

Name: \_\_\_\_\_

## CSE373 Winter 2014, Final Examination March 18, 2014

**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 **98 points** total, distributed **unevenly** among **10** questions (many with multiple parts):

Question	Max	Earned
1	7	
2	7	
3	8	
4	20	
5	8	
6	8	
7	16	
8	8	
9	7	
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. (7 points)

Given this code, write client code using `BankAccount` such that a call to `getOwnerAge` encounters a `NullPointerException`. Then explain briefly how to rewrite one of the methods in `BankAccount` to avoid this problem.

```
class Person {
    public String name;
    public Date birthdate; // has a Date(int year, int month, int day) constructor
}

class BankAccount {
    private Person owner;
    private float balance;
    public BankAccount(Person o, float b) {
        if (o == null || o.birthdate == null) {
            throw new IllegalArgumentException();
        }
        owner = o; balance = b;
    }

    public long getOwnerAge() {
        Date now = new Date(); // initialized to be the current date
        long millisecondsPerYear = 365 * 24 * 60 * 60 * 1000;
        // the Date getTime() method returns the date as the number
        // of milliseconds since January 1, 1970, 00:00:00 GMT
        return (now.getTime() - owner.birthdate.getTime()) / millisecondsPerYear;
    }
}
```

**Solution:**

```
Person p = new Person();
p.name = "Bob";
p.birthdate = new Date(1988, 10, 17);
BankAccount acct = new BankAccount(p, 10.0);
p.birthdate = null;
acct.getOwnerAge();
```

The constructor of `BankAccount` should do a deep copy of the `Person` object passed in.

Name: \_\_\_\_\_

2. (7 points)

Your friend wrote the code below to compute the maximum value from an array of integers. Unfortunately, your friend is not seeing the speedup they were expecting. Briefly explain why this is the case (i.e. identify the mistake they made). Then modify their code to take full advantage of the available parallelism. Indicate which lines of the original you are replacing.

```
class MaxThread extends java.lang.Thread {
    int lo; int hi; int[] arr; // arguments
    int ans = Integer.MIN_VALUE; // result
    MaxThread(int[] a, int l, int h) { ... }
    public void run() { // override
        if(hi - lo < SEQUENTIAL_CUTOFF)
            for(int i=lo; i < hi; i++)
                if (arr[i] > ans)
                    ans = arr[i];
        else {
            MaxThread left = new MaxThread(arr,lo,(hi+lo)/2);
            MaxThread right= new MaxThread(arr,(hi+lo)/2,hi);
            left.run();
            right.run();
            ans = Math.max(left.ans, right.ans);
        }
    }
}
```

```
int max(int[] arr){
    MaxThread t = new MaxThread(arr,0,arr.length);
    t.run();
    return t.ans;
}
```

**Solution:**

The current code is entirely sequential because a separate thread of execution is never created (i.e. `start()` is never called). Replace code in `else` block with:

```
MaxThread left = new MaxThread(arr,lo,(hi+lo)/2);
MaxThread right= new MaxThread(arr,(hi+lo)/2,hi);
left.start();
right.run();
left.join();
ans = Math.max(left.ans, right.ans);
```

Name: \_\_\_\_\_

3. (8 points)

You are at a summer internship working on a program that currently takes 12 minutes to run on a 2-processor machine. Two-thirds of the execution time (8 minutes) is spent running sequential code on one processor and the other third is spent running parallel code on both processors. Assume the parallel code enjoys perfectly linear speedup for any number of processors.

Note/hint/warning: This does *not* mean two-thirds of the work is sequential. Two-thirds of the *running time* is spent on sequential code.

Your manager has a budget of \$2,000 to speed up the program. She figures that is enough money to do only one of the following:

- Buy a 8-processor machine.
- Hire a CSE373 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.

**Solution:**

This is an application of Amdahl's Law. Let  $S$  be the running time for the sequential portion,  $L$  be the running time for the parallel portion on 1 processor and  $P$  be the number of processors. Then  $T_p = S + L/P$ . Initially  $T_p$  is 12 minutes,  $S$  is 8 minutes and  $P = 2$ , which means  $L$  is 8 minutes. So the total work is  $8 + 8 = 16$  minutes.

