

---

<b>Name</b>	
<b>Student ID</b>	

# UW CSE 351, Summer 2013

## Midterm Exam

### Instructions:

- Make sure that your exam is not missing any of the 12 pages, then write your full name and UW student ID on the front.
- Read over the entire exam before starting to work on the problems! The last page is a reference page that you may tear off to use during the exam; it does not have to be turned in.
- Feel free to use the backs of pages for working on problems, but *write your answers in the space provided below each problem*. If you make a mess, clearly indicate your final answer. Be sure to answer all parts of all questions.
- Do not spend too much time on a problem if there are other easy problems that you haven't solved yet.
- **No books, notes, or electronic devices may be used during the exam.** You may not communicate with other students during the exam, but please ask the instructor if you need clarification for some problem.
- If you read this far, smile and have fun!

<b>Section</b>	<b>Awarded</b>	<b>Possible</b>
<b>Section 1</b>		<b>5</b>
<b>Section 2</b>		<b>25</b>
<b>Section 3</b>		<b>40</b>
<b>Section 4</b>		<b>30</b>
<b>Total</b>		<b>100</b>

---

**Note:** This exam was tough and about 10-15% too long in retrospect. Some questions in section 4 were especially subtle. We graded by adding a buffer to everyone's absolute score to reflect the fact that the exam was too long. Answers within often go into more detail than needed for completing the exam.

---

## 1. Instruction Set Architectures (5 points)

- (a) The original x86 processor, the Intel 8086, had a *word size* of 16 bits. What does *word size* mean and what does it determine in a computer system? (one or two sentences)

**Word size defines register size; it is the largest piece of data a CPU operates on at once. In machines we have looked at, it is also the same as the size of an address/pointer and in all machines it puts a bound on the size of an address, and hence the addressable memory.**

Some people defined word size in terms of basic units of memory accesses, but recall that memory is still byte-addressable. Every byte of memory may be read and written individually, even if the system really moves around a full word (or more) on every access.

- (b) Circle all of the features below that are part of the *instruction set architecture* of a machine:
- The number of registers.
  - The number of machine cycles to execute a single instruction.
  - The effect that a certain instruction has on memory and registers when executed.
  - The condition codes and what causes them to be set.
  - The available memory addressing modes.
  - The machine word size.

**All except (ii).** The ISA defines the parts of the machine an assembly programmer needs to know about to understand *what* a program does. The number of cycles per instruction does not affect *what* the program does.

## 2. Numbers and Bits (25 points)

- (a) (3 points) Is there anything wrong with the code below? If so, name one thing. Assume that  $n$  is the number of elements in the `values` array. (Write one or two sentences.)

```
double product_except(int n, float values[], float skip) {
    float prod = 1.0;
    int i;
    for (i = 0; i < n; i++) {
        if (values[i] != skip) {
            prod = prod * values[i];
        }
    }
    return (double)prod;
}
```

**There are things that can go wrong: 1. floating point values should never be compared for equality, due to rounding imprecision; 2. the multiplication of floats could result in overflow (e.g., reaching positive or negative infinity).** Either of these answers received full credit. Rounding float to double is not a problem, since there's always enough space in a double to store the value of a float without rounding.

---

(b) (8 points) Consider a machine with 6-bit integers. (While we generally use powers of two today, 36-bit words with 6-bit characters were commonplace several decades ago.)

- i. What are the binary *and* decimal representations of the sum of these *two's complement signed* integers represented in binary? (Please write the sum both in binary and in decimal.)

$$100101 + 011101 = -27 + 29 \text{ (not required as part of answer)}$$

binary: **000010**

decimal: **2**

- ii. What are the binary *and* decimal value of the sum of the 6-bit *two's complement* representations of these signed integers (shown here in decimal)? (Please write the sum both in binary and in decimal.)

$$(-24) + (-12) = 101000 + 110100 \text{ (not required as part of answer, but this translation tripped up a few people)}$$

binary: **011100**

decimal: **28**

- iii. *Either* your answer for (i) *or* your answer for (ii) should be different than the answer for arithmetic on the equivalent mathematics integers. Which one? What is the name of the condition that causes it to be different?

