Name:_____

CSE373 Fall 2013, Final Examination
December 10, 2013

Please do not turn the page until the bell rings.

Rules:

- The exam is closed-book, closed-note, closed calculator, closed electronics.

- **Please stop promptly at 4:20.**

- There are **114 points** total, distributed **unevenly** among **10** questions (many with multiple parts):

Question	Max	Earned
1	28	
2	8	
3	11	
4	15	
5	7	
6	7	
7	12	
8	8	
9	9	
10	9	

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.** Make sure you get to all the problems.

- If you have questions, ask.

- Relax. You are here to learn.

Name:_____

1. (**28** points)   **Don't miss part (b) of this question.**

   (a) Fill in this table with the **worst-case** asymptotic running time of each operation when using the data structure listed. Assume the following:

   - Items are *comparable* (given two items, one is less than, equal, or greater than the other) in $O(1)$ time.
   - For insertions, it is the client's responsibility not to insert an item if there is already an equal item in the data structure (so the operations do not need to check this).
   - For insertions, assume the data structure has enough room (do *not* include any resizing costs).
   - For deletions, assume we do *not* use lazy deletion.

	insert	lookup	delete	getMin	getMax
	(take an item and add it to the structure)	(take an item and return if it is in the structure)	(take an item and remove it from the structure, if present in it)	(return smallest item in the structure)	(return largest item in the structure)
sorted array					
unsorted array					
array kept organized as a min-heap					
AVL tree					
Hashtable with chaining for collision resolution					

   (b) The entries in the table above are worst-case. For each one, circle the entry if we expect an asymptotically better run-time in practice. The correct answer for this question circles two entries.

2. (**8** points)   Suppose there is a class `C` that has a constructor that takes an argument of type `A` and stores the content of the `A` object for later use by other methods. For each definition of class `A` below, answer "yes copy" if the constructor in `C` would need to make a deep copy of the argument to preserve the abstraction defined by `C` or "no copy" if it is not necessary.

(a)
```
public class B {
      public int x;
      public int y;
      public B(int _x, int _y) { x = _x; y = _y; }
}
public class A {
      public B a;
      public B b;
      public A(B _a, B _b) { a = _a; b = _b; }
}
```

(b)
```
public class B {
      public final int x;
      public final int y;
      public B(int _x, int _y) { x = _x; y = _y; }
}
public class A {
      public final B a;
      public final B b;
      public A(B _a, B _b) { a = _a; b = _b; }
}
```

(c)
```
public class B {
      public final int x;
      public final int y;
      public B(int _x, int _y) { x = _x; y = _y; }
}
public class A {
      public B a;
      public B b;
      public A(B _a, B _b) { a = _a; b = _b; }
}
```

(d)
```
public class B {
      public int x;
      public int y;
      public B(int _x, int _y) { x = _x; y = _y; }
}
public class A {
      public final B a;
      public final B b;
      public A(B _a, B _b) { a = _a; b = _b; }
}
```

3. (**11** points)   For each of the six questions in parts (a)-(c), answer in terms of big-Oh and the number of vertices in the graph $|V|$.

   (a) Suppose a graph has no edges.

       i. What is the asymptotic space cost of storing the graph as an adjacency list?

       ii. What is the asymptotic space cost of storing the graph as an adjacency matrix?

   (b) Suppose a graph has every possible edge.

       i. What is the asymptotic space cost of storing the graph as an adjacency list?

       ii. What is the asymptotic space cost of storing the graph as an adjacency matrix?

   (c) Suppose an undirected graph has one node A that is connected to every other node and the graph has no other edges.

       i. What is the asymptotic space cost of storing the graph as an adjacency list?

       ii. What is the asymptotic space cost of storing the graph as an adjacency matrix?

   (d) Is an adjacency list faster or slower than an adjacency matrix for answering queries of the form, "is edge (u, v) in the graph"?

   (e) Is an adjacency list faster or slower than an adjacency matrix for answering queries of the form, "are there any directed edges with u as the source node"?

4. (**15** points)

   (a) Draw a weighted undirected graph with exactly 3 nodes that has exactly 0 minimum spanning trees.

   (b) Draw a weighted undirected graph with exactly 3 nodes that has exactly 1 minimum spanning tree.

   (c) Draw a weighted undirected graph with exactly 3 nodes that has exactly 2 minimum spanning trees.

   (d) Draw a weighted undirected graph with exactly 3 nodes that has exactly 3 minimum spanning trees.

   (e) Argue in roughly 2-3 English sentences that no weighted undirected graph with 3 nodes can have more than 3 minimum spanning trees.

