

Polynomial Time Approximation Schemes for Euclidean Traveling Salesman and Other Geometric Problems

SANJEEV ARORA

Princeton University, Princeton, New Jersey

Abstract. We present a polynomial time approximation scheme for Euclidean TSP in fixed dimensions. For every fixed $c > 1$ and given any n nodes in \mathcal{R}^2 , a randomized version of the scheme finds a $(1 + 1/c)$ -approximation to the optimum traveling salesman tour in $O(n(\log n)^{O(c)})$ time. When the nodes are in \mathcal{R}^d , the running time increases to $O(n(\log n)^{O(\sqrt{dc})^{d-1}})$. For every fixed c, d the running time is $n \cdot \text{poly}(\log n)$, that is *nearly linear* in n . The algorithm can be derandomized, but this increases the running time by a factor $O(n^d)$. The previous best approximation algorithm for the problem (due to Christofides) achieves a $3/2$ -approximation in polynomial time.

We also give similar approximation schemes for some other NP-hard Euclidean problems: Minimum Steiner Tree, k -TSP, and k -MST. (The running times of the algorithm for k -TSP and k -MST involve an additional multiplicative factor k .) The previous best approximation algorithms for all these problems achieved a constant-factor approximation. We also give efficient approximation schemes for Euclidean Min-Cost Matching, a problem that can be solved exactly in polynomial time.

All our algorithms also work, with almost no modification, when distance is measured using any geometric norm (such as ℓ_p for $p \geq 1$ or other Minkowski norms). They also have simple parallel (i.e., NC) implementations.

Categories and Subject Descriptors: F.2.2 [Analysis of Algorithms and Problem Complexity]: Geometrical problems and computations, Routing and layout; G.2.2 [Graph Theory]: Path and circuit problems, Trees

General Terms: Algorithms, Theory

Additional Key Words and Phrases: Approximation Algorithms, Traveling Salesman Problem, Steiner Problem, Network Design, Matching

1. Introduction

In the *Traveling Salesman Problem* (“TSP”), we are given n nodes and for each pair $\{i, j\}$ of distinct nodes, a distance $d_{i,j}$. We desire a closed path that visits each node exactly once (i.e., is a *salesman tour*) and incurs the least *cost*, which is the sum of the distances along the path. This classic problem has proved a rich

The work of S. Arora was supported by NSF CAREER award NSF CCR 95-02747, an Alfred Sloan Fellowship, and a Packard Fellowship.

Author’s address: Computer Science Department, Princeton University, 35 Odden Street, Princeton, NJ 08544; e-mail: arora@cs.princeton.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery (ACM), Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 1998 ACM 0004-5411/98/0900-0753 \$05.00

testing ground for most important algorithmic ideas during the past few decades, and influenced the emergence of fields such as operations research, polyhedral theory and complexity theory. For a fascinating history, see Lawler et al. [1985].

Since the 1970s, mounting evidence from complexity theory suggests that the problem is computationally difficult. Exact optimization is NP-hard (Karp [1972]). So is approximating the optimum within any constant factor (Sahni and Gonzalez [1976]). There are also other reasons to believe in the TSP's nastiness (cf. D^P completeness [Papadimitriou and Yannakakis 1984] and PLS-completeness [Johnson et al. 1988]).

But TSP instances arising in practice are usually quite special, so the hardness results may not necessarily apply to them. In *metric TSP* the nodes lie in a metric space (i.e., the distances satisfy the triangle inequality). In *Euclidean TSP* the nodes lie in \mathbb{R}^2 (or more generally, in \mathbb{R}^d for some d) and distance is defined using the ℓ_2 norm. Note that Euclidean TSP is a subcase of metric TSP.

Unfortunately, even Euclidean TSP is NP-hard (Papadimitriou [1977]; Garey et al. [1976]). Therefore, algorithm designers were left with no choice but to consider more modest notions of a “good” solution. Karp [1977], in a seminal work on *probabilistic analysis of algorithms*, showed that when the n nodes are picked uniformly and independently from the unit square, then the *fixed dissection heuristic* with high probability finds tours whose cost is within a factor $1 + 1/c$ of optimal (where $c > 1$ is arbitrarily large). Christofides [1976] designed an *approximation algorithm* that runs in polynomial time and for every instance of metric TSP computes a tour of cost at most $3/2$ times the optimum.

Two decades of research failed to improve upon Christofides' algorithm for metric TSP. The Held–Karp heuristic is conjectured to have an approximation ratio $4/3$ (some results of Goemans [1995] support this conjecture) but the best upperbound known is $3/2$ (Wolsey [1980], Shmoys and Williamson [1990]). Some researchers continued to hope that even a PTAS might exist. A PTAS or *Polynomial-Time Approximation Scheme* is a polynomial-time algorithm—or a family of such algorithms—that, for each fixed $c > 1$, can approximate the problem within a factor $1 + 1/c$. The running time could depend upon c , but for each fixed c has to be polynomial in the input size. PTASs are known for very few problems; two important ones are Subset-Sum (Ibarra and Kim [1975]) and Bin-Packing (Fernandez de la Vega and Lueker [1981]; see also Karmarkar and Karp [1982]). Recently Arora et al. [1992] showed that if $P \neq NP$, then metric TSP and many other problems do not have a PTAS. Their work relied upon the theory of *MAX-SNP-completeness* (Papadimitriou and Yannakakis [1991]), the notion of *probabilistically checkable proofs* or PCPs (Feige et al. [1991]; Arora and Safra [1992]), and the connection between PCPs and hardness of approximation [Feige et al. 1991].

The status of Euclidean TSP remained open, however. In this paper, we show that Euclidean TSP has a PTAS. For every fixed $c > 1$, a randomized version of this algorithm computes a $(1 + 1/c)$ -approximation to the optimal tour in $O(n(\log n)^{O(c)})$ time. When the nodes are in \mathbb{R}^d , the running time rises to $O(n(\log n)^{O(\sqrt{d}c)^{d-1}})$. Our algorithm can be derandomized, but this seems to multiply the running time by a factor $O(n^d)$ in \mathbb{R}^d . Our techniques also apply to many other geometric problems, which are described in Section 1.1.

We design the PTAS by showing that the plane can be recursively partitioned (using a randomized variant of the quadtree) such that some $(1 + 1/c)$ -

approximate salesman tour crosses each line of the partition at most $r = O(c)$ times (see Theorem 2.1.2). Such a tour can be found by dynamic programming. For each line in the partition the algorithm first “guesses” where the tour crosses this line and the order in which those crossings occur. Then the algorithm recurses independently on the two sides of the line. There are only $n \log n$ distinct regions in the partition. Furthermore, the “guess” can be fairly coarse, so the algorithm spends only $(O(\log n))^{O(r)} = (\log n)^{O(c)}$ time per region, for a total running time of $n \cdot (\log n)^{O(c)}$.

We remark that the idea of partitioning a TSP instance into smaller instances and dynamic programming has been used before, most famously in Karp [1977]. Smith [1988] showed how to solve the TSP in \mathbb{R}^2 to optimality in $2^{O(\sqrt{n})}$ time; the main idea is that the plane can be recursively partitioned such that an optimal tour crosses every partition only $O(\sqrt{n})$ times. Recently Grigni et al. [1995] designed an approximation scheme for planar graph TSP using similar ideas.

Finally, the inevitable question: Is our PTAS practical? A straightforward implementation (for even moderate values of c) is very slow, but we see no reason why a speedier, more subtle, implementation may not exist (see Section 4.1 in the appendix). At the very least, the Theorem gives a way of decomposing TSP instances into a large number of “independent” and smaller instances, and this may prove helpful in parallelizing existing programs. We note that classical local-exchange heuristics such as K -OPT or Lin–Kernighan [1973] are already known to compute very good tours on “real-life” TSP instances [Johnson and McGeoch 1997; Bentley 1992]. However, the performance of known heuristics for problems such as Steiner Tree or k -TSP is not as good.

We find it conceivable that our techniques, which show that Euclidean instances have near-optimal salesman tours with a very simple structure, may contribute to a theoretical understanding of local-exchange heuristics on Euclidean instances. (The “real-life” instances in TSPLIB [Reinelt 1991] that constitute the test-bed for many of these heuristics are often either Euclidean or derived from Euclidean instances.) For example, even our current dynamic programming algorithm can be viewed—after some twists in the definition of “local search”—as a local search algorithm that performs up to $O(c)$ edge exchanges per step (see Section 4.2). Note however that none of the known heuristics is believed to be a PTAS.¹

History. The current paper evolved out of preliminary results obtained in January 1996, culminating in a submission to *IEEE FOCS 1996* in April 1996 [Arora 1996]. A few weeks later, Mitchell [1996/1998] independently discovered an $n^{O(c)}$ time approximation scheme for points in \mathbb{R}^2 . His algorithm used ideas from his earlier constant-factor approximation algorithm for k -MST [Mitchell 1996]. It relies on the geometry of the plane and does not seem to generalize to

¹ The few published results in fact suggest the opposite. With an adversarially-chosen starting tour, 2-OPT may produce a tour whose cost is $\Omega(\log n / \log \log n)$ times the cost of the optimum tour, even when the n nodes lie in \mathbb{R}^2 [Chandra et al. 1994]. K -OPT can in the worst case produce tours whose cost is twice the optimum. In case of metric TSP, finding a locally-optimum tour for K -OPT (for $K \geq 8$) is PLS-complete [Krentel 1989]. This strongly suggests that no polynomial-time algorithm can find such a local optimum; see Johnson et al. [1988]. Many variants of Lin–Kernighan are also PLS-complete [Papadimitriou 1992].

higher dimensions. In January 1997, the author discovered the nearly-linear-time algorithm described in this paper. The key ingredient of this algorithm is Theorem 2.1.2, which the author had originally conjectured to be false. He is grateful to Karen Wang, whose inability [1996] to construct any counterexample to Theorem 2.1.2 motivated him to abandon his conjecture.

