## Andrew login ID:

$\qquad$
Full Name:

## CS 15-213, Spring 2004

## Exam 1

February 26, 2004

## Instructions:

- Make sure that your exam is not missing any sheets (there should be 15 ), then write your full name and Andrew login ID on the front.
- Write your answers in the space provided below the problem. If you make a mess, clearly indicate your final answer.
- The exam has a maximum score of 80 points.
- The problems are of varying difficulty. The point value of each problem is indicated. Pile up the easy points quickly and then come back to the harder problems.
- This exam is OPEN BOOK. You may use any books or notes you like. No electronic devices are allowed. Good luck!

| 1 (9): |
| ---: |
| $2(14):$ |
| 3 (12): |
| $4(8):$ |
| $5(12):$ |
| $6(12):$ |
| $7(6):$ |
| $8(7):$ |
| TOTAL (80): |

## Problem 1. (9 points):

Assume we are running code on a 10 -bit machine using two's complement arithmetic for signed integers. Short integers are encoded using 5 bits. Sign extension is performed whenever a short is casted to an int. For this problem, assume that all shift operations are arithmetic. Fill in the empty boxes in the table below.

```
int i = -42;
unsigned ui = i;
short s = -7;
unsigned short us = s;
```

Note: You need not fill in entries marked with "-". TMax denotes the largest positive two's complement number and TMin denotes the minimum negative two's complement number. Finally, you must use hexidecimal notation for your answers in the "Hex Representation" column, failure to do so will result in being marked incorrect on that portion of the question.

| Expression | Decimal Representation | Hex Representation |
| :---: | :---: | :---: |
| Zero | 0 | $\mathbf{0 0 0}$ |
| - | -9 | $\mathbf{3 f 7}$ |
| i | -42 | $\mathbf{3 d 6}$ |
| i >>5 | -2 | $\mathbf{3 f e}$ |
| ui | $\mathbf{9 8 2}$ | $\mathbf{3 d 6}$ |
| (int) s | -7 | $\mathbf{3 f 9}$ |
| (int)(s $\wedge-12)$ | $\mathbf{1 3}$ | $\mathbf{0 0 d}$ |
| (int) us | $\mathbf{2 5}$ | $\mathbf{0 1 9}$ |
| TMax | $\mathbf{5 1 1}$ | $\mathbf{1 f f}$ |
| TMin | $\mathbf{- 5 1 2}$ | $\mathbf{2 0 0}$ |

## Problem 2. (14 points):

Consider the following 11-bit floating point representation based on the IEEE floating point format:

- There is a sign bit in the most significant bit.
- The next $k=4$ bits are the exponent. The exponent bias is 7 .
- The last $n=6$ bits are the significand.

Numeric values are encoded in this format as a value of the form $V=(-1)^{s} \times M \times 2^{E}$, where $s$ is the sign bit, $E$ is exponent after biasing, and $M$ is the significand.

## Part I

How many FP numbers are in the following intervals $[a, b)$ ?
For each interval $[a, b)$, count the number of $x$ such that $a \leq x<b$.
A. Interval $[1,2)$ : $\qquad$
B. Interval $[2,2.5): \quad 2^{4}=\mathbf{1 6}$

## Part II

Answer the following problems using either decimal (e.g., 1.375) or fractional (e.g., 11/8) representations for numbers that are not integers.
A. For denormalized numbers:
(a) What is the smallest value $E$ of the exponent after biasing? -6
(b) What is the largest value $M$ of the significand? $\qquad$
B. For normalized numbers:
(a) What is the smallest value $E$ of the exponent after biasing? -6
(b) What is the largest value $E$ of the exponent after biasing? 7
(c) What is the smallest value $M$ of the significand? $\qquad$ 1
(d) What is the largest value $M$ of the significand? 127/64

## Part III

Fill in the blank entries in the following table giving the encodings for some interesting numbers.

| Description | $E$ | M | V | Binary Encoding |
| :---: | :---: | :---: | :---: | :---: |
| Zero | -6 | 0 | 0 | 00000000000 |
| Smallest Positive (nonzero) | -6 | 1/64 | 1/4096 | 00000000001 |
| Largest denormalized | -6 | 63/64 | 63/4096 | 00000111111 |
| Smallest positive normalized | -6 | 1 | 1/64 | 00001000000 |
| Positive Infinity | - | - | - | 01111000000 |
| Negative Infinity | - | - | - | 11111000000 |

## Problem 3. (12 points):

Consider the following assembly code:


The assembly on the previous page corresponds to the C code below. Fill in the blanks in the C code to match the operations done in the assembly.

```
unsigned int func(unsigned int a, unsigned int b)
{
    unsigned int result = 0;
    while (result < a*b ) {
        result += a*b
        result -= func(__(a+b)/2__ a-1 );
    }
    return result;
}
```


## Problem 4. ( 8 points):

Consider the following C declarations:

```
struct a_struct {
    char a;
    struct a_strcut *b;
};
struct b_struct {
    char c;
    int i;
    double * d;
    short e[3];
    struct a_struct m;
};
```

