

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

Email address (UWNetID): \_\_\_\_\_

## **CSE 332 Autumn 2018 Final Exam**

(closed book, closed notes, no calculators)

**Instructions:** Read the directions for each question carefully before answering. We may give partial credit based on the work you **write down**, so show your work! Use only the data structures and algorithms we have discussed in class so far. Writing after time has been called will result in a loss of points on your exam.

**Note:** For questions where you are drawing pictures, please circle your final answer.

You have 1 hour and 50 minutes, work quickly and good luck!

Total: Time: 1 hr and 50 minutes.

<b>Question</b>	<b>Max Points</b>	<b>Topic</b>
1	8	Hashing
2	7	Graphs
3	11	More Graphs
4	10	Prefix
5	14	ForkJoin
6	12	Concurrency
7	22	Code & Sorting
8	6	Speedup
9	10	P/NP
<b>Total</b>	100	

**1) [8 points total] Hash Tables**

- a) [4 pts] You are designing a **separate chaining hash table** and you are trying to decide between using a Binary Search Tree (BST) or a linked list as what each bucket should point to. List one pro and one con for each approach. Be specific and give answers as distinct from each other as possible.

BST:

Pro:

Con:

Linked List:

First describe how you would order the linked list. (circle one)

Ordered by increasing value OR most recently inserted item at the beginning

Pro:

Con:

- b) [2 pts] Insert the following elements in this order: 7, 5, 24, 57, 16, 46, 4 into the Quadratic Probing hash table below. You should use the primary hash function  $h(k) = k \% 10$ . If an item cannot be inserted into the table, please indicate this and continue inserting the remaining elements.

0	1	2	3	4	5	6	7	8	9

- c) [2 pts] The Linear Probing hash table below already contains several elements, including two that have been lazily deleted. Insert the elements 77 and 33 (in that order) into the table below. You should use the primary hash function  $h(k) = k \% 10$ .

0	1	2	3	4	5	6	7	8	9
	11		23	(deleted)	35	56	(deleted)	88	

## 2) [7 points total] Graphs!

After moving into your new apartment in the Graf district, you want to make a sandwich to celebrate, but you need to buy ingredients first. You have a map of nearby buildings and their phone numbers.

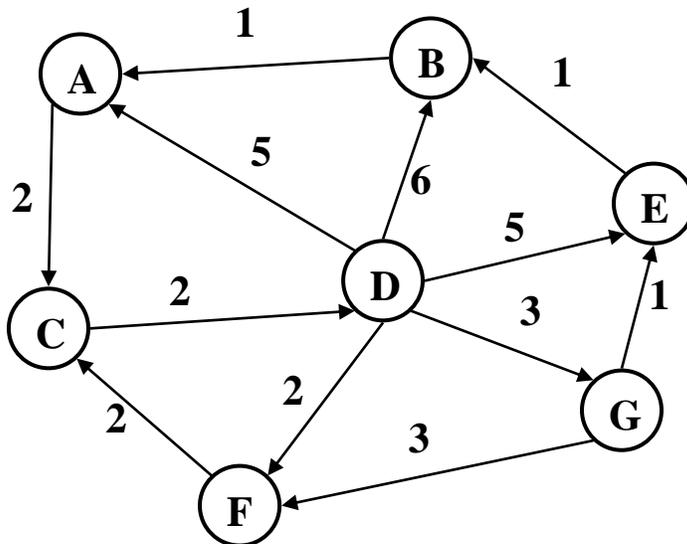
- a) [2 pts] First, you plan to call each building to find one that sells groceries. Given that edges between buildings in the Graf district are **all of length 1**, what simple algorithm should you use to efficiently find the nearest building that is a grocery store? (The name of the algorithm is sufficient. For full credit give the simplest algorithm possible.)
  
- b) [2 pts] It turns out that your phone doesn't work in the Graf district yet...so you plan to find the closest grocery store by first, just walk to all buildings 1 unit away, then walking to all buildings 2 units away, then all buildings 3 units away, and so on; going home after each iteration to rest, until you find a grocery store. What algorithm is this?
  
- c) [2 pts] As you prepare to start your journey, it occurs to you that maybe you could save time by planning a **single** route that visits each building in the Graf district **exactly once**, finally returning home! **Ignoring constant factors**, is this a good or bad idea, and **why**?

----- Below here not related to the Graf district -----

- d) [1 pts] In class we discussed two data structures to represent graphs. Give the name of the one that has a faster running time for the operation of determining if a particular edge exists.

