

# CS232 Final Exam

## May 5, 2006

Name: \_\_\_\_\_

- This exam has 10 exciting pages, including this handsome cover.
- There are five breathtaking questions, worth a total of 100 points.
- The final three pages are for your reference and can be detached from the exam for your convenience. They include: 1) a summary of the MIPS instruction set, 2) a datapath diagram of a pipelined MIPS processor, and 3) a reference on ASCII coding.
- No other written references or calculators are allowed.
- Write clearly and show your work. State any assumptions you make.
- You have ~2 fun-filled hours. Budget your time, and good luck!

Question	Maximum	Your Score
1	25	
2	15	
3	20	
4	15	
5	25	
Total	100	

### Question 1: MIPS programming (25 points)

Below is a function that reads a file and counts the number of times the string “cs” exists in the file. It does this with a simple two-state finite-state machine, where the first state (state == 0) is the state where the previous character was **not** ‘c’ and the second state (state == 1) is when the previous character was ‘c.’ The code uses the `fgetc` function (prototype shown to the right) to get the next character in the FILE; `fgetc` returns EOF (end of file) when there are no more characters in the file. Write MIPS code for the `count_sequence` function (your code should call `fgetc`, but not implement it). Be sure to observe the MIPS calling convention. Refer to the attached MIPS and ASCII references.

```
int count_sequence(FILE *fptr) {
    int count = 0, state = 0, c;

    while ((c = fgetc(fptr)) != EOF) {
        if ((state == 1) && (c == 's')) {
            count++;
            state = 0;
        } else if (c == 'c') {
            state = 1;
        } else {
            state = 0;
        }
    }
    return count;
}
```

```
int fgetc(FILE *);
#define EOF (-1)
```

```
## int count_sequence(FILE *fptr) {
count_sequence:
    sub    $sp, $sp, 16
    sw     $s0, 0($sp)
    sw     $s1, 4($sp)
    sw     $s2, 8($sp)
    sw     $ra, 12($sp)
    move   $s0, $a0

    ##    int count = 0, state = 0, c;
    li     $s1, 0
    li     $s2, 0

    ##    while ((c = fgetc(fptr)) != EOF) {
cs_loop:
    move   $a0, $s0
    jal    fgetc
    beq    $v0, -1, cs_done

    ##        if ((state == 1) && (c == 's')) {
    ##            count++;
    ##            state = 0;
    bne    $s2, 1, cs_case2
    bne    $v0, 115, cs_case2
    li     $s2, 0
    add    $s1, $s1, 1
    j      cs_loop

    ##        } else if (c == 'c') {
    ##            state = 1;
    cs_case2:
    bne    $v0, 99, cs_case3
    li     $s2, 1
    j      cs_loop

    ##        } else {
    ##            state = 0;
    cs_case3:
    li     $s2, 0
    j      cs_loop

    ##        }
    ##    }
    ##    return count;
cs_done:
    move   $v0, $s1

    lw     $s0, 0($sp)
    lw     $s1, 4($sp)
    lw     $s2, 8($sp)
    lw     $ra, 12($sp)
    add    $sp, $sp, 16
    jr     $ra

## }
}
```

## Question 2: Processor and Memory System Performance (15 points)

$$\begin{aligned}\text{Execution Time} &= \# \text{instructions} * \text{CPI} * \text{clock\_period} \\ \text{CPI} &= \text{Cycles Per Instruction} \\ \text{AMAT} &= \text{hit\_time} + (\text{miss\_rate} * \text{miss\_penalty}) \\ \text{Memory Stall Cycles} &= \# \text{ memory operations} * \text{miss\_rate} * \text{miss\_penalty}\end{aligned}$$

Consider a program whose execution consists of: 50% arithmetic instructions, 20% loads, 10% stores, and 20% branches (25% not-taken, 75% taken). For half of the arithmetic and load instructions, the subsequent instruction is a dependent arithmetic or load instruction.

