

GRADING/ANSWER SHEET

Fall 2010 CS61C Final

Your Name: _____

Your TA: Andrew Michael Conor Charles

Login: cs61c-__

This exam is worth 186 points, or about 20% of your total course grade. The exam contains 10 questions.

This booklet contains 14 numbered pages including the cover page, plus the 2 pages for the green card. Put all answers on these pages, please; don't hand in stray pieces of paper.

POSSIBLE MAX MAX ACTUAL AVERAGE MIN

R, Ch
M
M
D
D
Co
Ch
A
R/A
R

| Question | Points (Minutes) | Score |
|-----------------------------|--------------------------------|------------------|
| 1. Justified | 30 points (15 minutes) | 20.5 |
| 2. Virtuosity | 14 (7) | 9.8 |
| 3. Mo' Cache, Mo' Problems | 24 (12) | 18.1 |
| 4. My Secret Cache | 34 (17) | 20.5 |
| 5. Heaven's Gate | 8 (4) | 7.4 |
| 6. Off the Map | 20 (10) | 12.7 |
| 7. Pay It Forward | 12 (6) | 8.3 |
| 8. Bigger, Stronger, Faster | 10 (5) | 6.3 |
| 9. Altered States | 10 (5) | 5.6 |
| 10. Three's Company | 24 (12) | 17.0 |
| Total | 186 points (93 minutes) | 127.2 |

TOTAL 177 132 42
MAX AVG MIN

1. **Justified.** True or False with Justification

Circle whether each statement in questions a) to g) is True or False, and **give a one sentence justification**. We reserve the right to dock points for long-winded answers. Be concise! If you feel like you can give a clear justification in just one word, please do!

3 pts/question; -2 if explanation partially correct
-1 if explanation "close"

a. Designing a processor with dynamic multiple issue places extra burden on the compiler-author to ensure the correctness of emitted code.

TRUE / FALSE

Justification:

The instruction set doesn't change because of multiple issue

b. As a result of hitting the Power Wall, Moore's Law no longer holds.

TRUE / FALSE

Justification:

Moore's Law is about transistor count, not clock rates, which are what is limited by the Power Wall.

c. Increasing the depth of a pipeline tends to increase performance primarily by decreasing the amount of time it takes on average to execute a particular instruction.

TRUE / FALSE

Justification:

(Accepted either T/F due to ambiguity about "takes on average to execute")
Decreases time between instructions → increased throughput
but not decreased latency

d. In a single-cycle MIPS CPU during the execution of a jump instruction, the adder block inside the ALU always computes the sum of its inputs.

TRUE / FALSE

Justification:

This hardware is still powered and has inputs, even though its output is not used.

e. After enough time, a circuit composed of combinational logic without feedback will always show the same output for a given set of inputs.

TRUE / FALSE

Justification:

This is the definition of combinational logic.
(i.e. It is stateless)

- f. For pipelined execution, if multiple exceptions happen in the same clock cycle for different instructions, the hardware should set the EPC to address of the earliest issued instruction.

TRUE / FALSE

Justification:

(Assuming "earliest issued" means "earliest issued among those w/ exceptions")
(snorky out-of-order issue.)

Exception handlers should be invoked as if instructions are executed one at a time.

- g. To get good performance from future microprocessors in the next 5 to 10 years, SIMD via successors to SSE4 instructions could be as important to use as MIMD via more cores.

TRUE / FALSE

Justification:

SIMD widths are growing nearly as fast as core counts.

- h. A particular memory access results in a page fault but not a protection error. Which of the following statements about this access is always true? Circle the answer.

1 pt / each

Always true / Not always true There was a TLB miss.

Always true / Not always true The cache(s) will miss.

Always true / Not always true Space for the required page is allocated somewhere on disk.

→ (accepted "not always true" w/ explanation)

- i. Which of the following could be enabled by adding more ALUs to a processor design? (Check all that apply.)

More structural hazards

More issue slots

Wider SIMD width

Larger CPI

-1 for each wrong (max -3)

- j. Multiple Choice. (Pick one.) Virtual machines are used in Cloud Computer by Amazon to

I. Protect users from each other

II. Offer more price points

III. Simplify introduction of new and faster hardware

A. None of the above

B. I only

C. II only

D. III only

E. I and II

F. II and III

G. I and III

H. I, II, and III

(-1 for each of I, II, III not identified)

5 pts each.

