



## M1) *Hacker's Delight* (10 pts, 20 min)

We present `mystery`, a new helper routine for your C programming.  
In parts (a) and (b), you'll show us you know how to use `mystery` from C.  
In parts (c) and (d), you'll show us you understand its limitations.

```
mystery:  ori    $v0 $0 0x0
          beq    $a0 $0 done
          la     $t0 mystery
          lw     $t1 0($t0)
          addiu $t1 $t1 0x1
          sw     $t1 0($t0)
          lw     $v0 0($a0)
done:     jr     $ra

          main() {
            char A[4], char4 = '4';
            int pi[] = { 3, 1, 4, 1, 5, 9 }
            float float4 = 4;
            // part (a)
            ... // more code, function calls, etc.
          }
```

- a) If you're at "`part (a)`" in the C code, show a single call to `mystery` so that it returns 4.

```
printf("Here is mystery returning four ... %d\n", mystery(_____));
```

- b) Complete the documentation of `mystery` for a fellow programmer. Use good abstraction – don't tell the user *how* it does what it does, just tell them *what* it does and *how* it's to be used.

"When called with a *non-NULL* argument `arg`, the subroutine `mystery` ...

\_\_\_\_\_".

"When called with a *NULL* argument, the subroutine `mystery` ...

\_\_\_\_\_".

"Overall, `mystery` is a subroutine used to ...

\_\_\_\_\_".

- c) We'd like to know if there is a limit to the # of times `mystery` can be called with a *NULL* argument (so that it still does what you described in part b). If there *is*, state what the limit is and what happens if it's called one more time. If there *isn't* a limit, write N/A (not applicable) in both blanks.

"With a *NULL* argument, `mystery` may be called at most \_\_\_\_\_ times. Calling it once more...

\_\_\_\_\_".

- d) We'd like to know if there is a limit to the # of times `mystery` can be called with a *non-NULL* argument (so that it still does what you described in part b). If there *is*, state what the limit is and what happens if it's called one more time. If there *isn't* a limit, write N/A in both blanks.

"With a *non-NULL* argument, `mystery` may be called at most \_\_\_\_\_ times. Calling it once more...

\_\_\_\_\_".

## M2) Cache Money, y'all (10 pts, 20 min)

This C code runs on a 32-bit MIPS machine with 4 GiB of memory and a single L1 cache. Vectors **A**, **B** live in different places of memory, are of equal size (**n** is a power of 2 and a [natural number] multiple of the cache size), block aligned. The size of the cache is **C**, a power of 2 (and always bigger than the block size, obviously).

<pre>// sizeof(uint8_t) = 1 SwapLeft(uint8_t *A, uint8_t *B, int n) {     uint8_t tmp;     for (int i = 0; i &lt; n; i++) {         tmp = A[i];         A[i] = B[i];         B[i] = tmp;     } }</pre>	<pre>// sizeof(uint8_t) = 1 SwapRight(uint8_t *A, uint8_t *B, int n) {     uint8_t tmpA, tmpB;     for (int i = 0; i &lt; n; i++) {         _____         _____         _____         _____     } }</pre>
--	---

Let's first just consider the **SwapLeft** code for parts (a) and (b).

- a) If the cache is **direct mapped** and the *best* hit:miss ratio is "H:1", what is the block size in bytes? \_\_\_\_\_
  
- b) What is the *worst* hit:miss ratio? \_\_\_\_\_:\_\_\_\_\_
  
- c) Fill in the code for **SwapRight** so that it does the same thing as **SwapLeft** but improves the (b) hit:miss ratio. You may not need all the blanks.
  
- d) If the block size (in bytes) is *a*, what is the *worst* hit:miss ratio for **SwapRight**? \_\_\_\_\_:\_\_\_\_\_
  
- e) We next change the cache to be **2-way set-associative**, and let's go back to just considering **SwapLeft**. What is the **worst** hit:miss ratio for **SwapLeft** with the following replacement policies? The cache size is *C* (bytes), the block size is *a* (bytes), LRU = Least Recently Used, MRU = Most Recently Used.