Therefore,  $T_8 = 8 + 8/8$  is 9 minutes and that's the "buy an 8-processor" option.

Under the "hire an intern" option  $S$  becomes 1 and  $L$  becomes 15 minutes. So the total time is  $1 + 15/2$  which is 8.5 minutes.

So the better choice is to hire the intern. Hurrah for interns!



**Solution:**

(a) Hashtable

insert, lookup, and remove all have a  $O(1)$  running times (credit also given for  $O(n)$ , which accounts for worst-case hash collisions)

(b) AVL Tree

insert, find, and remove all have  $O(\log n)$  running times.

Getting all students in sorted order take  $O(n)$  time (doing an in-order traversal).

(c) Union-Find

Find has a  $O(\log^* n)$  running time (credit also given for  $O(\log n)$  or almost  $O(1)$ ).

Union has a  $O(1)$  running time.

(d) Hashtable

insert, lookup, and remove all have a  $O(1)$  running times (credit also given for  $O(n)$ , which accounts for worst-case hash collisions)

(e) Stack

Pop has a  $O(1)$  running time and push has a  $O(n)$  running time (if array-based) (or a  $O(1)$  amortized running time) or a  $O(1)$  running time (if linked-list based).

Name: \_\_\_\_\_

5. (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):

17, 45, 877, 31, 7, 222, 42, 43, 301, 2525, 9, 27, 64

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

**Solution:**

After first pass:

bucket:	0	1	2	3	4	5	6	7	8	9
contents:		31	222	43	64	45		17		9
		301	42			2525		877		
								7		
								27		

After second pass:

bucket:	0	1	2	3	4	5	6	7	8	9
contents:	301	17	222	31	42		64	877		
		7	2525		43					
		9	27		45					

Name: \_\_\_\_\_

6. (8 points)

For each of the following situations, name the best sorting algorithm from among those we studied. There may be more than one answer deserving full credit, but you only need to give one answer for each.

- (a) You need a very fast sort on average, and you can only use a constant amount of extra space.
- (b) The array is in perfect sorted order.
- (c) You have a large data set, but you know all the values are between 0 and 999.
- (d) Copying your data is very fast, but comparisons are relatively slow.

Please provide an example of each of the following (or say NONE if no example exists) from among the sorting algorithms we studied.

- (e) A stable comparison sort with a worst-case  $O(n^2)$  running time
- (f) A stable comparison sort with a worst-case  $O(n)$  running time
- (g) An efficient (i.e.  $O(n \log n)$ ) stable comparison sort
- (h) A stable sort with a worst-case running time linear in  $n$

**Solution:**

- (a) Quicksort
- (b) Insertion sort
- (c) Radix sort or bucket sort
- (d) Merge sort
- (e) Insertion sort
- (f) NONE
- (g) Merge sort
- (h) Radix sort or bucket sort



Name: \_\_\_\_\_

7. (16 points)

Below are four graph-based computational tasks. For each one, **specify the type of graph** most appropriate for the data in question in terms of undirected or directed, and unweighted or weighted. In addition, **choose the graph algorithm** from the following list best suited to computing a solution: Breadth-First Search, Minimum Spanning Tree (e.g. Prim's or Kruskal's), Dijkstra's Algorithm, Topological Sort, Depth-First Search. No additional explanation is required.

- (a) You have data for the rail transport (i.e. train track) network in North America, including the length of each section of track. You need to compute the shortest (in terms of distance) route for a train traveling from New York City to Atlanta, GA.
  
- (b) You need to lay water lines for a new housing development (i.e. pipes to carry water to and from each house). You have data on all the places it's possible to build pipes, and how much pipe would be required in each case. From the many potential pipeline options, you must compute which pipes to actually build such that you use the least amount of pipe overall.
  
- (c) You are compiling a program from source code and you know the dependencies between source files (i.e. for each file  $F$ , you know which other files, if any, can't be compiled until  $F$  is compiled). You need to compute a valid compilation order such that each file is compiled after all of its dependencies are compiled.
  
- (d) You have data about friend relationships from a social network, and you want to model the spread of a rumor based on these relationships. Your model is that a rumor starts with a single person, who then tells all of their friends. Then at each subsequent step, everyone who knows about the rumor spreads it to all of their friends who don't know it. You want to compute how many people know about the rumor after  $k$  steps.

