CS4234: Optimization Algorithms Lecture 4

The Traveling Salesman

Lecturer: Seth Gilbert *September 1, 2015*

Abstract

The goal of the Traveling Salesman Problem is to find a tour that connects all the nodes in a graph at a minimal cost. There are several variants, depending on whether repeated visits are allowed, and depending on whether the distances satisfy a metric. We discuss the relationship between these variants, and give a simple 2-approximation algorithm. We then develop a more involved 1.5-approximation algorithm that relates the travelling salesman problem to Eulerian tours.

1 The Travelling Salesman Problem

Today we consider he Traveling Salesman Problem, often abbreviated TSP. The "TSP" problem is perhaps one of the most famous (and most studied) problems in combinatorial optimization.

1.1 Problem Definition

Given a set of cities (i.e., points), the goal of the traveling salesman problem is to find a minimum cost circuit that visits all the points. More formally, the problem is stated as follows:

Definition 1 *Given a set V of n points and a distance function $d : V \times V \to \mathbb{R}$, find a cycle C of minimum cost that contains all the points in V. The cost of a cycle $C = (e_1, e_2, \ldots, e_m)$ is defined to be $\sum_{e \in C} d(e)$, and we assume that the distance function is non-negative (i.e., $d(x, y) \geq 0$).*

Notice that, unlike the Steiner tree problem, there are no Steiner nodes: you have to visit every city. The cities/points may be in some geometric space (e.g., the Euclidean plane), or they may not. If the points are in the Euclidean plane, then it is natural to define d to be the Euclidean distance. Otherwise, d can be arbitrary. See Figure 1 for an example of one instance of the traveling salesman problem.

1.2 Variants

As with the Steiner tree problem, there are several variants:

- *Metric vs. General:* In the *metric* version of the traveling salesman problem, the distance function d is a metric, i.e., it satisfies the triangle inequality. In the general version, d can assign any arbitrary weight to an edge.
- *Repeated visits vs. No-Repeats:* The goal of the TSP problem is to find a cycle that visits every node. Can the cycle contain repeated nodes (or does it have to be a simple cycle)? In the version with No-Repeats, the TSP cycle must visit each node *exactly* once. In the version with repeats, it is acceptable to visit each node more than once (if that results in a shorter route).

We will summarize these four variants as follows:

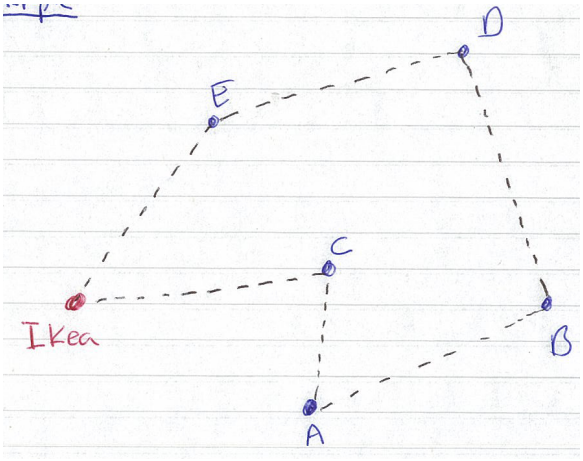

Figure 1: Example of the traveling salesman problem. Here, there are six locations. The problem is to find the shortest delivery circuit starting from Ikea and visiting each of the five houses, and then returning to Ikea.

	Repeats	**No-Repeats**
Metric	M-R TSP	M-NR TSP
General	G-R TSP	G-NR TSP

1.3 NP-Hard

All the variants of the traveling salesman problem are NP-hard. In fact, the problem is NP-hard even for planar graphs with maximum degree 3.

For the No-Repeats variants of the problem, this is easily seen by reduction from the Hamiltonian Cycle problem: a Hamiltonian cycle is a cycle that contains every node exactly once; the goal of the Hamiltonian Cycle problem is to determine whether a given graph has a Hamiltonian Cycle. Deciding whether a graph has a Hamiltonian Cycle is NP-hard (and this can be shown by reduction from 3SAT). Clearly, if we can solve TSP, then we can solve the Hamiltonian Cycle problem (by setting the distances properly). There are similar reductions for the other variants.

