

CIS 121—Data Structures and Algorithms with Java—Spring 2018

Due: Monday, March 26th at 11:59 PM on Gradescope

Important: There are no late days on this homework

6 Required Problems (75 points), Written Portion (10 pts), Style & Tests (15pts)

DO NOT modify methods or method headers that were already provided to you.

DO NOT add public or protected methods.

DO introduce helper methods that are **package-private** (i.e., do not have a privacy modifier) as you deem necessary.

Submission (to Gradescope)

12 Files to Submit: Graph.java, GraphTest.java, BFS.java, BFSTest.java, UnionFind.java, UnionFindTest.java, Kruskal.java, KruskalTest.java, WidestPath.java, WidestPathTest.java, Client.java, ClientTest.java *(please do not submit any extra files)*

For all of your source files **DO NOT** place any in a package, and ensure that no file has a package declaration e.g. `package src;`. When you submit, be sure that all files are located at the root level i.e. do not submit a folder that contains the source files. The easiest way to do this is to drag the files individually onto Gradescope.

Important: Be sure that your own test cases pass and do not stall. Otherwise, your submission will not be processed correctly.

Important: Do not wait until the last minute to try submitting. You are responsible for ensuring your submission goes through by the deadline.

You have unlimited submissions for this homework. Upon submission, our infrastructure will compile your code and run some basic test cases against the interface. These are here to test that your code compiles (i.e. you didn't forget to add a class etc.) when we run real tests. They should not be used to test your own codes correctness since the functionality they test is very very basic.

Motivation: Network Connectivity

In Homework 6, you considered lossless compression as a way to transmit more data over a fixed Internet link. Now, we'll ask questions about the links themselves. The Internet is inherently a graph where machines and routers are vertices, and the wires between them are edges. Various questions concern the implementation and structure of the Internet, such as “How many hops must a message take to reach a destination?” and “What is the cheapest way to connect and route all computers to each other?”

In this assignment, you will implement several data structures and algorithms to answer these questions. Note: using items from Java collections is okay (examples: `Collections.sort`, `Queue` interface)!

Part One: Graph Representation (10 points)

Files to submit: Graph.java, GraphTest.java

Recall the [graph lab](#), where we discussed graph representations. Here, we'll implement an undirected, optionally-weighted graph. The graph should be immutable with respect to its vertices. Note that for an undirected graph, the edge from a vertex u to v is the same as the edge from vertex v to u .

We call the graph “optionally weighted” because it can be used by algorithms that use weights (like Kruskal's) and by algorithms that do not (like BFS). An algorithm like BFS would simply ignore any weights present.

The implementation details are largely up to you, but be sure to consider the runtimes of each operation. In the `Graph.java` stub, we've provided target runtimes for each method you need to implement; be sure to keep these in mind! (Particularly, for anything that says expected or amortized, it is fair game to use a `HashSet`, `HashMap` or `ArrayList` on if you deem it to be necessary. These data structures offer expected-case $O(1)$ `find/insert/delete`). Also, make sure that your solution only uses $O(m+n)$ space - please don't use adjacency matrices!

Your solution does need to handle self-loop, but not parallel edges (although you are welcome to); however, it must support negative edge weights. Overall, the graph implementation is largely up to you for design as long as it meets the space and time requirements. This could include implementing an edge wrapper class if you chose to. The graph implementation also assumes that the vertices will be represented by integers that begin at 0.

Note that we've also provided a basic weighted edge class in `Edge.java` with `compareTo`, `equals`, and `hashCode` implementations to make it easy to use for hashing and sorting.

This `Edge` class is tangential to the graph representation—it is just a wrapper for passing around edges, just as `Graph.getEdges` does (side note: the `getEdges` method should run in $O(m)$ time). It will be useful in some later problems. It would probably be helpful for you to read through the `Edge` class to see how it works. *DO NOT MODIFY Edge.java!*

Something else to keep in mind is what data your methods expose. Remember, you should practice *defensive coding*. If you return the data structures containing your graph to a client, then they can modify them as they wish (which is very bad!). Thus, you should use `Collections.unmodifiableXXX` (where the XXX is `Map`, `Set`, etc.) to ensure that the user cannot modify the data structures you are passing them.

