

**15-213 Introduction to Computer Systems**  
**Final Exam**  
May 10, 2007

Name: \_\_\_\_\_  
Andrew User ID: \_\_\_\_\_  
Recitation Section: \_\_\_\_\_

- This is an open-book exam.
- Notes and calculators are permitted, but not computers.
- Write your answer legibly in the space provided.
- You have 180 minutes for this exam.

	<b>Problem</b>	<b>Max</b>	<b>Score</b>
Floating Point	1	20	
Assembly Language	2	20	
Optimization	3	20	
Cache Memory	4	20	
Signals	5	20	
Garbage Collection	6	20	
Threads	7	20	
Synchronization	8	20	
	<b>Total</b>	<b>150+10</b>	

## 1. Floating Point (20 points)

In this problem we consider properties of floating point operations. For each property state whether it is true or false. If false, give a counterexample as a (possibly negative) power of 2 within the range of precision for the variables. We assume that the variables on an x86\_64 architecture are declared as follows

```
float x, y, z;  
double d, e;
```

and initialized to some unknown value **different from NaN,  $+\infty$ , and  $-\infty$** . We have given the first answer as an example.

$(x + y) + z == x + (y + z)$	<b>false</b>	$x = 1, y = 2^{127}, z = -2^{127}$
If $x > 0$ then $x / 2 > 0$		
$(x + y) * z == x * z + y * z$		
If $x \geq y$ and $z \leq 0$ then $x * z \leq y * z$		
If $x > y$ then $(\text{double})x > (\text{double})y$		
If $d > e$ then $(\text{float})d > (\text{float})e$		
$x + 1 > x$		

## 2. Assembly Language (20 points)

In this problem we consider an illustrative program for multiplication of two unsigned int's, returning an unsigned long int holding the product.

```
unsigned long mult (unsigned i, unsigned k) {
    unsigned long p = 0;
    unsigned long q = k;
    while (i != 0) {
        if (i & 1)
            p = p + q;
        q = q << 1;
        i = i >> 1;
    }
    return p;
}
```

The following is the resulting machine code when compiled on an x86.64 machine with gcc -O2, omitting two instructions.

```
mult:
    xorl    %ecx, %ecx
    mov     %esi, %edx
    testl   %edi, %edi
    jmp     .L8

.L10:
    leaq   (%rcx,%rdx), %rax
    testb  $1, %dil

    _____ # missing conditional move
    addq   %rdx, %rdx
    shrl   %edi

.L8:
    jne    .L10

    _____ # missing move
    ret
```

1. (5 pts) For each register, give the value it holds during the iteration, expressed in terms of the C program.

Register	C expression
<code>%rcx</code>	
<code>%rdx</code>	
<code>%rax</code>	
<code>%edi</code>	
<code>%dil</code>	

2. (5 pts) Fill in the missing two instructions in the code.
3. (4 pts) Rewrite the loop to use a conditional jump instead of a conditional move.

4. (3 pts) Explain briefly why the compiler preferred to use a conditional move instruction.

5. (3 pts) Assume we declared and initialized

```
int i, k;  
long m;
```

and called

```
m = (long)mult((unsigned)i, (unsigned)k);
```

using the above definition of `mult`. Will `m` hold the correct value of the signed product of `i` and `k`? Circle the correct answer.

yes      no

Briefly explain your answer.

### 3. Optimization (20 points)

Consider the following code for calculating the dot product of two vectors of double precision floating point numbers.

```
double dot_prod(double A[], double B[], int n) {
    int i;
    double r = 0;
    for (i = 0; i < n; i++)
        r = r + A[i] * B[i];
    return r;
}
```

Assume that multiplication has a latency of 12 cycles and addition a latency of 7 cycles and load 4 cycles. Also assume that there are an unlimited number of functional units. **[Hint:** Under this assumption, theoretically optimal performance is dominated by the critical data dependency path.]

1. (5 points) What is the theoretically optimal CPE for this loop?
2. (10 points) Show the code for the loop unrolled by 2. You may apply associativity and commutativity of multiplication and addition, assuming that rounding errors are insignificant.

```
double dot_prod2(double A[], double B[], int n) {
    int i;
    double r = 0;

    return r;
}
```

3. (5 points) What is the theoretically optimal CPE for this loop?

#### 4. Cache Memory (20 points)

In this problem we explore the operation of a basic TLB as a cache. Assume the following

- Virtual addresses are 32 bits.
- The virtual page number (VPN) is 24 bits.
- The physical page number (PPN) is 32 bits.
- The TLB is 2-way set associative containing a total of 512 lines.

1. (6 points) Please fill in the following blanks by giving a bit range, such as “0–15”.

- (a) The VPO of a virtual address consists of bits \_\_\_\_\_ of the VA.
- (b) The VPN of a virtual address consists of bits \_\_\_\_\_ of the VA.
- (c) The PPO of a physical address consists of bits \_\_\_\_\_ of the PA.
- (d) The PPN of a physical address consists of bits \_\_\_\_\_ of the PA.
- (e) The TLB index (TLBI) consists of bits \_\_\_\_\_ of the VPN.
- (f) The TLB tag (TLBT) consists of bits \_\_\_\_\_ of the VPN.