A. Using the templates below (allowing a maximum of 32 bytes), indicate the allocation of data for struct b_struct. Mark off and label the areas for each individual element (arrays may be labeled as a single element). Cross hatch the parts that are allocated, but not used, and be sure to clearly indicate the end of the structure. Assume the Linux alignment rules discussed in class.

```
struct b_struct:
```


B. How would you define the struct b2_st ruct structure to minimize the number of bytes allocated for the structure using the same fields as the struct b_struct structure?
struct b2_struct \{

C. What is the value of sizeof(struct b2_struct)? 24

## Problem 5. (12 points):

Consider the following C code:

```
struct triple
{
    int x;
    char c;
    int y;
};
int mystery1(int x);
int mystery2(int x);
int mystery3(struct triple* t);
int main()
{
    struct triple t = {35, 'q', 10};
    int result1 = mystery1(42);
    int result2 = mystery2(19);
    int result3 = mystery3(&t);
    printf("result1 = %d\n", result1);
    printf("result2 = %d\n", result2);
    printf("result3 = %d\n", result3);
    return 0;
}
```

Using the assembly code for mystery1, mystery2, and mystery 3 on the next page, fill in the proper values in this program's output:

```
result1 =
```

$\qquad$

```
result2 =
```

$\qquad$

```
result3 \(=\)
``` \(\qquad\)
\begin{tabular}{|c|c|c|c|c|}
\hline \multicolumn{5}{|l|}{\(080483 d 0\) <mystery1>:} \\
\hline 80483d0: & 55 & & push & \%ebp \\
\hline 80483d1: & 89 e 5 & & mov & \%esp, \%ebp \\
\hline 80483d3: & 53 & & push & \%ebx \\
\hline 80483d4: & 8 b 45 & 08 & mov & \(0 \times 8\) (\%ebp), \%eax \\
\hline 80483d7: & 89 c & & mov & \%eax, \%ebx \\
\hline 80483d9: & 83 e & 01 & and & \$0x1, \%ebx \\
\hline 80483dc: & 85 c & & test & \%eax, \%eax \\
\hline 80483 de: & 74 0b & & je & 80483 eb <mystery1+0x1b> \\
\hline 80483e0: & c1 f8 & 01 & sar & \$0x1, \%eax \\
\hline 80483e3: & 50 & & push & \%eax \\
\hline 80483e4: & e8 e7 & \(f f\) ff ff & call & \(80483 d 0\) <mystery1> \\
\hline 80483e9: & 01 c & & add & \%eax, \%ebx \\
\hline 80483 eb : & 89 d8 & & mov & \%ebx, \%eax \\
\hline 80483 ed : & 8 b 5 & fc & mov & 0xfffffffc (\%ebp), \%ebx \\
\hline 80483f0: & c9 & & leav & \\
\hline 80483f1: & c3 & & ret & \\
\hline \multicolumn{5}{|l|}{080483f4 <mystery2>:} \\
\hline 80483f4: & 55 & & push & \%ebp \\
\hline 80483f5: & 89 e 5 & & mov & \%esp, \%ebp \\
\hline 80483f7: & 8 b 5 & 08 & mov & 0 x 8 (\%ebp), \%edx \\
\hline 80483fa: & 31 c & & xor & \%eax, \%eax \\
\hline 80483fc: & 85 d & & test & \%edx, \%edx \\
\hline 80483fe: & 7 e 0 & & jle & 8048406 <mystery2+0x12> \\
\hline 8048400 : & 40 & & inc & \%eax \\
\hline 8048401 : & 4 a & & dec & \%edx \\
\hline 8048402 : & 85 d & & test & \%edx, \%edx \\
\hline 8048404: & 7 f f & & jg & 8048400 <mystery2+0xc> \\
\hline 8048406 : & c9 & & leav & \\
\hline 8048407 : & c3 & & ret & \\
\hline \multicolumn{5}{|l|}{08048408 <mystery3>:} \\
\hline 8048408: & 55 & & push & \%ebp \\
\hline 8048409: & 89 e & & mov & \%esp, \%ebp \\
\hline 804840 b : & 8 b 45 & 08 & mov & \(0 \times 8\) (\%ebp), \%eax \\
\hline 804840e: & 8 b 10 & & mov & (\%eax), \%edx \\
\hline 8048410: & \(0 f a\) & 5008 & imul & 0 x 8 (\%eax), \%edx \\
\hline 8048414: & 89 d0 & & mov & \%edx, \%eax \\
\hline 8048416: & c9 & & leav & \\
\hline 8048417: & c3 & & ret & \\
\hline
\end{tabular}

\section*{Problem 6. ( 12 points):}

This problem tests your understanding of byte ordering and the stack discipline. The following program reads a string from standard input and prints an integer in it's hexadecimal format based on the input it was given.
```

\#include <stdio.h>
int get_key () {
int key;
scanf ("%s", \&key);
return key;
}
int main () {
printf ("0x%8x\n", get_key());
return 0;
}