1.1. DEFINITIONS AND RESULTS. The input to our algorithm is a set of n points in \mathcal{R}^d , given by their coordinates. For $p \geq 1$, the distance between two points $(x_1, \dots, x_d), (y_1, \dots, y_d) \in \mathcal{R}^d$ in the ℓ_p norm is defined as $(\sum_{i=1}^d |x_i - y_i|^p)^{1/p}$. When $p = 2$, this norm is called the *Euclidean norm*; this is the norm we will usually discuss in detail.

We will use the standard Real RAM model of computation, which assumes a unit cost for arithmetic on real numbers. Strictly speaking, we will need computations on real numbers only to “preprocess” the input in linear time. After this, we truncate coordinates and edge costs to their $2 \log n$ most significant digits. This affects all edge costs (and hence the optimum tour cost) by at most a multiplicative factor $(1 + 1/n^2)$ —which is negligible in the context of an approximation scheme.

Now we define the geometric problems for which we will design approximation schemes. Except for Euclidean Matching, all are NP-hard. Prior to our work, the best approximation algorithms for the NP-hard problems achieved a constant factor approximation in polynomial time (see the survey by Bern and Eppstein [1996]). These algorithms used problem-specific ideas and usually require at least $\Omega(n^2)$ time (sometimes a lot more). In contrast, our approximation schemes for the different problems rely on essentially the same idea.

Euclidean Traveling Salesman. Given n nodes in \mathcal{R}^d , find the shortest tour that visits all nodes.

Minimum Steiner Tree. Given n nodes in \mathcal{R}^d , find the minimum-cost tree connecting them.² In general, the minimum spanning tree is not an optimal solution. In \mathcal{R}^2 (with distances measured in ℓ_2 norm) the cost of the MST can be as far as a factor $2/\sqrt{3}$ from the optimum. Furthermore, the famous Gilbert–Pollak [1968] conjecture said it can’t be any further from the optimum; this conjecture was recently proved by Du and Hwang [1992]. A spate of research activity in recent years (starting with the work of Zelikovsky [1993]) has provided better algorithms, with an approximation ratio around 1.143 [Zelikovsky 1996]. The metric case is MAX-SNP-hard [Bern and Plassmann 1989].

k-TSP. Given n nodes in \mathcal{R}^d and an integer $k > 1$, find the shortest tour that visits at least k nodes. An approximation algorithm due to Mata and Mitchell [1995] achieves a constant factor approximation in \mathcal{R}^2 .

k-MST. Given n nodes in \mathcal{R}^d and an integer $k \geq 2$, find k nodes with the shortest Minimum Spanning Tree. The problem is NP-hard [Fischetti et al. 1994]. Blum et al. [1995] gave the first $O(1)$ -factor approximation algorithm for points in \mathcal{R}^2 and Mitchell [1996] improved this factor to $2\sqrt{2}$.

² It appears that this problem was first posed by Gauss in a letter to Schumacher (R. L. Graham, personal communication).

Euclidean Min-Cost Perfect Matching (EMCPM). Given $2n$ points in \mathbb{R}^2 (or \mathbb{R}^d in general), find the minimum cost set of nonadjacent edges that cover all vertices. This problem is in P (even for nongeometric instances). Vaidya [1988] shows how to solve it optimally in $\tilde{O}(n^{2.5})$ time, and to approximate it within a factor $(1 + 1/c)$ in $O(\text{poly}(c)n^{1.5} \log^3 n)$ time [Vaidya 1989].

1.1.1. *For each fixed d , the \mathbb{R}^d version of each of the above problems has a randomized PTAS. The algorithm computes a $(1 + 1/c)$ -approximation with probability at least $1/2$. In \mathbb{R}^2 the running time is $n(\log n)^{O(c)}$ for TSP, Steiner Tree, and Min-Cost Euclidean Matching, and $nk(\log n)^{O(c)}$ for k -TSP and k -MST. All running times on instances in \mathbb{R}^d are larger by a factor $(O(\log n))^{(O(\sqrt{dc}))^{d-1}}$. The above expressions for running times are unchanged when the problem is specified using any Minkowski norm instead of the Euclidean norm. Furthermore, all the above PTASs can be derandomized, by increasing the running time in \mathbb{R}^d by $O(n^d)$.*

The various parts of this theorem will be proved in different sections (titled appropriately). The TSP part will be proved in Section 2 and the parts corresponding to other problems will be proved in Section 3.

We remark that though our algorithms generalize to \mathbb{R}^d with no difficulties, this was not always the case with previous (constant factor) approximation algorithms for k -TSP and k -MST. Those algorithms relied on the geometry of the plane and broke down even in \mathbb{R}^3 . But a recent algorithm of Garg [1996]—discovered independently of our paper—works in any metric space.

Geometric versions of polynomial time problems have been studied for many years, especially MST and Euclidean Matching. Exploiting geometric structure for these problems is known to lead to faster algorithms than the corresponding algorithms for general graphs. The best MST algorithm requires $\tilde{O}(2^d n^{4/3})$ time in \mathbb{R}^d for $d > 2$ [1]. Note that treating the instance as a weighted graph with $\binom{n}{2}$ edges would lead to $\Omega(n^2)$ running time. Vaidya described a nearly-linear-time approximation scheme for geometric MST, which computes a $(1 + 1/c)$ -approximation in $\tilde{O}(cn2^{O(d)})$ time.

2. The TSP Algorithm

This section describes our approximation scheme for the TSP. Section 2.1 describes the algorithm for Euclidean TSP in \mathbb{R}^2 . The proof of the algorithm's correctness relies upon Theorem 2.1.2, which is proved in Section 2.2. Then Section 2.3 describes the algorithm for Euclidean TSP in \mathbb{R}^d and Section 2.4 describes the (trivial) extensions of our algorithm when distances are measured in some non-Euclidean norms.

2.1. EUCLIDEAN TSP IN \mathbb{R}^2 . As mentioned in the introduction, the main idea in the algorithm is to perform a (recursive) geometric partitioning of the instance. The geometric partitioning is very simple: it is a randomized variant of the familiar quadtree. Theorem 2 below says that with probability at least $1/2$, (over the choice of the randomness used to construct this partition), there exists a $(1 + 1/c)$ -approximate tour that crosses each line of the partition $O(c)$ times. Furthermore, these crossings happen at some prespecified points; see below for more details. Such a tour can be found easily using dynamic programming that runs in $n \cdot (\log n)^{O(c)}$ time.




FIG. 1. (a) The dissection. (b) The corresponding quadtree.

To simplify exposition we perform an $O(n \log n)$ -time “perturbation” of the instance at the start that makes all coordinates integral, and makes the minimum internode distance at least 8. (The perturbation is described later.) Let the bounding box of the perturbed instance be the smallest axis-aligned square that contains it, and let L be the length of each side of this square. We will refer to L as the size of the box, and will assume without loss of generalities that it is a power of 2.

A *dissection* of the bounding box is a recursive partitioning into smaller squares. We view it as a 4-ary tree whose root is the bounding box. Each square in the tree is partitioned into four equal squares, which are its children. We stop partitioning a square if it has size ≤ 1 (and therefore at most one node). Note that there are $O(L^2)$ squares in the dissection and its depth is $\log L$. A *quadtree* is defined similarly, except we stop the recursive partitioning as soon as the square has at most one node. The quadtree may in general have fewer squares than the dissection; see Figure 1. In fact, since each leaf either contains a node or is a sibling of a square that contains a node, the quadtree contains $O(n)$ leaves and thus $O(n \log L)$ squares in all.

If a, b are integers in $[0, L)$, then the (a, b) -shift of the dissection is defined by shifting the x - and y - coordinates of all lines by a and b respectively, and then reducing modulo L . (Of course, the nodes in the TSP instance do not move.) In other words, the middle vertical line of the dissection is moved from the x -coordinate $L/2$ to the x -coordinate $a + (L/2) \bmod L$, and the middle horizontal line from the y -coordinate $L/2$ to the y -coordinate $b + (L/2) \bmod L$. The rest of the dissection is then “wrapped-around,” so that the left edge of the dissection comes to rest at the x -coordinate a , and the lower edge of the dissection comes to rest at the y -coordinate b (see Figure 2). Note that we treat a “wrapped-around” square in the shifted dissection as a single region; this will greatly simplify notation later. The reader can also think of a wrapped-around square as a disjoint union of 2 (or 4) rectangles.

The *quadtree with shift* (a, b) is obtained from the corresponding shifted dissection by cutting off the partitioning at squares that contain only 1 node. It is easy to see that in general the shifted quadtree has a very different structure than the original quadtree.




FIG. 2. Dissection with shift (a, b) . Only the four children of the root are shown; each is identified by a distinctive shading. Note that three of the children are “wrapped-around” squares.

