

CS 7B - Spring 2019 - Final Exam Solutions

1. Write the number of the definition on the right next to the term it defines.

- | | |
|---------------------------------|---|
| (a) copy <u>3</u> | (1) (1) a value used to identify a typed object in memory; (2) a variable holding such a value. |
| (b) overload <u>4</u> | (2) an operation that transfers a value from one object to another, leaving behind a value representing “empty.” |
| (c) container <u>12</u> | (3) An operation that makes two objects have values that compare equal.. |
| (d) pointer <u>1</u> | (4) Define two functions or operators with the same name but different argument (operand) types. |
| (e) reference <u>8</u> | (5) A user-defined type that may contain data members, function members, and member types. |
| (f) class <u>5</u> | (6) An operation that initializes an object. Typically establishes an invariant and often acquires resources needed for an object to be used (which are then typically released by a destructor). |
| (g) invariant <u>10</u> | (7) The region of program text (source code) in which a name can be referred to. |
| (h) type <u>9</u> | (8) (1) a value describing the location of a typed value in memory; (2) a variable holding such a value. |
| (i) byte <u>11</u> | (9) Something that defines a set of possible values and a set of operations for an object. |
| (j) constructor <u>6</u> | (10) Something that must be always true at a given point (or points) of a program; typically used to describe the state (set of values) of an object or the state of a loop before entry into the repeated statement. |
| (k) scope <u>7</u> | (11) The basic unit of addressing in most computers. |
| (l) move <u>2</u> | (12) An object that holds elements (other objects). |

2. Write the output of the following C++ code, which uses several kinds of C++ parameter passing:

```
1 int mystery(int a, int* b, int& c) {
2     a++;
3     (*b)++;
4     c++;
5     return a;
6 }
7
8 int main() {
9     int a{0}, b{0}, c{0}, d{0};
10    mystery(a, &b, c);
11    cout << a << " " << b << " " << c << " " << d << endl;
12    mystery(c, &d, a);
13    cout << a << " " << b << " " << c << " " << d << endl;
14    c = mystery(b, &a, d);
15    cout << a << " " << b << " " << c << " " << d << endl;
16    a = mystery(a, &a, a);
17    cout << a << " " << b << " " << c << " " << d << endl;
18 }
```

SOLN:

Variables `a`, `b`, `c`, and `d` are initialized to zero and then `a`, `b`, and `c` are passed to the function `mystery()`, with `a` passed by value, `&b` is passed by its address (which is kind of like being passed by reference, but `mystery()` then interprets it as a pointer variable) and `c` is passed by reference. So `mystery()` will leave `a` unchanged (it increments a copy of `a`, and that incremented copy is returned by `mystery()`, but the function call on line 10, does nothing with the returned value). So the value of `a` in the scope of `main()` is still 0 on line 11. On the other hand, `mystery()` dereferences the pointer to `b` and increments that value, so in the scope of `main()`, `b` is now 1, and since `c` is passed by reference and incremented in `mystery()`, `c` will have the value 1 on line 11. Nothing was done with `d`, so it's still 0 and the output from line 11 is "0 1 1 0".

The call to `mystery()` on line 12 passes `c=1` by value, which `mystery()` makes a local copy of (renamed `a`) increments to 2 and returns, but the call on line 12 does nothing with that returned value. The address of `d=0` is passed by value to a pointer to variable `mystery()` names `b`, which is dereferenced and incremented on line 3, so on line 13, the value of `d=1`. The variable `a` is passed by reference and renamed `c`, which is incremented so that on line 13, `a=1`. Since `b=1` was not affected by that, the output on line 13 is "1 1 1 1".