We show a part of the TLB relevant to the next two questions.

Index	Valid?	Tag	Entry
3D	1	0x083F	0x0913ABDE
	1	0x083E	0xAB18ED24
3E	0	0xF3E9	0x0913ABDE
	1	0x083F	0xAB18ED24
3F	1	0x409A	0x0913ABDE
	1	0x083F	0xAB18ED24
40	0	0x083E	0x0913ABDE
	1	0x3E40	0xAB18ED24

2. (7 points) Assume the virtual address is  $0x083F3E9A$ . Fill in the following table in hexadecimal notation. Write **U** for any value that is unknown, that is, not determined from the parameters and the table above.

Parameter	Value
VPN	
VPO	
TLBI	
TLBT	
Cache Hit? (Y/N/U)	
PPN	
PA	

3. (7 points) Assume the virtual address is  $0x083E409B$ . Fill in the following table in hexadecimal notation. Write **U** for any value that is unknown, that is, not determined from the parameters and the table above.

Parameter	Value
VPN	
VPO	
TLBI	
TLBT	
Cache Hit? (Y/N/U)	
PPN	
PA	

## 5. Signals (20 points)

Consider the following program.

```
int counter = 0;

void handler (int sig) {
    counter++;
}

int main() {
    signal(SIGUSR1, handler);
    signal(SIGUSR2, handler);
    int parent = getpid();
    int child = fork();

    if (child == 0) {

        /* insert code here */

        exit(0);
    }

    sleep(1);
    waitpid(child, NULL, 0);
    printf("Received %d USR{1,2} signals\n", counter);
    return 0;
}
```

For each of the following four versions of the above code, list the possible outputs of this program, assuming that all function and system calls succeed and exit without error. You may also assume no externally issued signals are sent to either process.

1. (5 pts)

```
kill(parent, SIGUSR1);
kill(parent, SIGUSR1);
```

2. (5 pts)

```
kill(parent, SIGUSR1);  
kill(parent, SIGUSR1);  
kill(parent, SIGUSR1);
```

3. (5 pts)

```
kill(parent, SIGUSR1);  
kill(parent, SIGUSR2);
```

4. (5 pts)

```
kill(parent, SIGUSR1);  
kill(parent, SIGUSR2);  
kill(parent, SIGUSR1);  
kill(parent, SIGUSR2);
```

## 6. Garbage Collection (20 points)

In this problem we consider a tiny list processing machine in which each memory word consists of two bytes: the first byte is a pointer to the tail of the list and the second byte is a data element. The end of a list is marked by a pointer of  $0\times 00$ . We assume that the data element is never a pointer.

We start with the memory state on the left, where the range  $0\times 10-0\times 1F$  is the from-space and the range  $0\times 20-0\times 2F$  is the to-space. All addresses and values in the diagram are in hexadecimal.

Write in the state of memory after a copying collector is called with root pointers  $0\times 10$  and  $0\times 12$ , in this order. You may leave cells that remain unchanged blank.

Please be sure to use the proper breadth-first traversal algorithm covered in lecture.

Before GC			After GC		
Addr	Ptr	Data	Addr	Ptr	Data
10	14	A2	10		
12	1A	1F	12		
14	1E	02	14		
16	1E	20	16		
18	00	33	18		
1A	18	BC	1A		
1C	12	DF	1C		
1E	10	8F	1E		
			20		
			22		
			24		
			26		
			28		
			2A		
			2C		
			2E		

After garbage collection, free space starts at address \_\_\_\_\_

## 7. Threads (20 points)

Consider three concurrently executing threads in the same process using two semaphores  $s_1$  and  $s_2$ . Assume  $s_1$  has been initialized to 1, while  $s_2$  has been initialized to 0.

What are the possible values of the global variable  $x$ , initialized to 0, after all three threads have terminated?

```
/* thread A */
P(&s2);
P(&s1);
x = x*2;
V(&s1);

/* thread B */
P(&s1);
x = x*x;
V(&s1);

/* thread C */
P(&s1);
x = x+3;
V(&s2);
V(&s1);
```

## 8. Synchronization (20 points)

We explore the so-called *barbershop problem*. A barbershop consists of a  $n$  waiting chairs and the barber chair. If there are no customers, the barber waits. If a customer enters, and all the waiting chairs are occupied, then the customer leaves the shop. If the barber is busy, but waiting chairs are available, then the customer sits in one of the free chairs.

Here is the skeleton of the code, without synchronization.

```
extern int N;    /* initialized elsewhere to value > 0 */
int customers = 0;

void* customer() {

    if (customers > N) {

        return NULL;
    }

    customers += 1;

    getHairCut();

    customers -= 1;

    return NULL;
}

void* barber() {
    while(1) {

        cutHair();

    }
}
```

For the solution, we use three binary semaphores:

- `mutex` to control access to the global variable `customers`.
- `customer` to signal a customer is in the shop.
- `barber` to signal the barber is busy.

1. (5 points) Indicate the initial values for the three semaphores.

- `mutex`
- `customer`
- `barber`

2. (15 points) Complete the code above filling in as many copies of the following commands as you need, but no other code.

```
P (&mutex);  
V (&mutex);  
P (&customer);  
V (&customer);  
P (&barber);  
V (&barber);
```