

# COS 217, Spring 2022

## Final Exam

This exam consists of 7 questions, and you have 180 minutes – budget your time wisely. *Do all of your work on these pages (using provided blank pages for scratch space) and give your answers in the space provided.* Assume the ArmLab/Linux/C/gcc217 environment unless otherwise stated.

This is a closed-book, closed-note exam, with a 1-page study sheet *allowed*. Please place items that you will not need out of view in your bag or under your working space at this time. Electronic devices such as cell phones, laptops, music players, etc. may not be used during this exam.

Name:  NetID:

Precept (*circle one*): 1: MW 1:30 Christopher Moretti 4: TTh 1:30 Huihan Li  
2: MW 3:30 Christopher Moretti 7: TTh 3:30 Cedrick Argueta  
3: TTh 12:30 Maxine Perroni-Scharf

This examination is administered under the Princeton University Honor Code. Students should sit one seat apart from each other, and refrain from talking to other students during the exam. All suspected violations of the Honor Code must be reported to [honor@princeton.edu](mailto:honor@princeton.edu).

*In the box below, write out and sign the Honor Code pledge before turning in the test:*

*“I pledge my honor that I have not violated the Honor Code during this examination.”*

Pledge and Signature:

**Exam statistics:**  
**Mean 94.2/120 (78.5%) StdDev: 16.7/120**  
**Max 118/120. Percentiles: 25<sup>th</sup> – 86/120; 50<sup>th</sup> (median) – 98/120; 75<sup>th</sup> – 106/120.**

**For instructor use only:**

Question	Max Points	Points Earned	Question	Max Points	Points Earned
1	15		5	25	
2	20		6	5	
3	18		7	15	
4	22		TOTAL	120	

**(blank page)**

## Q1. Truthiness and Falsiness (15 points)

Circle **T** (True) or **F** (False) for each statement.

	Answer	Statement
<b>A</b>	<input checked="" type="radio"/> T <input type="radio"/> F	The pre-processor adds the contents of other files to a C file. <i>#include directives add header file contents to the file being preprocessed.</i>
<b>B</b>	<input checked="" type="radio"/> T <input type="radio"/> F	The compiler determines which variables go into the rodata, data, and bss sections.
<b>C</b>	<input type="radio"/> T <input checked="" type="radio"/> F	The compiler can detect memory leak errors in the code.
<b>D</b>	<input type="radio"/> T <input checked="" type="radio"/> F	The assembler generates assembly language code. <i>The compiler generates assembly code. The assembler generates machine code.</i>
<b>E</b>	<input type="radio"/> T <input checked="" type="radio"/> F	The linker creates relocation records. <i>The assembler creates relocation records. The linker processes relocation records.</i>
<b>F</b>	<input checked="" type="radio"/> T <input type="radio"/> F	The linker patches object code with virtual memory addresses.
<b>G</b>	<input type="radio"/> T <input checked="" type="radio"/> F	The expression $0xFF - 0xEE$ evaluates to 11 in decimal. <i><math>0xFF - 0xEE = 0x11</math>. <math>0x11</math> is 17 in decimal.</i>
<b>H</b>	<input checked="" type="radio"/> T <input type="radio"/> F	If <code>char s[5] = "mom";</code> then the expression <code>*(s+2) - *s</code> evaluates to 0.
<b>I</b>	<input type="radio"/> T <input checked="" type="radio"/> F	The expression <code>sizeof(7UL) &gt; sizeof(1L)</code> evaluates to 1 (true). <i>The sizes are equal. (Both 8 bytes on armlab)</i>
<b>J</b>	<input type="radio"/> T <input checked="" type="radio"/> F	AArch64 is a <i>Complex Instruction Set Computer</i> (CISC) architecture. <i>ARM is RISC: Reduced Instruction Set Computer</i>
<b>K</b>	<input checked="" type="radio"/> T <input type="radio"/> F	The <i>stack memory</i> can run out of space due to an unbounded number of function calls.
<b>L</b>	<input type="radio"/> T <input checked="" type="radio"/> F	<i>Statement</i> testing can require more test cases than <i>Path</i> testing.
<b>M</b>	<input type="radio"/> T <input checked="" type="radio"/> F	<i>Stress</i> testing is re-running all test cases after fixing a bug. <i>This is regression testing.</i>
<b>N</b>	<input type="radio"/> T <input checked="" type="radio"/> F	A <i>dangling pointer</i> is when you fail to free some chunk of allocated memory. <i>This is a memory leak.</i>
<b>O</b>	<input checked="" type="radio"/> T <input type="radio"/> F	An optimizing compiler can perform <i>function inlining</i> to improve performance.