In fact, the general no-repeats version of TSP (G-NR TSP) is NP-hard even to approximate! If $P \neq NP$, there is no constant factor approximation algorithm. In fact, something even stronger is true. Let $r(n)$ be any polynomial time computable function. (For example, perhaps $r(n) = 2^{100}2^{n^2}$.) Let A be an $r(n)$-approximation algorithm for G-NR TSP. If A is a polynomial time algorithm, then $P = NP$. Clearly, this is a problem that is very hard to approximate.

Even so, we are going to see how to (easily!) approximate the other three variants. In almost all common real-world cases (e.g., where the distance function satisfies the triangle inequality, or where repeats are acceptable), there are good approximation algorithms.

2 Equivalence

The first thing we are going to see is that the three other variants are equivalent. That is, $M-R-TSP$, $M-NR-TSP$, and $G-R-TSP$ are all equivalent: if we have a c-approximation algorithm for any one of these variants, then we also can construct a c-approximation algorithm for the other two.

We first focus on the issue of repeats: as long as the distance function is a metric, it does not matter whether or not we allow repeats. In many ways, this should be unsurprising (given our discussion of Steiner Trees last week): if we have repeated nodes on the cycle, we can always "short-cut" past them producing a cycle with no repeats at the same cost.

Lemma 2 *There exists a c-approximation algorithm for $M-NR-TSP$ if and only if there exists a c-approximation algorithm for $M-R-TSP$.*

Proof The proof has three claims. First, we show that both the version with repeats and no-repeats have the same optimal cost. We then show how to transform the solution from the NR variant into the R variant (which is trivial) and from the R variant into the NR variant (which involves skipping repeats).

Claim 1. Let (X, d) be an input for Metric TSP. Let $OPT(R)$ be the minimum cost TSP-cycle overall (i.e., with repeats), and let $OPT(NR)$ be the minimum cost TSP-cycle with no repeats. Then $OPT(R) = OPT(NR)$.

To show this, first we observe that any cycle with no repetitions is also a legal TSP-cycle overall (i.e., for the version that allows repeats). And so it follows immediately that $OPT(R) \leq OPT(NR)$. (Remember, when we minimize over a larger space of possible solutions, we always get a solution that is at least as good as the best solution of the smaller subspace.)

Next, let C be a cycle with repeats, for example:

$$C = A, B, C, D, C, E$$

We can now create cycle C' by skipping the repeated nodes. For example:

$$C' = A, B, C, D, E$$

By the triangle inequality, we know that $d(C') \leq d(C)$. For example, in constructing C' from C, we replace (D, C, E) with (D, E); we know that $d(D, C) + d(C, E) \leq d(D, E)$, by the triangle inequality. We can repeat this inductively for every skipped node.

If we assume that cycle $C = OPT(R)$, then we conclude that $OPT(NR) \leq d(C') \leq d(C) = OPT(R)$ Putting together the two inequalities, we have shown that $OPT(NR) \leq OPT(R)$ and $OPT(R) \leq OPT(NR)$, which yields our desired conclusion that $OPT(R) = OPT(NR)$.

Claim 2. If A is a c-approximation algorithm for $M - NR - TSP$, then A is a c-approximation algorithm for $M - R - TSP$.

Assume algorithm A outputs cycle C. By assumption, C has no repeats, and $d(C) \leq c \cdot OPT(NR) = OPT(R)$. Thus we see that C is also a valid solution for $M - R - TSP$ and provides a c-approximation.

Claim 3. If A is a c-approximation algorithm for $M - R - TSP$, then we can construct an algorithm A' which is a c-approximation for $M - NR - TSP$.

