

Homework 06: Analysis, sorting, and graphs

Due date: March 8, 11:59pm

Instructions:

Submit a typed or neatly handwritten scan of your responses on Gradescope here: <https://www.gradescope.com/courses/41121> (the assignment is “HW 6”). You will log in with your UW NetID.

For more details on how to submit, see https://courses.cs.washington.edu/courses/cse401/18au/hw/submitting_hw_guide.pdf.

Part 1. (10 points) Sorting algorithms

Answer each of the parts in English (not code or pseudocode). **Each subpart requires at most 2-4 sentences to answer.** Answers significantly longer than required will not receive full credit.

- (a) (3 points) Recall that merge sort works by taking an array, splitting it into two pieces, recursively sorting both pieces, then combining both pieces in $\mathcal{O}(n)$ time. Suppose we split the array into *three* equally sized pieces instead of two. And after we finish recursively sorting each of the three pieces, we first merge two of the three pieces together, then we merge that with the third piece to get the final sorted result.
- Write the recurrence for this variation of merge sort
 - Use the master theorem to give an asymptotic bound for this variation of merge sort.
 - Compare the recurrence and bound of this variation with that of the standard merge sort. Is this variation better or worse than the standard merge sort? Justify your answer.
- (b) (4 points) Suppose you are designing a search aggregator, which for a given query fetches search results from two different search engines and presents an intersection of the two search results. Here is a simplified version of this problem: Given two sorted integer arrays of lengths m and n , return a new array with elements that are present in both input arrays. The input array may contain duplicates, but there should be no duplicates in the output array. For example, if the input arrays are [17, 23, 23, 35, 43, 47, 69, 78, 80, 84, 84, 86] and [23, 35, 50], the output array should be [23, 35].
- Describe a brute force solution. What is the big- \mathcal{O} time complexity of this brute solution?
 - Describe a solution that has $\mathcal{O}(m \log n)$ time complexity.
 - Describe a solution that has $\mathcal{O}(m + n)$ time complexity.
 - Between subparts ii and iii above, is one solution always better than the other (in terms of runtime)? If yes, which one and why? If no, describe the situation (in terms of m and n) when one is better than the other.
- (c) (3 points) You're given an array where each element is an (age, name) pair representing guests at a DC-Marvel Universe party in Shire. You have been asked to print each guest at the party in ascending order of their ages, but if more than one guests have the same age, only the one who appears first in the array should be printed. For example, if the input array is
- [(23, Noah), (2000, Gandalf), (50, Frodo), (47, Emma), (23, Sophia), (83200, Superman), (23, Alice), (47, Madison), (47, Mike), (150, Dumbledore)]
- the output should be
- (23, Noah) (47, Emma) (50, Frodo) (150, Dumbledore) (2000, Gandalf) (83200, Superman)
- Describe a solution that takes $\mathcal{O}(1)$ space in addition to the provided input array. Your algorithm may modify the input array. This party has some very old guests and your solution should still work correctly for parties that have even older guests, so your algorithm can't assume the maximum age of the partygoers. Give the time complexity of your solution.

Part 2. (12 points) Running Dijkstra's algorithm

In this question, we will think about how to answer shortest path problems where we have more than just a single source and destination. Answer each of the parts in English (not code or pseudocode). **Each bullet point requires at most a few sentences to answer.** Answers significantly longer than required will not receive full credit.

- (a) (3 points) You and $k - 1$ of your friends have all decided to meet at Odegaard for a study session. Unfortunately you're scattered all over campus, and you'd like to start studying as soon as possible. You have a map of campus represented as the adjacency list for a weighted, directed graph (the graph has only positive edge weights), where the edge weights represent how much time it will take you to get from one way to another. In this version of UW, there are some one-way streets, which are represented in the graph as well. You have a list of the vertices representing the current locations of you and your friends. You also know which vertex corresponds to Odegaard. Figure 1 shows an example graph.

First, let's assume that you cannot alter the graph you have been given, and see what we can do.

- i. Describe how to find the quickest way for each of you and your friends to arrive at Odegaard. You should say which algorithm(s) you run and any inputs you would pass into the functions described in class.
Your answer should handle finding the path for each person, not just how long it will take them to get there.
 - ii. What is the big- \mathcal{O} running time for this process in terms of k , $|V|$ and $|E|$?
- (b) (3 points) Now, By modifying the graph representation, we're going to discover a more efficient algorithm for this problem. Suppose we wanted to reverse every edge in the graph (i.e. an edge pointing from u to v should end up pointing from v to u).
- i. How long would it take to make a reversed copy of the graph? Assume that the adjacency list uses hash tables, and all inserts will take constant time.
 - ii. What does a path from u to v in this new graph correspond to?
 - iii. Describe an algorithm to find the quickest way for all of your friends to get to Odegaard (again you should be able to return the full route, not just how long it takes).
 - iv. What is the big- \mathcal{O} running time to find all of the routes? Your answer in this part should be more efficient than the previous algorithm.