```

Here is the corresponding machine code on a Linux/x86 machine:
```

08048414 <get_key>:
8048414: 55 mell
8048417: 83 ec 18 sub \$0x18,%esp
804841a: 83 c4 f8 add \$0xffffffff8,%esp
804841d: 8d 45 fc lea 0xfffffffc(%ebp),%eax
8048420: 50
8048421: 68 b8 84 04 08 push \$0x80484b8 format string for scanf
8048426: e8 e1 fe ff ff call 804830c <_init+0x50> call scanf
804842b: 8b 45 fc mov 0xffffffffc(%ebp),%eax
804842e: 89 ec mov %ebp,%esp
8048430: 5d pop %ebp
8048431: c3 ret
08048434 <main>:
8048434: 55 push %ebp
8048435: 89 e5 mov %esp,%ebp
8048437: 83 ec 08 sub \$0x8,%esp
804843a: 83 c4 f8 add \$0xffffffff8,%esp
804843d: e8 d2 ff ff ff call 8048414 <get_key>
8048442: 50 push %eax integer arg of printf
8048443: 68 bb 84 04 08 push \$0x80484bb format string for printf
8048448: e8 ef fe ff ff call 804833c <_init+0x80> call printf
804844d: 31 c0 xor %eax,%eax
804844f: 89 ec mov %ebp,%esp
8048451: 5d pop %ebp
8048452: c3 ret

```

Here are a few notes to help you with the problem:
- scanf ( \({ }^{\prime} \% s^{\prime \prime}\), i) reads a string from the standard input stream and stores it at stores it at address \(i\) (including the terminating ' \(\backslash 0^{\prime}\) character. It does not check the size of the destination buffer.
- printf (" \(0 \mathrm{x} \% 8 \mathrm{x} \backslash \mathrm{n} ", \mathrm{j})\) prints 8 digits of the integer i in hexadecimal format as \(0 \times x x x x x x\)
- Recall that Linux/x86 machines are Little Endian.
- You will need to know the hex values for the following characters:
\begin{tabular}{|c|c|c|c|}
\hline Character & Hex Value & Character & Hex Value \\
\hline 'E' & 0x45 & 's' & 0x73 \\
\hline 'v' & 0×76 & 'h' & 0x68 \\
\hline 'i' & 0x69 & '!' & \(0 \times 21\) \\
\hline '1' & 0x6c & \(\prime \backslash 0 \prime\) & \(0 \times 00\) \\
\hline
\end{tabular}
A. Suppose we run this program on a Linux/x86 machine with the input string "Evil!".

Here is a template for the stack, showing the location of key. Indicate with a labelled arrow where \%ebp points to, and fill in the stack with the values that were just read in after the call to scanf (addresses increase from left to right).


What is the 4-byte integer (in hex) printed by print \(f\) inside main?
0x \(\qquad\)
\(6 c 697645\)
B. Suppose we instead gave it the input string "Evillish!".

For the remaining problems, each answer should be an unsigned 4-byte integer expressed as 8 hexadecimal digits.
(a) What is the value of (\&key) [1] just after scanf returns to get key?
(\&key) [1] = 0x_ 6873696c
(b) What is the value of \%ebp immediately before the execution of the ret instruction of get key? \(\% e b p=0 x \underline{6873696 c}\)

You can use the following template of the stack as scratch space. This will not be considered for credit.


\section*{Problem 7. (6 points):}

Consider the following code for a matrix multiplication function:
```

for (i=0; i<n; i++) {
for (j=0; j<n; j++) {
sum = 0.0;
for (k=0; k<m; k++)
sum += a[i][k] * b[k][j];
c[i][j] = sum;
}
}

```
with matricies \(a[n][m], b[m][n]\) and \(c[n, n]\).
1. Assume that \(m\) is twice as large as \(n\). Is the above loop optimally arranged for preserving locality? If not, state the optimal nesting of these loops.

No,
Nest them as
```

for ( ... i ...)
for ( ... k ... )
for ( ... j ... )

```
see class 13 , slides 30-40
2. Assume that n is twice as large as m . Is the above loop optimally arranged for preserving locality? If not, state the optimal nesting of these loops.

No,
Nest them as
```

for ( ... i ...)
for ( ... k ... )
for ( ... j ... )

```
see class 13 , slides 30-40

\section*{Problem 8. (7 points):}

Answer true or false for each of the statements below. For full credit your answer must be correct and you must write the entire word (either true or false in the answer space.
1. If it takes less time to do either a cache access or a main memory access than to perform a branch, then loop unrolling always improves performance.
2. On the Fish machines the most effective way to time a short routine is to use the cycle counter.
3. Loop invariant code motion can be applied to expressions involv-
false ing the loop induction variable

4a. typedef int (*a) (int *);
typedef a b[10];
typedef b* (*c) ();
c d;
int *(*(*) (int *)) [10] (*e) ();
d and e have the same type.

4b. e and c have the same type.
either
false
5. In C, the variable \(f\) is declared as: int \(f[12][17]\). The ad-
false dress for \(\mathrm{f}[3][8]\) can sometimes be greater than the address for f[8][3].
6. The largest possible finite denormalized IEEE floating number is
false greater than the smallest possible positive normalized IEEE floating number:```