### Part (a)

Consider the pipelined MIPS processor discussed in this class and whose pipeline is shown on page 9 of this exam. This machine predicts branches as not-taken and resolves branches in the ID stage of the pipeline, flushing any incorrect instructions on a branch misprediction. The processor includes forwarding from both the EX/MEM and MEM/WB registers to either input of the functional units. Assuming a perfect memory system (no stalls due to cache misses), compute the cycles per instruction (CPI) for this machine. (5 points)

Stalls when load followed by either arithmetic or load instructions:  $(20\% \text{ loads} * 50\%) * 1 \text{ cycle} = .1 \text{ CPI}$

Flushes due to taken branches:  $(20\% \text{ branches} * 75\% \text{ taken}) * 1 \text{ cycle} = .15 \text{ CPI}$

Total CPI = 1 + stalls + flushes =  $1 + .1 + .15 = 1.25 \text{ CPI}$

### Part (b)

Consider the following realistic memory hierarchy described below: *Note: the next level of the hierarchy is accessed only after a miss is detected in the current level.*

L1: two-way set-associative, write-back, write-allocate 16KB cache with 32B blocks and 1 cycle access time.  
L2: eight-way set-associative, write-back, write-allocate 1MB cache with 32B blocks and 8 cycle access time.  
Memory: 200 cycle access time.

Assuming the instruction mix above, when 1000 instructions are executed, the processor encounters, on average, 50 L1 data cache misses and 3 L2 data cache misses, what is the average memory access time (AMAT) in cycles? (5 points)

$L1\_miss\_rate = 50 \text{ misses} / (1000 \text{ inst} * (20\% \text{ loads} + 10\% \text{ stores})) = 50/300 = 1/6$

$L2\_miss\_rate = 3 \text{ misses} ./ 50 \text{ accesses} = 3/50$

$L2\_AMAT = L2\_hit\_time + L2\_miss\_rate * \text{memory\_access\_time} = 8 + 3/50 * 200 = 8 + 12 = 20$

$AMAT = L1\_hit\_time + L1\_miss\_rate * L2\_AMAT = 1 + 1/6 * 20 = 1 + 10/3 = 4 \frac{1}{3}$

### Part (c)

Compute a new CPI for the processor described in Part (a) using the memory system described in Part (b). (5 points)

$CPI\_with\_mem = CPI\_w/o\_mem + (AMAT - 1) = 1.25 + (4.3333 - 1) = 4.583333..$

### Question 3: Input/Output Systems (20 points)

Assume a hard drive that has a 6,000 rpm rotational speed, a 3ms average seek time, and that overhead is negligible. Each disk track has 64 sectors, and each sector holds 1 KB of data. The drive can read or write data as fast as it rotates. For this problem, to make the computations easier, you may assume that 1 KB = 1,000 bytes and 1 MB = 1,000 KB and 1 GB = 1,000 MB. For parts (a)-(c), compute the time to read a 1 GB file in the following three scenarios. In each case you are welcome to leave your result as an expression.

#### Part (a)

The file is randomly spread across disk sectors. (5 points)

$1\text{GB}/1\text{KB}/\text{sector} = 1\text{M sectors}$        $\text{rot\_period} = 1/6000\text{rpm} * 60\text{sec}/\text{m} = 1/100 \text{ s} = 10\text{ms}$   
 $\text{Time to read one sector} = \text{seek\_time} + \text{rot\_delay} + \text{transfer\_time} = 3\text{ms} + .5 * 10\text{ms} + (10\text{ms}/64) = 8.15625 \text{ ms}$   
 $\text{Total time} = 1\text{M} * 8.15625\text{ms} = 8156 \text{ sec} = 2.265 \text{ hours}$

#### Part (b)

The file is written sequentially onto tracks, but the tracks are spread randomly across the disk. (5 points)

$\# \text{ tracks} = 1\text{M sectors} / 64 \text{ sectors}/\text{track} =$   
 $\text{Time to read one track} = \text{seek\_time} + (1 \text{ or } 0) * \text{rot\_delay} + \text{rot\_period} = 13\text{ms or } 18\text{ms}$   
 $\text{Total time} = 1\text{M}/64 * (13\text{ms or } 18\text{ms}) = 203 \text{ or } 281 \text{ sec} = 3.4 \text{ or } 4.7 \text{ min}$

(We'll accept assumptions of either needing to wait for rotational delay or not)

