NAME: ________________________________ NETID: ________________________

CS2110 Spring 2013 Final
10 May 2013

Write your name and Cornell netid. *There are 5 questions plus one extra credit question on 7 numbered pages. Check now that you have all the pages. Write your answers in the boxes provided. Use the back of the pages for workspace. Ambiguous answers will be considered incorrect. The exam is closed book and closed notes. Do not begin until instructed. You have 90 minutes. Good luck! And have a nice summer!*

	1	2	3	4	5	ext	Total
Score	/20	/20	/20	/20	/20	/5	
Grader							

1. (20 points) Your team is writing a program that searches for asteroids in photos collected by the Hubble Telescope. The program reads in images, scans for possible objects that might need more study, and prints details for each one it finds that go into a list for later examination.

An initial version of the program was able to process 6 images per second, but then you changed it into a multi-threaded version with 3 threads: (1) one reads images from a big database, (2) one scans each image for possible unknown asteroids, and (3) one prints information for anything found. Two bounded buffers were used between the stages: one holds images that have been read but not yet processed, and one holds descriptions of objects for further study until they have been printed.

Everyone is excited because with this change, the program is now running at 23.7 images per second — almost four times faster. Your job is to explain why concurrency can give such a big benefit.

(a) [5 points for the explanation, 5 more points for the nice graphic]. Explain what a thread is, why we use the term "producer/consumer" for threads used in the manner described here, and what a bounded buffer does. Also draw a picture showing the 3 threads and the 2 bounded buffers (make them 5 elements each), labeling each with its name. Design the graphic as if you intend to use it when explaining the program to a teammate.

(b) [5 points]. With just 3 threads, you got nearly a 4x speedup. Explain why threads can speed up a program. Then explain why the speedup could be larger than the number of threads that were used. *Hint: think about a thread forced to wait, like when it reads data from a disk that needs to spin up and reach the right spot before the read can be performed.*

(c) [5 points]. The bounded-buffer class we saw in lecture uses the Java "synchronized" keyword and the "wait" and "notifyAll" methods. Explain "synchronized," "wait" and "notify".

2. (20 points) True or false?

a	T	F	If we think of the Internet as a graph with vertices corresponding to routers and weighted edges corresponding to network links and their delays, we could use Dijkstra's algorithm to compute a lower bound on the delay from a source to each of a set of destinations.
b	T	F	If a graph G is connected, then the minimum spanning tree for G will also be connected.
c	T	F	We can test whether a directed graph G is a DAG by repeatedly removing vertices with in-degree 0. G is a DAG if eventually every vertex is removed.
d	T	F	If your program is crashing because of deadlocks but your computer only has one core, running it on a multicore machine might guarantee that deadlocks can no longer occur.
e	T	F	With a quantum computer, operations always have O(1) complexity.
f	T	F	Technologies like JQL allow us to express complex operations on large collections of data very concisely. A disadvantage is that these approaches are so slow: by writing the same operations by hand, we would almost always get much better performance.
g	T	F	If we use Java to perform operations on a database <u>*stored on disk*</u>, we often would need to explicitly tell Java when to save updates back into the database.
h	T	F	Consider classes A and B, with neither being a subclass of the other. If you create an ArrayList<A>() and attempt to insert an object declared to have type B into it, **you will get a compile-time error**.
i	T	F	If a function includes a try { *something;* } catch(IOException e) { handle-IOExceptions; } then the function must return null in the event that an exception does occur.
j	T	F	Suppose method X contains an exception handler for NullReferenceException, and inside the try X calls method Y, which doesn't catch NullReferenceException. Now suppose that on line of 10 of method Y, a null reference occurs. *Then after X handles the exception, Y will resume on line 11: the line after the one that caused the exception.*
k	T	F	Dijkstra's algorithm computes the "all pairs" shortest paths in a directed graph with non-negative edge weights. That is, after running the algorithm a single time, starting from some single vertex, it will have computed *MinDist(x,y)* for all vertices x and y in the graph.
l	T	F	Suppose class A implements interface T. If you pass an object of type A to a method that expects an argument of type T, the code compiles but Java will throw a runtime exception.
m	T	F	Suppose B is a subclass of A and you create a new A object. Then B's constructor will be executed automatically by Java because when a superclass is instantiated, the constructors for all the subclasses are executed too.
n	T	F	Given graphs G and H, suppose that g_0 is a randomly picked vertex in G, and h_0 is a randomly picked vertex in H. Then *any* correct implementation of the .equals() method for graph vertices *must* return true for g_0.equals(h_0) if G and H represent the same graph, and false if G and H differ in any way (e.g. different vertices or edges, different weights, etc).
o	T	F	A hashcode method that often has collisions would be legal, but could cause terrible performance for Java utilities such as HashMap and HashSet.
p	T	F	If objects x, y, and z are of type A and x.equals(y), and x.equals(z) then y.equals(z). You may assume that the implementation of "equals" is correct.
q	T	F	An invariant is a property of a data structure that is always true.
r	T	F	A recurrence relation is an equation that defines a sequence. One or more base cases are given, and then each subsequent value is defined as a function of the preceding ones.
s	T	F	One important form of inductive proof starts by proving the base case(s) and then shows that if the desired property holds for all $N_0 < k < N$, it also holds for k=N.
t	T	F	Concurrent threads that both read and update some shared object can avoid race conditions by using Java synchronization correctly.