Specifically, construct algorithm A' as follows: Run algorithm A to get cycle C; then construct C' from C by skipping any repeated node. There are two things we have to show. First, the cost of cycle C is a good approximation of the optimal for $M - NR - OPT$:

$$\begin{aligned} d(C) &\leq c \cdot OPT(R) \\ &\leq c \cdot OPT(NR) \end{aligned}$$

Second, the cycle we have constructed in A' has cost bounded by the cost of C:

$$d(C') \leq d(C)$$

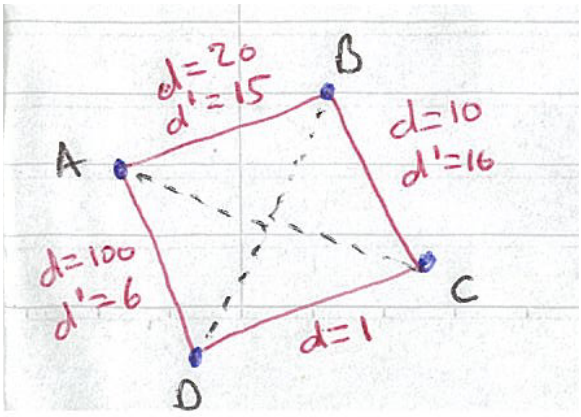

Figure 2: Example of constructing A', reducing $G - R - TSP$ to $M - R - TSP$.

This follows from the triangle inequality. Putting these two facts together, we see that the algorithm A' is a c-approximation algorithm for $M - NR - TSP$. □

We can show exactly the same type of equivalence for the general and metric versions, as long as we allow repeats. In this case, we cannot create "shortcuts" because the triangle inequality does not hold; however, we do not need to create shortcuts because repeats are allowed.

Lemma 3 *There exists a c-approximation algorithm for $M - R - TSP$ if and only if there exists a c-approximation algorithm for $G - R - TSP$.*

Proof In this case, the proof contains two claims: one transforming G into M (which is trivial), and the other M into G.

Claim 1. If A is a c-approximation algorithm for $G - R - TSP$, then A is a c-approximation algorithm for $M - R - TSP$.

This claim is immediately and obviously true. The two problems are identical, except for the restriction that the input to $M - R - TSP$ must be a metric. However, if a given input *is* a metric, then it clearly satisfies the requirements of $G - R - TSP$ and we can execute algorithm A on that instance. Moreover, the optimals solution will be the same for $G - R - TSP$ and $M - R - TSP$ (since they are optimizing over the exact same set of possible cycle).

Claim 2. If A is a c-approximation algorithm for $M - R - TSP$, then we can construct algorithm A' is a c-approximation algorithm for $G - R - TSP$.

Let (X, d_g) be the input to algorithm A', i.e., an input for $G - R - TSP$. We construct algorithm A' as follows.

First, we need to construct a distance metric d_m. For every pair of nodes (u, v), define $d_m(u, v)$ to be the shortest path from u to v according to the distances d_g. (Recall that d_g is not a metric, i.e., does not necessarily satisfy the triangle inequality.) We can find d_m by construct the complete graph on X, assign the weight of edge (u, v) to be $d_g(u, v)$, and executing an all-pairs-shortest-path algorithm. You will notice that d_m is a distance metric, i.e., it satisfies the triangle inequality. Recall that d_m is called the metric completion of the graph.

Second, execute algorithm A on (X, d_m), producing cycle C. We then need to construct a new cycle C' out of C. For each edge (u, v), we add the shortest path from u to v according to d_g to our output cycle C'. Notice this produces a cycle C' that may contain may repeated nodes.

We now argue that the result is a good approximation of the $OPT(X, d_g)$. First, observe that $d(C') = d(C)$. Each edge in C got replaced by a path in C' with the exact same cost, and hence we end up with the same cost cycle.

Second, we know by assumption that $d(C) \leq c \cdot OPT(X, d_m)$, by the assumption that A is a c-approximation algorithm.

Finally, we argue that $OPT(X, d_m) \leq OPT(X, d_g)$. In particular, if R is the optimal cycle for (X, d_g), then R is also a cycle in (X, d_m). Moreover, for every pair (u, v), we know that $d_m(u, v) \leq d_g(u, v)$ (because d_m is defined as the shortest path). Thus $d_m(R) \leq d_g(R)$. Since the optimal cycle with metric d_m has to be at least as good as $d_m(R)$, we conclude that $OPT(X, d_m) \leq d_m(R) \leq d_g(R) = OPT(X, d_g)$.