Part Two: Breadth-First Search (15 points)

Files to submit: `BFS.java`, `BFSTest.java`

Recall the discussion of breadth-first search in lecture and in recitation. We want to implement BFS as a single-source shortest path algorithm, which we will use as a tool later in this homework assignment. Your BFS solution must include the source and the target in the resulting path (if it does not reach the target return an empty path). Remember that even if the graph is weighted, the BFS algorithm does not consider the value of the edge weights. The `BFS.java` comments contain necessary clarifying information, and you already know the algorithm, so give it your best shot!

Important: Do not modify the private constructor.

Part Three: Union-Find (15 points)

Files to submit: `UnionFind.java`, `UnionFindTest.java`

Union-find is a data structure for tracking a set of elements partitioned into disjoint subsets. We can `union` these subsets together and then `find` which set an element belongs to. In particular, we can determine when two elements belong to the same set, which ends up being quite useful for Kruskal's algorithm.

We want to start with the pointer-based implementation, in which each element keeps track of a "parent" pointer. At the beginning, everyone is their own parent, and two elements are in the same set if they share an ancestor. This allows for fast `union` operations, only involving changing a parent pointer. Think carefully, though, about what parent pointer that should be.

As with everything in CIS 121, there's a tradeoff at play here. Fast `union` means slow `find` if we're not careful, since `find` requires us to follow the parent pointers up to a root node. If our tree gets very tall, this could take a while. To alleviate this, we'll implement the two optimizations discussed in lecture to help make `find` as fast as possible (and it turns out that that's pretty fast!).

We recommend implementing this without any optimizations at first. Donald Knuth, responsible for popularizing big-Oh notation and inventing `TeX`, once famously said "Premature optimization is the root of all evil." There are many ways to interpret this piece of wisdom, but here we take premature to mean "before it even works at all." It's also much easier to optimize working code when you have test cases to make sure you didn't break anything.

The two optimizations are described in the following sections. You must implement **both** optimizations in the same method, or your runtime will be sub-optimal.

Weighting

When we union two sets, one set essentially absorbs the other; the absorbing set continues to exist, but it's the end of the road for the absorbed set. How, then, should we decide which set does the absorbing, and which set gets absorbed?

You'll need to start keeping track of some property of your sets, and use that information to make a decision here. Remember that we're trying to make find as fast as possible, so we want to keep our tree heights low; take this into consideration when deciding what property to track.

If you implement this optimization correctly, `find()` and `union()` will run in $O(\log n)$ time.

Path Compression

Another technique we can use to keep our traversals back to the root node from getting too long is by making them take just one hop. We call this path compression; whenever we traverse the tree to perform a `find` operation, we can just point the nodes along the traversal to the root. This should only take one or two lines of code to implement!

The amortized running time of `find()` and `union()` is very hard to reason about (take CIS 502 if you would like to grapple with that proof, or look [here](#)). Unbelievably, the path compression optimization makes the running time of those operations “effectively constant.” In reality, they are $O(\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann function. $\alpha(n)$ grows so slowly that for almost any practical value of n , $\alpha(n) < 5$. Hence, “effectively constant”! That's pretty cool.

Remember to implement both optimizations in the same method. Also, when implementing path compression keep in mind that rank is not necessarily equal to depth. Finally for consistency, during the union step for u and v , if i is the representative of u and j is the representative for v make i the parent of j if both i and j have the same rank.

Part Four: Kruskal's Algorithm (10 points)

Files to submit: `Kruskal.java`, `KruskalTest.java`

Kruskal's algorithm computes the minimum spanning tree of a graph by using the union-find data structure. You are to implement this algorithm, which is further described by the comments in `Kruskal.java`. For this algorithm you are allowed to use `Collections.sort` to sort the edge weights. Furthermore, if the edge weights are not distinct you can return any valid MST.

Important: Do not modify the private constructor.