Our Structure Theorem below will show that if the shift (a, b) is picked randomly, then with probability at least $1/2$ (over the choice of the random shift), there exists a $(1 + 1/c)$ -approximate tour that crosses the boundary of each square in the shifted quadtree $O(c)$ times. In fact, the Structure Theorem says something stronger: these crossings happen at one of a small set of prespecified points (called *portals*). The previous sentence may mystify the reader: how can the crossing points be prespecified?

The reason is that we will allow the salesman to deviate from the straight-line path while traveling between nodes; these digressions allow it to pass through any desired prespecified point. Such a salesman tour with digressions will be called a *salesman path*.³

Definition 1. Let m, r be positive integers. An m -regular set of portals for a shifted dissection is a set of points on the edges of the squares in it. Each square has a portal at each of its 4 corners and m other equally-spaced portals on each edge.

A *salesman path* is a path in \mathbb{R}^2 that visits all the input nodes, and some subset of portals. It may visit a portal more than once.

The salesman path is (m, r) -light with respect to the shifted dissection if it crosses each edge of each square in the dissection at most r times and always at a portal.

Notes

- (i) Between visits to two successive input nodes i and j , the tour may cross region boundaries and therefore have to pass through a sequence of portals P_1, P_2, \dots , which may not lie on the straight line connecting nodes i, j . Thus the “edge” from i to j consists of line segments $(i, P_1), (P_1, P_2), (P_2, P_3) \dots$. We think of edge (i, j) as being “bent” at P_1, P_2, \dots . The bent edges will also play an important role in the proof of Theorem 2.1.2.
- (ii) We will assume that no node lies on the boundary of any region in the dissection. We can ensure this (among other ways) by scaling distances by a factor 2 and then ensuring that nodes have odd coordinates and the lines of the dissection have even coordinates.

³ This usage should not be confused with another common use of the term *salesman path*, namely, a tour that does not return to the starting point.

In the rest of this section, we will let OPT denote the cost of the optimal salesman tour (with no bent edges). Our algorithm will return a salesman path of cost at most $(1 + \epsilon)OPT$. Of course, we can straighten the bent edges at the end to get a salesman tour, without raising the cost. (This follows from the triangle inequality.) Now we state our main theorem; the proof appears in Section 2.2.

THEOREM 2 (STRUCTURE THEOREM). *Let $c > 0$ be any constant. Let the minimum nonzero internode distance in a TSP instance be 8 and let L be the size of its bounding box. Let shifts $0 \leq a, b \leq L$ be picked randomly. Then with probability at least $1/2$, there is a salesman path of cost at most $(1 + 1/c)OPT$ that is (m, r) -light with respect to the dissection with shift (a, b) , where $m = O(c \log L)$ and $r = O(c)$.*

Remark. Currently, the best constant we know how to achieve in the $O(c)$ is around 10, but that can probably be improved.

2.1.1. THE ALGORITHM. Assuming the truth of the Structure Theorem, we describe our PTAS.

STEP 1: PERTURBATION. As mentioned earlier, the algorithm first “perturbs” the instance to make it *well-rounded*. This means: (i) all nodes have integral coordinates, (ii) each (nonzero) internode distance is at least 8 units, and (iii) the maximum internode distance is $O(n)$. The instance is made well-rounded as follows: Let L_0 be the size of the bounding box of the given instance and OPT be the optimum tour cost. Note that $OPT \geq L_0$. We place a grid of granularity $L_0/8nc$ in the plane and move each node to its nearest gridpoint. (More than one node may map to the same gridpoint, which is why we phrased Theorem 2.1.2 so that only the minimum nonzero internode distance is 8.) Note that for any fixed order of visiting the nodes, the cost of that tour in the two instances differs by at most $2n \cdot L_0/8nc < OPT/4c$. It follows that the perturbation affected the cost of the optimum by at most $OPT/4c$. Now we divide distances by $L_0/64nc$. Thus, all coordinates become integral and the minimum (nonzero) internode distance is ≥ 8 . Furthermore, L , the size of the bounding box of the perturbed instance, is $O(nc)$, which is $O(n)$ since c is a constant. Thus, the instance has become well-rounded. Note that we now need to compute a $(1 + 3/4c)$ -approximation in this new instance instead of a $(1 + 1/c)$ -approximation, but that doesn’t matter since $c > 1$ represents an arbitrary constant.

STEP 2: CONSTRUCTING A SHIFTED QUADTREE. We pick a shift (a, b) randomly and compute a quadtree with these shifts. Since the bounding box has size $L = O(n)$, the depth of the shifted quadtree is $O(\log n)$, and the number of squares in it is

$$T = O(\text{number of leaves with a node} \times \text{depth}) = O(n \log n).$$

The shifted quadtree can be constructed easily using a sorting-based algorithm in $O(n \log^2 n)$ time; faster algorithms also exist [Bern et al. 1993].

STEP 3: DYNAMIC PROGRAMMING. Next, we use dynamic programming to find the optimal (m, r) -light salesman path with respect to the shifted quadtree computed in Step 2, where $m = O(c \log n)$ and $r = O(c)$. With probability at least $1/2$ (over the choice of the shift used to construct the quadtree), this

salesman path has cost at most $(1 + 3/4c)OPT$. The running time of this step is $O(T \cdot (m)^{O(r)})$, which in our context is $O(n \cdot (\log n)^{O(c)})$.

The dynamic programming uses the following observation. Suppose S is a square of the shifted quadtree and the optimal (m, r) -light salesman path crosses the boundary of S a total of $2p \leq 4r$ times. Let a_1, a_2, \dots, a_{2p} be the sequence of portals where these crossings occur. (The portals have been numbered in the order they are traversed by the salesman path.)

Then the portion of the optimal salesman path inside S is a sequence of p paths such that (i) for $i = 1, \dots, p$, the i th path connects a_{2i-1} to a_{2i} , (ii) together the paths visit all nodes that lie inside S , and (iii) the collection of p paths is (m, r) -light; in other words, they collectively cross each edge of each square in the quadtree at most r times, and these crossings always happen at portals.

Since the salesman path is optimal, the above sequence of p paths must be the set of paths that have lowest cost among all paths with properties (i), (ii), and (iii). This observation motivates us to define the (m, r) -multipath problem. An instance of this problem is specified by the following inputs:

- (a) A nonempty square in the shifted quadtree.
- (b) A multiset of $\leq r$ portals on each of the four edges of this square such that the sum of the sizes of these multisets is an even number $2p \leq 4r$.
- (c) A pairing $\{a_1, a_2\}, \{a_3, a_4\}, \dots, \{a_{2p-1}, a_{2p}\}$ between the $2p$ portals specified in (b).

The goal in the (m, r) -multipath problem is to find a minimum cost collection of p paths in the square that is (m, r) -light. The i th path connects a_{2i-1} to a_{2i} , and the p paths together visit all the nodes in the square. (If $p = 0$, then the goal is to find the optimum (m, r) -light salesman path for the nodes in the square.)⁴

The dynamic programming builds a lookup table containing the costs of the optimal solutions to *all* instances of the (m, r) -multipath problem arising in the quadtree. Once this table is built the algorithm is done, since the optimal (m, r) -light salesman path occurs as one of the entries in this table—the one corresponding to the root of the quadtree and $p = 0$.

The number of entries in the lookup table is just the number of different instances of the (m, r) -multipath problem in the shifted quadtree. In a quadtree with T nonempty squares, this number is $O(T \cdot (m + 4)^{4r} \cdot (4r)!)$. (For each of the T squares, there are at most $(m + 4)^{4r}$ ways to choose the multiset of portals and at most $4r!$ pairings among those portals.)

The table is built up in a bottom-up fashion. Instances at the leaves of the quadtree contain at most 1 node and $O(r)$ selected portals, so they are solved optimally in $O(r)$ time, by trying all r ways of placing the single node in $O(r)$ paths. Inductively, suppose the algorithm has solved all (m, r) -multipath problems for squares at depth $> i$ and let S be any other square at depth i . Let S_1, S_2, S_3, S_4 be its four children in the quadtree. For every choice in (b), (c) for S , the algorithm enumerates all possible ways in which an (m, r) -multipath could cross

⁴ Note that the (m, r) -multipath problem may also be viewed as a *multiple traveling salesmen* problem in which a team of salesmen have to visit a set of clients. Each client has to be visited by some salesman, and each salesman has a designated starting and stopping point, which is a portal on the boundary.

the edges of S_1, \dots, S_4 . This involves enumerating all choices for the following: (a') a multiset of $\leq r$ portals on the four inner edges of the children (note that the outer edges are part of the edges of S and so we already know where they are crossed and in what order) (b') an order in which the portals in (a') are traversed by the optimum (m, r) -multipath. The number of choices in (a') is at most $((m + 4)^r)^4$ and the number of choices in (b') is at most $(4r)^{4r}(4r)!$ (where the term $(4r)^{4r}$ upperbounds the number of ways of choosing, for each of the portals chosen in (a'), one of the $\leq 4r$ paths in which it lies). Each choice in (a') and (b') leads to a (m, r) -multipath problem in the four children, whose optimal solutions—by induction—already exist in the lookup table. Adding the cost of these four optimal solutions, the algorithm determines the cost of the pair of choices made in (a') and (b'). Doing this for each pair of choices shows the algorithm the optimal choice in (a'), (b').