Putting the pieces together, we conclude that $d(C') = d(C) \leq c \cdot OPT(X, d_m) \leq c \cdot OPT(X, d_g)$, and hence algorithm A' is a c-approximation algorithm for $G - R - TSP$. □

3 2-Approximation Algorithm

We now present a 2-approximation algorithm for $G - R - TSP$. We already know that this yields a 2-approximation algorithm for the other two variants. Assume that the input is (X, d) where X is a set of points and d is a distance function (but not necessarily a metric). Consider the following algorithm:

1. Construct the complete graph $G = (X, E)$ with weights w where E contains every pair $(u, v) \in X \times X$ and $w(u, v) = d(u, v)$.
2. Let T be the minimum spanning tree of G.
3. Let C be the cycle constructed by performing a depth-first search of T.

To analyze this algorithm, as usual, we start with the optimal solution and work backwards. Let C^* be the optimal (minimum cost) TSP cycle for input (X, d) and let E^* be the edges in C^*. Notice that the graph $G^* = (X, E^*)$ is connected, because C^* is a cycle that includes every point. Let T^* be the minimum spanning tree of $G = (X, E^*)$. At this point, we know that:

$$d(T^*) \leq d(C^*) = OPT$$

Since the tree T is a minimum spanning tree of $G = (X, E)$ and $E^* \subseteq E$, we know that:

$$d(T) \leq d(T^*)$$

Finally, since C is construct by a depth-first-search traversal of T, we know that C includes each edge in T exactly twice, and so:

$$d(C) = 2d(T)$$

Putting these inequalities together, we conclude that:

$$d(C) = 2d(T) \leq 2d(T^*) \leq 2d(C^*) \leq 2OPT$$

That is, the cost of C is at most twice the cost of OPT, and hence the algorithm is a 2-approximation algorithm.

4 Eulerian Cycles

We now want to develop a better approximation algorithm, one that will have a better approximation ratio than 2. In order to do that, we are going to need a few additional tools. Let's begin with a famous problem, known as the *Bridges of Konigsberg*.

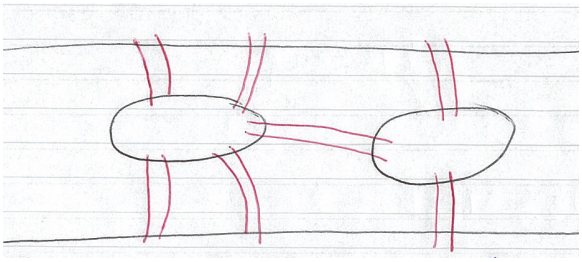

Figure 3: Illustration of the Bridges of Konigsberg.

In Konigsberg, there is a river with two islands. Historically, these two islands were connected by seven bridges. The bridges were beautiful and famous, and so the question that everyone wanted to answer was whether it was possible to cross each bridge exactly once—and end up back where you started. With a little bit of thought, it becomes obvious that you cannot. But how would you prove it? Euler solved this problem, and his solution is often seen as marking the beginning of graph theory as a mathematical field. Notably, he realized that the problem could be represented as a graph: Once the problem was represented abstractly as a graph, we can apply more powerful techniques to solve the

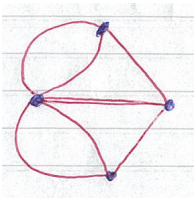

Figure 4: Illustration of the Bridges of Konigsberg as a graph.

problem.

For the purpose of this section, we will be talking about *multigraphs* rather than simple graphs. A *multigraph* is a graph $G = (V, E)$ where E is a multiset, rather than a set, of edges. This means that each edge can appear in the graph more than once. For example, a given edge (u, v) might appear in graph G four times.

Definition 4 *A **Eulerian Cycle** in a multigraph G is a cycle that crosses each edge exactly once.*

Notice that, unlike a TSP-cycle, an Eulerian Cycle focuses on edges, rather than nodes. The goal is to cross each edge exactly once—though that may mean visiting a given node many times. Clearly the problem of the Bridges of Konigsberg is simply asking whether there is an Eulerian Cycle in the associated graph.

The key claim is to characterize when a graph has an Eulerian Cycle. Amazingly, there is a very simple characterization:

Lemma 5 *A multigraph $G = (V, E)$ has a Eulerian Cycle if and only if:*

- *It is connected (except for nodes of degree 0).*

- *Every node has even degree.*

Proof Notice that one direction of this claim is easy. Assume graph G has an Eulerian Cycle. In that case, since the cycle cross every edge, clearly the cycle visits every node (except those nodes with degree zero) and hence graph is connected—except for nodes with degree zero. In addition, the cycle enters and exits each node the same number of times. For example, if the cycle enters a node v some k times, then it also exits node v the same k times. From this we conclude that every node has even degree. (In our example, node v has degree $2k$, which is even.)