LRU and an empty cache	MRU and a full cache
_____:	_____:

**M3) What is that Funky Smell? Oh, it's just Potpourri... (10 pts, 20 mins)**

- a) How many non-negative floats are  $< 2$ ? \_\_\_\_\_ (you must show your work above for credit)
- b) What's the biggest change to the PC as the result of a jump on a 32-bit MIPS system? \_\_\_\_\_ (answer in IEC format, like 16 kibibytes or 128 gibibytes)

c) Fellow 61C student Ben Bitdiddle was told to write a function `count_az` that takes an input string of lower-case letters (only 'a' through 'z') and returns an array of the number of occurrences of all letters; a histogram if you will. The returned array will be zero-indexed and the indices will correspond to their respective order in the alphabet (i.e. a = 0, b = 1, ..., z = 25). E.g., if the input `str` is "baaadd", the output array will look like the right column of the table on the right. Fix all the errors; we should be able to call it like this: `myAZ = count_az(str); yourAZ = count_az(str);`

index	count[index]
0 (for 'a')	3
1 (for 'b')	1
2 (for 'c')	0
3 (for 'd')	2
...	...
25 (for 'z')	0

```

1 int count_az(char *str) {
2
3     int count[26];           // Create the count array
4
5     while(*str) {           // Go through the whole string
6
7         int index = &str - 0x97; // The 97 is from the MIPS green sheet..
8
9         count[index]++;      // Increment the appropriate bucket
10
11        str++;               // Go to the next character
12
13    }
14
15    free(str);               // Free the string storage
16
17 }
```

Line #	Add Change Remove	Additions / Changes / Removals



## F2) V(I/O)rtual Potpourri ... (23 pts, 30 mins)

For the following questions, assume the following:

- 32-bit virtual addresses
- 1 MiB pages
- 512 MiB of physical memory with LRU page replacement policy
- Fully associative TLB with 32 entries and an LRU replacement policy

a) How many entries does a page table contain? \_\_\_\_\_

b) How wide is the *page table base register*? \_\_\_\_\_

```
int histogram[MAX_SCORE];
void update_hist(int *scores, int num_scores) {
    for (int i = 0; i < num_scores; i++)
        histogram[scores[i]] += 1;
}
```

Assume that only the code and the two arrays take up memory, ALL of code fits in 1 page, the arrays are page-aligned (start on page boundary), and this is the only process running.

c) If `update_hist` were called with `num_scores = 10`, how many page faults can occur in the worst-case scenario? \_\_\_\_\_

d) In the best-case scenario, how many iterations of the loop can occur before a TLB miss? You can leave your answer as a product of two numbers. \_\_\_\_\_

e) For a particular data set, you know the scores are clustered around fifty different values, but you still observe a high number of TLB misses during `update_hist`. What pre-processing step could help reduce the number of TLB misses?

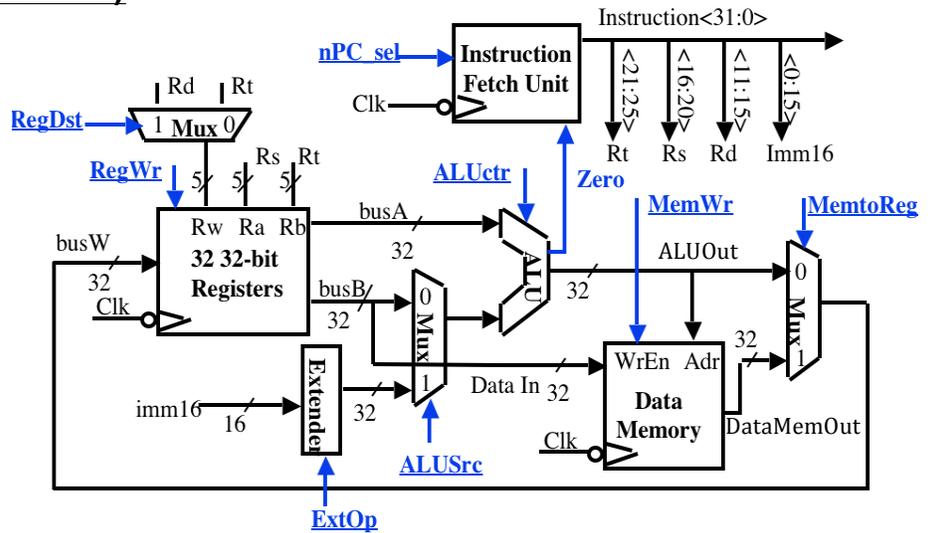
### F3) Datapathology ... (22 pts, 30 mins)

Consider the single cycle datapath as it relates to a new MIPS instruction, memory add:

```
madd rd, rs, rt
```

The instruction does the following:

- 1) Reads the value of *memory* at the address stored in *rs*.
- 2) Adds the value in the register specified by *rt* to the memory value and stores the resulting value in *rd*.



