


- 1 Given a directed graph $G = (V, E)$ with non-negative edge lengths (denoted by $\ell(e)$, for an edge $e \in E$), and a node $s \in V$, describe an algorithm to find the length of a shortest cycle containing the node s .

Solution:

Let C be the shortest such cycle. Consider the “last” edge $v \rightarrow s$ in C . The idea is to compute the length of the path from s to all the vertices in G including v , and then add to it the cost of the last edge $v \rightarrow s$.



The algorithm. Run Dijkstra’s algorithm to find the shortest distances $d(s, v)$ from s to all the vertices $v \in V$. The length of the shortest cycle containing s is

$$\alpha = \min_{v \in V, v \rightarrow s \in E} (d(s, v) + \ell(v \rightarrow s)).$$

To this end, let v^* be the vertex realizing this minimum, and let π be the shortest path from s to v^* in G . The desired cycle is $\pi + (v^* \rightarrow s)$ [here $+$ denote concatenation].

Correctness. The algorithm outputs a cycle π that goes through s , and its length is α . A shortest cycle containing s must be composed of some path from s to a vertex x , followed by an edge $x \rightarrow s$. Namely, the length of the shortest desired cycle is at least α (since it is one of the solutions considered by the min above). As such, the computed cycle is a shortest cycle through s .

Running time analysis. Running Dijkstra takes $O(n \log n + m)$ time. Computing α , and thus v^* can be done in $O(n + m)$ time (this can be done in $O(n)$ time if the graph presentation lists all the incoming edges into a vertex). Extracting π from the shortest path tree can be done in $O(n)$ time. As such, the overall running time is $O(n \log n + m)$.

Comment on notations. One can use either (u, v) or $u \rightarrow v$, or even $(u \rightarrow v)$ to denote the directed edge going from u to v . It is fine to use either notation, but it is a good practice to use the same convention throughout. In the rest of this writeup, we use the notation (u, v) .

2 Suppose we have a collection of cities and different airlines offer flights between various pairs of cities. Some airlines only fly between some pairs of cities. Some pairs of cities are served by many airlines. Each airline charges perhaps different amounts for their one-way tickets.

2.A. Suppose you'd like to get from City A to City B at the least total cost. Describe an efficient solution. (Your solution may change planes to a different airline as needed.)

Solution:

This is a direct application of shortest path, where the graph $G = (V, E)$, where V is the set of airports, and

$$E = \{(u, v) \mid \text{some airline flies directly from city } u \text{ to city } v\}.$$


But what is the length $\ell(u, v)$ of edge (u, v) ? We could include multi edges in which for each airline flying directly from u to v we include an edge with length the cost of the airfare, or simpler, we use a single edge with length $\ell(u, v)$ defined to be the minimum fare offered by any airline that flies directly from u to v .

Finding a shortest path from city A to city B in G thus gives us the minimum airfare from A to B in the stated problem.


2.B. It turns out that airports charge usage taxes. Different airports may charge different amounts in tax. Your cost of traveling from A to B now includes all of the flight costs, plus all of the taxes of the airports that you stopover along the way from A to B . Model this as a graph problem and give an efficient solution to find the least cost way to get from A to B .

Solution:

This is a bit trickier because of the airport taxes. We employ a trick typically used when a graph has costs associated with vertices as well as costs on edges: we split each vertex into two pieces, and insert an edge between them with weight equal to the cost of the vertex. Each vertex will have an “in” version, and an “out” version, with all incoming edges directed to the “in” version, and all outgoing edges emanating from the “out” version. Refer to the picture for an example.



original graph
with vertex weights



new graph
with only edge weights

The solution is then to create the graph as in part (a), but then split nodes as described and insert edges with lengths corresponding to the airport taxes. The minimum cost path from A_{in} to B_{out} will

give the least cost of traveling from A to B , paying all airport taxes, including those at airports A and B . (What if taxes at airports A and B were not required to be paid?)

Formally, the graph we create is $G = (V, E)$, where V and E and the edge weights are defined as follows:

- $V = V_{in} \cup V_{out}$, where $V_{in} = \{v_{in} \mid v \text{ is a city}\}$, and $V_{out} = \{v_{out} \mid v \text{ is a city}\}$.
- $E = E_1 \cup E_2$, where $E_1 = \{(u_{out}, v_{in}) \mid \text{there is an airline that flies from city } u \text{ to city } v\}$, and $E_2 = \{(v_{in}, v_{out}) \mid v \text{ is a city}\}$.
- If $e = (u_{out}, v_{in}) \in E_1$, then $\ell(e) =$ the least cost airfare for flying directly from city u to city v .
- If $e = (v_{in}, v_{out}) \in E_2$, then $\ell(e) =$ the airport tax for the airport in city v .

The time to create the graph G is linear in the number of cities and number of flights specified, and the running time of the algorithm is thus dominated by running the shortest path algorithm which takes time $O(m + n \log n)$, where m is the number of flights, and n the number of cities.

3 There are n galaxies connected by m intergalactic teleport-ways. Each teleport-way joins two galaxies and can be traversed in both directions. However, the company that runs the teleport-ways has established an extremely lucrative cost structure: Anyone can teleport *further* from their home galaxy at no cost whatsoever, but teleporting *toward* their home galaxy is prohibitively expensive.

Judy has decided to take a sabbatical tour of the universe by visiting as many galaxies as possible, starting at her home galaxy. To save on travel expenses, she wants to teleport away from her home galaxy at every step, except for the very last teleport home.