3. (20 points) In class we learned about graph *colorings*: an assignment of "colors" to the vertices in a graph such that no two adjacent vertices have the same color. Suppose a vertex is defined by the class shown below. Your job is to provide code for some methods that would be used to find out how many colors are in use in a graph and whether or not the assignment of colors is a graph coloring. We are *not* asking you to write the code that computes the coloring: you can assume each node already has a color when these methods are called. Moreover, you may assume that the methods are *never* called with null arguments, and that the graph itself is undirected and fully connected.

```
public class Vertex {
    public enum Color { red, yellow, blue, black, green, … };    // A long list of colors
    public Color vColor;                                          // Color of this node
    public Set<Vertex> neighbors;                                 // Neighbors of this node
}
```

a. [5 points] Write the body of the method shown below

```
/** Return true if the neighbors of v have colors different from v's color.
 *   Return false if some neighbor of v has the same color as v. */
public static boolean colorCheck(Vertex v) {










}
```

b. [5 points] Write the body of the method shown below.

```
/** Return true if the current vertex coloring of graph g is a legal graph coloring, false otherwise */
private static boolean isALegalGraphColoring(Set<Vertex> g) {            // You can use colorCheck.

}
```

c. [5 points] Write the body of the recursive method shown below

```
/** Precondition: v is not in set setV of visited vertexes. Add to setC the colors in use by v and
      vertexes reachable from v along paths of vertexes not in setV. Add to setV all visited vertexes. */
private static void setOfColors(Vertex v, Set<Vertex> setV, Set<Color> setC) {

}
```

d. [5 points] Write the body of the method shown below

```
/** Return the number of distinct colors in use in v and vertexes of the graph that are reachable from v.
 *   (You are expected to use function setOfColors, above). */
public static int numberOfColors(Vertex v) {

}
```

4. (20 points)
You and your CS2110 friends have decided to get rich by selling the ultimate iPhone and Android app for rock climbers. A rock-climbing location might have a number of possible routes. A route starts at some point on the cliff, then has a series of pitches (sections), each rated by difficulty and by the amount of time required, and finally ends at the top of the cliff. *A pitch goes from some point on the cliff to some other point on the cliff; it may involve climbing sideways or even back downward, not just up, and there are routes that climb up, circle around, and then return to a point previously visited.*

a. [5 points] How do you represent a great cliff for rock climbing as a Java data structure? Give a short and clear description of the data structure you would use. If your data structure makes use of some standard Java utility or interface (FIFO queue, priority queue, set, stack, list, set, etc) you don't need to tell us how that standard Java utility works. Just tell us what you would use and why.

b. [5 points] Define: *directed acyclic graph* (*DAG*). Would the data structure you recommended in part 4a be a DAG? Explain why or why not.

c. [5 points] Define: *connected graph*. Would the data structure you recommended in part 4a be connected? Explain why or why not.

d. [5 points] A climber might want the app to recommend the fastest way to the top of the cliff that doesn't involve any pitches with difficulty greater than some limit. Tell us what problem is solved by Dijkstra's algorithm. Then explain how we can adapt this algorithm to solve the problem of recommending a route for this climber. *Note: Recall that a given cliff could have many possible starting points. The recommended route would need to include a recommended starting point and ending point.*