---

3) [11 points total] **More Graphs!** Use the following graph for this problem:



a) [2 pts] List a valid **topological ordering** of the nodes in the graph above (if there are no valid orderings, state why not).

b) [5 pts] Step through Dijkstra's Algorithm to calculate the single source shortest path from A to every other vertex. For full credit, you **must show all of your steps** in the table below **by crossing through Distance and Path values that are replaced by a new value**. Break ties by choosing the lowest letter first; ex. if B and C were tied, you would explore B first. *Note that the next question asks you to recall what order vertices were declared known.*

Vertex	Known	Distance	Path
A			
B			
C			
D			
E			
F			
G			

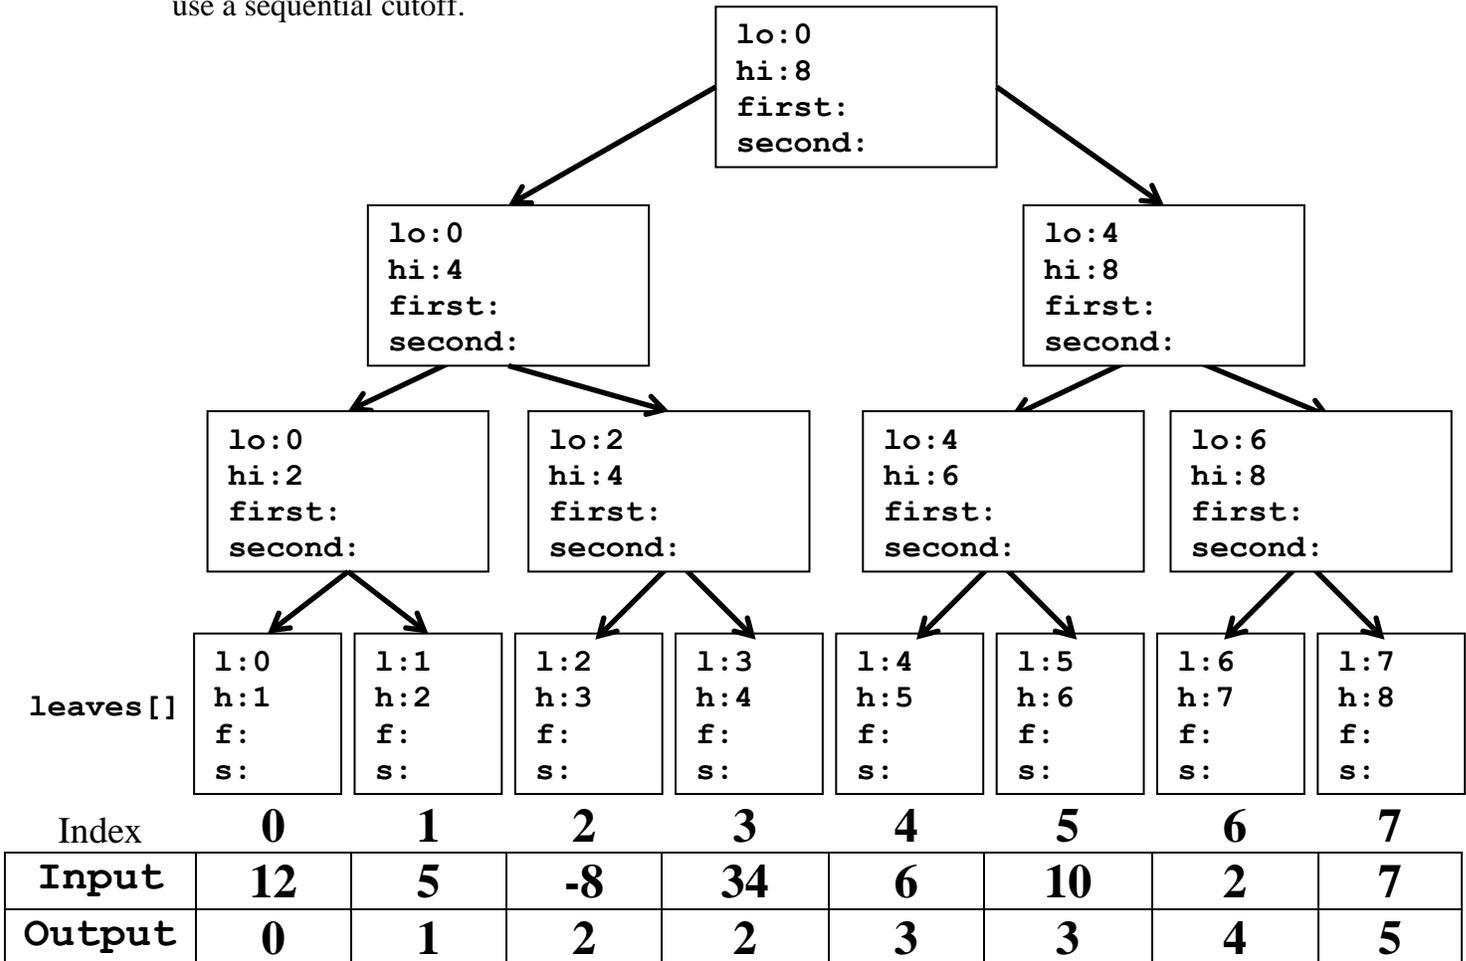
c) [1 pt] In what order would Dijkstra's algorithm mark each node as *known*?