**(ii) is different. Modular arithmetic gives us 28 instead of -36 due to overflow.**

(c) (4 points) What are two disadvantages to *sign and magnitude* integer representation as compared to *two's complement*? (one sentence or phrase each)

**There are two representations of 0.**

**The addition algorithm for signed numbers is different than that for unsigned (also positive vs. negative), and is more cumbersome.**

**On the other hand, two's complement has a single representation of 0, and a single addition algorithm**

(d) (10 points) Design a space-efficient encoding of the latitude, longitude, and altitude of a satellite's position relative to Earth using only a single 32-bit C int to store *all* of this information at once. It must be easy to retrieve the individual dimensions from the encoded position. Signed latitudes range from -90 to +90 in increments of 1 and signed longitudes range from -180 to +179 in increments of 1. The satellite altitude (in your favorite units) ranges from zero to some maximum in increments of 1.

- i. Allocate the 32 available bits below by bracketing and labeling which bits belong to which dimensions. *Put the latitude in the highest order bits, followed by the longitude, followed by the altitude in the lowest order bits.* Be sure to maximize the altitude you can represent.

**latitude (signed), 8 bits:**  $-(2^{8-1}) = -128 \leq -90 < -(2^{7-1}) = -64$

**longitude (signed), 9 bits:**  $-(2^{9-1}) = -256 \leq -180 < -(2^{8-1}) = -128$

**altitude (unsigned), 15 bits: 32 - 8 - 9**

```
latitude: 31 - 24   | longitude: 23 - 15   | altitude: 14 - 0
31 30 29 28 27 26 25 24|23 22 21 20 19 18 17 16 15|14 13 12 11 10  9  8  7  6  5  4  3  2  1  0
```

- ii. What is the maximum altitude your encoding can represent? You may write an expression involving exponents and/or arithmetic.  
 $2^{15} - 1$ , **due to 15 bits unsigned (altitudes are non-negative...)**

- iii. Implement a function to retrieve the longitude as a signed 32-bit int given an int encoded using your answer to (i). Only bitwise operators, shifts, and addition/subtraction are allowed, for efficiency. You may use local variables, but they are not necessary.

```
int get_longitude(int position) {
    return (position << 8) >> (8 + 15);
}
```

**It was necessary to shift left before shifting right so that the sign bit of longitude would be preserved by arithmetic shift right .**

---

### 3. Reverse Engineering Stacks and Procedures (40 points)

*There are questions that use the code below on the following pages. You may tear this page out carefully if you want to avoid flipping back for reference. You must turn this page in, but it does not need to be attached to the rest of your exam.*

Running short on sleep after a research paper deadline, your instructor accidentally removed several important files for a C program! Fortunately, the compiled machine code survived, but he needs your help reverse-engineering it. He remembers that the code uses a linked list of integers, with each node in the list represented by the following type:

```
typedef struct list_node {
    int value;
    struct list_node * next;
} list_node;
```

Each `list_node` stores an `int` `value` and a pointer (`next`) to the next `list_node` in the list. The last element in a list stores the address `0x0` in its `next` field, showing that there is no next element after the last element.

Additionally, we know that `mystery` looks like this:

```
int mystery(list_node* current) { ... }
```

The following machine code has been recovered and disassembled:

```
0x080483a0 <mystery>:
  0x080483a0 <+0>:      push   %ebp
  0x080483a1 <+1>:      mov    %esp,%ebp
  0x080483a3 <+3>:      sub   $0x8,%esp
  0x080483a6 <+6>:      mov   0x8(%ebp),%edx
  0x080483a9 <+9>:      mov   0x4(%edx),%eax
  0x080483ac <+12>:     test  %eax,%eax      <-- Pause for (c).
  0x080483ae <+14>:     jne   0x80483b4 <mystery+20>
  0x080483b0 <+16>:     mov   (%edx),%eax
  0x080483b2 <+18>:     jmp   0x80483bc <mystery+28>
  0x080483b4 <+20>:     mov   %eax,(%esp)
  0x080483b7 <+23>:     call  0x80483a0 <mystery>
  0x080483bc <+28>:     leave                               <-- Stop for (d).
  0x080483bd <+29>:     ret

0x080483bb <intrigue>:
  ...
  ...                               <-- Start.
  0x080483eb <+48>:     call  0x80483a0 <mystery>
  0x080483f0 <+53>:     mov   %eax,0x2c(%esp)
  ...
```