Thus, the running time is $O(T \cdot (m + 4)^{8r}(4r)^{4r}(4r!)^2)$, which is $O(n(\log n)^{O(c)})$.

Thus far, we have omitted details on how to solve the multipath problem for “wrapped-around” squares. Note however that dynamic programming over such squares is if anything easier than in “normal” squares (i.e., the number of choices explored by the algorithm is smaller), since the multipath cannot go between the two (or four) portions of the wrapped-around square.

This completes the description of our PTAS.

Remarks

- (1) The dynamic programming as described above only computes the *cost* of the optimal (m, r) -light salesman path, and not the path itself. But the path can be easily reconstructed from the lookup table at the end, by looking at the decisions made at each step of the dynamic programming.
- (2) It should be clear now why we insisted that the salesman path enter and leave the regions of the quad tree only at portals. This aids efficiency in our dynamic programming, which has to enumerate all possible ways in which the path could enter and leave the region. If we had allowed the salesman tour to use normal edges while entering and exiting, then this enumeration could have required as much as $\binom{n^2/2}{O(r)}$ time (since the number of edges among n nodes is $\binom{n}{2} \approx n^2/2$), instead of $m^{O(r)}$.
- (3) Our algorithm may return a salesman path that is self-crossing. It is well-known how to remove self-crossings at the end without raising the cost [Lawler et al. 1985].
- (4) In the above description, we simplified the expression for the running time by noticing that for $c < \log n$, $O(c^{O(c)}n(\log n)^{O(c)})$ is $O(n(\log n)^{O(c)})$. We use similar simplifications elsewhere in the paper.

Derandomization. Note that only Step 2 of the algorithm uses randomization, namely to pick a pair of random numbers $a, b < L$. Hence, we can derandomize the algorithm by going through all choices for the pair (a, b) and running Step 3 for each choice (and at the end returning the lowest cost salesman tour ever found). This multiplies the running time by $L^2 = O(n^2)$.

2.2. PROOF OF THE STRUCTURE THEOREM. In this section, we prove Theorem 2. Lemmas 3 and 4 will be important ingredients of the proof.

Lemma 3 is implicit in prior work on Euclidean TSP [Beardwood et al. 1959; Karp 1977] and can safely be called a “folk theorem.” However, we have never seen it stated precisely as it is stated here. When we later use this lemma, the “closed path” π of the hypothesis will be a salesman path on some set of nodes. The closed path π' of the conclusion of the lemma will be a new salesman path.

LEMMA 3 (PATCHING LEMMA). *There is a constant $g > 0$ such that the following is true. Let S be any line segment of length s and π be a closed path that crosses S at least thrice. Then there exist line segments on S whose total length is at most $g \cdot s$ and whose addition to π changes it into a closed path π' that crosses S at most twice.*

Remark. Note that we strongly use the fact that a salesman path is allowed to have “bent” edges, since otherwise it would not be possible in general to add line segments from S to π .

PROOF. Suppose π crosses S a total of t times. Let M_1, \dots, M_t be the points on which π crosses S . Break π at those points, thus causing it to fall apart into t paths P_1, P_2, \dots, P_t . In what follows, we will need two copies of each M_i , one for each side of S . Let M'_i and M''_i denote these copies.

Let $2k$ be the largest even number less than t . Let J be the multiset of line segments consisting of the following: (i) A minimum cost salesman tour through M_1, \dots, M_t . (ii) A minimum cost perfect matching among M_1, \dots, M_{2k} . Note that the line segments of J lie on S and their total length is at most $3s$. We take two copies J' and J'' of J and add them to π . We think of J' as lying on the left of S and J'' as lying on the right of S .

Now if $t = 2k + 1$ (i.e., t is odd), then we add an edge between M'_{2k+1} and M''_{2k+1} . If $t = 2k + 2$, then we add an edge between M'_{2k+1} and M''_{2k+1} and an edge between M'_{2k+2} and M''_{2k+2} . (Note that these edges have length 0.)

Together with the paths P_1, \dots, P_{2k} , these added segments and edges define a connected 4-regular graph on

$$\{M'_1, \dots, M'_t\} \cup \{M''_1, \dots, M''_t\}.$$

An Eulerian traversal of this graph is a closed path that contains P_1, \dots, P_t and crosses S at most twice. Hence we have proved the theorem for $g = 6$. \square

Remarks. (i) A more careful argument using Christofides’ technique shows $g = 3$ suffices. (ii) A similar patching can be done for paths in \mathbb{R}^{d+1} whose endpoints lie inside a d -dimensional cube of side s . By a well-known upperbound on the length of the shortest tour (and hence also the length of the minimum matching) on k nodes in a cube (see Proposition A1 in the appendix), the patching cost is $O(k^{1-1/d}s)$.

The next lemma will be useful in the algorithm’s analysis. It uses a simple argument that often appears in geometric probability.

Let us grid the bounding box by putting a sequence of vertical and horizontal lines at unit distance from one another. The lemma relates the cost of a tour to the total number of times it crosses the grid lines. If l is one of these lines and π is a salesman tour, then let $t(\pi, l)$ denote the number of times that π crosses l .

LEMMA 4. *If the minimum internode distance is at least 4, and T is the length of π , then*

$$\sum_{l:\text{vertical}} t(\pi, l) + \sum_{l:\text{horizontal}} t(\pi, l) \leq 2T.$$

PROOF. The main observation is that the left hand side roughly measures the ℓ_1 length of the tour, which is at most $\sqrt{2}$ times T .

Specifically, we show that an edge of π that has length s contributes at most $2s$ to the left-hand side. Suppose u and v are the lengths of the horizontal and vertical projections of the edge; thus $u^2 + v^2 = s^2$. Then it contributes at most $(u + 1) + (v + 1)$ to the left-hand side, and

$$u + v + 2 \leq \sqrt{2(u^2 + v^2)} + 2 \leq \sqrt{2s^2} + 2.$$

Finally, since $s \geq 4$, we have $\sqrt{2s^2} + 2 \leq 2s$. \square

Now we are ready to prove the Structure Theorem.

PROOF (STRUCTURE THEOREM). Let $s = 12gc$, where g is the constant appearing in the Patching Lemma, and let $r = s + 4$, $m \geq 2s \log L$. Let π be the optimum salesman tour and suppose shift (a, b) is picked randomly. We prove the Structure Theorem by modifying π over many steps (using a deterministic procedure) into a salesman path that is (m, r) -light with respect to the randomly-shifted dissection. This may increase the tour cost slightly, which we upperbound (in the expectation) as follows. For accounting purposes, we place a grid of unit granularity in the bounding box. We “charge” any cost increase to some (horizontal or vertical) line of the grid. We will show that for each line l of the grid,

$$E_{a,b}[\text{charge to line } l \text{ when shift is } (a, b)] \leq \frac{3gt(\pi, l)}{s}, \tag{1}$$

where g is the constant appearing in the Patching Lemma. By linearity of expectations, it then follows that the expected increase in the cost of the tour is

$$\sum_{l:\text{vertical}} \frac{3gt(\pi, l)}{s} + \sum_{l:\text{horizontal}} \frac{3gt(\pi, l)}{s},$$

which is $\leq 6g \text{ OPT}/s$ by Lemma 4.

Since $s \geq 12gc$, the expected increase in the tour cost is at most $\text{OPT}/2c$. Markov’s inequality implies that with probability at least $1/2$ this increase is no more than OPT/c . We conclude that with probability at least $1/2$ the cost of the best (m, r) -light salesman path for the shifted dissection is at most $(1 + 1/c)\text{OPT}$.

Thus, to prove the theorem it suffices to describe how we modify the optimum tour π and charge resulting cost increases. Assume without loss of generality that the size of the bounding box L is a power of 2. Thus, all lines used in the dissection are grid lines. Recall how the dissection with shift (a, b) is obtained from the dissection. The middle vertical line of the dissection is moved from the x -coordinate $L/2$ to the x -coordinate $a + (L/2) \bmod L$, and the middle




FIG. 3. If a vertical line l is at level i , then for $p = 0, 1, 2, \dots, 2^i - 1$, the segment of the line between the y -coordinates $b + p \cdot (L/2^i) \bmod L$ and $b + (p + 1) \cdot (L/2^i) \bmod L$ is the edge of a level i square. It contains a portal at each end and m other equally spaced portals in between. An (m, s) -light salesman path must cross this segment at most s times, and always at one of the portals.

horizontal line from the y -coordinate $L/2$ to the y -coordinate $b + (L/2) \bmod L$. Then, the rest of the dissection is “wrapped-around,” so that the left edge of the dissection comes to rest at the x -coordinate a , and the lower edge of the dissection comes to rest at the y -coordinate b (see Figure 2). Note that since a, b are integers, the lines of the shifted dissection still lie on grid lines.

Recall that squares in a dissection form a hierarchy, and have a natural notion of “level” (the bounding box is at level 0, its four children are the squares at level 1, and so on). We say that a grid line l has level i in the shifted dissection if it contains the edge of some level i square. Note that the edge of a level i square gets subdivided to yield the edges of two level $i + 1$ squares, so a line that is at level i is also at level j for all $j > i$. For each $i \geq 1$ there are 2^i horizontal lines and 2^i vertical lines at level i . (To simplify notation, we do not include the boundaries of the original bounding box in this calculation, since they do not shift and are also not crossed by the tour.) The vertical lines have x -coordinates $a + p \cdot (L/2^i) \bmod L$, where $p = 0, \dots, 2^i - 1$ and the horizontal lines have y -coordinates $b + p \cdot (L/2^i) \bmod L$, where $p = 0, \dots, 2^i - 1$. The maximal level of a line is the highest level it is at.