d) [1 pt] List the shortest path from A to G? (Give the actual path NOT the cost.)

e) [2 pts] Is this graph strongly connected? (circle one)                      YES                      NO  
For any credit, briefly explain your answer.

4) [10 points] **Parallel CountUnderTen**: Given the following array as input, perform the parallel prefix algorithm to fill the **output** array with the count of only the values less than 10 contained in all of the cells to the left (including the value contained in that cell) in the input array. Values  $\geq 10$  in the input array should not contribute to the count.

a) Fill in the values for **first** and **second** in the tree below. Use these variables as you wish: *in part b) you will explain how you used them.* The output array has been filled in for you. Do not use a sequential cutoff.



b) Give formulas for the following values where **p** is a reference to a tree node *other than a leaf node* and **leaves[i]** refers to the leaf node in the tree visible just above the corresponding location in the **input** and **output** arrays in the picture above.

**p.first** =

**p.left.second** =

**p.right.second** =

**output[i]** =

c) Describe how you assigned a value to: **leaves[i].first**

5) [14 points] In Java using the ForkJoin Framework, write code to solve the following problem:

- **Input:** An array of Strings
- **Output:** Print the sum of the lengths of the strings AND the position of the longest string.

For example, if the input array is { "", "hi", "abcdef", "a"}, the program would print 9 for the sum and 2 for the position of the string "abcdef".

- Do **not** employ a sequential cut-off: **the base case should process one element.** (You can assume the input array will contain at least one string.)
- Give a class definition, LengthSumTask, **along with any other code or classes needed.**
- Fill in the function printSumAndPosition **below.**

**You may not use any global data structures or synchronization primitives (locks).**

- a) Write the code.
  - b) Answer this: Is this a **map** or a **reduction (circle one)? Why?**

```
import java.util.concurrent.ForkJoinPool;
import java.util.concurrent.RecursiveTask;
import java.util.concurrent.RecursiveAction;

class Main{
    public static final ForkJoinPool fjPool = new ForkJoinPool();

    // Print the sum of the lengths of the strings in input.
    // Also prints the position of the longest string.
    public static void printSumAndPosition (String[] input) {
        int sum, pos;
        // Your code here:

        System.out.println("Sum of lengths is: " + sum);
        System.out.println("Position of longest string is:" + pos);
    }
}
```

**Please fill in the function above and write your class(es) on the next page.**

Don't forget to answer  
b) on the previous page!

5) (Continued) Write your class(es) on this page.

**6) [12 points] Concurrency:** The `BubbleTea` class manages a bubble tea order assembled by multiple workers. Multiple threads could be accessing the same `BubbleTea` object. Assume the `Stack` objects ARE THREAD-SAFE, have enough space, and operations on them will not throw an exception.

```
public class BubbleTea {
    private Stack<String> drink = new Stack<String>();
    private Stack<String> toppings = new Stack<String>();
    private final int maxDrinkAmount = 8;

    // Checks if drink has capacity
    public boolean hasCapacity() {

        return drink.size() < maxDrinkAmount;

    }

    // Adds liquid to drink
    public void addLiquid(String liquid) {

        if (hasCapacity()) {

            if (liquid.equals("Milk")) {

                while (hasCapacity()) {

                    drink.push("Milk");

                }

            } else {

                drink.push(liquid);

            }

        }

    }

    // Adds newTop to list of toppings to add to drink
    public void addTopping(String newTop) {

        if (newTop.equals("Boba") || newTop.equals("Tapioca")) {

            toppings.push("Bubbles");

        } else {

            toppings.push(newTop);

        }

    }

}
```

**6) (Continued)**

a) Does the `BubbleTea` class above have (circle all that apply):

a race condition,          potential for deadlock,          a data race,          none of these

If there are any problems, give an example of when those problems could occur. Be specific!

b) Suppose we made the `addTopping` method **synchronized**, and changed nothing else in the code. Does this modified `BubbleTea` class above have (circle all that apply):

a race condition,          potential for deadlock,          a data race,          none of these