2. Virtuosity

A computer has a 32-bit virtual address with 8 KB pages and a 16 KB cache with 32 Byte blocks that is 2-way set associative

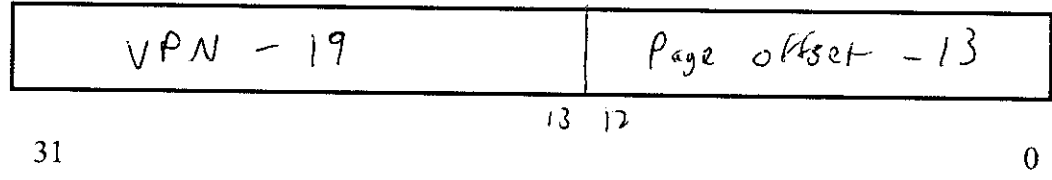
-2 pts: Off by one error.

-2 pts: Wrong ordering of fields

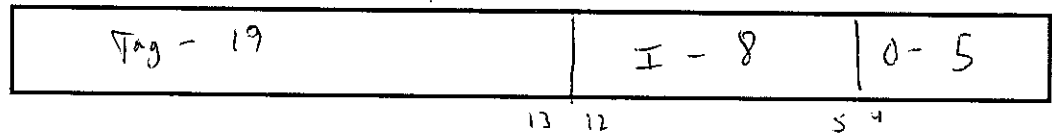
-2/3 points: Field completely wrong.

Only deduct one for common error between parts.

IV. Show how the 32-bit address is divided into the virtual page number and page offset, showing the number of bits in each field:



V. Show how the 32-bit address could be divided into the cache tag, cache index, and block offset, showing the number of bits in each field:



31 → One pt per part for at least getting fields in right order.

VI. Given these layouts, must the lookup in the TLB to translate from virtual address to physical address be done sequentially before access the cache? Why or why not?

4 pts.

3 pts - Minor misstatement
2 pts - mostly correct, some inaccurate statements
1 pt - partly correct.

No. ~~The cache index field~~ The cache index field is contained within the page offset field, so we can retrieve the relevant blocks from the cache while we look up the translation for the VPN in the TLB. We can verify hit/miss on the blocks by using the PPN as the tag.

3. Mo' Cache, Mo' Problems

a) A naive hardware developer designs the cache protocol for a 2 CPU system, each CPU with a 2 block direct mapped cache, 1 byte blocks. Memory addresses are 32 bits (byte addressed). The cache policies are Allocate on write, and Write through.

Assume that both caches start out with all blocks invalid. Which of the following access patterns, executed independently from one another, will always yield correctly updated results? Assume that each Read/Write finishes completely during its time period.

4 pts each

| | CPU | Time 1 | Time 2 | Time 3 | Correct? |
|------|-----|-----------|-----------|----------|----------|
| i. | 1 | Write 0x0 | --- | Read 0x0 | N |
| | 2 | --- | Write 0x0 | --- | |
| ii. | 1 | Write 0x0 | --- | Read 0x1 | Y |
| | 2 | --- | Write 0x1 | --- | |
| iii. | 1 | Write 0x0 | --- | Read 0x0 | Y |
| | 2 | --- | Write 0x2 | --- | |
| iv. | 1 | Read 0x0 | --- | Read 0x2 | Y |
| | 2 | --- | Write 0x2 | --- | |
| v. | 1 | Read 0x0 | --- | Read 0x0 | N |
| | 2 | --- | Write 0x0 | --- | |

Incorrect when CPU 1 has a read hit, AND CPU 2 had written to

that mem's address earlier. The read hit will return stale data. (inconsistent with the most up to date copy in memory)

b) What feature could you add to this cache protocol to make it work in all cases? Be specific; your revision should describe a particular behavior.

Invalidate on write.

4 pts,

Partial credit possible.