Since the horizontal shift a is chosen randomly, we have for each vertical line l in the grid, and each $i \leq \log L$,

$$\Pr_a[l \text{ is at level } i] = \frac{2^i}{L}. \tag{2}$$

Of course, a similar statement is true for horizontal lines.

First, we try to make the salesman path (m, s) -light. We will almost succeed in this; however, at the end we will need to allow 4 more crossings on each square edge in the quadtree and thus finish with a (m, r) -light salesman path.

Recall what it means for the path to be (m, s) -light. First, for each vertical grid line l , if i is the maximal level of this line, then for $p = 0, 1, 2, \dots, 2^i - 1$, the segment of this line lying between the y -coordinates $b + p \cdot (L/2^i) \bmod L$ and $b + (p + 1) \cdot (L/2^i) \bmod L$ is crossed by the salesman path at most s times. Second, all these crossings happen at portals. (See Figure 3.) Of course, an analogous statement holds for all horizontal grid lines.

How can we modify the optimum tour to satisfy the first condition? An obvious idea suggests itself. Go through all lines l in the grid, and if its maximal level is i ,

then go through each of its 2^i segments of length $L/2^i$ and apply the Patching Lemma whenever one of them is crossed by the tour more than s times. This certainly works, but unfortunately the resulting cost increase may be too high (the reader may wish to check this). A better idea is to call the procedure $\text{MODIFY}(l, i, b)$, which does the patching “bottom up” for all levels $j \geq i$ and therefore is more miserly with cost increases. (We describe the procedure only for vertical lines; the description for horizontal lines is identical.)

$\text{MODIFY}(l, i, b)$

(l is a vertical grid line, b is the vertical shift of the dissection, and i is the maximal level of line l)

For $j = \log L$ down to i do:

For $p = 0, 1, \dots, 2^j - 1$, if the segment of l between the y -coordinates $(b + p \cdot (L/2^j) \bmod L)$ and $(b + (p + 1) \cdot (L/2^j) \bmod L)$ is crossed by the current salesman path more than s times, then use the Patching Lemma to reduce the number of crossings to 4.

Remarks on MODIFY: (i) The careful reader might note that we assume that the number of crossings after patching is 4, whereas the statement of the Patching Lemma seems to ensure this number is 2. The reason for the discrepancy is that the segment could be “wrapped-around”, and the patching has to be done separately for its two parts. (ii) The salesman path is updated iteratively, so the modifications made for the $j = k$ iteration of the loop are influenced by the modifications for $j = k + 1, \dots, \log L$. (iii) The patching on a vertical line adds to the cost of the salesman path and could increase the number of times the path crosses a horizontal line. We ignore this effect for now, and explain at the end of the proof that this costs us only an additional 4 crossings. (iv) The structure of the loop (i.e., the fact that j decreases from $\log L$ to i) is important and will be used in the argument that follows.

If $j \geq i$, let $c_{l,j}(b)$ denote the number of segments to which we apply the patching lemma in the iteration corresponding to j in the “for” loop in $\text{MODIFY}(l, i, b)$. Note that $c_{l,j}(b)$ is independent of i since the loop computation for j does not depend on i . We claim furthermore that

$$\sum_{j \geq i} c_{l,j}(b) \leq \frac{t(\pi, l)}{s - 3}. \tag{3}$$

The reason is that the optimum tour π crossed line l only $t(\pi, l)$ times, and each of the applications of the Patching Lemma counted on the left hand side replaces at least $s + 1$ crossings by at most 4.

Furthermore, the cost increase can be estimated using the Patching Lemma as follows:

$$\text{Increase in tour cost due to } \text{MODIFY}(l, i, b) \leq \sum_{j \geq i} c_{l,j}(b) \cdot g \cdot \frac{L}{2^j}, \tag{4}$$

where g is the constant appearing in the Patching Lemma. We charge this cost to l . Of course, whether or not this charge occurs depends on whether or not i is the

maximal level of line l , which by (2) happens with probability at most $2^i/L$ (over the choice of the horizontal shift a). Thus, for every vertical line l and every $0 \leq b \leq L - 1$,

$$\begin{aligned}
 & E_a[\text{charge to } l \text{ when horizontal shift is } a] \\
 &= \sum_{i \geq 1} \frac{2^i}{L} \cdot \text{cost increase due to } \text{MODIFY}(l, i, b) \\
 &\leq \sum_{i \geq 1} \frac{2^i}{L} \cdot \sum_{j \geq i} c_{l,j}(b) \cdot g \cdot \frac{L}{2^j} \\
 &= g \cdot \sum_{j \geq 1} \frac{c_{l,j}(b)}{2^j} \cdot \sum_{i \leq j} 2^i \\
 &\leq g \cdot \sum_{j \geq 1} 2 \cdot c_{l,j}(b) \\
 &\leq \frac{2gt(\pi, l)}{s - 3}.
 \end{aligned}$$

The next modification to the salesman path consists in moving each crossing to the nearest portal. If a line l has maximal level i , then each of the $t(\pi, l)$ crossings might have to be displaced by $L/2^{i+1}m$ – i.e., half the interportal distance – to reach the nearest portal. Instead of actually deflecting the edge, we break it at the point where it crosses l and add to it two line segments (one on each side of l) of length at most $L/2^i m$, so that the new edge crosses l at a portal. Thus, the expected increase in the length of the salesman path when we move every crossing in l to its nearest portal is at most

$$\sum_{i=1}^{\log L} \frac{2^i}{L} \cdot t(\pi, l) \cdot \frac{L}{2^i m} = \frac{t(\pi, l) \log L}{m},$$

which is at most $t(\pi, l)/2s$ when $m \geq 2s \log L$.

Thus, the expected cost (over the choice of the random shift a) of making the final salesman path (m, r) -light at line l is

$$\frac{2gt(\pi, l)}{s - 3} + \frac{t(\pi, l)}{2s} \leq \frac{3gt(\pi, l)}{s}.$$

(Where the last calculation assumes $s > 15$.) This is what we set out to prove.

To finish our proof it only remains to explain our remark (iii) on the `MODIFY` procedure. This concerned the following situation. Whenever we apply `MODIFY` on a vertical line l , we use the Patching Lemma and augment the salesman path with some segments lying on l . These segments could cause the path to cross some horizontal line l' much more than the $t(\pi, l')$ times it was crossing earlier. However, our analysis assumed that the number of crossings remains constant at $t(\pi, l')$ throughout the modification, and this requires an explanation. The explanation is simple: we can show that without loss of generality the increase in

the number of crossings at l' due to the patching on l is at most 2. The reason is that if the increase were more than 2, we could just use the Patching Lemma to reduce it to 2. Furthermore, since the Patching Lemma is being invoked for segments lying on l , these have zero horizontal separation (that is, they lie on top of each other) and therefore the tour cost does not increase! Also, we apply the patching separately on both sides of l , so the number of crossings on l does not change. Arguing similarly about all pairs of grid lines, we can ensure that at the end of all our modifications, each side of each square in the shifted dissection is crossed by the modified tour up to $s + 4$ times. So, as promised at the start of the proof, we ended up with an (m, r) -light salesman path.

A similar accounting trick explains why we assumed that the $t(\pi, l)$'s are not affected when we move each edge-crossing to its nearest portal. \square

2.3. ALGORITHM FOR EUCLIDEAN TSP IN \mathbb{R}^d . Suppose d is a constant independent of n . Our methodology for \mathbb{R}^d mimics that for \mathbb{R}^2 . As before, we start with a rescaling and perturbation to make the instance well-rounded. If L is the length of the smallest axis-aligned cube enclosing the instance (implying $OPT \geq L$), we place a grid of granularity $L/8cn\sqrt{d}$ in \mathbb{R}^d and move each node to its nearest gridpoint. This changes the cost of the optimal tour by at most $2n \cdot \sqrt{d}L/8cn\sqrt{d}$, which is at most $L/4c$. Then we rescale distances by $L/64cn\sqrt{d}$, so that the minimum nonzero internode distance is 8 units and the size of the bounding cube is $O(c\sqrt{d}n)$. Since c, d are constant, we think of the size of the cube as $O(n)$. Thus, the instance is now well-rounded. (Note that this procedure just requires d rounds of sorting, and thus takes $O(nd \log n)$ time.)

The “quad” trees in \mathbb{R}^d are 2^d -ary trees. The shifted 2^d -ary tree is defined in the obvious way. Note that a region of this tree is a cube (possibly “wrapped around”) in \mathbb{R}^d , so the boundaries of the region are cubes of dimension $d - 1$. An m -regular set of portals on a dimension $d - 1$ cube is an orthogonal lattice of m points in the cube. Thus, if the cube has length W , then the spacing between the portals is $W/m^{1/(d-1)}$.