**Section 3 continues on the next page.**

### 3. [Continued] Reverse Engineering Stacks and Procedures (40 points)

We have provided you with the state of the heap and the stack *just before executing the call instruction in `intrigue` at `0x080483eb`. **Questions (a) through (e) on the next page will guide you through filling the blanks in this diagram.** Feel free to make notes in the margins, but make sure all labels requested in (a) through (e) are clearly visible. (Hints: you should not need more stack space than we have drawn and the heap contents will not change.)*

**Stack**                      **initial %ebp = 0xffffffff58**                      **initial %esp = 0xffffffff4c**

	Address	Value	Description of value
	0xffffffff4c	0x0b000000	(b) argument to <code>mystery</code> ( <code>list_node*</code> )
	0xffffffff48	0x080403f0	(b) return address in <code>intrigue</code>
	0xffffffff44	0xffffffff5b	(b) saved <code>%ebp</code>
	0xffffffff40		(b) unused
	0xffffffff3c	0x0b000018	(d) argument to <code>mystery</code> ( <code>list_node*</code> )
	0xffffffff38	0x080403bc	(d) return address in <code>mystery</code>
(d) <code>%ebp</code> →	0xffffffff34	0xffffffff44	(d) saved <code>%ebp</code>
	0xffffffff30		(d) unused
(d) <code>%esp</code> →	0xffffffff2c		(d) unused

#### Heap

Address	Value	Description of these 8 bytes
0x0b00001c	0x0	(next) <code>list_node</code> (a)
0x0b000018	351	(value)
...	...	...
0x0b000004	0x0b000018	<code>list_node</code> (a)
0x0b000000	341	

**Section 3 continues on the next page.**

---

### 3. [Continued] Reverse Engineering Stacks and Procedures (40 points)

These questions refer to the diagrams and code on the previous two pages.

- (a) (3 points) Based on `mystery`'s signature and the initial state of the stack, the heap, and the `%ebp` and `%esp` registers immediately before the call instruction at `0x080483eb` is executed, label the *type* of what should be stored in the 8 bytes of memory at `0x0b000000` and the 8 bytes of memory at `0x0b000018`. **See previous page.**
- (b) (10 points) Next, simulate the execution of the program *starting just before the call instruction at `0x080483eb`. Stop just before you reach the test instruction at `0x080483ac`. Fill in the stack diagram as your simulation progresses, writing each value stored on the stack and a description of what this value represents. Write "unused" if the location is not used. You may find it helpful to look ahead to part (e) of this problem to consider as you simulate the program. See previous page.*
- (c) (4 points) What is the combined purpose of the `test` instruction at `0x080483ac` and the following `jne` instruction at `0x080483ae`? (one sentence or phrase)

**It tests to see whether the next pointer in the current list\_node is 0 (end of list). You also got credit if you described this at the assembly level since the question did not specify which...**

- (d) (10 points) Continue simulating execution and filling in the stack diagram until just before you execute a `leave` instruction, and then label where the frame pointer ("`←%ebp`") and stack pointer ("`←%esp`") point. **See previous page.**
- (e) (10 points) We just heard that the original C code had the following structure! Using what you now know about `mystery`, fill in the blank lines with C code, using no variables except `current` in your expressions. Do not write register names.