Name:_____

5. (**7** points)   Complete this Java method so that it properly implements **insertion sort**. Make sure the result is in-place (do not use another array).

```java
void insertionSort(int [] array) {
    for(int i=1; i < array.length; i++) {
        // YOUR CODE HERE




















    } // end of the for-loop
}
```

6. (**7** points)

Complete this Java method so that it properly implements **selection sort**. Make sure the result is in-place (do not use another array). Notice there are two places to add code.

```java
void selectionSort(int [] array) {
    for(int i=0; i < array.length; i++) {
        int index_of_least = i;
        int j;
        for(j=i+1; j < array.length; j++) {
            // YOUR CODE HERE #1




        } // end of the inner for-loop
        // YOUR CODE HERE #2





    } // end of the outer for-loop
}
```

Name:_____

7. (**12** points)   Sorting short answer: In all questions about quick-sort, assume n is large enough that any use of a cut-off is not relevant.

   (a) What is the worst-case asymptotic running time of heap-sort?

   (b) What is the worst-case asymptotic running time of merge-sort?

   (c) What is the worst-case asymptotic running time of quick-sort?

   (d) Can heap-sort be done in-place?

   (e) Can merge-sort be done in-place?

   (f) Can quick-sort be done in-place?

   (g) Consider *one* item in an array that is sorted with mergesort. In asymptotic (big-Oh) terms, how many times can that *one* item move to a different location?

   (h) What is the asymptotic running time of quick-sort if the array is already sorted (or almost sorted) and the pivot-selection strategy picks the leftmost element in the range-to-be-sorted?

   (i) What is the asymptotic running time of quick-sort if the array is already sorted (or almost sorted) and the pivot-selection strategy picks the rightmost element in the range-to-be-sorted?

   (j) What is the asymptotic running time of quick-sort if the array is already sorted (or almost sorted) and the pivot-selection strategy picks the middle element in the range-to-be-sorted?

   (k) Before starting a comparison sort for an array with n elements, how many possible orders are there for the elements in the final result?

   (l) If at some point during a comparison sort, the algorithm has enough information to narrow the final result down to k possibilities, what is the least number of possibilities that remain after performing an additional comparison?

8. (**8** points)

Suppose we use radix sort to sort the numbers below, using a radix of 10. Show the state of the sort *after each of the first two passes*, *not* after the sorting is complete. If multiple numbers are in a bucket, put the numbers that "come first" closer to the top.

Numbers to sort (in their initial order):

13, 105, 1492, 411, 1776, 1, 6, 2014, 99, 98, 96, 11, 21

Put the result after the first pass here:

0	1	2	3	4	5	6	7	8	9

Put the result after the second pass here:

0	1	2	3	4	5	6	7	8	9

9. (**9** points)

(a) Suppose you have a (large) linked list and you want to perform some operation `f` on every element and some operation `g` on every element. Your first version of the code first does all the `f` operations in one list traversal and then all `g` operations in a second traversal. Your second version does one traversal, doing `f` and then `g` for each element. Why might the second version be faster?

    i. temporal locality
    ii. spatial locality
    iii. it is faster asymptotically
    iv. none of the above

(b) Suppose you have a large array of n data items that are in a random order and you want to return a sample of 1% of them. Your first version of the code returns (a copy of) every 100^{th} item (elements at index 0, 100, 200, etc.). Your second version returns (a copy of) the first 1% of the items (elements at index 0 up to $n/100$). Why might the second version be faster?

    i. temporal locality
    ii. spatial locality
    iii. it is faster asymptotically
    iv. none of the above

(c) Suppose you have a large linked list of n integers and you want to print them in reverse order (the numbers closer to the end of the list first). The first version of your code follows this algorithm:

    • Traverse the list from the beginning to determine what n is.
    • For $i = n, n-1, n-2, ..., 1$, traverse the list from the beginning to the i^{th} element and print it

The second version of your code looks like this, calling `printReverse` on the first node in the list.

```
class ListNode {
    int x;
    ListNode next;
    void printReverse() {
      if(next != null)
        next.printReverse();
      System.out.println(x);
    }
}
```

Why might the second version be faster?

    i. temporal locality
    ii. spatial locality
    iii. it is faster asymptotically
    iv. none of the above

10. (**9** points)

(a) If you take a correct parallel program and comment out the lines that call a `join` method, which **one** of the following is most likely to happen?

    i. It could slow down because we will never have two threads execute at the same time.

    ii. It could slow down because the operating system will only assign one processor to run the program.

    iii. It could produce the wrong answer because the code that is supposed to combine answers from other threads will never run.

    iv. It could produce the wrong answer because the code that is supposed to combine answers from other threads will run before the locations it reads hold the correct data.

(b) For **each** of the following, indicate whether it can be computed with:

- A parallel map operation
- A parallel reduce operation
- Neither

    i. Replace each string in an array with an all-capitals version of the same string

    ii. The number of strings in an array that are not the empty string

    iii. The median of an array of numbers

    iv. The largest index in an array of numbers that holds a negative number

(c) According to Amdahl's Law, if 25% of a program's execution time is spent on operations that cannot be parallelized, then what is the maximum speed-up parallelism can bring to the overall program execution time?