Figure 1: An example graph where $k = 3$ (the highlighted vertices). Odegaard library is L, and Q, F, R are other locations on campus. The shortest path from 1, 2, and 3 to L are 1-L, 2-Q-L, and 3-2-Q-L respectively.

- (c) (1 point) Give another real-world path finding example that you could solve using the strategy of reversing edges.

- (d) (4 points) You are in charge of routing ambulances to emergency calls. You have three ambulances in your fleet that are parked at different locations. There is an emergency and you have to dispatch the ambulance closest to the emergency to help as soon as possible.

You have a map of Seattle represented as the adjacency list for a weighted, directed graph (the graph has only positive edge weights). You also have a list of vertices representing the locations of each of your ambulances and the vertex representing where the new emergency is located.

You could run Dijkstra's from each of the three ambulance vertices, but since every second counts, you decide you'd like an algorithm that has better constant factors (even if you can't improve the big- \mathcal{O}). Your goal is to modify the graph so that you only need to run Dijkstra's once.

- i. Add a "dummy node" (i.e. a new vertex that doesn't represent any real location) to the graph. We want to run Dijkstra's with our dummy node as the source. How should we connect it to the existing graph so Dijkstra's will still return paths for the ambulances to follow in real life, and so we only have to run Dijkstra's once?

Be sure to describe both the direction and weight of any edges you add.

- ii. In your modified graph, once you have run Dijkstra's, how do you tell which ambulance to send?
 - iii. How long will it take to add the dummy node and any edges you added? Assume the adjacency list is implemented as a hash table of hash tables, and you won't get any collisions when you insert.
- (e) (1 point) Give another real-world path finding example that you could solve using the strategy of inserting a dummy vertex.

Part 3. (10 points) Modifying Dijkstra's algorithm

- (a) (4 points) Suppose you are given a graph that has negative-cost edges, but no negative-cost cycles. You want to find the shortest path between two points and decide to do so using the following algorithm:
- i. Add every edge by some constant k such that there are no more negative-cost edges.
 - ii. Run Dijkstra's algorithm to find the shortest path.

Unfortunately, this algorithm does not work. Give an example of a graph where this algorithm fails to return the correct answer. Be sure to explain why the algorithm fails on this graph.

Note: your example does not need to be complicated – you should need to use at most three or four nodes.

- (b) (3 points) Suppose we have a version of Dijkstra's algorithm implemented using binary heaps with a $\mathcal{O}(\log(n))$ decreaseKey method. As we discussed in lecture, this version of Dijkstra's algorithm will have a worst-case runtime of $\mathcal{O}(|V| \log(|V|) + |E| \log(|V|))$.

Now, suppose we know we will always run this algorithm on a *simple* graph. Given this information, rewrite the worst-case runtime of Dijkstra's algorithm so it only uses the variable $|V|$. Briefly justify your answer.

- (c) (3 points) Why does your answer to part (b) apply only to simple graphs? Why is it an inaccurate worst-case bound for non-simple graphs?

Part 4. (10 points) Modeling with graphs

Suppose we are trying to implement a program that finds a shortest path between any two locations within Seattle.

In case of a tie, the program should find us the route with the *fewest vertices on the path*. For example, suppose we are considering two different routes from UW to the Space Needle. Both routes are 3 miles long, but route one has 3 vertices and route two has five vertices. In this case, our algorithm should pick route one.

- (a) (7 points) Explain how you would model this scenario as a graph. Answer the following questions in bullet points with short sentences, then give an overall description on anything else that is relevant:
- i. What are your vertices and edges?
 - ii. What information do you store for each vertex and edge?
 - iii. Is this a weighted graph or an unweighted one?
 - iv. Is this a directed or undirected graph?
 - v. Do you permit self-loops? Parallel edges?
- (b) (3 points) At a high-level, how you would modify Dijkstra's algorithm to find us the best route according to the specifications listed above. In particular, be sure to explain:
- i. How you combine or use different factors like road length and number of vertices in the path while finding the best route.
 - ii. How you would modify Dijkstra's algorithm to break ties in the manner described above. (State which lines you would modify, if any, and/or lines you would add to the following pseudocode; state how and when you would check and resolve ties, e.g., during the BFS search or after the BFS search)

Use the following pseudocode for Dijkstra's algorithm as a base for your response to this question; it will be helpful to list specific lines that you will be modifying to fit this specific problem. Be sure your algorithm makes sense when combined with the graph representation you chose in part a.

Answer in English (not code or pseudocode) and in at most a few sentences.

Answers significantly longer than required will not receive full credit.

```

1: function Dijkstra(Graph G, Vertex source)
2:   initialize distances to  $\infty$ 
3:   mark source as distance 0
4:   mark all vertices unprocessed
5:   while there are unprocessed vertices do
6:     let u be the closest unprocessed vertex
7:     for each edge (u, v) leaving u do
8:       if u.dist + w(u, v) < v.dist then
9:         v.dist = u.dist + w(u, v)
10:        v.predecessor = u
11:    mark u as processed

```