If there are any **FIXED** problems, describe why they are **FIXED**. If there are any **NEW** problems, give an example of when those problems could occur. Be specific!

c) Modify the code on the previous page to use locks to *allow the most concurrent access* and to avoid all of the potential problems listed above. **For full credit you must allow the most concurrent access possible without introducing any errors.** Create locks as needed. Use any reasonable names for the locking methods you call. **DO NOT use synchronized.** You should create re-entrant lock objects as follows:

```
ReentrantLock lock = new ReentrantLock();
```

d) Clearly circle all of the critical sections in your code on the previous page. **DON'T FORGET**

**7) [22 points] Code Analysis & Sorting** - A friend of yours has collected an array of delicious recipes to try baking. Each line of the recipe contains either an "ingredient" (e.g. "ingredient: eggs") or an actual step in the baking process (e.g. "step: crack the eggs into the bowl"). Recipes always have at least one ingredient and at least one step. A sample recipe with a total of 7 lines and 4 steps might look like this:

```
ingredient: butter
ingredient: sugar
step: mix the butter and sugar
ingredient: flour
step: Add flour
step: Mix well
step: Bake
```

Your friend wrote some code to sort the recipes by how many **steps** they have, so that they can try baking the recipes in order of difficulty, from easiest (fewest number of steps) to hardest (largest number of steps). Their (pseudo)-code is shown below:

```
// sorts the provided array of recipes in ascending order of difficulty
void confectionSort(recipes):
    for i from 0 to recipes.length - 1:
        for j from i to recipes.length - 1:
            if isEasier(recipes[j], recipes[i]):
                swap(recipes[i], recipes[j])

// helper function that returns true iff recipeA has fewer steps than recipeB
boolean isEasier(recipeA, recipeB):
    stepsInA = 0
    stepsInB = 0

    for line in recipeA.lines:
        if line is a step:
            stepsInA++
    for line in recipeB.lines:
        if line is a step:
            stepsInB++
    return stepsInA < stepsInB
```

Your friend wants your help analyzing and improving the code. **For each of the following four options** give the worst case Big-Oh runtime of the `confectionSort` algorithm, in terms of  $R$  (the number of recipes), and  $L$  (the maximum number of lines in any single recipe). Keep all terms of  $R$  and  $L$  in your final answers (e.g. do not assume  $R > L$ ).

a) **ORIGINAL SORT**: [Code shown above]  
**Running time:**

b) **CHANGE #1**: [Original `isEasier` function, NEW `confectionSort` function that uses a **merge-sort strategy**.]  
**Running time:**

7) (Continued)

- c) **CHANGE #2:** [NEW Modified `isEasier` function, Original `confectionSort` function.]

Describe how you would modify the `isEasier` function to improve overall `confectionSort` running time. You are allowed to add additional fields to the `recipe` objects if needed. You do not need to show code.

**Description of Change:**

**Running time:**

- d) **CHANGE #3:** [Throw everything out and start over!]

After talking some more, your friend mentions that although recipes sometimes contain many ingredients, every recipe has **at most 10 steps**. You are overjoyed! Describe what sorting algorithm your friend **should** be using.

**Description of Change/New algorithm:**

**Running time:**

- e) Is ORIGINAL `confectionSort` an in-place sort? For any credit, briefly explain why.

YES                      NO

Explanation:

- f) Is ORIGINAL `confectionSort` a stable sort? For any credit, briefly explain why.

YES                      NO

Explanation:

----- Below here not related to `confectionSort` -----

- g) [3 points] Give the **recurrence** for Quicksort (parallel sort & sequential partition) – worst case span: (Note: We are NOT asking for the closed form.)

**T(n) =**

- h) [3 points] Quicksort's partition step can also be parallelized. **What** is the big-O span of a single parallel partition? For full credit, **explain why** it has that span?

**8) [6 points] Speedup**

What *fraction of a program must be parallelizable* in order to get 5x speedup on 15 processors?

**You must show your work for any credit.** For full credit give your answer as a number or a simplified fraction (not a formula).

9) [10 points] **P, NP, NP-Complete**

a) [1 point] “NP” stands for \_\_\_\_\_

b) [2 points] What does it mean for a problem to be NP-complete?

c) [4 points] Draw a diagram demonstrating how (we are pretty sure) the sets P, NP, and EXP overlap/don't overlap with each other. Label each set clearly.

d) [3 points] Provide an example of ONE (if you list more than one you will get zero points) problem from each of the three sets below. Give a DIFFERENT problem for each set.

P: \_\_\_\_\_

NP: \_\_\_\_\_

EXP: \_\_\_\_\_

----- SCRATCH PAPER -----