## Q2. Is this in scope? (20 points)

Consider the following C code in `file1.c`:

```
int a;
int b = 0;
static int c = 1;

static int DoubleIt(int input) {
    int d = 2;
    static int e = 3;
    static int f;
    int g = d * input;
    return g;
}
```

(a) Fill in the blanks in the following table: (15 points)

variable	Scope	Linkage	Duration
a	file	external	process
b	file	external	process
c	file	internal	process
d	block	internal	temporary
e	block	internal	process
f	block	internal	process
g	block	internal	temporary

(b) Which section of memory is `f` in? (2 points)

BSS, because it is a process duration variable that was not initialized in its declaration.

We now create a second C file, `file2.c`, which includes the following declarations outside any function, plus a main function (not shown):

```
extern int a;  
int b = 0;  
static int c = 2;
```

(c) When we build `file1.c` and `file2.c` (and no other files) together into a program, which of the three lines from `file2.c` will cause an error? (1 point)

`int b = 0;`

(d) What is the problem? (1 point)

This would result in two different externally linked definitions of `b`.  
The other two are fine: `a` is a declaration that there is an externally linked definition that exists somewhere (it does, in `file1.c`); `c` is internally linked as is `file1.c`'s `c`, so there is no conflict.

(e) Which stage of the build process will report the error? (1 point)

Linker.

### Q3. Modular Table (18 points)

Consider the following excerpts from an interface (top box) and implementation of that interface (bottom box) for a variant of the symbol table module you wrote in Assignment 3:

Interface:

```
/* A SymTable_T is a set of key-value pairs */
typedef struct SymTable* SymTable_T;

/* instantiate and return a new empty
SymTable_T object with k buckets, or NULL
if there is insufficient memory */
SymTable_T ST_new(size_t k);
```

Implementation:

```
typedef struct binding* Binding_T;
struct binding {
    char* key;
    void* val;
    Binding_T next;
};
struct SymTable {
    Binding_T* buckets;
    size_t count;
    size_t numBuckets;
};
SymTable_T ST_new(size_t k) {
    /* implementation not shown */
}
```

- (a) The interface (top box) above violates an ADT design principle from the Modularity lecture. What is the problem found in the interface above? (1 point)

The function comment for `ST_new` partially reveals the underlying implementation of the object (that it maintains the object's contents in `k` buckets)

- (b) The following lines of code make up the body of a new interface function `ST_maxValue`, which takes an instance pointer `st` and a comparison function (with `strcmp` semantics) `comp` and returns the value (not the key!) in the symbol table with the maximal value based on the comparison function. These lines, however, have become scrambled and have lost their indentation. In the 13 blanks below, write the line numbers in their proper order that results in a correct implementation of this function. (13 points)

```
1    } /* end while */
2    if(max) return max->val;
3    size_t i;
4    Binding_T c = st->buckets[i];
5    return NULL;
6    if(!max || (*comp)(c->val, max->val) > 0)
7    } /* end for */
8    while(c != NULL) {
9    Binding_T max = NULL;
10   c = c->next;
11   for(i = 0; i < st->numBuckets; i++) {
12   max = c;
13   assert(st != NULL && comp != NULL);
```

3 or 9 9 or 3 13 11 4 8 6 12 10 1 7 2 5

\_\_\_\_\_

- (c) What is the type of the `comp` parameter in the `ST_maxValue` function? (2 points)

`int (*)(void*, void*)`

- (d) Now consider a representation invariant for the symbol table that each bucket is stored in descending order (i.e., for any valid symbol table, the largest value in each bucket is at the head of that list and the smallest value is at the tail of that list). Describe in two sentences or less a simplification that can be made to the `ST_maxValue` function above by taking advantage of this constraint. (2 points)

We would not need the while loop through all nodes in each bucket, because the head of each bucket is our only candidate for the max element.