THEOREM 5 (STRUCTURE THEOREM FOR EUCLIDEAN TSP IN \mathbb{R}^d). *Let $c > 0$ be any constant. Let the minimum (nonzero) internode distance in an instance be \mathbb{R}^d and let L be the size of the bounding box. Let $0 \leq a_1, \dots, a_d < L$ be picked randomly. Then with probability at least $1/2$ the dissection with shift (a_1, \dots, a_d) has an associated $(1 + 1/c)$ -approximate salesman path that is (m, r) -light, where $m = (O(\sqrt{dc} \log L))^{d-1}$, $r = (O(\sqrt{dc}))^{d-1}$.*

The dynamic programming algorithm for Euclidean TSP in \mathbb{R}^d is basically the same as for \mathbb{R}^2 . The 2^d -ary tree has $O(2^d n \log n)$ regions, and each region has $2d$ facets. Hence, the amount of time needed to “guess” the portals at which the salesman path enters and leaves a region is $m^{O(2dr)}$. Hence, the total running time is $O(2^d m^{O(2dr)} n \log n)$, which is $O(n(\log n)^{O(\sqrt{dc})^{d-1}})$.

The proof of Theorem 5 is analogous to that of Theorem 2, and we only list the calculations that need to be modified.

Patching Lemma. If a salesman path crosses a dimension $d - 1$ cube k times (where we assume that the cube contains no nodes itself), then the number of crossings can be reduced to at most 2, by augmenting the salesman path with segments in the cube whose total length is $O(k^{1-(1/(d-1))}W)$, where W is the length of the cube’s side.

The proof of the lemma uses remark (ii) following Lemma 3.

Modification to Lemma 4. The grid “lines” are $d - 1$ dimensional hyperplanes in \mathbb{R}^d . The Lemma now says: If the minimum internode distance is at least 4, then the sum of the number of times a salesman tour crosses the hyperplanes of the unit grid is at most $5\sqrt{d}/4$ times the cost of the tour.

The proof of the modified lemma is the same as in the planar case, except for the fact that \sqrt{d} is the only available upperbound on the ratio of the ℓ_1 and ℓ_2 costs of the tour.

Proof of Structure Theorem. As before, we start with an optimal salesman tour and use the Patching Lemma repeatedly to make it (m, r) -light. We charge resulting cost increases to grid hyperplanes and show below that the expected charge to hyperplane h is at most $t(\pi, h)/2c\sqrt{d}$, where $t(\pi, h)$ denotes the number of times the optimum tour π crosses hyperplane h . Then the theorem follows by linearity of expectations and the modified Lemma 4.

The *maximal level* of a grid hyperplane h is the level of the largest subcube of the shifted 2^d -ary tree that has its boundary on h . Let $a_1, a_2, \dots, a_d \in [0, L - 1]$ be the random shifts used to construct the shifted 2^d -ary tree. In the discussion below, let a grid hyperplane h be perpendicular to the j th coordinate axis. Then, the maximal level of h only depends upon a_j , and

$$\Pr_{a_j}[\text{maximal level of } h \text{ is } i] = \frac{2^i}{L}. \tag{5}$$

Note that if the maximal level of h is i , then 2^{id} level- i cubes (each of side $L/2^i$) have their boundaries on h . An (m, r) -light salesman path can cross each side of these cubes at most r times, and always through a portal. We use the Patching Lemma to modify the optimum tour and make it satisfy this condition. Furthermore we do this modification in a bottom-up fashion: we invoke a MODIFY procedure that ensures for $j = \log L, \log L - 1, \dots, i$, that the tour crosses the boundaries of each level j cube adjacent to h at most r times.

Let $c_j(h)$ (or just c_j since h is clear from the context) be the number of level j cubes for which the MODIFY procedure invokes the Patching Lemma. For $k \leq c_j$, suppose the k th invocation of the Lemma involved replacing $t_{jk} \geq r + 1$ crossings by at most $2 \cdot 2^d$ crossings (where the reason for the $2 \cdot 2^d$ is that the cube may be “wrapped around” and thus actually be a union of up to 2^d smaller regions, for all of which the Patching Lemma may have to be invoked separately). Thus, we have

$$\sum_{j=1}^{\log L} \sum_{k=1}^{c_j} (t_{jk} - 2^{d+1}) \leq t(\pi, h). \tag{6}$$

Furthermore, the cost of the MODIFY procedure if the maximal level of h is i is

$$\sum_{j=1}^{\log L} \frac{L}{2^j} \sum_{k=1}^{c_j} g \cdot t_{jk}^{1-(1/(d-1))},$$

where g is the constant appearing in the Patching Lemma for \mathcal{R}^{d-1} . By (5), the probability that the maximal level of h is i is at most $2^i/L$, so we may upperbound the expected cost increase at hyperplane h by

$$\sum_{i=1}^{\log L} \frac{2^i}{L} \sum_{j=1}^{\log L} \frac{L}{2^j} \sum_{k=1}^{c_j} g \cdot t_{jk}^{1-(1/(d-1))}.$$

Upon rearranging the order of summation, this becomes

$$\sum_{j=1}^{\log L} \frac{1}{2^j} \sum_{k=1}^{c_j} g \cdot t_{jk}^{1-(1/(d-1))} \sum_{i=1}^j 2^i.$$

Since $\sum_{i=1}^j 2^i = 2^{j+1} - 1$, this is upperbounded by

$$2g \cdot \sum_{j=1}^{\log L} \sum_{k=1}^{c_j} t_{jk}^{1-(1/(d-1))}. \tag{7}$$

Since the t_{jk} 's satisfy (6) and are $\geq r + 1$, this cost is maximized when each $t_{jk} = r + 1$. In this case, (6) simplifies to

$$\sum_{j=1}^{\log L} c_j \leq \frac{t(\pi, h)}{r + 1 - 2^{d+1}}.$$

Similarly, the upperbound on the expected cost in (7) simplifies to

$$2g \cdot (r + 1)^{1-(1/(d-1))} \sum_{j=1}^{\log L} c_j,$$

which is at most

$$2g(r + 1)^{1-(1/(d-1))} \cdot \frac{t(\pi, h)}{r + 1 - 2^{d+1}}.$$

When $r = (O(\sqrt{dc}))^{d-1}$, this upperbound is at most $t(\pi, h)/4c\sqrt{d}$.

The analysis of the cost of moving all crossings to their nearest portal is essentially unchanged from the planar case. Namely, on a grid hyperplane h , at most $t(\pi, h)$ crossings will ever need to be moved to their nearest portal. If the maximal level of the hyperplane is i , then each crossing needs to move at most a distance $\sqrt{d}L/2^i m^{1-(1/(d-1))}$ (since the portals on the boundary of a level i cube form a grid of granularity $L/2^i m^{1-(1/(d-1))}$). The probability that i is the maximal level of h is at most $2^i/L$, so an upperbound on the cost of moving all crossings on h to their nearest portals is

$$\sum_{i=1}^{\log L} \frac{2^i}{L} \cdot t(\pi, h) \cdot \frac{\sqrt{d}L}{2^i m^{(1/(d-1))}}.$$

When $m = (O(\sqrt{dc} \log L))^{d-1}$, this cost is at most $t(\pi, h)/4c\sqrt{d}$.

2.4. TSP IN OTHER GEOMETRIC NORMS. A Minkowski norm in \mathcal{R}^d is defined using a convex body C that is symmetric around the origin. The length of $x \in \mathcal{R}^d$ under this norm is defined to be $|x|_2/|y|_2$, where y is the intersection of the surface of C with the line connecting x to the origin. It is easy to see that this definition generalizes the ℓ_p norm for $p \geq 1$ ($C =$ the ℓ_p -unit ball centered at the origin).

Our algorithm generalizes to the case where distances are measured using any fixed Minkowski norm. The only change is that the value of “ m ” and “ r ” changes by some constant factor depending upon the norm. The reason for this change is that distances under this norm are within some constant factor (depending upon the convex body) of the distance under the ℓ_2 norm. Hence, the proofs of Lemma 4 and the Patching Lemma are still valid (except the constant “2” in Lemma 4 has to be replaced by some other constant and the constant “ g ” in the Patching Lemma may also be different) and so the proof of the Structure Theorem also goes through after we multiply the old values of m , r by appropriate constants.

2.5. PARALLEL IMPLEMENTATIONS. The algorithm presented in the sections above has an NC implementation (i.e., an implementation that runs in $\text{poly}(\log n)$ time using $\text{poly}(n)$ processors; we do not attempt to investigate the exact complexity). This is obvious for Steps 1 and 2, since they can be implemented using sorting. Step 3, consisting of dynamic programming, can also be implemented efficiently in parallel because the depth of the dynamic programming is $O(\log n)$. In fact, the steps involving enumeration of all possible portal combinations seems quite ideal for a multiprocessor.

3. Approximation Schemes for Other Geometric Problems

To obtain approximation schemes for the geometric problems mentioned earlier, we imitate our TSP methodology. First, we rescale and perturb the instance to make it “well rounded.” A well-rounded instance is one in which all coordinates are integral, the minimum nonzero internode distance is 8, and the size of the bounding square is $O(n^3)$. The perturbation procedure for the different problems are similar to the one for the TSP, although some care is occasionally needed (see below). For all the problems we prove an analogue of the Structure Theorem, whose statement is essentially the same as that of Theorem 2 except the words “salesman path” have to be replaced by problem-specific words such as “Steiner Tree.” The theorem is proved just as in the TSP case using a charging argument. The dynamic programming algorithm to find the optimum (m, r) -light solution is straightforward.

