we can drive our system with? (Write your answer as an expression involving the spec variables.) Feel free to draw timing diagrams if you wish.

CLK																			
PS																			
I																			
NS																			



F3) On a final exam, a little cache can boost your performance...

No, we can't be bribed, but we *can* improve our memory access time...let's take a look!

- a) For the purposes of this question, assume our MIPS processor has two levels of data caches with the capabilities shown in the table on the right. Also assume:
- It takes **100 cycles** to go to memory
 - 1GiB of physical address space
- Fill in the table on the right.

	L1	L2
Cache Data Size	32 KiB	512 KiB
Block Size	8 B	32 B
Associativity	4-way	Direct-mapped
Hit Time	1 cycle	33 cycles
Miss Rate	10%	2%
Write Policy	Write-through	Write-through
Replacement Policy	LRU	LRU
Tag		
Index		
Offset		
AMAT (in clock cycles)		

- b) In software, you decide to live on the wild side so you turn OFF your L2 cache and clear your L1 cache. Here are other features:
- 4 GiB of virtual address space
 - 2 KiB page size
 - 8-entry TLB, LRU replacement
 - 32 MiB ARRAY_SIZE
 - char A[] starts at a block & page boundary (i.e., block, page aligned)
 - The following code is run on the system with no other users and process switching turned off.

```
main() {
    int i,j;
    char *A = (char *) malloc (ARRAY_SIZE * sizeof(char));           // block,page aligned
    for (i = 0 ; i < (ARRAY_SIZE/STRETCH) ; i++) {                   // # of STRETCHes
        for (j = 0 ; j < STRETCH ; j++) sum += A[i*STRETCH + j];    // go up to STRETCH
        for (j = STRETCH-1 ; j >=0 ; j--) prod *= A[i*STRETCH + j]; // down from STRETCH
    }
}
```

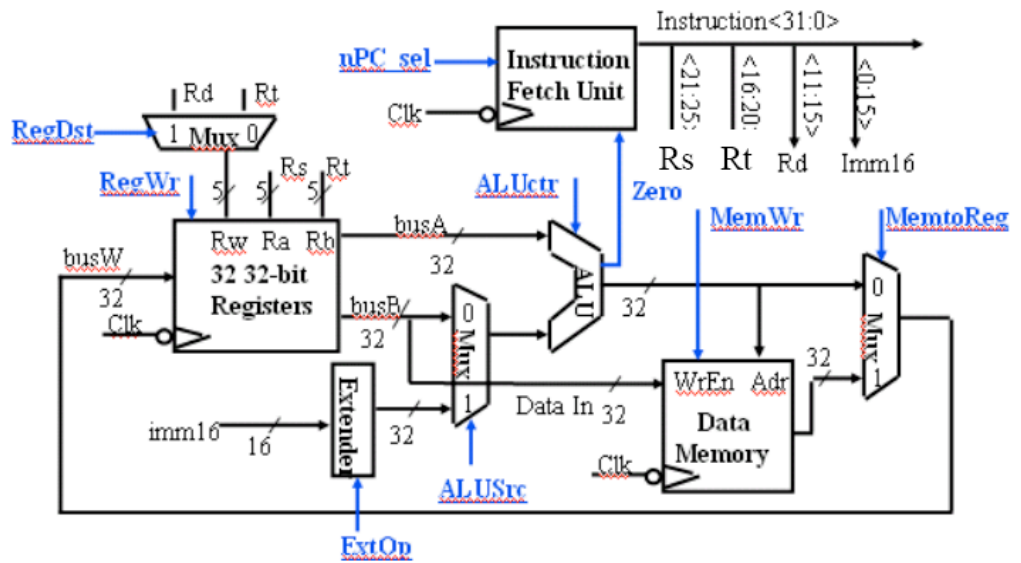
- i. What is number of bits used for the VPN (assume byte addressing)? _____
- ii. What is number of bits used for the PPN (assume byte addressing)? _____
- iii. As we double our STRETCH from 1 to 2 to 4 (...etc), we notice the number of cache misses doesn't change! What is the largest value of STRETCH before cache misses changes? (Use IEC terms, like 64 TiB, 128 GiB, etc.) _____
- iv. If we double the STRETCH from (iii), what is ratio of cache hits to misses? _____
- v. As we double our STRETCH from 1 to 2 to 4 (...etc), we notice the number of TLB misses doesn't change! What is the largest value of STRETCH before TLB misses changes? (Use IEC terms, like 64 TiB, 128 GiB, etc.) _____
- vi. For any value of STRETCH what is the fewest number of page faults we could ever generate? (Use IEC terms, like 64 TiFaults, 128 GiFaults, etc.) _____

F4) Don't just sit and wait for another datapath you by!