#### Q4. Buggy Parents (22 points)

The function `getParentPath` should return a defensive copy of a string representing the *parent* of a `path` parameter like in Assignment 4. If `path` is the root, the string returned should be the empty string.

You should assume that all relevant header files have been `#included`. You may assume the `path` passed in as an argument is well formed by the standards of Assignment 4: it is a non-empty `'\0'`-terminated string that neither begins nor ends with a `'/'` character and does not contain consecutive `'/'` delimiters.

Here are some concrete examples of the function argument and resulting return values:

`path: "cos"`, i.e. `{'c', 'o', 's', '\0'}` should return a new string with contents: `""`, i.e. `{'\0'}`

`path: "cos/217"` should return a new string with contents: `"cos"`

`path: "cos/217/forever!/Yay:)"` should return a new string with contents: `"cos/217/forever!"`

Now consider this implementation:

```
1 char* getParentPath(const char* path) {
2     char* parent;
3     char* end = path;
4     assert(path != NULL);
5     end += strlen(path);
6     while(*end != '/') end--;
7     parent = malloc(end - path);
8     strncpy(parent, path, end - path);
9     return parent;
10 }
```

As a reminder, the `string.h` interface function `strncpy(dest, src, n)` copies the first `n` bytes of `src` to `dest`. If `src`'s length is less than `n`, `strncpy` writes null bytes to `dest` until `n` bytes have been written.

(a) Unfortunately, the implementation above is buggy!

For this problem, the following count as bugs: warnings or errors from `gcc217`, runtime crashes, incorrect behavior, inefficiency as judged in Assignments 2 and 3 (i.e., extra traversals beyond what is required), or dynamic memory management issues observable by `MemInfo` or `Valgrind`.

In the spaces below, identify 3 *distinct* bugs (there may be more than that in the code, and they could be bugs of any kind, but you should provide only 3). To identify each bug, list the line number(s) where the bug occurs and a one-sentence description of the bug. (9 points)

**Bug 1:** Here are several correct answers:

1. `path` is a `const char*`, `end` is a `char*`, so line 3 will result in a compiler warning.
2. if `path` is the root, it has no `/` characters, so line 6 will traverse off the front of the string

**Bug 2:** 3. line 7 allocates one too few bytes of memory (the new string must end in a nullbyte)

4. Between lines 7-8, there is no check that `malloc` succeeded
5. `strncpy` in line 8 does not copy a nullbyte into `parent`, so we must do so manually after

**Bug 3:** 6. the algorithm on lines 5, 6, 8 will traverse the entirety of `path` twice. It is possible to keep track of the last `/` on the first traversal and thus only have to traverse the parent portion of `path` in the second traversal.