#### Part (c)

The file is written sequentially onto sequential tracks. You can assume it takes only 1ms to seek to a neighboring track. (5 points)

$\# \text{ tracks} = 1\text{M sectors} / 64 \text{ sectors}/\text{track} =$   
 $\text{Time to read one track} = \text{seek\_time} + (1 \text{ or } 0) * \text{rot\_delay} + \text{rot\_period} = 11\text{ms or } 16\text{ms}$   
 $\text{Total time} = 1\text{M}/64 * (11\text{ms or } 16\text{ms}) = 2.9 \text{ or } 4.2 \text{ min}$

(We'll accept assumptions of either needing to wait for rotational delay or not)

#### Part (d)

Explain the difference between latency and throughput and give examples of when each would be an appropriate metric. (5 points)

Latency is the time to do one thing (units = time) , throughput is the rate which many things are done (units = things/time).

At a bank you would consider the time it took for one customer to be served a latency, but the rate at which the bank served the customers would be a throughput.

#### Question 4: Pipeline Implementation (15 points)

##### Part (a)

In lecture, we discussed how the pipeline handled conditional branches. These aren't the only kinds of control flow. To other kinds of important control flow are the jump-and-link (`jal`) instructions used for function calls and jump-register (`jr`) instructions used for return. In what pipeline stage would you expect `jal` and `jr` to be resolved (i.e., when can the next PC be known for sure)? (Note: They may be resolved in different pipeline stages.) Justify your answer. You can ignore dependences with other instructions in the pipeline. (5 points)

`jal` would likely be resolved in the ID stage because there is no condition to compute and the branch target can be computed when the instruction is decoded.

`jr` would likely be resolved in ID also, because at the end of the ID stage we've read the branch target out of the register file.

##### Part (b)

In the pipeline datapath at the end of the exam, there is no forwarding datapath present for the following dependence.

<code>lw</code>	<code>\$12, 0(\$11)</code>
<code>sw</code>	<code>\$12, 0(\$13)</code>

How many stall cycles are necessary to handle this hazard? Explain. (5 points)

A single cycle stall would enable forwarding from the mem/wb latch to the ALU in EX using the existing bypass datapath (through the `rt_mux` and `imm_mux`).

It would also be correct to stall the machine for two cycles, keeping the `sw` instruction in the ID stage so that the register write of the `lw` and the register read of the `sw` are done in the same cycle.

##### Part (c)

Write an equation for the hazard detection logic for the above stall condition. (5 points)

```
if (IF/ID.mem_write && ID/EX.mem_read &&
    ID/EX.rt == IF/ID.rt) {
    stall;
}
```

It is also alright to use EX/MEM and ID/EX for ID/EX and IF/ID (respectively) above. Also, accept ID/EX.rd for `rt`.

### Question 5: Conceptual Questions (25 points)

#### Part (a)

A coding scheme has two valid code words (1010 and 0101) and uses a error detection/correction scheme. Below, we've shown how this ECC scheme handles a sample of errors.

1000 - Error Corrected to 1010

0110 - Error Detected (cannot correct)

1111 - Error Detected (cannot correct)

0111 - Error Corrected to 0101

Based on the above sample, determine how the following errors will be handled. State any assumptions.

1001 - Error Detected (cannot correct): Hamming distance 2 from both code words

1110 - Error Corrected to 1010: Hamming distance 1

0001 - Error Corrected to 0101: Hamming distance 1

0000 - Error Detected (cannot correct): Hamming distance 2 from both code words

This is 1 bit correction, 2 bit detection scheme. *Fill in the blanks.* (5 points)

#### Part (b)

Consider a machine with a byte-addressable memory that uses 32b addresses for both virtual and physical addresses. Compute the size of a single-level page table if 4KB pages are used. (5 points)

Page table size = # pages \* storage/page =  $2^{20} * 4 \text{ bytes} = 4\text{MB}$

# pages = virtual address space size / page size =  $2^{32} / 4\text{KB} = 2^{32} / 2^{10} = 2^{20}$  page

PPNsize = PA\_bits - page\_offset\_bits =  $32 - 10 = 22\text{b}$

Storage per page = round up (PPN size + 1 (dirty) + 1 (valid)) = round up ( $22\text{b} + 1 + 1$ ) =  $32\text{b} = 4\text{B}$