On the right is the **single-cycle** MIPS datapath presented during lecture. **Ignore pipelining for the question.** Your job is to modify the diagram to accommodate a new MIPS instruction.

Your modification may use simple adders, shifters, mux chips, wires, and new control signals. If necessary, you may replace original labels.

We want to add a new MIPS instruction (we'll call it `addpr` for "add to pointed reg") that is almost identical to `addi` but with a twist. Instead of storing into the `rt` register the sum of the constant and the value of the register specified by `rs`, it stores into the `rt` register the sum of the constant and the value of the register specified by the lowest 5 bits in memory at the address specified by the pointer stored in the `rt` register. Said another way, first get the pointer stored in the `rt` register. Follow that pointer to get its value from memory. Take the lowest 5 bits of that value, treat it as a register number, and find out what is stored in that register. Add that to the immediate, and store it in the `rt` register.



a) Make up the syntax for the I-type MAL MIPS instruction that does it (show an example if the pointer lives in `$v0`, and the constant is 5). On the right, show the register transfer language (RTL) description of `addpr`.

Syntax	RTL

b) Change as *little as possible* in the datapath above and list all changes below. You may not need all boxes.

(i)	
(ii)	
(iii)	

c) We now want to set all the control lines appropriately. List what each signal should be, an intuitive name or {0, 1, x - don't care}. Include any new control signals you added.

RegDst	RegWr	nPC_sel	ExtOp	ALUSrc	ALUctr	MemWr	MemtoReg			

d) In the context of a single-cycle CPU, `lw` used to be the "critical path" instruction that really determined our fastest clock period, since it utilized the most components of our datapath. Using the terms below to create an expression that will determine *how much slower our clock period will be if we also consider the `addpr` instruction*: `PCRegClkToQ`, `InstMemAccess`, `ControlLogicDelay`, `RegFileAccess`, `ALUdelay`, `DataMemAccess`, `RegSetup`, `RegHold`, `MemSetup`, `MemHold`.

F5) Whose pipeline is it anyway? 500 points if you can guess...

You are the owner of a company that manufactures computers on an assembly line 24 hrs/day. You have 5 stages in your assembly line, and each hour (adjustable if need be) the partially-finished computer moves to the next stage. With 3 shifts of 5 workers working 8 hrs/day, you've been in business for weeks. The 5 stages:

- **PARTS:** Lay out all the parts to be used in the later stages (takes 1 full hour of operator time)
- **HW:** Plug in all the boards and wire it up (takes 1 full hour of operator time)
- **SW:** Install the software (takes 1 full hour of operator time clicking “Ok”, “Yes, install everything”, etc)
- **TEST:** Test that the machine works (the test always takes 1 full hour, but setting it up only takes a moment of operator time, the rest of the hour the operator sits idle waiting for the test to complete). In the rare case of failure, the test software makes a loud annoying sound and the machine is placed in the parts recycling bin; most of the time the test is successful and the computer moves to the next stage.
- **BOX:** Put the machine in a box, affix a shipping label & load it on a truck (takes 1 full hour of operator time)

- a) What is the latency of your assembly line?
- b) What is the throughput of your assembly line?
- c) How much faster are you than your competitor who manufactures the same computer but only hires 1 worker per each 8 hr shift? (use Nx, not M% notation)
- d) All was going fine until yesterday. Congress passes a law requiring you run a more rigorous *5-hour test* on every machine! This requirement stands to throw a major wrench in your well-crafted manufacturing process. Without hiring anyone else, and with minimal changes to what you already have in place, what do you do to be maximally productive? Draw a diagram to explain yourself.

- e) Ok, now back to MIPS. You buy a pimped-up single-user, single-process (aside from the OS), single-CPU, single-core MIPS machine. Miraculously, all 5 pipeline stages *a/ways* complete in exactly 1ns (you spared no expense buying the most cutting-edge parts, especially memory – no cache of any sort was needed since the memory was so fast). A thief breaks into your dorm room and swaps your awesome 1ns memory with cheap 5ns memory. Both memories have the same standard, simple interface (in terms of input and outputs) you saw when we built the MIPS datapath. Your system clock (initially set to 1ns) is adjustable. Does the technique you used in (d) work here? Why or why not?

- f) How many cycles would it take to execute the first loop iteration for the code on the left on the 5-stage MIPS machines on the right? Note: memory and registers CAN be written and read in the same cycle, branch compares occur in the 2nd stage, we stall on ALL hazards, & there is no out-of-order execution.

Code	Non-delayed branch No forwarding	Delayed branch Forwarding
<pre> Loop: lw \$t0, 42(\$s1) addu \$t1, \$t0, \$s2 addu \$v0, \$t1, \$t1 beq \$t0, \$0, Loop sll \$0, \$0, 0 </pre>		