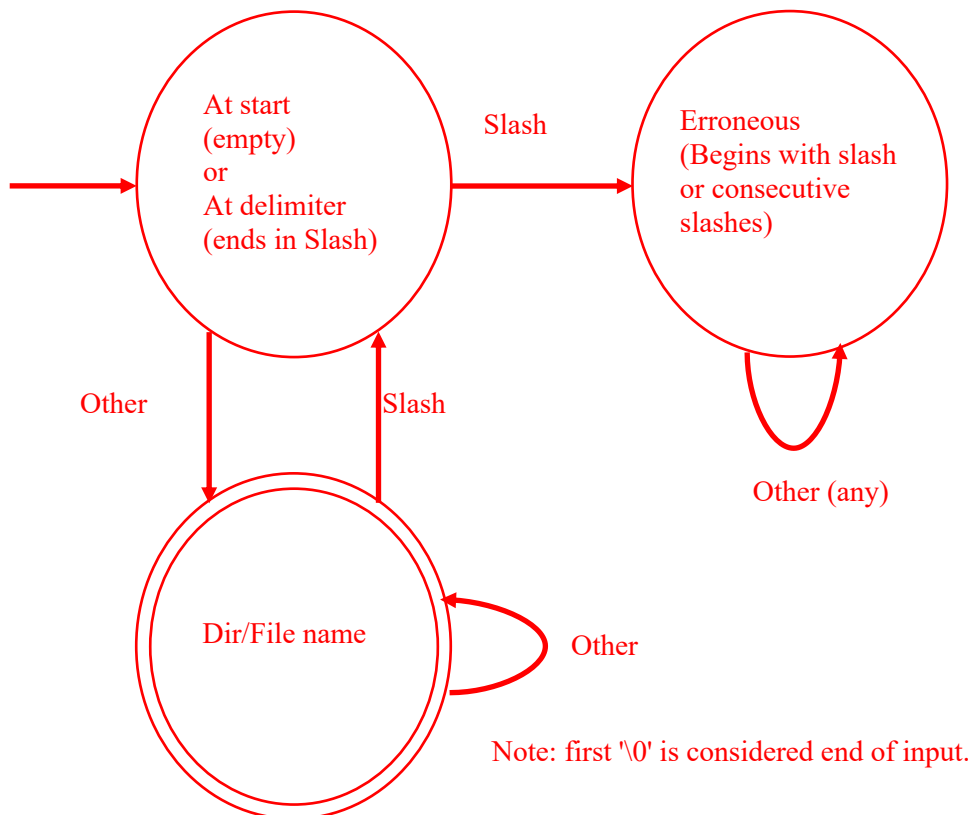
- (b) Write a function comment for `getParentPath` that accords to the COS 217 assignment standards for establishing a contract with clients who will call your function. (5 points)

Return a string with a copy of the path parameter's parent's path. If path is the root (i.e., it does not contain a slash), return the empty string. If memory for the new string cannot be allocated, return `NULL`. The caller owns the memory returned.

- (c) Define a DFA (in either the visual format from lecture or the textual format from Assignment 1) that accepts only paths that are well formed as specified above. Here is that description, repeated for convenience:

A non-empty `'\0'`-terminated string that neither begins nor ends with a `'/'` character and does not contain consecutive `'/'` characters.

Remember to identify your start state, classify each state as accept or reject, label all state transitions, and name each state with a semantically meaningful name. (8 points)



## Q5. Mysterious Assembly (25 pts)

An AArch64 reference is provided at the end of this question. Consider the following AArch64 program:

```
.section .rodata
scanFormat:
.string "%d"
printFormat:
.string "Answer is %d\n"

.section .text
.global mystery
mystery:
    sub    sp, sp, 48
    str    x30, [sp]
    str    x19, [sp, 8]
    str    x20, [sp, 16]
    str    x21, [sp, 24]

    mov    w19, wzr
    add    w20, w19, 1
    add    w21, w19, 2
loop1:
    cmp    w20, w0
    bgt    endloop1
    mul    w20, w20, w21
    add    w19, w19, 1
    b      loop1
endloop1:
    sub    w0, w19, 1

    ldr    x19, [sp, 8]
    ldr    x20, [sp, 16]
    ldr    x21, [sp, 24]
    ldr    x30, [sp]
    add    sp, sp, 48
    ret

.size    mystery, .-mystery

.global main
main:
    /* prologue */
    sub    sp, sp, 16
    str    x30, [sp]
    /* call to scanf */
    adr    x0, scanFormat
    add    x1, sp, 8
    bl     scanf
    /* call to mystery */
    ldr    w0, [sp, 8]
    bl     mystery
    /* call to printf */
    mov    w1, w0
    adr    x0, printFormat
    bl     printf
    /* return 0 and epilogue */
    mov    w0, 0
    ldr    x30, [sp]
    add    sp, sp, 16
    ret

.size    main, .-main
```

(a) Consider first the function `mystery`. How many parameters does it take as input? (1 point)

It takes one parameter (of type `int`).

(b) Does it return a value? (Yes, or No?) (1 point)

Yes, it returns a value (of type `int`).

