Polynomial Time Approximation Schemes for the Euclidean Traveling Salesman Problem

Shanshan Wu Vatsal Shah

October 20, 2015

Abstract

In this report, we aim to understand the key ideas and major techniques used in the assigned paper "Polynomial Time Approximation Schemes for Euclidean Traveling Salesman and Other Geometric Problems" by Sanjeev Arora. We also provide a review of related literature with an emphasis on the concurrent work by Joseph S. B. Mitchell. One presentation topic is selected from Arora's paper and the other topic is from Mitchell's paper.

1 INTRODUCTION

The *traveling salesman problem* (or TSP) is a very classic and well-known combinatorial problem. It considers the following optimization problem:

Given a complete graph $G = (V, E)$ with non-negative costs for each edge, find a Hamiltonian tour (i.e., a tour that visits each node exactly once) of minimum total cost.

In class we have learned that the corresponding decision problem of TSP is NP-complete [4]. In other words, for a given value $k > 0$, deciding whether G has a Hamiltonian tour of cost at most k cannot be done in polynomial time, unless P=NP. Further, it is also hard to get an approximate solution in polynomial time, because the existence of any constant factor approximation algorithm of TSP can be used as a solver for the Hamiltonian-cycle problem, and hence implies P=NP.

Although the general TSP is hard to solve and hard to approximate, by restricting our attention to special cases of the TSP inputs, it is possible to obtain a good approximation algorithm. For example, the *Metric TSP* restricts the cost function to satisfy triangle inequality. There have been constant-factor approximation algorithms developed for Metric TSP (We have seen in class a 2-approximation algorithm [8]; please refer to Section 3 for more related results). However, no *polynomial-time approximation scheme* (PTAS) exists for Metric TSP, unless P=NP [2]. Recall that a PTAS for TSP is defined as follows:

Let OPT be the cost of the optimum traveling salesman tour. A PTAS is a family of algorithms $\{A_\epsilon\}$*, where* $\forall\epsilon > 0$*, a (1+*ϵ*)-approximate algorithm exists whose running time is polynomial in* n *(where* $n = |V|$*).*

Another special class of the TSP instances is *Euclidean TSP*, which assumes that the input graph G lies in an Euclidean plane and the costs between any two vertices are their Euclidean distances. For example, in $\mathbb{R}^2$, the cost between two points (x_i, y_i) and (x_j, y_j) is $c_{ij} = \sqrt{(x_i - x_j)^2 + (y_i - y_j)^2}$. Obviously, Euclidean TSP is a sub-class of Metric TSP. It remained an open question as to whether Euclidean TSP has a PTAS. A breakthrough result was achieved by Arora in 1996 [1]. In this seminal paper, a PTAS was developed for any $d-$dimensional Euclidean TSP, which runs in time $O(n(\log n)^{(O(\sqrt{d}/\epsilon))^{d-1}})$. Around the same time, Mitchell [6] independently proposed a PTAS for 2-dimensional Euclidean TSP with running time $n^{O(1/\epsilon)}$. The goal of this project is to understand the key ideas as well as the main techniques used in these two PTASs.

2 Results and Techniques

2.1 Results

As mentioned in the introduction section, the major contribution of Arora's paper [1] as well as Mitchell's paper [6] is the discovery of a PTAS for Euclidean TSP. Prior to their results, the best known polynomial time algorithm is a 3/2-approximation algorithm proposed by Christofides in 1976 [3].

At a high level, Arora's PTAS works in four steps. First, we perturb the given instance to a new input so that this perturbed instance has certain "nice" properties. This perturbation is mainly done by rescaling all the points, and then moving each point to the nearest grid point. Second, for the perturbed nice instance, we partition its bounding box recursively into smaller squares. Since each square has 4 children squares, this recursive partition results in a 4-ary tree (also called a *dissection*). Third, we allow tours to have "bent" edges and restrict our attention to tours with the following properties: each edge enters and leaves the squares only at specified locations (called "portals", we place m equally-spaced portals along each boundary of the squares); the number of times that a tour can cross each boundary of the squares are bounded by an integer r (hence these tours are called *(m, r)-light portal-respecting tours*). Then we use dynamic programming to find the optimal (m, r)-light portal-respecting tour. This is done in a bottom-up manner over the 4-ary tree: find the cheapest (m, r)-light portal-respecting tour for each of the four children squares, and then combine the results to find the cheapest tour for the parent square. In the fourth step, we transform the optimal (m, r)-light portal-respecting tour back to a traveling salesman tour with respect to the original instance.