Describe and analyze an algorithm to compute the maximum number of galaxies that Judy can visit. Your input consists of an undirected graph G with n vertices and m edges describing the teleport-way network, an integer $1 \leq s \leq n$ identifying Judy's home galaxy, and an array $D[1 \dots n]$ containing the distances of each galaxy from s .

Solution:

We reduce this problem to finding the length of the longest path in a DAG $G = (V, E)$ as follows:

- V is the set of n galaxies, plus an artificial target node t . Let s denote Judy's home galaxy.
- E contains a directed edge (u, v) if either of the following conditions is satisfied
 - There is a teleport-way between galaxy u and galaxy v , and v is farther from s than u .
 - $v = t$ and there is a teleport-way between galaxy u and s .
- We need to compute the length of the longest path in G from s to t .
- We can compute this length using dynamic programming as described in the lecture (and in the lecture notes).
- The algorithm runs in $O(|V| + |E|) = O(n + m)$ time.

To think about later: You probably heard of the phrase “six degrees of separation” and the “small world” phenomenon; see https://en.wikipedia.org/wiki/Six_degrees_of_separation. The idea is that in many interesting networks people or objects are within a small distance of each other. At the same time we believe in “locality” in that each person may only a small number of people compared to the total population. The next two problems explore the tradeoffs between diameter and degree in a graph to explore this in a more quantitative fashion.

- 4 Suppose G is a connected graph with maximum degree d . The diameter of the graph is $\max_{u,v} \text{dist}(u, v)$. Prove that the diameter of the graph is $\Omega(\log_d n)$ where n is the number of nodes (here, the edges are not weighted). It is easier to consider $d = 5$ or some other small constant for simplicity. *Hint:* Consider the BFS layers starting at any vertex v .

The point of the problem is to show that if all degrees are small then the diameter must grow with the number of nodes.

Solution:

Suppose the maximum degree is d . Let the diameter be $D = d(s, t)$ for some pair of nodes s and t . Then the number of nodes that can be in layer 1 of a BFS search starting at s (distance 1 away from s) is, by definition of degree, at most d . The number of nodes that can then be at layer 2 is at most d^2 , since each node at layer 1 is connected to at most d other nodes. In general, the number of nodes at layer i is at most d^i . Since the diameter is D , the total number of nodes summing across all of the D layers in a BFS starting from s is at most $\sum_{i=0}^D d^i \leq d^{D+1}$. But if there are n nodes, this means $n \leq d^{D+1}$, and taking logs base d , we have $\log_d n \leq D + 1$, so diameter $D = \Omega(\log_d n)$ as was to be shown.

- 5 Suppose the diameter of an undirected simple graph is d . Prove that there is a node with degree at most $3n/d$. *Hint:* Consider the BFS layers for the pair defining the diameter. It is easier to prove a bound such as $9n/d$.

This problem is to show you that if the diameter is small then there must be a large degree node.

Solution:

Note that the graph is connected if it has finite diameter. If the diameter is d then there is a pair of nodes s, t such that $d(s, t) = d$. Consider the BFS layers starting from s . Hence we have $L_0 = \{s\}$, L_1, L_2, \dots, L_d where $t \in L_d$. Consider any node u in a layer L_i where $1 \leq i \leq d-1$. From the properties of the BFS layers u can be connected only to nodes in L_{i-1}, L_i, L_{i+1} . Thus $\text{deg}(u) \leq |L_{i-1}| + |L_i| + |L_{i+1}|$. For $i = 0$ we can say that $\text{deg}(s) \leq |L_1|$. For $i = d$ we can say that $\text{deg}(t) \leq |L_{d-1}| + |L_d|$. Suppose the claim is false and hence $\text{deg}(v) > 3n/d$ for every $v \in V$. For $1 \leq i \leq d-1$ pick an arbitrary node u_i from L_i (each layer is non-empty so we can do this). We have $\text{deg}(u_i) \leq |L_{i-1}| + |L_i| + |L_{i+1}|$ and since we assumed that the claim is false, $3n/d < |L_{i-1}| + |L_i| + |L_{i+1}|$ for each $1 \leq i \leq d-1$. For $i = 0$ and $i = d$ we have $|L_1| > 3n/d$ and $|L_{d-1}| + |L_d| > 3n/d$. Summing these inequalities from $i = 0$ to d we obtain

$$|L_0| + 3(|L_1| + |L_2| + \dots + |L_{d-1}|) + 2|L_d| > (d+1) \cdot 3n/d > 3n.$$

The above implies that

$$|L_0| + |L_1| + \dots + |L_d| > n$$

which is a contradiction since the every node is in exactly one of the layers. Thus there must be a node with degree at most $3n/d$.

A more elegant solution.

Solution:

Let $P = v_0, v_1, \dots, v_d$ be the diameter-defining path of the graph. Here, unless $|i - j| = 1$, there cannot be an edge joining v_i and v_j , for otherwise the distance between v_0 and v_d would be smaller than d . In addition, any vertex u not in P can have at most three neighbors in P for the same reason.

Therefore, the sum of degrees of the vertices in P is

$$\sum_{i=0}^d \deg(v_i) \leq 2d + 3(n - d - 1) = 3n - d - 3$$

By the Pigeonhole Principle, there is a node in P with degree at most

$$\frac{1}{d+1} \sum_{i=0}^d \deg(v_i) \leq \frac{3n - d - 3}{d+1} \leq \frac{3n}{d}.$$