The call to `mystery()` on line 14 takes `b=1` by value, renames it `a`, increments that value and returns it, thereby assigning 2 to `c`. Meanwhile, the address of `a=1` is passed by value and renamed `b`, a pointer to an `int`, which is then dereferenced and incremented to 2, so on line 15, `b = 2`. Also, the call on line 14 passes `d=1` by references, renames is `c` and increments that value to 2, so that `d = 2`. The only variable unchanged then is `b=1`, which was passed in by value. So the output on line 15 is "2 1 2 2".

On line 16, the only thing that really matters is that `a=2` is passed in by value and the local copy is incremented and returned so that the value assigned to `a` is 3. This happens after `a` is incremented to 4 in the body of `mystery()`, so none of that matters and the output on line 17 is "3 1 2 2", giving the total output of

```
0 1 1 0
1 1 1 1
2 1 2 2
3 1 2 2
```

3. Consider the following complete program

```

1 #include <iostream>
2 #include <vector>
3 #include <chrono>
4 using namespace std;

6 int f(int);
  int f(int num, vector<int>& memoizeF);
8 int64_t f(int64_t num, int64_t *memoizeF);

10 int main() {
    int num;
12    int64_t num64{ 0 };
    vector<int> memoizeV{ 0,1 };
14    // -1 indicates that the value is not within the array
    while (cin >> num) {
16        int64_t* memoizeP = new int64_t[num + 1]{ 0, 1 };
        for (int i = 2; i <= num; i++) {
18            memoizeV.push_back(-1); // -1 means not done

```

```

20     memoizeP[i] = -1;
21 }
22 num64 = num; //promotion
23 auto t1 = std::chrono::system_clock::now();
24 cout << "The time to compute fibonacci(" << num << ") = "
25     << f(num) << " with dumb recursion: ";
26 auto t2 = std::chrono::system_clock::now();
27 cout << (t2 - t1).count() << endl;
28 t1 = std::chrono::system_clock::now();
29 cout << "The time to compute fibonacci(" << num << ") = "
30     << f(num, memoizeV) << " with memoization vector: ";
31 t2 = std::chrono::system_clock::now();
32 cout << (t2 - t1).count() << endl;
33 t1 = std::chrono::system_clock::now();
34 cout << "The time to compute fibonacci(" << num << ") = "
35     << f(num64, memoizeP) << " with memoization pointer: ";
36 t2 = std::chrono::system_clock::now();
37 cout << (t2 - t1).count() << endl;
38 delete [] memoizeP;
39 }
40 }
41
42 int64_t f(int64_t num, int64_t* memoizeP) {
43     if (memoizeP[num] == -1) {
44         memoizeP[num] = f(num-1, memoizeP) + f(num-2, memoizeP);
45     }
46     return memoizeP[num];
47 }
48
49 int f(int num, vector<int>& memoizeV) {
50     if (memoizeV[num] == -1) {
51         memoizeV[num] = f(num-1, memoizeV) + f(num-2, memoizeV);
52     }
53     return memoizeV[num];
54 }
55
56 int f(int n) {
57     if (n == 0 || n == 1) return n;
58     else return f(n - 1) + f(n - 2);
59 }

```

- (a) Describe the parameters that distinguish the three overloaded versions of `fibonacci()`.

SOLN: The first one just takes a single `int`. The second takes an `int` and reference to a `vector` of `int`—the third a 64-bit `int` and a pointer to a 64-bit `int`.

- (b) Describe what happens on line 13.

SOLN: A `vector` of two `ints` is created and initialized to contain `{0,1}`.

- (c) Describe, in detail, the effect of the command on line 16. **SOLN:** Enough contiguous memory is allocated on the heap to contain `num+1` 64-bit `ints` and the memory address of the first one is assigned to the variable `memoizeP`. Also, the first two values are initialized as 0 and 1.

- (d) Why does line 16 need to be inside the `while`-loop but line 13 does not?

SOLN: To avoid a memory leak, the memory allocated for the `memoizeP` array is freed at the end of the `while` loop and so needs to be reallocated and initialized. Note that the same issue kind of exists for the `memoizeV` `vector`, except that it's not dynamically allocated (at least not in a way that's transparent to the coder), so we can just let the accumulated Fibonacci values take a ride for the next cycle.

