

Name: _____

**CSE 332, Spring 2010, Final Examination
8 June 2010**

Please do not turn the page until the bell rings.

Rules:

- The exam is closed-book, closed-note. You may use a calculator *for basic arithmetic only*.
- **Please stop promptly at 4:20.**
- You can rip apart the pages, but please staple them back together before you leave.
- There are **10** questions (many with multiple parts) worth **10 points each** for a total of 100 points.

Advice:

- Read questions carefully. Understand a question before you start writing.
- **Write down thoughts and intermediate steps so you can get partial credit. But clearly circle your final answer.**
- The questions are not necessarily in order of difficulty. **Skip around.**
- If you have questions, ask.
- Relax. You are here to learn.

Name: _____

1. Suppose we sort an array of numbers, but it turns out every element of the array is the same, e.g., $\{17, 17, 17, \dots, 17\}$. (So, in hindsight, the sorting is useless.)
 - (a) What is the asymptotic running time of insertion sort in this case?
 - (b) What is the asymptotic running time of selection sort in this case?
 - (c) What is the asymptotic running time of merge sort in this case?
 - (d) What is the asymptotic running time of quick sort in this case?

Name: _____

2. Consider this code. Assume any client passes an appropriate `Comparator` object for comparing elements. These are just the assumptions we made throughout the course when computing with `Comparator` objects:

- It returns 0 for equal elements.
- It returns a negative number if the first element is “less than” the second and a positive number if the first element is “greater than” the second.
- Elements are totally ordered (transitive, anti-symmetric).


```
interface Comparator<E> {
    int compare(E e1, E e2);
}

class C {
    static <E> boolean mystery(E[] array,
                               Comparator<E> c) {
        for(int i=0; i < array.length-1; i++) {
            for(int j=i+1; j < array.length; j++) {
                if(c.compare(array[i],array[j])!=0)
                    return false;
            }
        }
        return true;
    }
}
```

- (a) In English, *what* does `mystery` compute? Do not explain *how* it computes it.
- (b) What is the worst-case running time of `mystery` in terms of n , the length of `array`?
- (c) Describe in English and/or pseudocode an algorithm for the same problem with running time $O(n \log n)$. You do not need to explain the implementation of any standard algorithms you use as part of your solution. Assume it is acceptable to mutate the contents of the array.
- (d) Now assume it is *not* acceptable to mutate the contents of the array. How can you modify your algorithm to use extra space but not change the asymptotic running time?

Name: _____

3. Consider the following undirected, weighted graph:



Step through Dijkstra's algorithm to calculate the single-source shortest paths from A to every other vertex. Show your steps in the table below. Cross out old values and write in new ones, from left to right within each cell, as the algorithm proceeds. Also list the vertices in the order which you marked them known. Finally, indicate the lowest-cost path from node A to node F.

Known vertices (in order marked known): _ _ _ _ _ _ _

Vertex	Known	Distance	Path
A			
B			
C			
D			
E			
F			
G			

Lowest-cost path from A to F: _____

Name: _____

4. In Java using the ForkJoin Framework, write code to solve the following problem:

- Input: A `String[]`
- Output: A `Pair` of (a) the number of words starting with the letter 'c' and (b) the length of the longest word starting with the letter 'c'. (If no words start with 'c', the length is irrelevant.)

Your solution should have $O(n)$ work and $O(\log n)$ span where n is the array length. Do *not* employ a sequential cut-off: the base case should process one `String`.

We have provided some of the code for you. You only need to provide one class definition:

```
import java.util.concurrent.ForkJoinPool;
import java.util.concurrent.RecursiveTask;
class Pair {
    int count;
    int longest;
    Pair(int c, int l) { count = c; longest = l; }
}
class Main{
    static final ForkJoinPool fjPool = new ForkJoinPool();
    Pair processCWords(String[] array) {
        return fjPool.invoke(new ProcessCWords(array,0,array.length));
    }
}
```

Name: _____

5. You are at a summer internship working on a program that currently takes 10 minutes to run on a 4-processor machine. Half the execution time (5 minutes) is spent running sequential code on one processor and the other half is spent running parallel code on all 4 processors. Assume the parallel code enjoys perfectly linear speedup for any number of processors.

Note/hint/warning: This does *not* mean half the work is sequential. Half the *running time* is spent on sequential code.

Your manager has a budget of \$6,000 to speed up the program. She figures that is enough money to do only one of the following:

- Buy a 16-processor machine.
- Hire a CSE332 graduate to parallelize more of the program under the highly dubious assumptions that:
 - Doing so introduces no additional overhead
 - Afterwards, there will be 1 minute of sequential work to do and the rest will enjoy perfect linear speedup.

Which approach will produce a faster program? Show your calculations, including the total *work* done by the program and the expected running time for both approaches.