Part Five: Widest Path Algorithm (15 points)

Files to submit: `WidestPath.java`, `WidestPathTest.java`

Consider a graph and a path P . The path P is a *widest path* when the weight of smallest edge in P is maximal. This maximal minimum-weight edge is known as a *bottleneck edge*. In other words, if you wanted to transfer a large amount of something (internet packets, car traffic, etc.) along a graph, and you considered edge weights to be capacities, you would want to find a path with the largest overall capacity. If we think about Internet packets, then the overall capacity of a route would be limited by the lowest-bandwidth connection in that route, as we would not be able to send more information along that path. Thus, the widest path is the path along which we can send the most stuff. The constraining edge on that path, or in this example, the lowest-bandwidth connection, would be the bottleneck edge of the path.

How can we algorithmically find a graph's widest path? That's up to you to devise and implement. There are multiple ways to solve this problem, so consider how you can do it with what you've already built (namely, BFS and Kruskal's). Please do this without modifying the behavior of the BFS and Kruskal's implementations you have already built. Note that within a widest path, edge weights are not necessarily

distinct, and those weights take precedence over the length of the path. To clarify, if there are multiple widest paths with the same bottleneck but different lengths then any of those paths are a valid solution. Runtime for this algorithm should be $O(m \lg n)$.

Important: Do not modify the private constructor.

Part Six: Client Application (10 points)

Files to submit: `Client.java`, `ClientTest.java`.

Let us take a step back. We have spent a lot of time in this class implementing APIs that do cool things. In this homework, we will briefly use an API we have developed as a code client.

A brief aside: **API** stands for **Application Program Interface**. You've probably heard the initialism used before. In the context of the web, it often refers to a set of URLs and requests that one can use to interact with a third-party service; for example, Candy Crush integrates with the Facebook API. More generally, it refers to the points of contact between any two independent pieces of a software system. In our context, this means the public methods in a Java class or interface that you are allowed to call. For example, the `WidestPath.java` API is simply `WidestPath.getWidestPath()`.

Now, imagine that you work for an Internet service provider (ISP). Imagine also that you have a list of all direct links between routers on your company's network, and you know the bandwidth capacity of each link. As an ISP, you're concerned about effectively routing Internet traffic. Note: IP addresses are in the standard format of IP addresses (i.e. `XXX.XXX.XXX.XXX`, where the number of digits in each section is from 1 to 3) which can be stored as strings.

Next, consider strategies one might use in choosing a path between two routers. Some common metrics include geographic distance (weighted shortest path) and number of hops (unweighted shortest path). Your company, however, wants to consider widest paths between routers in order to maximize available bandwidth. You have the `WidestPath.java` API available, so you do not need to implement a solution; instead, you need to marshal your own data into the format of the API and convert the result of the API call into a usable answer. Figuring out how to do this is the challenge of part six.

Style & Tests (15 pts)

The above parts together are worth a total of 75 points. 10 pts are awarded for the attached written homework, and the rest of the 15 points are awarded for style & tests.

Style is worth 5 of the 15 points and you will be graded according to the [121 style guide](#).

You will need to write comprehensive test cases for each of the methods you implement. This assignment in particular is highly compositional, with later parts building heavily on earlier parts, so we'd like to stress the importance of doing *bottom-up testing for each component before using it elsewhere*. We also encourage you to work out small examples of BFS, Kruskal's, and Widest Path by hand in order to write unit tests.

Make sure you consider edge cases, i.e., the more "interesting" inputs and situations. Use multiple methods instead of cramming a bunch of asserts into a single test method. Be sure to demonstrate that you considered "bad" inputs such as null inputs. Be sure to also demonstrate that you've tested for inputs to methods that should throw exceptions! Your test cases will be auto-graded for code coverage, and are worth 10 of the 15 points. You will have to thoroughly test your code to get full points! This includes testing any helper methods you have written. **Note:** you will not be able to write JUnit test cases for any private methods. Instead, make them package-private by leaving off any privacy modifier (i.e., do not write `public` or `private`).