CMSC 132A, Final exam (Practice)

Spring 2018

NAME:_____

UID: _____

Question	Points
1	15
2	15
3	15
4	15
5	15
6	15
Total:	90

This test is open-book, open-notes, but you may not use any computing device other than your brain and may not communicate with anyone. You have 120 minutes to complete the test.

The phrase "design a program" means follow the steps of the design recipe. Unless specifically asked for, you do not need to provide intermediate products like templates or stubs, though they may be useful to help you construct correct solutions.

You may use any of the data definitions given to you within this exam and do not need to repeat their definitions.

Unless specifically instructed otherwise, you may use any Java language features we have seen in class.

You do not need to write visibility modifiers such as `public` and `private`.

When writing tests, you may use a shorthand for writing check-expects by drawing an arrow between two expressions to mean you expect the first to evaluate to same result as the second. For example, you may write `"foo".length()` $\rightarrow$ 3 instead of `t.checkExpect("foo".length(), 3)`. You do not have to place tests inside a test class.

Problem 1 (15 points). Succintly answer each of the following questions:

If a (concrete, non-abstract) class `C` implements interface `I` but doesn't actually define all of the methods listed in `I` within the definition of `C`, what can you conclude about `C` (assuming the Java type-checker accepts this program)?

If you have a method `Integer m(String s)`, what do you know about possible inputs to the method when the program is run?

Suppose you have a program that contains the following interface and class definitions (and potentially others):

```
interface I {
  Boolean m1();
  Boolean m2();
}
abstract class A implements I {
  public Boolean m1() { return this.m2(); }
}
class B extends A {
  public Boolean m2() { return true; }
}
```

If you have an object o of type `A`, what does `o.m1()` produce?

Besides when multiple classes have identical method definitions, when (if ever) does it make sense to define methods in an abstract class?

List three *conceptual* differences between Java and ISL. (Examples of non-conceptual differences: the syntax, the IDE, the number of programmers, etc.)

Problem 2 (15 points). Here's an idea for a sorting method that, given a comparator, sorts a list into ascending order according the comparator. This method is implemented in the `AListof<X>` abstract class:

```
// Sort this list in ascending order according to c
public Listof<X> sort(Comparator<X> c) {
    // Convert this list into a list of singleton lists
    // For example [1, 2, 3] --> [[1], [2], [3]]
    // Note: each singleton list is a sorted list.
    Listof<Listof<X>> singletons =
      this.map(x -> new Cons<>(x, new Empty<>()));

    // Combining function to merge two sorted lists into one.
    BiFunction<Listof<X>, Listof<X>, Listof<X>> combine =
            (s1, s2) -> s1.merge(s2, c);

    // Fold over singleton lists, merging into a sorted list
    return singletons.foldr(combine, new Empty<X>());
}
```

This method works (although it's not a very efficient algorithm!), assuming the following `merge` method in the `Listof<X>` interface:

```
// Merge all the elements of this list and that list into a sorted list
// ASSUME: this and that are already sorted according to c.
Listof<X> merge(Listof<X> that, Comparator<X> c);
```

Some examples (written in exam shorthand notation):

```
Comparison<Integer> lt = (i, j) -> i - j;
[1,2,3].merge([4,5,6], lt) --> [1,2,3,4,5,6]
[1,3,5].merge([2,3,6], lt) --> [1,2,3,4,5,6]
[4,5,6].merge([1,2,3], lt) --> [1,2,3,4,5,6]
```

Design the `merge` method so that `sort` works. You may design and add any needed helper methods for the `Listof<X>` interface, but you cannot change the original `sort` method (nor use it!). For full credit, your `merge` method should take time linear in the length of the resulting list, i.e. the sum of the length of this and that.

[Space for problem 2.]

Problem 3 (15 points). Here is the usual definition for binary trees:

```
interface BT<X> {}

class Leaf<X> implements BT<X> {}
class Node<X> implements BT<X> {
  X val;
  BT<X> left;
  BT<X> right;
  Node(X val, BT<X> left, BT<X> right) {
    this.val = val;
    this.left = left;
    this.right = right;
  }
}
```

Define the following methods to implement structural equality and respect the law of hash codes (your hashCode method should distinguish at least some nodes):

```
// Is this tree structurally equal to that tree?
boolean equals(BT<X> that);

// Hash code for this tree
int hashCode();
```

[Space for problem 3.]

Problem 4 (15 points). Here is a data representation for (a small subset of) JSON documents:

```
interface Json {}

class JStr implements Json {
  String s;
  JStr(String s) { this.s = s; }
}

class JInt implements Json {
  Integer i;
  JInt(Integer i) { this.i = i; }
}

interface LoJson extends Json {}

class JMt implements LoJson {}

class JCons implements LoJson {
  Json first;
  LoJson rest;
  JCons(Json first, LoJson rest) {
    this.first = first;
    this.rest = rest;
  }
}
```

Implement the visitor pattern for Json documents. (You do not need to write tests for this problem.)

[Space for problem 4.]

Problem 5 (15 points). Here is a representation for directed graphs:

```
// A node in a directed graph
class GNode {
    String name;
    Listof<GNode> neighbors;

    GNode(String name) {
        this.name = name;
        this.neighbors = new Empty<>();
    }

    // EFFECT: add the given node to this node's neighbor
    public void addNeighbor(GNode n) {
        this.neighbors = new Cons<>(n, this.neighbors);
    }

    // Are there any cycles in this graph reachable from this node?
    public Boolean hasCycle() { ... }
}
```

Design the hasCycle method (examples on the next page).

```
void testGraph(Tester t) {
    // +---------+
    // v         |
    // A--->B--->C-->D
    //      +--->E-->F
    //
    GNode a = new GNode("A");
    GNode b = new GNode("B");
    GNode c = new GNode("C");
    GNode d = new GNode("D");
    GNode e = new GNode("E");
    GNode f = new GNode("F");

    a.addNeighbor(b);
    b.addNeighbor(c);
    c.addNeighbor(d);
    b.addNeighbor(e);
    e.addNeighbor(f);
    c.addNeighbor(a);

    t.checkExpect(a.hasCycle(), true);
    t.checkExpect(b.hasCycle(), true);
    t.checkExpect(c.hasCycle(), true);
    t.checkExpect(d.hasCycle(), false);
    t.checkExpect(e.hasCycle(), false);
    t.checkExpect(f.hasCycle(), false);
}
```

Problem 6 (15 points). Design the follow method which is given a function produces doubles and a (non-empty) array list of elements. It produces an element of the list such that applying f to that element produces a number less than or equal to applying f to any other element in the list.

```
class Algorithm {
  // Find an element of xs that minimizes f.
  // ASSUME: xs has at least one element
  static <T> T argmin(Function<T,Double> f, ArrayList<T> xs) { ... }
}
```