**There are a few equivalent versions. Here's one.**

```
int mystery(list_node* current) {
    if( current->next == 0 ) {
        return current->value;
    }
    return mystery( current->next );
}
```

- (f) (3 points) Congratulations, you recovered the lost code! Explain what `mystery` does with its linked list argument. (one or two sentences)

**`mystery` returns the value in the last node in a linked list.**



---

#### 4. Translating C Arrays to Memory and Assembly (30 points)

The following two C functions appear to perform the same operations: setting all elements in the diagonal of a matrix to their index within the diagonal. For example, `matrix[3][3]` gets set to 3. **However, they are subtly different.**

```
void diagonalA(int n, int* matrix[]) {
    int i;
    for (i = 0; i < n; i++) {
        matrix[i][i] = i;
    }
}

void diagonalB(int n, int matrix[4][4]) {
    int i;
    for (i = 0; i < n; i++) {
        matrix[i][i] = i;
    }
}
```

For each statement below, create a true statement by filling the blank with one of:

- “Only A” if the statement is only true of `diagonalA`;
- “Only B” if the statement is only true of `diagonalB`;
- “Neither” if it is true of neither; or
- “Both” if it is true of both.

Additionally, where asked, explain how or why this is the case *in a brief sentence or phrase*.

- (a) (5 points) **Only B** uses multi-dimensional (2D) arrays, while **Only A** uses multi-level (2-level) arrays. What is the difference and how can you tell from the type?

**A uses an array of int pointers (which, in this case, are also pointers to int arrays). A’s array of pointers is contiguous in memory, but the sub-arrays they point may not be. B uses a fixed-size 2D array: an array of int arrays, allocated contiguously in row-major form.**

- (b) (2 points) **Both** could lead to a segmentation fault. How/why?

**Even if we assume the `matrix` argument well-formed, given the wrong `n`, both can write outside the bounds of the arrays they are given.**

- (c) (5 points) **Only B** could generate memory accesses to this sequence of addresses if the base address of the matrix array is `0x4000` and `n = 4`:

`0x4000, 0x4014, 0x4028, 0x403c`. How/why?

**These addresses represent accesses of the form  $0x4000 + 4 * \text{sizeof}(\text{int}) * i + 4 * i$  (for  $i$  from 0 to  $n - 1$ ) to a contiguous block of memory, with one access per iteration.**

Even a two-level array `a` with `a[0] = &a[1]`, would still need 8 accesses instead of 4.

- 
- (d) (5 points) **Only A** could generate memory accesses to this sequence of addresses if the base address of the `matrix` array is `0x4000` and `n = 4`:  
`0x4000, 0x4800, 0x4004, 0x568c, 0x4008, 0x4600, 0x400c, 0x4800`  
How/why?

**These are accesses to pointers stored in the top-level array, a contiguous block of memory (of the form *base address + sizeof(int\*)\*i*: `0x4000, 0x4004, 0x4008, 0x400c`), interleaved with accesses to sub-arrays elsewhere in memory. (`0x4800, 0x568c, 0x4600, 0x4800`).**

If you chose “Only A” and you mentioned something about non-contiguous memory or 2 accesses for every iteration, you got credit.

- (e) (3 points) **Neither** could lead to a buffer overflow in the current stack frame, corrupting the return address so that when `diagonalA/B` returns, it starts executing an attacker’s code instead. How/why?

**The arrays used by these functions are not allocated in the stack frames of these functions, so writing past the ends of the arrays cannot write in these stack frames or immediate return addresses.**

This was admittedly tricky, and I wish I had asked a more straightforward question. You got 1 point if you explained that there could be a buffer overflow by writing past the end of the arrays. But it turns out it is impossible to write in `diagonalA/B`’s stack frame or return address because of this.

Neither `diagonalA` nor `diagonalB` is operating on an array allocated *in its own stack frame*, so assuming the array is allocated in *some* stack frame, it must be allocated higher in memory than the return address for `diagonalA/B` (in the caller of `diagonalA/B` or its caller or ...). Thus, even if we write past the end of the array, we are overwriting the stack frame in which the array was allocated, not `diagonalA/B`’s stack frame or the return address it should use. Thus, there’s no way to overwrite that return address, and when `diagonalA/B` returns, it always returns to its caller. Perhaps its caller (or its caller or ...) will have a corrupted return value, but not

Getting even more far-fetched: if the array was allocated on the *heap* and the attacker chose a large enough *n*, the return address for `diagonalB` might eventually get overwritten. Even then, it would take great luck to be at the right alignment and *i* to write when reaching the return address slot on the stack. Overwriting past the size of the top-level array in `diagonalA` would likely cause a segmentation fault when treating garbage as a pointer to a sub-array.