The surprising result is the opposite direction. Assume graph G is connected (except for nodes of degree 0) and that every node has even degree. We need to construct a Eulerian Cycle. We will proceed by induction on the number of edges, proving a stronger invariant as we go: If every node in G has even degree, then each connected component with > 1 nodes has an Eulerian Cycle.

We will begin (in the base case) with an empty graph containing only the nodes, and add one edge at a time back to the graph. At every step, we will construct an Eulerian Cycle for each connected component.

Base case: In this case, $m = 0$, i.e., there are no edges, and hence trivially the claim holds.

Inductive step: Assume we have m edges. Let $V_1, V_2, \ldots, V_k$ be the connected components of G that contain > 1 nodes. We consider two cases:

Case 1: $k \geq 2$: There are at least two connected components. Each connected component has $< m$ edges. Hence, by induction, each component has an Eulerian Cycle.

Case 2: $k = 1$: V_1 is the only connected component with > 1 nodes. Recall that all the nodes in V_1 have even degree, and every node note in V_1 has degree zero.

First, we claim that there must be a cycle in V_1. Let $n_1 = |V_1|$. If there are no cycles in V_1, then V_1 is a tree and contains exactly $n_1 - 1$ edges. However, the degree of every node in V_1 is at least 2 (since the degree is even and > 0). Thus there are at least $2V_1$ endpoints of edges in V_1, and hence at least V_1 edges in V_1. (Notice in general that if $d(v)$ is the degree of node v, then a graph has exactly $\sum_{u \in V} d(u)/2$ edges.) This contradicts the claim that V_1 is a tree.

Let C be any cycle in V_1. Let G' be the the graph G where the edges in C are removed. Let $V'_1, V'_2, \ldots, V'_{k'}$ be the new connected components of G'. (By removing the cycle C, we may have split V_1 into several distinct connected components; or V_1 may remain connected.) Each of these components now has $< m$ edges. Hence, by induction, each of these components has an Eulerian Cycle. Let $C_1, C_2, \ldots, C_k$ be the Eulerian Cycles for connected components $V'_1, V'_2, \ldots, V'_{k'}$.

It remains to stitch these Eulerian Cycles back together. Notice that these connected components are all connected by the initial cycle C. That is, if we add back the cycle C, we will have a single connected component. Here's how we do that:

Let $C = (v_1, v_2, \ldots, v_\ell)$. Begin at node v_1. We know that node v_1 is part of one of the Eulerian Cycles V_j. Follow the Eulerian Cycle V_j until it returns to node v_1. Then proceed to node v_2. If v_2 is part of a new Eulerian Cycle (i.e., not V_j), then follow the new Eulerian Cycle until it returns to v_2. Otherwise move on to v_3. And so on. At each step, we move to the next node v_i in the cycle C. If that node v_i is part of a new Eulerian Cycle V_i that has not yet been traversed, then we traverse the cycle until it returns to v_i. Then we continue to v_{i+1}. This continues until we return to node v_1.

Notice that the new cycle we have constructed traverses every edge in the graph G, as it traverses all the edges in the connected components $V'_1, \ldots, V'_{k'}$ as well as all the edges in C. Hence we have constructed an Eulerian Cycle for G. □

The other interesting fact is that this Eulerian Cycle can be easily constructed in polynomial time, exactly by following the steps of the algorithm:

1. Find a cycle in G (e.g., via DFS). If there is none, then skip the next step.
2. Remove the cycle.
3. Find the connected components in the residual graph.
4. Recurse: find an Eulerian Cycle in each connected component.
5. If a cycle was found in Step 1, paste the cycles back together.

This can readily be implemented in $O(m^2)$ time. Can you optimize it further?

5 A Better TSP Approximation

We now consider the Metric Traveling Salesman Problem, with repeats. (As before, if we can approximate this variant, then we can also approximate the other two tractable variants.) Assume (X, d) is the input to our problem.