(c) Write flattened-C corresponding to the function `mystery`. Assume that registers `w19`, `w20`, `w21` correspond to local variables `r`, `p`, `m`, respectively. Use the same label names as the assembly language code. (12 points)

C code:

```
int mystery (int n) {
    int r = 0;
    int p = 1;
    int m = 2;

loop1:
    /* while (p <= n) { */
    if (p > n)
        goto endloop1;

    p *= m;
    r++;
    goto loop1;
    /* } */

endloop1:
    return (r-1);
}
```

- (d) Now consider the function `main`. Note that it has calls to `scanf`, `mystery`, and `printf`. What does this program print if the user provides input 16? (1 point)

Answer is 4 <newline>

- (e) What does this program print if the user provides input 15? (1 point)

Answer is 3 <newline>

- (f) What does this program print if the user provides an input that is 0 or a negative number? (1 point)

Answer is -1 <newline>

- (g) What does the function `mystery` do? (*Hint*: Try out the program on a few other positive numbers if you like.) Show its specification in the form of a function comment that meets the requirements from your programming assignments in this course. (4 points)

If the argument `n` is a positive number, `mystery` returns the floor of  $\log_2 n$ .  
If `n` is 0 or a negative number, it returns -1.

**Alternate acceptable explanations for positive case:**

If `n` is a positive number, it returns the integral part of  $\log_2 n$ .

If `n` is a positive number, it returns the highest power of 2, such that 2 raised to that power is smaller than or equal to `n`.

- (h) List *two distinct* optimizations (i.e., not the same optimization on two different variables) that you can perform on the assembly code for *mystery* to make the program run faster, and state how they would help. (4 points)

You *don't need to show* the optimized assembly code, just briefly describe the optimizations for the specific given code (not just a general technique).

Here are examples of correct responses:

1. Change the multiply (`mul`) instruction to a logical shift left (`lsl`).  
A left shift is a faster instruction than a multiply instruction.
2. Instead of using callee-saved registers (`w19`, `w20`, `w21`), use caller-saved registers, since *mystery* does not call any other function. Then, those registers do not need to be saved on the stack nor restored from the stack.
3. A more extreme extension of 2.: since *mystery* does not call any other function, there is no need for a prologue and epilogue at all.
4. Use the guarded-loop pattern to save one branch instruction

<b>AArch64 quick reference</b>	
<p><b>add/sub/mul dst, src1, src2</b> Add/subtract/multiply <code>src1</code> and <code>src2</code>, put result in <code>dst</code>.</p> <p><b>mov dst, src</b> Copy <code>src</code> to <code>dst</code></p> <p><b>ldr dst, [src]</b> Load word or quad from memory pointed to by <code>src</code> into <code>dst</code></p> <p><b>str src, [dst]</b> Store word or quad from <code>src</code> to memory pointed to by <code>dst</code></p> <p><b>cmp src1, src2</b> Compare two registers, set condition flags</p> <p><b>b label</b> Unconditional branch to <code>label</code></p> <p><b>bgt label</b> Branch to <code>label</code> if greater than (signed)</p> <p><b>bl label</b> Branch to <code>label</code> and save return address in <code>x30</code></p> <p><b>adr dst, var</b> Put address of <code>var</code> in <code>dst</code></p> <p><b>ret</b> Return, i.e., branch to address in <code>x30</code></p>	<p><b>x0..x30, xzr</b> 64-bit (quad) registers</p> <p><b>w0..w30, wzr</b> 32-bit (word) registers</p> <p><b>r0..r7</b> Parameters, scratch space, caller-saved</p> <p><b>r0</b> Return value from function</p> <p><b>r9..r15</b> Scratch, caller-saved</p> <p><b>r19..r28</b> Scratch, callee-saved</p> <p><b>x30</b> Link register (return address)</p> <p><b>zr</b> Always holds zero</p>

## Q6. Missing Branch (5 points)

Consider again the assembly language program from the previous question. This time we will focus on a small part of the objdump output shown below, which shows the disassembled instructions for the loop in the `mystery` function.

...

```
0000000000000020 <loop1>:
 20:6b00029f  cmp  w20, w0
 24:5400008c  b.gt 34 <endloop1>
 28:1b157e94  mul  w20, w20, w21
 2c:11000673  add  w19, w19, #0x1
 30:MISSING  b   20 <loop1>
```

**MISSING:**      17FFFFFFC

```
0000000000000034 <endloop1>:
 34:51000660  sub  w0, w19, #0x1
```

...