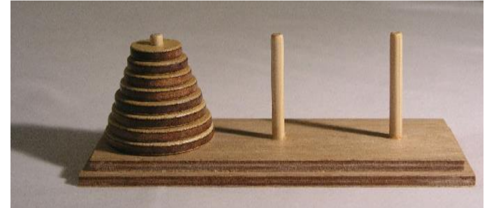

5. (20 points) To the right is a picture of a game called “Towers of Hanoi”. The game starts with a pile of disks on one of the three rods, as seen in the picture. Notice that the disks are arranged with the smallest on top and the largest on the bottom. The game involves moving the disks from the rod on the left to the rod on the right, one at a time, but without ever putting a large disk on top of a smaller one. Let’s name the rods from left to right: A, B and C.

Towers of Hanoi can be solved in a very elegant way using recursion: To move N disks from rod A to rod C, first move N-1 disks from rod A to rod B, then move one disk (this will be the biggest of them) from rod A to rod C, and finally move N-1 disks from rod B to rod A. The nice thing about this method is that because the biggest disk was moved last, it automatically satisfies the rule about never putting a big disk on a smaller disk. Moreover, as seen below, the instructions don’t need to say “which” disk to move, since the one to move is always on top. Your job is to write a method that implements this recursion, printing instructions as it does so. Here’s what it should print for a few values of N. *Hint: notice how the “move N-1 disks to the middle tower” step is performed in these three examples. Draw little pictures if that might help you build up intuition into how this works.*

With N=1: *tower(“A”, “B”, “C”, 1) prints*

Move one disk from A to C

With N=2: *tower(“A”, “B”, “C”, 2) prints*

Move one disk from A to B
Move one disk from A to C
Move one disk from B to C

With N=3: *tower(“A”, “B”, “C”, 3) prints*

Move one disk from A to C
Move one disk from A to B
Move one disk from C to B
Move one disk from A to C
Move one disk from B to A
Move one disk from B to C
Move one disk from A to C

a. [5 points].

```
/** Print instructions to move N disks from rod startRod to rod targetRod using rod otherRod as
 *  as the "middle one". (Note: the base case is when N = 0; there is nothing to print.)
 *  Precondition: No disk is on top of a smaller one. The N disks to be moved are smaller than the
 *  ones on rods targetRod and otherRod. */
public static void tower(String startRod, String otherRod, String targetRod, int N) {

}
```

b. [5 points]. We wish to figure out the cost of solving this problem, using a recurrence relation T(N). What would be the base case for T(N)?

c. [5 points]. Write T(N) as a more general recurrence function that expresses the cost of moving N disks in terms of the cost of moving N-1 disks.

d. [5 points]. Show how this recurrence relation can be solved to obtain a simple equation for the cost of moving N disks using method tower(). You can solve it in any way you wish –by algebraic manipulation, by looking at examples and inferring from them, by drawing a tree of recursive calls, or anything else.

ext. [5 points extra credit]
Suppose that an Internet provider has a graph representing all the network routers and the network links that it owns. The provider uses Prim's algorithm to compute a minimum spanning tree and provides you with this tree. The original class TreeNode is shown below. Now the provider obtains a list of (source, destination, traffic) triples, but the list includes many edges that aren't part of the minimum spanning tree. The class is also shown below. Each object of type Triple lists a node in the graph (the source for some flow of network traffic), a second node (the destination), and an integer giving the data rate, measured in megabits/second.

Write the body of the static method given below. The method is defined within the class TreeNode, and you can extend TreeNode with additional fields if you wish, use helper methods, etc.

```
public class Triple {
      public TreeNode source, destination;
      public int traffic;
}

public class TreeNode {
     public  Set<TreeNode> neighbors;      // The vertices adjacent to this one in the spanning tree


   /** v is a node of a spanning tree of a network and tL a list of traffic triples. Print
     *  the traffic load on each network link, sorted from maximum to minimum traffic.  Hint: if vertex
    *   v has a neighbor x, then x will list v as a neighbor too.  Your algorithm will need to be careful
    *  not to get stuck in a loop, visiting v, then x, then v, then x….
    */
   public static void printLinkLoads(TreeNode v, Triple[] tL) {










    }
}
```