5.1 Basic Idea

Imagine that we built a multigraph $G = (X, E)$ that contained an Eulerian Cycle. Then that cycle would give us a feasible solution to the TSP problem. Conversely, if had a solution to the TSP problem, we could build a multigraph consisting of only the edges in the TSP-cycle, which would result in a graph with a Eulerian Cycle. Hence solving TSP is equivalent to finding a minimum cost set of edges E such that $G = (X, E)$ has an Eulerian Cycle.

For example, consider the following proposed algorithm:

1. Find the MST T of the complete graph on X with weights defined by d.
2. Add every edge in T *twice* to the set E. This ensures that each node in X has even degree.
3. Since every node in multigraph $G = (V, E)$ has even degree degree, we can find a Eulerian Cycle C for G.
4. Return C.

This is a correct algorithm for solving the TSP problem. Unfortunately, it will not yield any improvement over the earlier 2-approximation algorithm because we are still adding each edge in the MST twice, leading to a cost that is at most twice optimal.

Notice that it is quite inefficient to add each edge twice! Our only goal is to ensure that each node has even degree—if a node already has even degree, why are we doubling all of its outgoing edges, and thus increasing our cost? Consider the example in Figure 5. Thus the question we must answer is: how do we add the minimal number of edges to a multigraph to ensure that each node has even degree? Thus, in summary, the resulting algorithm should look like:

1. Find the MST T of the complete graph on X with weights defined by d.
2. Add every edge in T to the set E.
3. Add the minimum cost set of edges E' to E such that in $E \cup E'$, every node has even degree.
4. Since every node in multigraph $G = (V, E \cup E')$ has even degree degree, we can find a Eulerian Cycle C for G.

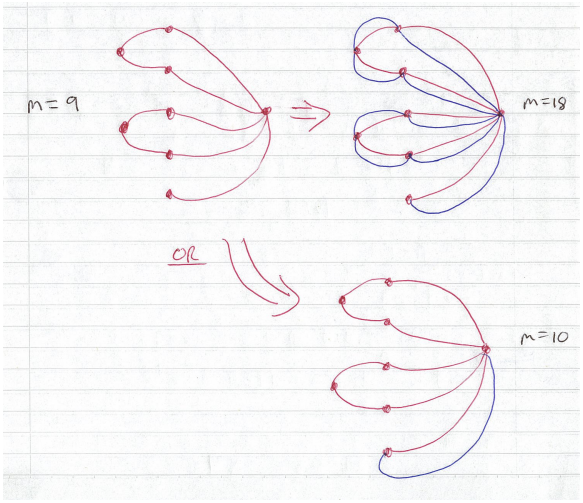

Figure 5: Example of finding a TSP by creating nodes with even degree. Initially, we have a graph with nine edges. If we double every edges, we have a graph with 18 edges. However, if we want to ensure that eachnode has even degree, it suffices to add just one edge to the graph, yielding a graph with only 10 edges.

5. Return C.

We already know that $cost(T) \leq cost(OPT)$, since the optimal TSP-cycle is also a spanning tree (once you have removed a single edge). The key is to prove that $cost(E') \leq OPT/2$.

5.2 Matchings

The last tool we need is a perfect matching:

Definition 6 *We say that (V, M) is a **matching** if no two edges in M share an endpoint, i.e., all nodes in V have degree ≤ 1. We say that (V, M) is a perfect matching if every node in V has exactly degree 1, i.e., every node is matched.*

Notice that a perfect matching only exists if $|V|$ is even: if there are an odd number of nodes, it is clearly impossible to match them all.

If each edge has a weight, then we may want to find the minimum cost perfect matching. One of the most beautiful results in combinatorial optimization yields exactly that:

Theorem 7 *There is a polynomial time algorithm for finding the minimum cost perfect matching.*

This finds the optimal solution, and it does it efficiently! Unfortunately, this algorithm is beyond the scope of today's class. (See Edmond's algorithm. The solution involves his famous *Blossom Algorithm* for finding matchings, along with the primal-dual method and linear programming.) For now, we are just going to assume that we can find a minimum cost perfect matching.