We emphasize that the main reason that the TSP methodology can be applied to these other problems is that the Patching Lemma and Lemma 4 generalize to these problems.

LEMMA 6 (GENERAL PATCHING LEMMA). *For each of the problems Min-Steiner Tree, k -TSP, k -MST and Min-Cost-Perfect-Matching, there is a constant $g > 0$ such that the following is true. Suppose π is any solution and S is any line segment of length s that π crosses least thrice. Then all but at most 2 of these crossings can be broken and instead some line segments on S can be added to the solution so that the new solution has cost at most $g \cdot s$ more than the old solution.*

Note. As in the TSP case, we are allowing solutions to contain “bent” edges. When we augment the original solution π with line segments on S , some of the new edges are “bent” edges. Solutions with bent edges pose no problems for our dynamic programming, and at the end we can straighten bent edges without increasing the cost.

PROOF. For k -TSP, the proof is the same as for the TSP. For k -MST and Min-Steiner-Tree, the proof is also trivial: we just break the tree at all but one of the points where it crosses S , and add straight line segments along S between all connected components of the tree. This makes the portion of the tree on each side of S connected, and the length of the added segments is clearly $\leq 2s$. The same method works for Min-Cost-Perfect-Matching. \square

LEMMA 7. *Let S be a collection of line segments in \mathbb{R}^2 each of length ≥ 4 . For a line l in the unit grid let $t(S, l)$ be the number of times these segments cross l . Then*

$$\sum_{l:\text{vertical}} t(S, l) + \sum_{l:\text{horizontal}} t(S, l) \leq 2 \cdot \text{cost}(S).$$

PROOF. As in the proof of Lemma 3, we just use the observation that the left hand side roughly measures the ℓ_1 cost of segments in S . \square

How to Make the Instance Well Rounded. The procedure to make TSP instances well-rounded does not generalize to k -TSP, k -MST and Euclidean matching. Instead, we outline a more general procedure. (The problem-specific details are given later.)

The TSP procedure does not work for these problems because it relies on the fact that $OPT \geq L$ for the TSP, where L is the size of the bounding box. This may not hold for our problems (since $OPT \ll L$ is possible).

Let us recall why $OPT \geq L$ was needed in the TSP case. If d_{\min} denotes the minimum nonzero internode distance, then the depth of the quadtree is $O(\log(L/d_{\min}))$. A priori this depth could be large but by coalescing nodes that are “too close,” we ensured $d_{\min} \geq L/8cn$. Hence, the quadtree has depth $O(\log n)$, which is crucial for efficiency. Furthermore, the coalescing affects the cost of the optimal salesman tour by at most $2n \cdot L/8cn$, which is at most $OPT/4c$ since $OPT \geq L$.

The rounding procedure for our problems first computes in nearly linear time a crude approximation to OPT that is correct within a factor n . (Algorithms to compute such approximations were known before our work.) Let $A \leq OPT \cdot n$ be this approximation. The procedure lays a grid of granularity $A/8cn^2$ in the plane and moves every node to its nearest gridpoint. This ensures $d_{\min} \geq A/8cn^2$ in the new instance, and triangle inequality implies that the change in the cost of the optimum solution is at most $2n \cdot A/8cn^2 \leq OPT/4c$.

If $A \geq L/n^2$, say, then we’re done because the depth of the quadtree is $O(\log(L/d_{\min})) = O(\log n)$ in that case. Now suppose $A < L/n^2$. We pick random shifts as before and construct a shifted quadtree (this can be done in $O(n \log n)$ time [Bern et al. 1993]). Then we treat all squares of size more than $A \log n$ in the shifted quadtree as independent instances of the problem. (In case of EMCPM, this means finding a matching within each square; the final matching is the union of such matchings.) Note that each of these squares has a quadtree of depth $O(\log 8cn^2) = O(\log n)$, where n is the original number of nodes.

In order to justify the correctness of the above algorithm, we show that if $A \leq L/n^2$ then with probability at least $1 - 2/\log n$ (over the choice of the shifts), no edge of the optimum solution crosses the boundary of any square of size $A \log n$ in the shifted quadtree. (In case of EMCPM, this implies in particular that with probability at least $1 - 2/\log n$, every square of size $A \log n$ in the quadtree has an even number of nodes.) We recall from Lemma 7 that at most $2A$ lines of the unit grid are crossed by edges of the optimum solution; let's call these the good lines. For any fixed grid line, the probability that the random shift causes it to become the boundary of a quadtree square of length $A \log n$ is at most $1/A \log n$. (To see that this is true, note that a quadtree square of length $A \log n$ contains $A \log n$ horizontal and vertical grid lines.) Hence the probability that any of the $2A$ good grid lines is on the boundary of a square of size $> A \log n$ in the shifted quadtree is at most $2/\log n$. Hence, with probability $> 1 - 2/\log n$, no square of size more than $A \log n$ will be crossed by any edge of the optimum solution, in which case our decision to treat those squares as independent instances is justified!

This finishes our outline of the perturbation procedure. We note that the use of randomness in the previous step can be removed by trying all possible shifts between 1 and $A \log n/d_{\min}$, which multiplies the running time by a factor $O(n^6)$ (by using a better quality approximation at the start, this factor can be reduced to $O(n^2 \log^2 n)$ in case of Min Cost Matching and to $O(k^3)$ in case of k -TSP and k -MST).

3.1. STEINER TREE. A *well-rounded* instance of the Steiner Tree problem is one in which all nodes and Steiner nodes have integer coordinates, the minimum (nonzero) internode distance is at least 8 and the smallest square bounding the instance has size at most $O(n)$.

In order to make the instance well-rounded, we use the fact that a Steiner node has degree ≥ 3 (this is also true in \mathcal{R}^d and also if distances are measured in any fixed geometric norm), so an optimum solution on n nodes has at most n Steiner nodes.

Let L_0 be the size of the bounding box of the instance, so $OPT \geq L_0$. We place a grid of granularity $L_0/16nc$ and move each node to its nearest grid point. We thereafter require all Steiner nodes also to lie on the grid. Since the optimum solution contains only $2n$ nodes (including Steiner nodes) and each moved by a distance at most $L_0/16nc$, the cost of the optimal tree changed by at most $L_0/4c$, which is at most $OPT/4c$. Rescaling distances by $L_0/128nc$, the instance becomes well rounded.

Note that by requiring the Steiner nodes to lie on the grid, we have forestalled any precision issues which would otherwise have arisen when the algorithm “guesses” the location of a Steiner node.

Recall that in the TSP case we had introduced portals in the quadtree and (somewhat unnaturally) forced the salesman path to cross region boundaries only at portals. In the Steiner Tree problem portals can be viewed more naturally: they are “Steiner” nodes. We define an (m, r) -light Steiner tree by analogy to an (m, r) -light salesman path.

THEOREM 8 (STRUCTURE THEOREM FOR STEINER TREE). *Let $c \geq 0$ be any constant. The following is true for every instance of Steiner tree in which the minimum (nonzero) distance between any pair of nodes (including Steiner nodes) is*

8 and the size of the bounding box is L . Let shifts $0 \leq a, b < L$ be picked randomly. Then with probability at least $1/2$, the dissection with shift (a, b) has an associated Steiner tree of cost at most $(1 + 1/c)$ -OPT that is (m, r) -light, where $m = O(c \log L)$ and $r = O(c)$.

PROOF. The proof is essentially the same as that of Theorem 2. We start with an optimal Steiner tree and refine it (using a procedure analogous to MODIFY) until it becomes (m, r) -light with respect to the shifted dissection. The cost increase is estimated using our old charging technique. \square

The dynamic programming algorithm to compute the optimal (m, r) -light Steiner tree is similar to that for the TSP, except for two changes. First, the basic subproblem in the dynamic programming is the (m, r) -Steiner forest problem. An instance of this problem is specified by (a) a square of the quadtree, (b) a multiset B containing $\leq r$ portals on each of the four sides of the square, and (c) a partition (B_1, B_2, \dots, B_p) of B . (In the multipath problem, the instance was specified by a pairing of the portals.) The goal is to find an optimum (m, r) -light collection of p Steiner trees, where the i th tree contains all the portals in B_i , and the trees together contain every node inside the square.

The other change from the TSP case is that a square that does not contain any input nodes may nevertheless contain Steiner nodes, which have to be “guessed” by the algorithm. However, since the square is entered and left only $4r$ times, we can treat it as an instance of the Steiner Tree problem of size at most $4r$, which can be solved in constant time since r is constant [Melzak 1961]. Note that this also implies that the number of leaves in the shifted quadtree is $O(n)$, since a leaf must contain either an input node or a Steiner node. Thus the size of the shifted quadtree is $O(n \log n)$ and the running time of the dynamic programming is similar to that in the TSP case.

Finally, we observe that the proof of Theorem 8 actually proves something about the structure of a near-optimal Steiner tree that may be of independent interest.

COROLLARY 9. *In the $(1 + 1/c)$ -approximate Steiner tree whose existence is guaranteed by Theorem 8, every Steiner node used in the tree is either a portal of the shifted dissection or contained in a square at the leaf of the dissection.*

PROOF. In the proof of Theorem 8, as the MODIFY procedure does a cut-and-patch operation on the optimum Steiner tree, the tree edges get “bent” to make them pass through portals. Thus, those portals turn into Steiner nodes. But no other Steiner nodes are introduced. The only other Steiner nodes in the final tree lie in squares at the leaves of the dissection. \square

3.1.1. *Steiner Tree Problem in \mathcal{R}^d and with non-Euclidean Norms.* The description of the approximation scheme for Minimum Steiner Tree in \mathcal{R}^d can also be derived easily from the corresponding TSP algorithm. The same is true for the case of Minkowski norms.

