
Problem Set 5

This problem set is divided into two parts: Part A problems are theory questions, and Part B problems are programming tasks.

Both Part A and Part B questions are due Monday, April 11th at 11:59PM.

Solutions should be turned in through the course website. Your solution to Part A should be in PDF format using \LaTeX . Your solution to Part B should be a valid Python file which runs from the command line. A template for writing up solutions in \LaTeX is available on the course website. Remember, your goal is to communicate. Full credit will be given only to a correct solution which is described clearly. Convolved and obtuse descriptions might receive low marks, even when they are correct. Also, aim for concise solutions, as it will save you time spent on write-ups, and also help you conceptualize the key idea of the problem.

Part A

Problem 5-1. [17 points] Arbitrage

Arbitrage is the use of discrepancies in currency exchange rates to transform one unit of a currency into more than one unit of the same currency. For example, suppose that 1 U.S. dollar buys 0.82 Euro, 1 Euro buys 129.7 Japanese Yen, 1 Japanese Yen buys 12 Turkish Lira, and one Turkish Lira buys 0.0008 U.S. Dollars. Then, by converting currencies, a trader can start with 1 U.S. dollar and buy $0.82 \times 129.7 \times 12 \times 0.0008 \approx 1.02$ U.S. dollars, thus turning a profit of 2 percent.

Suppose that we are given n currencies c_1, c_2, \dots, c_n and an $n \times n$ table of R exchange rates, such that one unit of currency c_i buys $R[i, j]$ units of currency c_j . Give an efficient algorithm to determine whether there exists a sequence of currencies $\langle c_{i_1}, c_{i_2}, \dots, c_{i_n} \rangle$ such that

$$R[i_1, i_2] \cdot R[i_2, i_3] \cdots R[i_{k-1}, i_k] \cdot R[i_k, i_1] > 1$$

Analyze the running time of your algorithm.

Problem 5-2. [17 points] Negative Weight Edges

Let $G = (V, E, w)$ be a weighted directed graph with exactly two negative-weight edges and no negative-weight cycles. Give an algorithm to find the shortest path weights $\delta(s, v)$ from a given $s \in V$ to all other vertices $v \in V$ that has the same running time as Dijkstra.

Problem 5-3. [16 points] Integer Weights

Let $G = (V, E)$ be a weighted, directed graph with weight function $w : E \rightarrow \{0, 1, \dots, W\}$ for some nonnegative integer W . Modify Dijkstra's algorithm to compute the shortest path weights from a given source vertex $s \in V$ in $O(WV + E)$ time.

Part B

Problem 5-4. [50 points] Speeding Up Dijkstra

The Howe & Ser Moving Company is transporting the Caltech Cannon from Caltech's campus to MIT's and wants to do so most efficiently. Fortunately, you have at your disposal the National Highway Planning Network (NHPN), packaged for you along with the code. You can learn more about the NHPN at <http://www.fhwa.dot.gov/planning/nhpn/>. This data includes node and link text files from the NHPN. Open `nhpn.nod` and `nhpn.lnk` in a text editor to get a sense of how the data is stored (`datadict.txt` has a more precise description of the data fields and their meanings). To save you the trouble of parsing these structures from a file, we have provided you with a Python module `nhpn.py` containing code to load the text files into Node and Link objects. Read `nhpn.py` to understand the format of the Node and Link objects you will be given.

Your goal in this problem is to implement and test several techniques for speeding up Dijkstra's algorithm in order to compute shortest paths between various pairs of locations.

An implementation of Dijkstra's algorithm is already provided. Function

```
dijkstra(nodes, edges, weight, source)
```

is given a graph with non-negative edges (represented as a list of Node objects and a list of undirected Edge objects), a function `weight(node1, node2)` that returns the weight of any edge between `node1` and `node2`, and a source node `source`. This function updates the `node.visited` field for all nodes, which indicates whether a shortest path to `node` has been found, as well as `node.distance` and `node.parent` for visited nodes, which are the length of the shortest path from the source node to `node` and the previous node on that path respectively. The function returns the number of nodes visited during the execution of the algorithm.

The links you are given do not include weights, so instead we use the geographical distance between their endpoints. Function `distance(node1, node2)` returns the distance between two NHPN nodes. Nodes come with latitude and longitude (in millionths of degrees). For simplicity, we treat these as (x, y) coordinates on a flat surface, where the distance between two points can be calculated using the Pythagorean Theorem.

Dijkstra's algorithm uses a priority queue, but this priority queue has one subtle requirement. Dijkstra's algorithm calls `decrease_key`, but `decrease_key` requires the index of an item in the heap, and Dijkstra's algorithm would have no way of knowing the current index corresponding to a particular Node. To solve this problem, the course staff have written an augmented heap object, `heap_id`, with the following extra features:

1. `insert(key)` returns a unique ID.
2. `decrease_key_using_id(ID, key)` takes an ID instead of an index.
3. `extract_min_with_id()` extracts the minimum element and returns a pair `key, ID`.

Additionally, we have provided some tools to help you visualize the output from your algorithms. You can use the `Visualizer` class to produce a KML (Google Earth) file. To view such a file on Google Maps, place it in a web-accessible location, such as your `Athena Public` directory, and then search for its URL on Google Maps.