5.3 Evens and Odds

We want to use the idea of perfect matchings to pair up the nodes in the multigraph that have odd degree. If we can assign each of them a partner, then they will have even degree. But there is one potential problem: what if there are an

odd number of nodes that need a partern? Luckily, that is impossible due to a nice counting argument:

Lemma 8 *In every graph $G = (V, E)$, there are an even number of nodes with odd degree.*

Proof Let O be the set of nodes with odd degree, and $V - O$ the set of nodes with even degree. Let $deg(u)$ be the degree of u.

Recall that we can count the number of edges in the graph by summing up the total degree and diving by 2, i.e.:

$$\sum_{u \in V} deg(u) = 2|E|$$

That is, the sum of the degrees of all the nodes is even.

Now, consider just the nodes in $V - O$: since each node has even degree, the sum of their degrees must also be even:

$$\sum_{u \in V-O} deg(u) = 2k$$

Summing everything, we conclude that:

$$\begin{aligned} \sum_{u \in O} deg(u) + \sum_{u \in V-O} deg(u) &= \sum_{u \in V} deg(u) \\ \sum_{u \in O} deg(u) + 2k &= 2|E| \end{aligned}$$

Since the right-hand-side of the expression is even, the left-hand side of the expression must also be even. That is:

$$\sum_{u \in O} deg(u) = 2\ell$$

But each node in O has odd degree. Hence the only way that the sum of their degrees can be even is if there are an even number of odd nodes, i.e., $|O|$ is even. □

5.4 Christofides Algorithm

Putting the pieces together, we now have the following algorithm:

1. Find the MST T of the complete graph on X with weights defined by d.
2. Add every edge in T to the set E.
3. Let O be the nodes in T with odd degree. Notice that $|O|$ is even.
4. Let M be the minimum cost perfect matching for O.
5. Construct the multigraph $G = (X, E \cup M)$.
6. Since every node in multigraph $G = (V, E \cup M)$ has even degree degree, we can find a Eulerian Cycle E for G.
7. Return E.

To analyze this, we notice that the cost consists of the following:

$$cost(E) = \sum_{e \in E} d(e) + \sum_{e \in M} d(e)$$

The first part of the expression comes from the MST, and hence we already know that:

$$\sum_{e \in E} d(e) \leq OPT$$

Since OPT forms a cycle, we can remove any edge from the cycle to get a spanning tree; the MST T must have cost no greater than this cycle.

We now need to argue that $2 \sum_{e \in M} d(e) \leq OPT$. Let C be the cycle for OPT. Let C' be the same cycle as C where we skip all the nodes in $X - O$, and we skip repeats. That is, C' is a cycle on the odd nodes with no repeats. Notice that $cost(C') \leq cost(C)$, since we have only skipped nodes (and the triangle inequality holds).

Also, notice that cycle C' has an even number of nodes (because there are an even number of odd nodes) and an even number of edges. Assume that $C' = (v_1, v_2, \ldots, v_{2k})$.

We can now construct two different perfect matchings M_1 and M_2. We define them as follows:

$$\begin{aligned} M_1 &= (v_1, v_2), (v_3, v_4), (v_5, v_6), \ldots, (v_{2k-1}, v_{2k}) \\ M_1 &= (v_2, v_3), (v_4, v_5), (v_6, v_7), \ldots, (v_{2k}, v_1) \end{aligned}$$

Notice that each of these perfect matchings has $k = |O|/2$ edges, and both are valid perfect matchings for the set O. Since M is the minimum cost perfect matching, we conclude that:

$$\begin{aligned} cost(M) &\leq cost(M_1) \\ cost(M) &\leq cost(M_2) \end{aligned}$$

Notice, though, that $cost(M_1) + cost(M_2) = cost(C') \leq cost(C)$. So, if we sum the matchings, we get:

$$2cost(M) \leq cost(M_1) + cost(M_2) \leq cost(C) = OPT$$

Thus we have shown that $cost(M) \leq OPT/2$.

Putting the pieces together, we see that the cost of the cycle output by the algorithm is:

$$\begin{aligned} cost(E) = \sum_{e \in E} d(e) + \sum_{e \in M} d(e) &\leq cost(T) + cost(M) \\ &\leq OPT + OPT/2 \\ &\leq 1.5 \cdot OPT \end{aligned}$$

Thus, we have discovered a 1.5-approximation algorithm for the traveling salesman problem.