Ignore pipelining for parts (a)-(c).

a) Write the Register Transfer Language (RTL) corresponding to `madd rd, rs, rt`

b) Change as *little as possible* in the datapath above (**draw your changes right in the figure**) to enable `madd`. List all your changes below. Your modification may use muxes, wires, constants, and new control signals, but nothing else. (You may not need all the provided boxes.)

(i)	.
(ii)	
(iii)	

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

RegDst	RegWr	nPC_sel	ExtOp	ALUSrc	ALUctr	MemWr	MemtoReg			

d) Briefly (one sentence) explain why `madd` CANNOT be run on the standard 5-stage MIPS pipeline.

e) Let's fix our datapath so we can pipeline `madd`! If each stage currently takes 100 ps and an additional ADDER will have a delay of 50 ps, fill in the following table for two possible solutions: (1) a 6-stage pipeline and (2) a 5-stage pipeline with an extended MEM stage. Assume a memory access takes the full 100 ps. Leave your throughput answers as fractions.

	6-stage pipeline	Extended MEM
Latency:		
Throughput:		

#### F4) What do you call two L's that go together? (22 pts, 30 mins)

Suppose we have `int *A` that points to the head of an array of length `len`. Below are 3 different attempts to set each element to its index (i.e. `A[i]=i`) using OpenMP with `n>1` threads. Determine which statement (A)-(E) correctly describes the code execution and provide a one or two sentence justification.

**Answers without a one or two sentence justification will receive NO credit.**

a) 

```
#pragma omp parallel for
for (int x = 0; x < len; x++){
    *A = x;
    A++;
}
```

- A) Always **Incorrect**
- B) Sometimes **Incorrect**
- C) Always **Correct**, Slower than Serial
- D) Always **Correct**, Speed relative to Serial depends on Caching Scheme
- E) Always **Correct**, Faster than Serial

Justification:

b) 

```
#pragma omp parallel
{
    for (int x = 0; x < len; x++){
        *(A+x) = x;
    }
}
```

- A) Always **Incorrect**
- B) Sometimes **Incorrect**
- C) Always **Correct**, Slower than Serial
- D) Always **Correct**, Speed relative to Serial depends on Caching Scheme
- E) Always **Correct**, Faster than Serial

Justification:

c) 

```
#pragma omp parallel
{
    for(int x = omp_get_thread_num(); x < len; x += omp_get_num_threads()) {
        A[x] = x;
    }
}
```

- A) Always **Incorrect**
- B) Sometimes **Incorrect**
- C) Always **Correct**, Slower than Serial
- D) Always **Correct**, Speed relative to Serial depends on Caching Scheme
- E) Always **Correct**, Faster than Serial

Justification:

#### F4) What do you call two L's that go together? (Continued)

The moving average (a type of low pass filter) is an operation commonly used to smooth noisy data. Here we compute a *centered moving average* of width **WIDTH** on an array of data of size **SIZE**, where each element in our output array is the average of the *current* element, the *previous*  $(\text{WIDTH}-1)/2$  elements, and the *next*  $(\text{WIDTH}-1)/2$  elements. **Assume that WIDTH is odd for simplicity** and use zeros where “required” elements do not exist.

Example Input of `SIZE`: `float[] A = [ 7, 2, 3, 4, 8, 6 ]`  
Output for `WIDTH=3`: `float[] result = [ 3, 4, 3, 5, 6, 4.6666667 ]`

Fit this problem to the MapReduce paradigm using a single map and reduce by filling in the blanks below. You may assume that you have access to the global variables **WIDTH** and **SIZE**. We expect you to use C syntax with the addition of a few java-like pseudocode elements (e.g. arrays have `.length`).

```
// receives data one element at a time
// Inputs: (key) is index i, (value) is A[i]
map(int key, float value){
    _____ {
        context.write( x, value);
    }
}

// outputs elements of centered moving average
// Outputs MUST be of the form:
// (key) is index i
// (value) is moving average of width WIDTH centered at i
reduce(int key, float[] values){
    float total = 0;

    // do not emit keys that do not exist in output array
    if ( (key >= 0) && (key < SIZE) ) {
        _____
        _____

        context.write( _____ , _____ );
    }
}
```