- (e) Describe what the containers `memoizeV` and `memoizeP` do that makes the computation of `f()` more efficient.

SOLN: Lines 17, 18 and 19 fill the next `num-2` elements of the vector `memoizeV` and the array `memoizeP` with the flag value “-1”, to indicate that these Fibonacci values have not yet been computed, then when the recursive functions `f()` are called, a new Fibonacci value is computed from previous values if, and only if, the Fibonacci value has not yet been computed. Since this leaves many fewer function calls on the call stack, and since looking up stored results and returning the cached result when the same inputs occur again is much less expensive in both time and space than expensive function calls, memoization is a great optimization.

- (f) A sampling of output for this program as the user enters 20,30,40 and 50 is shown below. Note the use of the `chrono` library to calculate the times of executions (measured in ticks).

20

```
The time to compute fibonacci(20) = 6765 with dumb recursion: 69501
The time to compute fibonacci(20) = 6765 with memoization vector: 46258
The time to compute fibonacci(20) = 6765 with memoization pointer: 40211
```

30

```
The time to compute fibonacci(30) = 832040 with dumb recursion: 974679
The time to compute fibonacci(30) = 832040 with memoization vector: 60426
The time to compute fibonacci(30) = 832040 with memoization pointer: 61168
```

40

```
The time to compute fibonacci(40) = 102334155 with dumb recursion: 121545885
The time to compute fibonacci(40) = 102334155 with memoization vector: 64019
The time to compute fibonacci(40) = 102334155 with memoization pointer: 68222
```

50

```
The time to compute fibonacci(50) = -298632863 with dumb recursion: 15006249361
The time to compute fibonacci(50) = -298632863 with memoization vector: 50684
The time to compute fibonacci(50) = 12586269025 with memoization pointer: 47437
```

Compare the growth in the time of computation for the three overloaded versions of `f()`.

SOLN: With memoization, the time does not seem to grow with the size of the input, first going up by about 50%, then more like 10%, and then, strangely...going down! Without the memoization, however, the ratios of time are $\frac{974679}{69501} \approx 1400\%$, $\frac{121545885}{974679} \approx 12500\%$ and $\frac{15006249361}{121545885} \approx 12500\%$. Those are huuuge leaps!

- (g) Why do the “dumb recursion” and “memoization vector” versions give an incorrect value for `f(50)` but the “memoization point” version gives the correct value?

SOLN: That’s a memory overflow error, since $F_{50} = 12586269025 \geq 2^{31}$, a 32-bit `int` can’t contain the value.

4. Suppose the following C++ classes have been declared:

<pre> 1 class C1 { 2 public: 3 void m1() { 4 cout << "C1 m1 "; 5 } 6 void m2() { 7 cout << "C1 m2 "; 8 m3(); 9 } 10 void m3() { 11 cout << "C1 m3 "; 12 m1(); 13 } 14 }; </pre>	<pre> 1 class C2 : public C1 { 2 public: 3 void m1() { 4 cout << "C2 m1 "; 5 } 6 void m2() { 7 C1::m3(); 8 cout << "C2 m2 "; 9 } 10 void m3() { 11 cout << "C2 m3 "; 12 m1(); 13 } 14 }; </pre>
--	--

and that a client program declares the following variables:

```

C1* var1 = new C2();
C2* var2 = new C2();

```

Write the output that would be produced by each of the following calls, as it would appear on the console.

```

var1->m1();      C1 m1
var1->m2();      C1 m2
                 C1 m3
                 C1 m1
var1->m3();      C1 m3
                 C1 m1
var2->m1();      C2 m1
var2->m2();      C1 m3
                 C1 m1
                 C2 m2
var2->m3();      C2 m3
                 C2 m1