RUBRIC

4 PTS FOR (a)
2 PTS / QUESTION IN (b)

4. My Secret Cache

The average memory access time (AMAT) is:

$$\text{AMAT} = \text{Hit Time}_{L1} + \text{Miss Rate}_{L1} \times \text{Miss Penalty}_{L1}$$

4 POINTS

(a) Write down the AMAT equation for a three level cache (use HT for Hit Time, MR for Miss Rate, and MP for Miss Penalty):

$$\text{HT}_{L1} + \text{MR}_{L1} \times (\text{HT}_{L2} + \text{MR}_{L2} \times (\text{HT}_{L3} + \text{MR}_{L3} \times \text{MP}_{L3}))$$

(4)

or
Memory access time

(b) Given the following specification:

For every 1000 CPU-to-memory references

40 will miss in L1\$;

20 will miss in L2\$;

10 will miss in L3\$;

L1\$ hits in 1 clock cycle;

L2\$ hits in 10 clock cycles;

L3\$ hits in 100 clock cycles;

Main memory access is 400 clock cycles;

There are 1.25 memory references per instruction; and

The ideal CPI is 1.

TRIED TO GIVE POINTS EVEN IF GAVE WRONG ANSWER ON EARLIER PARTS

Answer the following questions:

2/PARTS

(i) What is the local miss rate in the L2\$?

$$20/40 = 50\%$$

(2)

(ii) What is the global miss rate in the L2\$?

$$20/1000 = 2\% \text{ or } 4\% \times 50\% = 2\%$$

(2)

(iii) What is the local miss rate in the L3\$?

$$10/20 = 50\%$$

(2)

(iv) What is the global miss rate in the L3\$?

$$10/1000 = 1\% \text{ or } 50\% \times 2\% \times 50\% = 1\%$$

(2)

(v) What is the AMAT with all three levels of cache?

$$1 + 4\% (10 + 50\% \times (100 + 50\% \times 400)) = 1 + .4 + 2\% (100 + 200) = 1.4 + 6 = 7.4$$

(2)

(vi) What is the AMAT for a two-level cache without L3\$?

$$1 + 4\% (10 + 50\% \times 400) = 1 + .4 + 4\% \times 200 = 9.4$$

(2)

16

RUBERIC

(vii) What is the AMAT for a one-level cache without L2\$ and L3\$?

$$1 + 4\% (400) = 17 \quad 2$$

(viii) What is the average memory stalls per reference in the system of question (v)?

$$7.4 - 1 = 6.4 \quad 2$$

(ix) What is the average memory stalls per reference in the system of question (vi)?

$$9.4 - 1 = 8.4 \quad 2$$

(x) What is the average memory stalls per reference in the system of question (vii)?

$$17 - 1 = 16 \quad 2$$

(xi) What is the average memory stalls per instruction in the system of question (v)?

$$6.4 \times 1.25 = 8 \quad 2$$

5/4 4 | 32

(xii) What is the average memory stalls per instruction in the system of question (vi)?

$$8.4 \times 1.25 = 10.5 \quad 2$$

5/4 4 | 42

(xiii) What is the average memory stalls per instruction in the system of question (vii)?

$$16 \times 1.25 = 20 \quad 2$$

5/4

(xiv) What is the performance of the system of question (vii) versus that of question (v)?
(Long-division challenged can give just the ratio.)

$$\frac{16}{6.4} \approx 2.5 \quad 2$$

$$7.4 \overline{) 170} \begin{array}{r} 23 \\ 148 \\ \hline 220 \end{array} \quad 2$$

(xv) What is the performance of the system of question (vi) versus that of question (v)?
(Long-division challenged can give just the ratio.)

$$\frac{8.4}{6.4} = 1.3 \quad 2$$

$$8.4 \overline{) 109.2} \begin{array}{r} 13 \\ 84 \\ \hline 252 \\ \hline 200 \end{array} \quad 2$$

(18)

RUBEN C

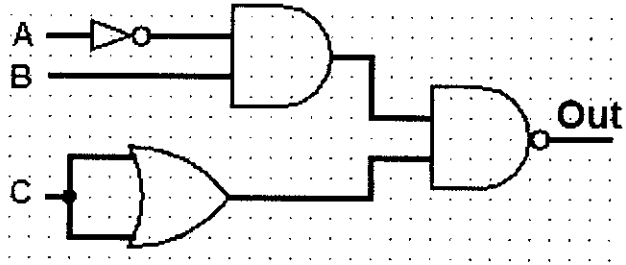
5. Heaven's Gate

What is the truth table for the following circuit?
(It might be simpler to first write the equation)