Name: _____

6. (a) On our midterm we saw a somewhat complicated algorithm for finding the second-smallest element in a binary search tree. It descended to the correct node without making any modifications to the tree. Suppose this algorithm is implemented in a method `secondSmallestMidterm`. Would it be correct for two threads to execute `secondSmallestMidterm` concurrently? Explain briefly.
- (b) Here is a simpler algorithm for finding the second smallest element in a binary search tree in terms of some other operations:

```
synchronized E deleteMin() { ... }
synchronized E findMin() { ... }
synchronized void insert(E x) { ... }
E secondSmallestFinal() {
    E min = this.deleteMin();
    E ans = this.findMin();
    this.insert(min);
    return ans;
}
```

Notice `deleteMin`, `findMin`, and `insert` are synchronized methods.

- i. Suppose two threads call `secondSmallestFinal` concurrently. Demonstrate how one of them can get the wrong answer.
- ii. Does the code above have any *data races*? Explain briefly.
- iii. What is the easiest way to fix `secondSmallestFinal`?

Name: _____

7. You are designing a new social-networking site to take over the world. To handle all the volume you expect, you want to support multiple threads with a fine-grained locking strategy in which each user's profile is protected with a different lock. At the core of your system is this simple class definition:

```
class UserProfile {
    static int id_counter;
    int id; // unique for each account
    int[] friends = new int[9999]; // horrible style
    int numFriends;
    Image[] embarrassingPhotos = new Image[9999];
    UserProfile() { // constructor for new profiles
        id = id_counter++;
        numFriends = 0;
    }
    synchronized void makeFriends(UserProfile newFriend) {
        synchronized(newFriend) {
            if(numFriends == friends.length
                || newFriend.numFriends == newFriend.friends.length)
                throw new TooManyFriendsException();
            friends[numFriends++] = newFriend.id;
            newFriend.friends[newFriend.numFriends++] = id;
        }
    }
    synchronized void removeFriend(UserProfile frenemy) {
        ...
    }
}
```

- (a) The constructor has a concurrency error. What is it and how would you fix it? A short English answer is enough – no code or details required.
- (b) The `makeFriends` method has a concurrency error. What is it and how would you fix it? A short English answer is enough – no code or details required.
- (c) Rather than throwing an exception in `makeFriends` if an array is full, give two alternatives. Describe them only at a high level – a sentence or two is enough – without getting into any code details. One alternative should be easy and have nothing to do with concurrency. The other should involve concurrency.

Name: _____


8. Note: This problem is doable before or after the previous problem. This problem does not involve concurrency.

Consider again the `UserProfile` class from the previous problem. Suppose we have an array of type `UserProfile[]` that contains every user profile. Let n be the length of the array and m be the number of “friendships” in the entire system.

- (a) This array is essentially a representation of a graph.
 - i. What are the nodes and what are the edges?
 - ii. Is it more like an adjacency-list or an adjacency-matrix (e.g., in how efficiently one can perform graph operations on it)?
- (b) Suppose we want to implement a method `isConnected` that takes two numbers and determines if the accounts with these id numbers are connected by any sequence of friendships. What would be the asymptotic worst-case running time of an efficient algorithm for this problem?
- (c) Suppose you are looking over some profile information and you realize that almost all `isConnected` queries are for profiles that are actually very close to each other, meaning only a small number of friendships are needed to connect them. How should this information guide your choice for how to traverse the data such that your algorithm is more efficient in practice (even though it will not improve the asymptotic worst-case)? Be specific about how you would traverse the data.

Name: _____

9. Consider the following undirected, weighted graph (seen previously in the exam):



Apply Kruskal's algorithm to compute a minimum spanning tree (MST). In the designated spaces below, write down the edges in the order they are considered by the algorithm. If the edge is part of the MST found by the algorithm, write it down in the first list of edges that form the MST. Write down the other edges considered in the order you considered them. Assume that the algorithm terminates as soon as an MST has been found. Don't forget part (c).

(a) Edges that form part of the MST, in the order considered:

(b) Other edges considered, but not included in the MST, in the order considered:

(c) Is there another MST in addition to the one you listed in (a)? Explain.

Name: _____

10. Suppose we have an undirected, connected, weighted graph with n vertices such that:

- Every weight is an integer greater than zero.
- No two edges have the same weight.

(a) In terms of n , what is the lowest cost a minimum spanning tree (MST) could have for a graph with n nodes that meets the description above? (The cost of an MST for some n -node graphs meeting the description will be higher. The question here is what the lowest possible MST cost could be — you get to pick a graph that minimizes MST cost subject to the criteria.)

Give a closed-form solution, which just means that, for full credit, your answer should not be in terms of a series.

(b) Prove your answer to part (a) is correct for all $n \geq 1$.