```

None of these are terribly surprising. What if, however we make `m3()` a virtual function by modifying it in both class definitions. In `C1`:

```

virtual void m3() {
    cout << "C1 m3  " << endl;;
    m1();
}

```

...and in `C2`:

```

virtual void m3() {
    cout << "C2 m3  " << endl;;
    m1();
}

```

Then, since C2's definition of `m3()` overrides C1's definition, we get

```
var1->m3();    C2 m3
              C2 m1
```

5. Assume that an integer and a pointer each takes 4 bytes. Also, assume that there is no alignment in objects.

Predict the output following program. Why?

```
#include <iostream>
2 using namespace std;

4 class Test {
    int x;
6     int *ptr;
    int y;
8 };

10 int main() {
    Test t;
12     cout << sizeof(t) << " ";
    cout << sizeof(Test *) << " ";
14 }
```

SOLN: The output is

12 4

The "12" indicates that an object of type `Test` has three member variables, each one is an `int` requiring 4 bytes, so, in all, 12 bytes is the size of a variable of type `Test`. A pointer to a variable of type `Test` is the address, which on a 64-bit machine, takes 8 bytes. On a 32-bit machine, or with a compiler using a 32-bit platform, this would be "4."

6. Write a program to print the area and perimeter of a triangle having sides of 3, 4 and 5 units by
- creating a class named 'Triangle' with a function to print the area and perimeter.

SOLN:

```
#include <iostream>
2 #include <string>
#include <cmath>
4 using namespace std;

6 class Triangle
{
8 public:
    void print_area(int s1, int s2, int s3)    {
10         double s = (s1+s2+s3)/2.0;
        cout << "Area is " << sqrt((s-s1)*(s-s2)*(s-s3)*s) << endl; //Heron's
        formula
12         cout << "Perimeter is " << s << endl;
    }
14 };

16 int main() {
    Triangle t;
18     t.print_area(3,4,5);
    return 0;
20 }
```

- (b) creating a class named 'Triangle' with the constructor having the three sides as its parameters.

SOLN:

```

1 #include <iostream>
2 #include <string>
3 #include <cmath>
4 using namespace std;

6 class Triangle {
7 public:
8     int s1,s2,s3;
9     Triangle(int a,int b,int c) : s1(a), s2(b), s3(c) { }
10    void print_area() {
11        double s = (s1+s2+s3)/2.0;
12        cout << "Area is " << sqrt((s-s1)*(s-s2)*(s-s3)*s) << endl; //Heron's
13        formula
14        cout << "Perimeter is " << s << endl;
15    }
16 };
17 int main() {
18     Triangle t(3,4,5);
19     t.print_area();
20     return 0;
21 }

```

7. Print the sum, difference and product of two complex numbers by creating a class named 'Complex' with separate functions for each operation whose real and imaginary parts are entered by the user.

SOLN: This is pretty bare bones—your fractal project contained more robust code!

```

1 #include <iostream>
2 using namespace std;
3 class Complex {
4 private:
5     int real;
6     int imag;
7 public:
8     Complex(int r, int i) real(r), imag(i) { }
9     int get_real() {
10        return real;
11    }
12    int get_imag() {
13        return imag;
14    }
15    void add(Complex c1) {
16        cout << c1.get_real() + real << "+i" << c1.get_imag() + imag << endl;
17    }
18    void difference(Complex c1) {
19        cout << real - c1.get_real() << "+i" << imag - c1.get_imag() << endl;
20    }
21    void multiply(Complex c1) {
22        cout << ((real*c1.get_real()) - (imag*c1.get_imag())) << "+i" <<
23        ((real*c1.get_imag()) + (imag*c1.get_real())) << endl;
24    }
25 };
26 int main() {
27     Complex c1(4, 5);
28     Complex c2(2, 3);
29     c1.add(c2);
30     c1.difference(c2);
31     c1.multiply(c2);
32 }

```