8 POINTS

$$\overline{(\overline{A} \cdot B) \cdot C} =$$

$$A + \overline{B} + \overline{C}$$



| A | B | C | Out |
|---|---|---|-----|
| 0 | 0 | 0 | 1 |
| 0 | 0 | 1 | 1 |
| 0 | 1 | 0 | 1 |
| 0 | 1 | 1 | 0 |
| 1 | 0 | 0 | 1 |
| 1 | 0 | 1 | 1 |
| 1 | 1 | 0 | 1 |
| 1 | 1 | 1 | 1 |

~~8 POINTS~~
1 POINT PER CORRECT ANSWER, UNLESS (MISTAKEN) (DID AND US AND), THEN GET 5 POINTS OUT OF 8

6. Off the Map

Given:

```
typedef struct llnode {
    struct llnode *next;
    int contents;
} llnode;
```

Implement map in C. map should take a pointer to an llnode and a map_fun, which is just a pointer to a function taking one int argument and returning int. It should return a pointer to a **new** linked list, where the contents of each of the new nodes is equal to the value returned by the map function called on the contents of the corresponding old node. Don't forget to fill in the return value correctly (pointers to function are tricky so we've done the arguments for you). Assume llnode lists are terminated by a NULL pointer. You should **not** modify any part of the list identified by head. You may assume map_fun has been properly declared by a typedef somewhere else.

```
llnode map(llnode *head, map_fun f)
{
    if ( !head ) {
        return NULL;
    }
    llnode *new = malloc(sizeof(llnode));
    new->contents = f(head->contents);
    new->next = map(head->next, f);
    return new;
}
```

(Note that to call a function pointer, you just use it like any other function. So, in the code template given above, saying `int result = f(3);` would correctly call the function pointed to by `f` with the value 3 and put the return value into `result`.)

Return type - 2

Base case - 4

Malloc - 6

Setting contents - 2

Recursive call - 4

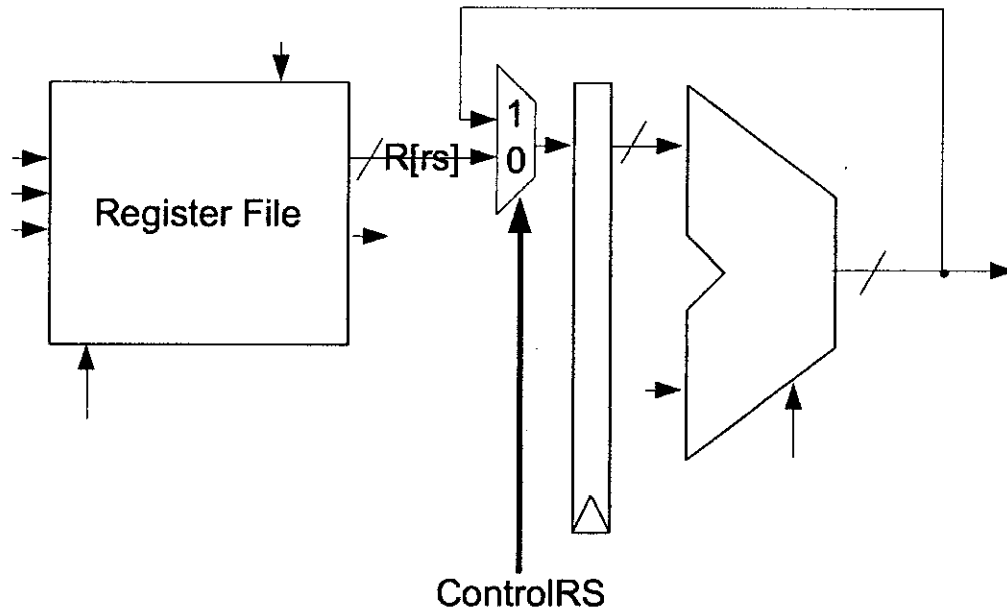
Return value - 2

iterative OK

(2 points max if
no malloc)

7. Pay It Forward

Consider the *excerpt* below of a 5-stage pipelined MIPS datapath.



a. Consider the following sequence of instructions

5 pts -

- [A] srl \$zero, \$zero, 0
- [B] addu \$t0, \$t1, \$t2
- [C] addu \$t0, \$t0, \$t2
- [D] lw \$s0, 0(\$t3)
- [E] subu \$t3, \$s0, \$t0

During which of these instructions' decode stages in the sequence above should ControlRS be 1 to avoid pipeline stalls? Use the labels [A], [B], ...