For this problem, you will modify the file `dijkstra.py`. As you solve each part of the problem, check your work by running the appropriate test functions. We have provided several test functions that test each part separately or perform comparison tests of several methods. You should follow the instructions for each part of the problem, perform appropriate tests and draw conclusions. Please submit the modified `dijkstra.py` file with the code and `ps5-username.pdf` file with proofs and short answers. Keep them short.

- (a) [3 points] Examine the code provided in `nhpn.py`, `heap_id.py` and `dijkstra.py` to learn the structure of the `Node` and `Link` classes and the implementation of Dijkstra's algorithm. Run `test_a()`. Is there a significant difference in the execution time for different pairs of nodes? Explain your observation.
- (b) [7 points] One way to speed up Dijkstra's algorithm is to terminate the algorithm early once a shortest path to the destination has been found.

Implement function

```
dijkstra_early_stop(nodes, edges, weight, source, dest)
```

that performs this optimization. As with the function `dijkstra()`, this function should update the `node.visited`, `node.distance` and `node.parent` fields, and return the number of nodes visited during its execution. Run `test_b()`. What characterizes pairs of nodes for which there is a significant speed-up using this optimized version of Dijkstra's algorithm?

Hint: Reuse the implementation of Dijkstra's algorithm provided, making the required changes to allow for early termination.

- (c) [20 points] We will apply the potentials method with a landmark node to obtain a faster shortest-path algorithm. For a given landmark node l , we denote the potential of a node u with respect to a destination node t by $\lambda_t^l(u)$. The potential is defined as $\lambda_t^l(u) = \delta(u, l) - \delta(t, l)$ if there exists a path from u to t through l , and $\lambda_t^l(u) = C$ where C is some fixed constant if no such path exists. (Here, $\delta(u, v)$ denotes the length of a shortest path from u to v .)

1. Prove that this potential function is feasible, i.e. the modified weight of every edge is non-negative.
2. Implement function

```
compute_landmark_distances(nodes, edges, weight,
                           landmark)
```

that computes shortest paths from all nodes to the given landmark node `landmark`. For each node `node`, `node.land_distance` should be set to

the value of the shortest path distance from `node` to `landmark` or to some constant C if no such path exists (e.g., $C = 10^9$). Why is it more useful to precompute distances to the landmark node than precomputing the potentials themselves?

3. Implement function

```
dijkstra_with_potentials(nodes, edges, weight, source,
                        dest)
```

that performs Dijkstra's algorithm using edge weights modified according to the potentials method (i.e., $w'(u, v) = w(u, v) - \lambda_t^l(u) + \lambda_t^l(v)$ for a landmark l and destination t) and terminates as soon as a shortest path to the destination node `dest` has been found. This function assumes that `node.land_distance` is already set to the proper value (no need to call

`compute_landmark_distances()` from it). As before, this function should update `node.visited`, `node.distance` and `node.parent` fields, and return the number of nodes visited during its execution. Run `test_d()`. In which scenarios is the speed-up most significant (compare to both Dijkstra's algorithm and Dijkstra's algorithm with early termination)?

Hint: Reuse the implementation of Dijkstra's algorithm provided, making the necessary changes.

(d) [20 points] We will now describe a potentials method where multiple landmarks are used. For a given set of landmarks L , the potential of a node u with respect to a destination node t is $\lambda_t^L(u) = \max_{l \in L} \lambda_t^l(u)$, where $\lambda_t^l(u)$ is defined as before.

1. Prove that this potential function is also valid, i.e. the modified weight of every edge is non-negative.
2. How does the potential function $\lambda_t^L(u)$ compare to a potential function $\lambda_t^l(u)$ for any single landmark vertex $l \in L$ in terms of the number of visited nodes when used in Dijkstra's algorithm with early termination? Explain which is better and why.
3. Implement function

```
compute_multi_landmark_distances(nodes, edges, weight,
                                landmarks)
```

that computes shortest paths from all nodes to all given landmark nodes in `landmarks`. For each node `node`, `node.land_distances` should be set to a list of values, such that `node.land_distances[i]` is the shortest path length from `node` to `landmarks[i]` or some constant C if no such path exists (e.g., $C = 10^9$).

4. Implement function

```
dijkstra_with_max_potentials(nodes, edges, weight,
                             source, dest)
```

that performs Dijkstra's algorithm using edge weights modified according to the potentials method with multiple landmarks (i.e., $w'(u, v) = w(u, v) - \lambda_t^L(u) +$

$\lambda_t^L(v)$ for a set of landmarks L and destination t), and terminates as soon as a shortest path to the destination node `dest` is found. This function assumes that `node.land_distances` is already set to the proper list of values. As before, this function should update `node.visited`, `node.distance`, and `node.parent` fields, and return the number of nodes visited during its execution. Run `test_f()`. Does the performance of your algorithm match your assertions from part 2?

- (e) **(Optional)** Included in `nhpn.py` is a method to convert a list of nodes to a `.kml` file. `.kml` files can be viewed using Google Maps, by putting the file in a web-accessible location (such as your Athena Public directory), going to <http://maps.google.com> and putting the URL in the search box.

Run `visualize_path.py`. This will create two files, `path_flat.kml` and `path_curved.kml`. Both should be paths from Pasadena CA to Cambridge MA. `path_flat.kml` was created using the distance function you wrote in part (b), and `path_curved.kml` was created using a distance function that does not assume that the Earth is flat. Can you explain the differences? Also, try asking Google Maps for driving directions from Caltech to MIT to get a sense of how similar their answer is.