In the box above on the right, provide the machine language encoding which is **MISSING** for only the branch (b) instruction highlighted in red, for which the machine language format is shown here:



### 26-bit signed PC-relative branch offset variant

B <label>

### Decode for this encoding

bits(64) offset = `SignExtend(imm26:'00', 64)`;

### Assembler symbols

<label>      Is the program label to be unconditionally branched to. Its offset from the address of this instruction, in the range +/-128MB, is encoded as "imm26" times 4.

### Operation

`BranchTo(PC[] + offset, BranchType_JMP);`

- Note that the first column lists the offset of each instruction (in hexadecimal), which you can use to compute the displacement that is needed to fill the correct value in the instruction format.
- Recall also that each machine language instruction is 32-bits, so a machine language instruction can be written as 8 hexadecimal digits (hexits), as shown in the objdump output above. You will likely want to work in binary, however your answer should be in hex.

**(blank page)**

Explanation for Q6:

1. Bits 31-26 are 000101 as shown in the format.
2. You have to compute the imm26 part of the machine language instruction.
3. This is a relative displacement, i.e., offset of branch target – offset of current instruction.
  - a. Offset of branch target instruction: 0x0000 0000 0000 0020
  - b. Offset of current instruction: 0x0000 0000 0000 0030
  - c. Displacement:  
negative 0000 0000 0000 0000 0000 0001 0000 (28 bits)
  - d. Displacement shifted right by 2 bits:  
negative 00 0000 0000 0000 0000 0000 0100 (26 bits)
  - e. Displacement in 2's complement:  
11 1111 1111 1111 1111 1111 1100 (26 bits)
4. Putting together bits 31-26 and bits 25-0 (imm26), we get machine language instruction:  
0001 0111 1111 1111 1111 1111 1111 1100 (32 bits)
5. In hex, the machine language instruction is:  
1 7 F F F F F C (8 hexits = 32 bits)

### Q7. Finally, some fun! (15 points)

For each code snippet, give a short (at most one line) description of what the program does. Assume all necessary header files are included.

(a) What does the following function print to `stdout` for a non-negative number `a`? (5 points)

```
void funA(unsigned int a) {
    if (a & (~1))
        funA(a>>1);
    putchar('0'+(a & 1));
}
```

It prints the binary representation of `a`.

Explanation:

1. `a & ~1` will be true when `a` has any non-zero bits in bit positions 1-31 (i.e., positions other than the rightmost bit).
2. In the above case, `funA` is called recursively on input `a` shifted right by 1 bit, which will print the binary representation of `a` divided by 2.
3. Note that the recursive call does not change `a` itself. After the recursive call returns, this call prints the character '0' or '1', depending on the 0<sup>th</sup> bit of `a`.

(b) How does `funB()` change the string passed to it? Recall that the function `toupper` converts a lower-case letter to the corresponding upper-case letter. (5 points)

```
char* funB(char *pStr) {
    int i, n = strlen(pStr);

    for (i = 3; i < n; i += 3)
        pStr[i] = toupper(pStr[i]);

    return pStr;
}
```

It changes every third character of the string from a lower-case to an upper-case letter, starting from the fourth character.

Alternate description: it capitalizes characters in indices of the string that are positive multiples of 3.

(If any such third character is not a lower-case letter, then that character is not changed.)



(c) For what kind of non-negative numbers  $n$  does the following function return 1? (5 points)

```
int funC(unsigned int n) {
    if (n & (n-1))
        return 0;
    if (n & 0x55555555)
        return 1;
    else
        return 0;
}
```

It returns 1 for any number that is a power of 4.

(Note that this corresponds to *even* powers of 2, which received full credit whether expressed in words or as  $2^{2n}$ . The response “every other power of 2”, received 4 points credit, as this description could just as well describe odd powers of 2, which would be wrong.)

Explanation:

1. The condition  $(n \ \& \ (n-1))$  will be true if  $n$  is *not* a power of 2. For such numbers, `funC` returns a 0 early.
2. If  $n$  is a power of 2, it will have a 1 in some position and all other bits are 0.
3. The next condition  $(n \ \& \ 0x55555555)$  checks if the 1 is in a position such that bitwise and-ing with 5 (i.e.,  $0x0101$ ) in every hexit results in true, in which case `funC` returns 1.

**(blank page)**