[C] - 3 if [C] omitted
 - 2 if others deleted
 - 1 for [B] alone (wrong stage)

b. Which fields of which instructions from part a does the control logic need to compute the value of ControlRS?

7 pts -

opcode of all instructions (to detect destination, R/I-type)
 rd of R-type instructions
 rs of instructions

- 1 pt - rd of R-type instructions

- 2 pts - no rd

- 2 pts - no rs

- 3 pts - no opcode / selecting by instruction type
 (2 pts if selection by instr. type w/o mentioning field's used)

8. Bigger, Stronger, Faster:

Suppose that you are running an algorithm for various problem sizes, and have obtained the data below. Sketch a weak scaling plot of parallel code performance that shows speedup over the serial implementation. Be sure to label the Y-axis.

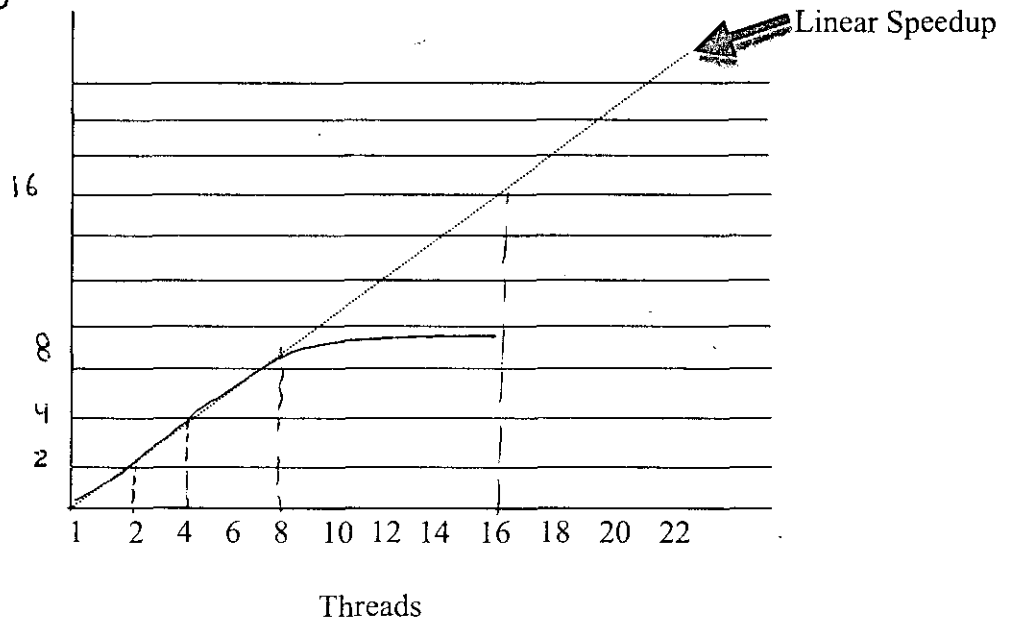
| Problem Size | Gflop/s (serial) | Threads | Gflop/s (parallel) |
|--------------|------------------|---------|--------------------|
| 100 | 5 | 1 | 5 |
| 200 | 5 | 2 | 10 |
| 400 | 5 | 4 | 19 |
| 600 | 5 | 6 | 25 |
| 800 | 5 | 8 | 35 |
| 1000 | 5 | 10 | 36 |
| 1200 | 5 | 12 | 37 |
| 1400 | 5 | 14 | 37 |
| 1600 | 5 | 16 | 38 |