#### Part (c)

Describe how a multi-level page table works and explain its advantages relative to a single-level page table. (5 points)

For the multi-level page table, the virtual page number (VPN) is broken into pieces ( $\text{VPN}_1 \dots \text{VPN}_n$ ). The first piece is used to look up in the first-level page table (pointed to by the page table base register), the entry found is a pointer to the next level page table. This pointer is used with the next piece of the VPN until the last level of the page table is found which holds the PPN.

The main advantage of using the multi-level page table is that you only need to allocate memory for those subsets of the virtual address space that are being used. Since there is spatial locality generally the parts of the virtual address space being used are close together.

### Question 5: Conceptual Questions, cont.



#### Part (d)

The above figure shows average miss rates for caches of different sizes and block sizes, measured on a variety of programs. The miss rate for the 1KB cache initially decreases with increased block size, but then begins to increase again; explain why this occurs. (5 points)

Initially since the blocks are small, we aren't getting much benefit from spatial locality resulting in extra misses when spatial locality exists. As we increase the block size the number of blocks decreases (at 256B blocks there are only 4 blocks); eventually the number of conflict misses outweigh the benefit of additional temporal locality.

#### Part (e)

When parallelizing a program for fork-join parallelism (like OpenMP) what factors would prevent a two-core machine from doubling the performance of a single-core machine? Explain at least two. (5 points)

1. Any portion of the program that can not be parallelized (serial portion, Amdahl's law)
2. Overhead resulting from forking and joining
3. Overhead due to synchronization (critical sections, etc.)
4. Overhead due to true/ false sharing in the cache

## MIPS instructions

These are some of the most common MIPS instructions and pseudo-instructions, and should be all you need. However, you are free to use *any* valid MIPS instructions or pseudo-instruction in your programs.

Category	Example Instruction	Meaning
Arithmetic / Logical	add    \$t0, \$t1, \$t2 sub    \$t0, \$t1, \$t2 and    \$t0, \$t1, \$t2 mul    \$t0, \$t1, \$t2 xor    \$t0, \$t1, \$t2 srl    \$t0, \$t1, \$t2	$\$t0 = \$t1 + \$t2$ $\$t0 = \$t1 - \$t2$ $\$t0 = \$t1 \& \$t2$ $\$t0 = \$t1 \times \$t2$ $\$t0 = \$t1 \wedge \$t2$ $\$t0 = \$t1 \ll \$t2$
Register Setting	move   \$t0, \$t1 li     \$t0, 100	$\$t0 = \$t1$ $\$t0 = 100$
Data Transfer	la     \$t0, label lw     \$t0, 100(\$t1) lh     \$t0, 100(\$t1) lb     \$t0, 100(\$t1) sw     \$t0, 100(\$t1) sh     \$t0, 100(\$t1) sb     \$t0, 100(\$t1)	$\$t0 = \text{address of data at label}$ $\$t0 = \text{Mem}[100 + \$t1] \quad (32 \text{ bits})$ $\$t0 = \text{Mem}[100 + \$t1] \quad (16 \text{ bits})$ $\$t0 = \text{Mem}[100 + \$t1] \quad (8 \text{ bits})$ $\text{Mem}[100 + \$t1] = \$t0 \quad (32 \text{ bits})$ $\text{Mem}[100 + \$t1] = \$t0 \quad (16 \text{ bits})$ $\text{Mem}[100 + \$t1] = \$t0 \quad (8 \text{ bits})$
Branch	beq    \$t0, \$t1, Label bne    \$t0, \$t1, Label bge    \$t0, \$t1, Label bgt    \$t0, \$t1, Label ble    \$t0, \$t1, Label blt    \$t0, \$t1, Label	if ( $\$t0 = \$t1$ ) go to Label if ( $\$t0 \neq \$t1$ ) go to Label if ( $\$t0 \geq \$t1$ ) go to Label if ( $\$t0 > \$t1$ ) go to Label if ( $\$t0 \leq \$t1$ ) go to Label if ( $\$t0 < \$t1$ ) go to Label
Set	slt    \$t0, \$t1, \$t2 slti   \$t0, \$t1, 100	if ( $\$t1 < \$t2$ ) then $\$t0 = 1$ else $\$t0 = 0$ if ( $\$t1 < 100$ ) then $\$t0 = 1$ else $\$t0 = 0$
Jump	j      Label jr     \$ra jal    Label	go to Label go to address in \$ra $\$ra = \text{PC} + 4$ ; go to Label

The second source operand of the arithmetic and branch instructions may be a constant.