The main theorem proved in Arora's paper is the *Structure Theorem*, which provides a probabilistic guarantee for the existentce of a (m,r)-light portal-respecting tour with cost at most $(1+\epsilon)\cdot$OPT. Before introducing this theorem, we need a definition for *shifted dissection*. Suppose that the perturbed instance is bounded by a $L \times L$ square (assume that L is a power of 2), with four corners located at $(0,0)$, $(0, L)$, $(L, 0)$, and (L, L). Then a *dissection* is generated by dividing the bounding box into four equally-sized squares (i.e., placing lines at $x = L/2$ and $y = L/2$), and then dividing each of the small squares into four equally-sized squares (i.e., placing lines at $x = L/4$, $x = 3L/4$, $y = L/4$, and $y = 3L/4$), and repeat until all squares are of unit length. Let $0 \le a, b \le L$ be integers, then a dissection with shift (a, b) is obtained by moving each of the grid points in the dissection, say (i, j), to $(i + a, b + j) mod L$. Now we are ready to state the Structure Theorem for the perturbed Euclidean TSP instance: let $m = O(c \cdot \log L)$ and $r = O(1/\epsilon)$, pick a shift $0 \le a, b \le L$ randomly, then with probability at least 1/2, there exists a (m, r)-light portal-respecting tour with cost at most $(1+\epsilon)$OPT. It is the central theorem that supports Arora's PTAS, and is also a major result in our assigned paper, so we choose this theorem as one presentation topic.

The main result in Mitchell's paper [6] is the proof of existence of a PTAS for 2-dimensional Euclidean TSP. More specifically, Mitchell proposed an $O(n^{20m+5})$ algorithm that computes an approximate tour for the TSP whose length is within a factor of $\left(1+2\sqrt{2}/m\right)$ of the length of the optimal tour.

A subdivision is *m-guillotine* if the polygonal subdivision has the property that there exists a (special) line whose intersection with the edge set consists of $O(m)$ connected components (where $O(m)$ is small) and the subdivisions on either side of the line are also *m-guillotine*. This paper proves the existence of this (special) line for any polygonal subdivision.

Mitchell [6] states that for any rectangular subdivision, there exists *m-guillotine* rectangular subdivision obtained by adding edges such that the total cost of the new edge set is less than $(1+1/m)$ times the original cost. He then extends this result for polygonal subdivision when the new edge set has a total cost of at most $\left(1+\sqrt{2}/m\right)$ times the cost of the edges in the original subdivision. The factor $\sqrt{2}$ comes as a result of the inclusion of each inclined edge in both the vertical and horizontal cuts. The cost of added edges is doubled (or another edge is included over the added edge) during the transformation to *m-guillotine* subdivision to ensure the existence of the Eulerian subgraph in the resultant solution. The total cost of the set of edges including the added edges is now less than $\left(1+2\sqrt{2}/m\right)$ times the initial cost of the edges in that subdivision.

2.2 Techniques

According to the previous subsection, Arora's PTAS involves two key techniques: (m,r)-light portal-respecting tour and shifted dissection. Recall that m is the number of equally-spaced portals placed along each side of squares in the dissection, and r is the upper bound that a tour can cross the boundary of each square. The reason why we only search for the optimal (m,r)-light portal-respecting tours is due to the running time control. Note that for a square S, the running time of enumerating all possible ways that a (m,r)-light portal-respecting tour would occur inside S is bounded by $O(m^{4r}\cdot(4r)!)$, because there are at most m^{4r} ways to select the portals that a tour can enter, and at most $(4r)!$ ways to pair the enter-leave portals. Hence, by restricting our attention to (m,r)-light portal-respecting tour, we can obtain a polynomial-time algorithm by carefully choosing m and r. However, to obtain a PTAS, we still need to show that the optimal (m,r)-light portal-respecting tour has a cost at most $(1+\epsilon)$OPT. This is achieved by adding randomization from shifted dissections. As mentioned in the previous subsection, the Structure Theorem ensures that with proper choice of m and r, the algorithm would find a tour of cost at most $(1+\epsilon)$OPT with probability at least $1/2$.

Similar to Arora's PTAS, Mitchell's PTAS also recursively partitions the planar Euclidean TSP into sub-partitions that can be solved using dynamic programming. Specifically, Mitchell's PTAS involves three major steps: (1) Transforming the optimal TSP tour into *m-guillotine* subdivisions by adding segments (in both directions) such that the total length of all the edges is less than $\left(1+2\sqrt{2}/m\right)$ times the length of the initial edge set, (2) Using dynamic programming to recursively find the minimum length *m-guillotine* subdivision that satisfies certain boundary and connectivity conditions and finally, (3) Utilizing the dynamic programming solution to form a Eulerian cycle which is used to generate a Hamiltonian tour.