Even the most likely option is unlikely: causing the program to pass just the right bogus array address argument, where that address is *below* the addresses of the `diagonalA/B` stack frame (in unallocated memory). In this case, the “attacker” still needs to cause your program to pass that right bogus address, perhaps via another buffer overflow.

- (f) (10 points) **Only A** could be correctly compiled to this assembly code (question below):

```
080483f4 <diagonalMystery>:
  80483f4: 55                push   %ebp
  80483f5: 89 e5            mov    %esp, %ebp
```

---

```
80483f7: 53          push   %ebx
80483f8: 8b 4d 08    mov    0x8(%ebp),%ecx
80483fb: 8b 5d 0c    mov    0xc(%ebp),%ebx
80483fe: 85 c9      test   %ecx,%ecx
8048400: 7e 12      jle   8048414 <diagonalMystery+0x20>
8048402: b8 00 00 00 00  mov    $0x0,%eax
8048407: 8b 14 83    mov    (%ebx,%eax,4),%edx
804840a: 89 04 82    mov    %eax,(%edx,%eax,4)
804840d: 83 c0 01    add    $0x1,%eax
8048410: 39 c8      cmp    %ecx,%eax
8048412: 75 f3      jne   8048407 <diagonalMystery+0x13>
8048414: 5b        pop    %ebx
8048415: 5d        pop    %ebp
8048416: c3        ret
```

***What was your key observation?***

There are many ways to answer this. One is: **There are two memory accesses every time around the loop: one to read the pointer (0x8048407) to the sub-array and one to write into the sub-array using an offset from its base address (0x804840a).**

Full credit was given for picking A and backing it up with a feasible explanation. Partial credit was give for picking A but noting a false reason or for picking something other than A and noting something true and relevant about the code.

---

## Reference

### 1. Hex Digits

0	1	2	3	4	5	6	7	8	9	A	B	C	D	E	F
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15

### 2. Powers of Two

$2^0 = 0x0001 = 1$	$2^6 = 0x0040 = 64$
$2^1 = 0x0002 = 2$	$2^7 = 0x0080 = 128$
$2^2 = 0x0004 = 4$	$2^8 = 0x0100 = 256$
$2^3 = 0x0008 = 8$	$2^9 = 0x0200 = 512$
$2^4 = 0x0010 = 16$	$2^{10} = 0x0400 = 1024$
$2^5 = 0x0020 = 32$	

### 3. Assembly Code Instructions

<code>pushl</code>	push a 4-byte value onto the stack
<code>pushq</code>	push a 8-byte value onto the stack
<code>popq</code>	pop a 8-byte value from the stack
<code>call</code>	push the address of the next instruction onto the stack and jump to target
<code>leave</code>	restore <code>ebp</code> from the stack
<code>ret</code>	pop return address from stack and jump there
<code>leal</code>	compute an address in first operand, put result in second
<code>movl</code>	move 4 bytes between values, registers and memory
<code>movq</code>	move 8 bytes between values, registers and memory
<code>movzbl</code>	move zero-extended value to long
<code>incl</code>	increment operand by 1
<code>decl</code>	decrement operand by 1
<code>addl</code>	(4 bytes) add first operand to second, put result in second
<code>addq</code>	(8 bytes) add first operand to second, put result in second
<code>subl</code>	subtract first operand from second, put result in second
<code>imull</code>	signed multiply of first operand and second, put result in second
<code>andl</code>	logical AND of first operand with second, put result in second
<code>sall</code>	arithmetic left shift of first operand, put result in second
<code>sarl</code>	arithmetic right shift of first operand
<code>cmpl</code>	subtract first operand from second; set flags
<code>testl</code>	logical AND of first operand with second; set flags (4 bytes)
<code>testb</code>	logical AND of first operand with second; set flags (1 byte)
<code>jmp</code>	jump to address
<code>jg</code>	conditional jump to address if last comparison was greater than
<code>je</code>	conditional jump to address if result of last comparison was zero
<code>jne</code>	conditional jump to address if result of last comparison was not zero
<code>jle</code>	conditional jump to address if result of last comparison was zero or negative