## Register Conventions

The *caller* is responsible for saving any of the following registers that it needs, before invoking a function.


\$t0-\$t9          \$a0-\$a3          \$v0-\$v1

The *callee* is responsible for saving and restoring any of the following registers that it uses.

\$s0-\$s7          \$ra



## Pipelined Datapath



## ASCII reference

Dec	Hx	Oct	Char	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr	Dec	Hx	Oct	Html	Chr
0	0	000	<b>NUL</b> (null)	32	20	040	&#32;	<b>Space</b>	64	40	100	&#64;	<b>@</b>	96	60	140	&#96;	<b>`</b>
1	1	001	<b>SOH</b> (start of heading)	33	21	041	&#33;	<b>!</b>	65	41	101	&#65;	<b>A</b>	97	61	141	&#97;	<b>a</b>
2	2	002	<b>STX</b> (start of text)	34	22	042	&#34;	<b>"</b>	66	42	102	&#66;	<b>B</b>	98	62	142	&#98;	<b>b</b>
3	3	003	<b>ETX</b> (end of text)	35	23	043	&#35;	<b>#</b>	67	43	103	&#67;	<b>C</b>	99	63	143	&#99;	<b>c</b>
4	4	004	<b>EOT</b> (end of transmission)	36	24	044	&#36;	<b>\$</b>	68	44	104	&#68;	<b>D</b>	100	64	144	&#100;	<b>d</b>
5	5	005	<b>ENQ</b> (enquiry)	37	25	045	&#37;	<b>%</b>	69	45	105	&#69;	<b>E</b>	101	65	145	&#101;	<b>e</b>
6	6	006	<b>ACK</b> (acknowledge)	38	26	046	&#38;	<b>&amp;</b>	70	46	106	&#70;	<b>F</b>	102	66	146	&#102;	<b>f</b>
7	7	007	<b>BEL</b> (bell)	39	27	047	&#39;	<b>'</b>	71	47	107	&#71;	<b>G</b>	103	67	147	&#103;	<b>g</b>
8	8	010	<b>BS</b> (backspace)	40	28	050	&#40;	<b>(</b>	72	48	110	&#72;	<b>H</b>	104	68	150	&#104;	<b>h</b>
9	9	011	<b>TAB</b> (horizontal tab)	41	29	051	&#41;	<b>)</b>	73	49	111	&#73;	<b>I</b>	105	69	151	&#105;	<b>i</b>
10	A	012	<b>LF</b> (NL line feed, new line)	42	2A	052	&#42;	<b>*</b>	74	4A	112	&#74;	<b>J</b>	106	6A	152	&#106;	<b>j</b>
11	B	013	<b>VT</b> (vertical tab)	43	2B	053	&#43;	<b>+</b>	75	4B	113	&#75;	<b>K</b>	107	6B	153	&#107;	<b>k</b>
12	C	014	<b>FF</b> (NP form feed, new page)	44	2C	054	&#44;	<b>,</b>	76	4C	114	&#76;	<b>L</b>	108	6C	154	&#108;	<b>l</b>
13	D	015	<b>CR</b> (carriage return)	45	2D	055	&#45;	<b>-</b>	77	4D	115	&#77;	<b>M</b>	109	6D	155	&#109;	<b>m</b>
14	E	016	<b>SO</b> (shift out)	46	2E	056	&#46;	<b>.</b>	78	4E	116	&#78;	<b>N</b>	110	6E	156	&#110;	<b>n</b>
15	F	017	<b>SI</b> (shift in)	47	2F	057	&#47;	<b>/</b>	79	4F	117	&#79;	<b>O</b>	111	6F	157	&#111;	<b>o</b>
16	10	020	<b>DLE</b> (data link escape)	48	30	060	&#48;	<b>0</b>	80	50	120	&#80;	<b>P</b>	112	70	160	&#112;	<b>p</b>
17	11	021	<b>DC1</b> (device control 1)	49	31	061	&#49;	<b>1</b>	81	51	121	&#81;	<b>Q</b>	113	71	161	&#113;	<b>q</b>
18	12	022	<b>DC2</b> (device control 2)	50	32	062	&#50;	<b>2</b>	82	52	122	&#82;	<b>R</b>	114	72	162	&#114;	<b>r</b>