3 Literature Survey

As stated in the introduction, the Travelling Salesman Problem is one of the oldest combinatorial problems around, conceived almost two centuries ago by W. Hamilton and Thomas Kirkman in the early 1800s. Consequently, this problem has been studied in great detail over the years and there is extensive literature available on this problem. One of the major breakthroughs came in 1972 when Richard Karp first showed that the Hamiltonian cycle was NP-complete which implies the NP hardness of the TSP. Since the general TSP does not admit a PTAS, we will only focus on the Euclidean TSP in the literature survey for a general TSP instead of looking at particular graphs. Recall that OPT indicates the cost of the optimal tour of the TSP and n is the number of distinct points on the graph.

3.1 Prior Results

The most common approaches to solve the Euclidean TSP using approximation algorithms involved the use of the Minimum Spanning Tree (MST) algorithm. 2-approximation algorithms existed for the TSP including the nearest addition algorithm in [8]. The breakthrough result was achieved by Christofides in [3] which gives an algorithm to generate a Hamiltonian tour that achieves a 3/2-approximation with respect to the optimal tour. Christofides' method involves finding a minimum cost perfect matching on the set of vertices with odd degree in the MST and then completing a Eulerian cycle by combining the edges in MST and the minimum cost perfect matching.

3.2 Concurrent Results

As mentioned in the introduction, Arora [1] and Mitchell [6] independently came up with a PTAS for the Euclidean TSP around the same time. Both these methods involve: (1) dividing the Euclidean plane, albeit in different ways, [6] involves using Guillotine subdivisions while [1] uses quadtrees, and (2) using dynamic programming to obtain a solution for the subdivisions. [1]

3.3 Followup Results

Rao and Smith [7] modified Arora's algorithm [1] to give a PTAS for 2-dimensional Euclidean TSP that runs in $O(n \log n)$ time. The major difference is that [7] begins by computing a $(1+\epsilon)$-spanner tree, followed by applying Arora's techniques [1] on spanner graph instead of the original graph. Unlike in Arora's PTAS, where we need to place m equally-spaced portals along each boundary of squares in the dissection, the $(1+\epsilon)$-spanner tree can be modified so that every small piece of it interacts outside at most certain number of times, and hence provides a more tailored location for placing the portals.

Dumitrescu and Mitchell [5] discussed a modified problem of finding approximation algorithms for TSP with neighborhoods (TSPN),where each neighborhood intersects with the edge set, E. The aim is to find the shortest tour that passes through all the n neighborhoods. The paper gives approximation algorithms for a variety of cases of the TSPN: (1) Constant factor approximation algorithms for connected neighborhoods with equal diameters/unequal diameters (with restrictions), (2) PTAS algorithms described for disjoint neighborhoods with equal diameters are extensions of work in [6], (3) Constant factor approximation algorithm for neighborhoods with a special line structure.

[1] Both Sanjeev Arora and Joseph Mitchell won the Godel prize for discovering the PTAS for the Euclidean TSP.

3.4 Connection to bibliography in the Assigned Paper

The assigned paper cites [3] and also recognizes the concurrent result obtained by [6] and the follow up result by [7]. These papers have been cited exactly in the assigned paper.

4 Presentation Topics

1. *Shanshan Wu* will present the main idea behind Arora's PTAS and the proof of the Structure Theorem (which corresponds to Theorem 2 in the paper). As mentioned in the previous section, it is the central theorem in obtaining a PTAS. In the assigned paper [1], the PTAS is explained in Section 2 (pp. 760-762). The Structure Theorem (for 2-dimensional Euclidean TSP) is stated in pp. 760. Its proof is elaborated in pp. 762-768.
2. *Vatsal Shah* will discuss the central theme (Theorem 3.1 and Lemma 3.2) behind Mitchell's PTAS [6]. The presentation will involve some preliminary results about the existence of *m-gullotine* subdivision for any rectangular subdivision and then describe the main result involving the TSP which is presented in Corollary 4.2 on pp. 1306-1308.

References

[1] S. Arora. Polynomial time approximation schemes for euclidean traveling salesman and other geometric problems. *Journal of the ACM (JACM)*, 45(5):753–782, 1998.

[2] S. Arora, C. Lund, R. Motwani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM (JACM)*, 45(3):501–555, 1998.

[3] N. Christofides. Worst-case analysis of a new heuristic for the travelling salesman problem. Technical report, DTIC Document, 1976.

[4] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. The MIT Press, 2009.

[5] A. Dumitrescu and J. S. B. Mitchell. Approximation algorithms for tsp with neighborhoods in the plane. *Journal of Algorithms*, 48(1):135–159, 2003.

[6] J. S. B. Mitchell. Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric tsp, k-mst, and related problems. *SIAM Journal on Computing*, 28(4):1298–1309, 1999.

[7] S. B. Rao and W. D. Smith. Approximating geometrical graphs via spanners and banyans. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 540–550. ACM, 1998.

[8] D. J. Rosenkrantz, R. E. Stearns, and P. M. Lewis, II. An analysis of several heuristics for the traveling salesman problem. *SIAM journal on computing*, 6(3):563–581, 1977.