**Solution:**

- (a) Undirected, weighted  
Dijkstra's algorithm
- (b) Directed, weighted  
Minimum spanning tree
- (c) Directed, unweighted  
Topological sort
- (d) Undirected, unweighted  
Breadth-first search for  $k$  steps or depth-first search to a depth of  $k$

Name: \_\_\_\_\_

8. (8 points)

- (a) Why might it be faster to access all the elements of an array-based queue than to access all the elements of a linked-list-based queue?
- temporal locality
  - spatial locality
  - it is faster asymptotically
  - none of the above

- (b) Why might the second version of `f` be faster than the first?

```
public void f(String[] strings) {
    for (int i = 0; i < strings.length; i++) {
        strings[i] = strings[i].trim(); // omits trailing/leading whitespace
    }
    for (int i = 0; i < strings.length; i++) {
        strings[i] = strings[i].toUpperCase();
    }
}
```

```
public void f(String[] strings) {
    for (int i = 0; i < strings.length; i++) {
        strings[i] = strings[i].trim(); // omits trailing/leading whitespace
        strings[i] = strings[i].toUpperCase();
    }
}
```

- temporal locality
  - spatial locality
  - it is faster asymptotically
  - none of the above
- (c) In any graph, why might checking if an edge exists be faster when using an adjacency matrix to represent the graph instead of an adjacency list?
- temporal locality
  - spatial locality
  - it is faster asymptotically
  - none of the above
- (d) In a dense graph, why might finding all the vertices adjacent to a particular vertex be faster when using an adjacency matrix to represent the graph instead of an adjacency list?
- temporal locality
  - spatial locality
  - it is faster asymptotically
  - none of the above

**Solution:**

- (a) i. spatial locality, (b) ii. temporal locality, (c) iii. it is faster asymptotically, (d) ii. spatial locality

Name: \_\_\_\_\_

9. (7 points)

(a) A *data race* is situation where the behavior of an application depends on the sequence or timing of parallel threads. That is, the application behavior may depend on which thread "gets there first." Which of the Java `Thread` class methods can be used to help prevent data races?

- i. `run`
- ii. `join`
- iii. `hashCode`
- iv. `start`

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

- A parallel map operation
  - A parallel reduce operation
  - Neither
- i. Increment each integer in an array.
  - ii. Compute the mode (i.e. the most frequently occurring value) of an array of integers.
  - iii. Convert an array of strings to an array of integers by parsing each string as an integer.
  - iv. Compute the average of an array of numbers.
  - v. Find the minimum element in an array of numbers.

**Solution:**

- (a) ii. `join`
- (b) i. `map`
  - ii. `neither`
  - iii. `map`
  - iv. `reduce`
  - v. `reduce`

Name: \_\_\_\_\_

10. (9 points)

- (a) Which of the following would generally be the best choice for computing the hashCode of an object with multiple fields?
  - i. Use the product of the hashcodes for each field.
  - ii. Take the maximum of the hashcodes for all the fields and multiply it by a large prime number.
  - iii. Sum the hashcodes for all the fields, multiplying by a prime number before each addition.
  - iv. Randomly choose a field and use its hashCode.
- (b) Java expects a particular relationship between the equals and hashCode methods of any object. What is that relationship?
- (c) Briefly describe a good approach for generating a hashCode from a string.
- (d) Why is quicksort not a good choice for an external sorting algorithm given data stored on traditional hard drives?
- (e) Suppose we can randomize the order of an array of elements in  $O(n)$  time. For which sorting algorithms (if any) could it be worthwhile to randomize the input array before performing the sort?

**Solution:**

- (a) iii. Sum the hashcodes for all the fields, multiplying by a prime number before each addition.
- (b) if `a.equals(b)` then `a.hashCode() == b.hashCode()`.
- (c) Sum integer values corresponding to each character, multiplying by a prime before every addition.
- (d) Quicksort requires lots of non-sequential accesses to the input array (i.e. jumps around the array), and on traditional hard drives non-sequential access is much slower than sequential access.
- (e) Quicksort