Rubric

- 5 y-axis nonexistent/ wrong
- 3 plot is above linear line

Speedup
 $\left(\frac{\text{parallel}}{\text{serial}} \right)$

Weak Scaling of Speedup over Serial

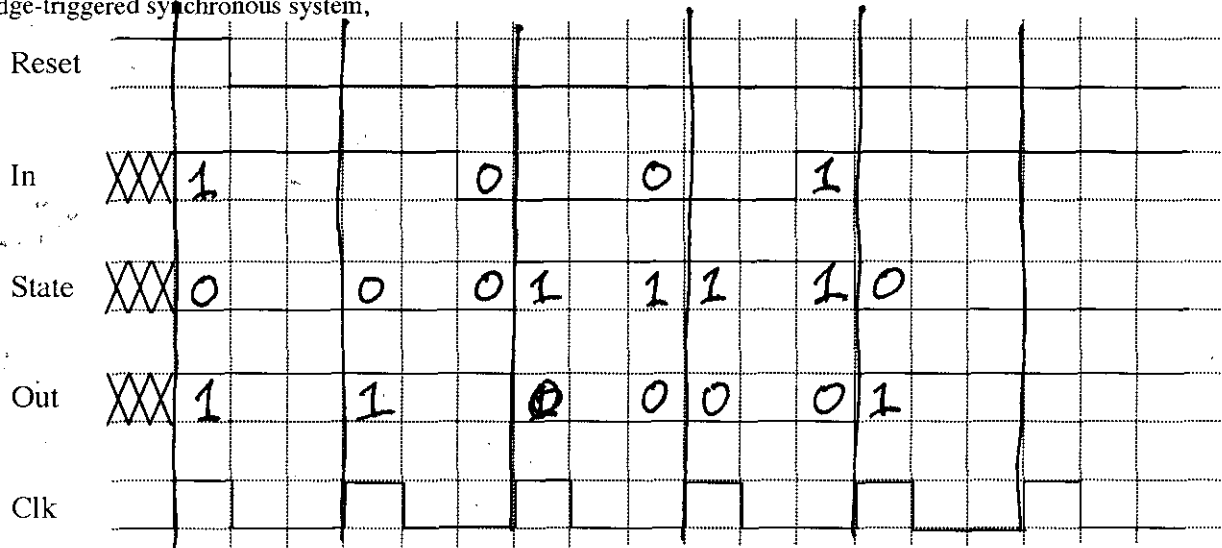


~~3 axes~~
~~5 axes~~

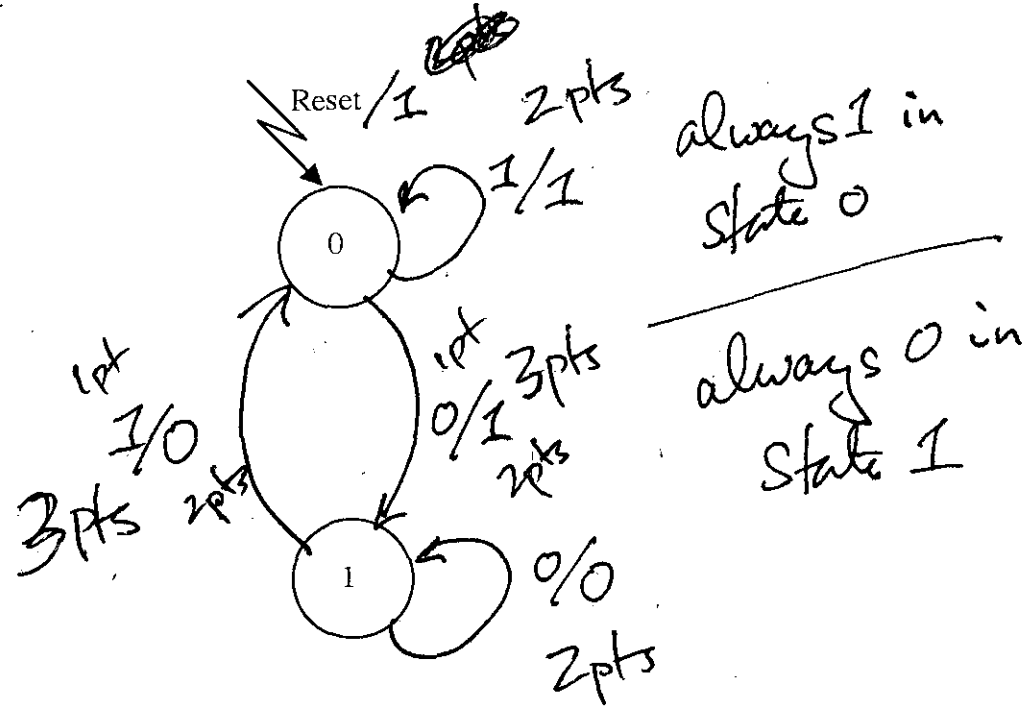
9. ALTERED STATES

Question State Diagram 10 pts

Given the following timing chart, and assuming 0 set-up, hold, and propagation times, as well as a positive edge-triggered synchronous system,



Fill-in the transitions and outputs of the following State Diagram so it has the identical behavior to the timing diagram above:



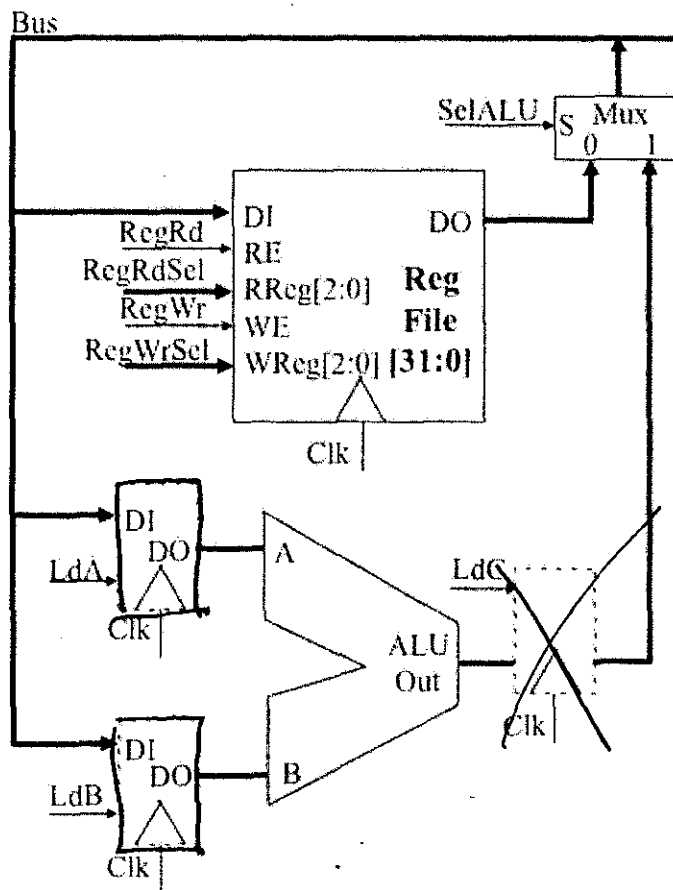
THREE'S COMPANY

Question Threesy 24 pts

Consider the following datapath with an Arithmetic Logic Unit (ALU) and an eight-register register file organized around a single bus. The ALU is to apply add, subtract, etc. operations to its two input operands to generate an output result. The register file has an asynchronous read and a synchronous write. That is, as soon as the Read Enable (RE) is asserted, the register file selects the indicated 32-bit register and presents its value on the Data Out (DO). On the other hand, the Write Enable (WE) is sampled only on the rising edge of the clock, and only writes the indicated register from the Data In (DI) lines on the same edge that WE is asserted. The ALU and Register File share the Bus via a 2:1 multiplexer. When SelALU is set to 1, the ALU path is connected to the Bus. Otherwise the Register File path is connected to the Bus.

The datapath must support three-address instructions of the form $R_z \leftarrow R_x \text{ <op> } R_y$.

To make use of a single bus architecture, the ALU can be surrounded by one, two, or three 32-bit temporary registers, labeled A, B, and C, as shown below (the temporary registers are shown as dotted lines – the correct solution requires at least one and possibly all three of the registers):



Using the fewest of the A/B/C registers and possible clock cycles, what is the fewest number of each to implement the register transfer for the instructions of the three-address type (circle one for each):

| | | | | |
|------------------------------|-----------|---|--------------|---|
| 1 4 | Registers | 1 | ② | 3 |
| | Cycles | 1 | 2 | ③ |

For your answer, on the previous cross out the registers you don't need, and fill-in the outline of the registers that you do. For each cycle you need, write below the control signals that must be asserted to implement the register transfers for the three-address instructions:

4
 Cycle 1: $SEL_{ALU} \leftarrow 0, \cancel{RegWr} \leftarrow 1, RegRd \leftarrow 1, RegRdSel \leftarrow X, LD_A,$
 (during cycle 1, $RegFile_x \rightarrow TempA$)

4
 Cycle 2: $SEL_{ALU} \leftarrow 0, RegRd \leftarrow 1, RegRdSel \leftarrow Y, LD_B,$
 (during cycle 2, $RegFile_y \rightarrow TempB$)

4
 Cycle 3:
 $SEL_{ALU} \leftarrow 1, RegWr \leftarrow 1, RegWrSel \leftarrow Z,$
 (ACU $\leftarrow OP$ Not shown on diagram
 (during Cycle 3, $ALU \rightarrow RegFile_z$)