19	13	023	<b>DC3</b> (device control 3)	51	33	063	&#51;	<b>3</b>	83	53	123	&#83;	<b>S</b>	115	73	163	&#115;	<b>s</b>
20	14	024	<b>DC4</b> (device control 4)	52	34	064	&#52;	<b>4</b>	84	54	124	&#84;	<b>T</b>	116	74	164	&#116;	<b>t</b>
21	15	025	<b>NAK</b> (negative acknowledge)	53	35	065	&#53;	<b>5</b>	85	55	125	&#85;	<b>U</b>	117	75	165	&#117;	<b>u</b>
22	16	026	<b>SYN</b> (synchronous idle)	54	36	066	&#54;	<b>6</b>	86	56	126	&#86;	<b>V</b>	118	76	166	&#118;	<b>v</b>
23	17	027	<b>ETB</b> (end of trans. block)	55	37	067	&#55;	<b>7</b>	87	57	127	&#87;	<b>W</b>	119	77	167	&#119;	<b>w</b>
24	18	030	<b>CAN</b> (cancel)	56	38	070	&#56;	<b>8</b>	88	58	130	&#88;	<b>X</b>	120	78	170	&#120;	<b>x</b>
25	19	031	<b>EM</b> (end of medium)	57	39	071	&#57;	<b>9</b>	89	59	131	&#89;	<b>Y</b>	121	79	171	&#121;	<b>y</b>
26	1A	032	<b>SUB</b> (substitute)	58	3A	072	&#58;	<b>:</b>	90	5A	132	&#90;	<b>Z</b>	122	7A	172	&#122;	<b>z</b>
27	1B	033	<b>ESC</b> (escape)	59	3B	073	&#59;	<b>;</b>	91	5B	133	&#91;	<b>[</b>	123	7B	173	&#123;	<b>{</b>
28	1C	034	<b>FS</b> (file separator)	60	3C	074	&#60;	<b>&lt;</b>	92	5C	134	&#92;	<b>\</b>	124	7C	174	&#124;	<b> </b>
29	1D	035	<b>GS</b> (group separator)	61	3D	075	&#61;	<b>=</b>	93	5D	135	&#93;	<b>]</b>	125	7D	175	&#125;	<b>}</b>
30	1E	036	<b>RS</b> (record separator)	62	3E	076	&#62;	<b>&gt;</b>	94	5E	136	&#94;	<b>^</b>	126	7E	176	&#126;	<b>~</b>
31	1F	037	<b>US</b> (unit separator)	63	3F	077	&#63;	<b>?</b>	95	5F	137	&#95;	<b>_</b>	127	7F	177	&#127;	<b>DEL</b>

Source: [www.LookupTables.com](http://www.LookupTables.com)

ASCII stands for American Standard Code for Information Interchange. Computers can only understand numbers, so an ASCII code is the numerical representation of a character such as 'a' or '@' or an action of some sort. ASCII was developed a long time ago and now the non-printing characters are rarely used for their original purpose. Below is the ASCII character table and this includes descriptions of the first 32 non-printing characters. ASCII was actually designed for use with teletypes and so the descriptions are somewhat obscure. If someone says they want your CV however in ASCII format, all this means is they want 'plain' text with no formatting such as tabs, bold or underscoring - the raw format that any computer can understand. This is usually so they can easily import the file into their own applications without issues. Notepad.exe creates ASCII text, or in MS Word you can save a file as 'text only'