3.2. k -TSP AND k -MST. The description of the algorithm for these two problems is almost the same as for TSP and Steiner tree respectively. We define (m, r) -lightness for k -TSP and k -MST by analogy with the TSP. The following is the Structure Theorem for them.

THEOREM 10. *Let $c > 0$ be any constant. The following is true for every instance of k -TSP (respectively, k -MST) in which the minimum nonzero internode distance is ≥ 8 and the size of the bounding box is L . Let shifts $0 \leq a, b < L$ be picked randomly. Then with probability at least $1/2$, the dissection with shift (a, b) has an associated (m, r) -light k -salesman path (respectively, k -minimum spanning tree with bent edges) that has cost $(1 + 1/c)OPT$, where $m = O(c \log L)$ and $r = O(c)$.*

Now we describe the k -TSP algorithm (the k -MST algorithm is analogous). A well-rounded instance for this problem is one in which all coordinates are integral, the minimum nonzero internode distance is 8, and the size of the bounding square is $O(k^2)$. We describe a simple $O(nk \log n)$ time perturbation that partitions the given instance into well-rounded instances. This relies on the observation (see Proposition 12 in the Appendix) that the cost of the optimum solution to k -TSP is at most $2\sqrt{k}$ times larger than the size of the smallest square containing at least k points (in \mathcal{R}^d it is at most $dk^{1-(1/d)}$ times the size of such a cube). Eppstein and Erickson [1994] show how to approximate the size of this square within a factor 2 by computing for each of the n nodes its k nearest neighbors. This takes $O(nk \log n)$ time.⁵

Once we have an approximation A such that $OPT \leq A \leq 2\sqrt{k} OPT$, we place a grid of granularity $A/16ck^{3/2}$ and move each node to its nearest gridpoint. By the triangle inequality, the cost of the optimum changed by at most $2k \cdot A/16ck^{3/2} \leq OPT/4c$. Then we pick random shifts in $[0, L)$ and construct a quadtree with these shifts. From then on, we treat all squares of size $A \log k$ in the quadtree as independent instances of the problem (i.e., in each of these squares we run a k -TSP algorithm). The general argument at the start of Section 3 shows that with probability $> 1 - 2/\log k$, no edge of the optimum solution crosses the boundaries of any of these squares, which means that the optimum lies entirely inside one of the squares. In each square, the ratio of the size of the square to the minimum internode distance is $\log k \cdot 16ck^{3/2} = O(k^2)$, so the instance inside the square is well-rounded. For each square, we construct a quadtree with random shifts (these shifts are picked randomly from 0 to L_1 , where L_1 is the size of the square) and use dynamic programming to find the optimum (m, r) -light k -salesman path for the points inside the square, where $m = O(c \log k)$ and $r = O(c)$. Having done this for each square of size $A \log k$, we output the lowest cost k -salesman path found.

The running time is analyzed as follows: For a well-rounded instance with n_1 nodes, let $S(n_1, k)$ and $T(n_1, k)$ be, respectively, the size of the lookup table and the running time of the dynamic programming. We will prove the following Claim below.

CLAIM. $S(n_1, k) = O(n_1 m^{O(r)} \log k)$ and $T(n_1, k) = O(n_1 k m^{O(r)} \log k)$.

The claim implies that the contribution of each node to the running time is $O(km^{O(r)} \cdot \log k)$, hence the overall running time (i.e., the time required to do the dynamic programming for all the independent instances we identified above) is $O(nkm^{O(r)}) = O(nk(\log k)^{O(c)})$, as claimed.

PROOF (OF CLAIM). The dynamic programming for k -TSP differs from the corresponding TSP algorithm in only the definition of the subproblem that is

⁵ Eppstein and Erickson have also described a $npoly(\log n)$ time algorithm to us.

solved for each region of the shifted quadtree. In the TSP case, this subproblem was the (m, r) -multipath problem, which involves finding a set of disjoint paths that visit all the nodes inside the given square. In the k -TSP case, the set of paths have to visit only k_1 nodes inside the given square, where $k_1 \leq k$. Since k_1 is not a priori known to the algorithm, it has to try all possible values for it. This is the new twist.

In a well-rounded instance, the size of the bounding box is $O(k^2)$, so a quadtree has $O(\log k)$ levels. If i is a level and $j \leq n$, then let t_{ij} be the number of squares at level i of the quadtree that contain exactly j input nodes. Note that every node is in exactly one square at level i . Since the total number of nodes is n_1 , we have for all i ,

$$\sum_{j \leq n} j \cdot t_{ij} = n_1. \tag{8}$$

Suppose a square has j input nodes and it contains $k_1 \leq \min\{k, j\}$ nodes from the optimum (m, r) -light solution. Not knowing k_1 , the algorithm tries all values for it. For each value of k_1 , the lookup table contains $m^{O(r)}$ subproblems (obtained by trying all possible ways in which the optimum (m, r) -light solution could enter/leave the square). Thus the table needs space for $m^{O(r)} \cdot \min\{k, j\}$ subproblems for this square. The table's size is therefore

$$S(n_1, k) = \sum_{i \leq O(\log k)} \sum_{j \leq n} \min\{k, j\} \cdot m^{O(r)} t_{ij} \leq m^{O(r)} \cdot \sum_i \sum_j j t_{ij},$$

which is $O(m^{O(r)} n_1 \log k)$ by (8).

Now we analyze the running time. First, we modify the dynamic programming a little. Namely, instead of storing solutions to subproblems corresponding only to squares, we also store solutions to subproblems corresponding to halves of squares. We will refer to these halves as *rectangles*. (Viewed in terms of the quadtree, our modification involves introducing additional nodes to reduce the degree of the tree from 4 to 2.) This modification only doubles the size of the lookup table.

Now suppose a square/rectangle has j input nodes. The lookup table contains $m^{O(r)} \cdot \min\{k, j\}$ subproblems for it, where each subproblem involves an input $k_1 \leq \min\{k, j\}$ that is the algorithm's "guess" for the number of nodes from the optimum (m, r) -light solution that lie in the square/rectangle. When the algorithm partitions the square/rectangle into two equal regions, it enumerates as before all the $m^{O(r)}$ choices for the number of ways in which the salesman path could cross between its two halves. In addition, it tries $k_1 + 1$ guesses for the number of nodes from the optimum solution in the two halves (these correspond to the number of ways of expressing the integer k_1 as a sum of two smaller integers). Thus, the running time for solving all the subproblems corresponding to such a square/rectangle is $m^{O(r)} \cdot \min\{k, j\}^2$, and the overall running time is

$$T(n_1, k) = \sum_{i \leq O(\log k)} \sum_{j \leq n} t_{ij} \cdot m^{O(r)} \cdot (\min\{k, j\})^2,$$

which is maximized when the only nonzero t_{ij} 's correspond to $j = k$. Thus, the running time is at most $O(n_1 k m^{O(r)} \log k)$. \square

The description of the algorithm in \mathcal{R}^d and for non-Euclidean norms is also straightforward.

3.3. MIN COST PERFECT MATCHING. We state the Structure Theorem for this problem.

THEOREM 11. *Let $c > 0$ be any constant. The following is true for every instance of EMCPM in which the minimum nonzero internode distance is ≥ 8 and the size of the bounding box is L . Let shifts $0 \leq a, b < L$ be picked randomly. Then with probability at least $1/2$, the dissection with shift (a, b) has an associated (m, r) -light perfect matching that has cost $(1 + 1/c)OPT$, where $m = O(c \log L)$ and $r = O(c)$.*

A well-rounded instance of the Euclidean Min-Cost Perfect Matching problem is one in which all coordinates are integral, the minimum nonzero distance is 8, and the size of the bounding square is $O(n^2)$. We already outlined a procedure that makes the instance well-rounded. Now we describe it in more detail. The procedure uses a crude approximation algorithm due to Vaidya [1989] that runs in $O(n \log^3 n)$ time and computes an $O(\log n)$ approximation to the cost of the optimum matching.⁶

Let $A \leq O(OPT \log n)$ be the cost of this solution. We place a grid of granularity $A/n \log^2 n$ and move each node to its nearest gridpoint. If a gridpoint receives an even number of nodes, we remove them; if it receives an odd number, we replace them by one node. (At the end of the algorithm we will put a trivial matching among each set of removed nodes, which will add at most $O(n \cdot A/n \log^2 n) \leq OPT/4c$ to the cost.) Using triangle inequality it is easy to see that the cost of the min-cost perfect matching in the new instance is at most $OPT + O(n \cdot A/cn \log n) \leq OPT(1 + 1/2c)$. Clearly, $d_{\min} \geq A/n \log^2 n$ in the new instance.

Finally, we construct a quadtree with random shifts and treat all squares of size more than $A \log n$ as independent instances. Our general argument at the start of Section 3 shows that with probability $> 1 - 2/\log n$, no edge of the optimum solution crosses the boundaries of any of these squares, which means that the optimum perfect matching is a disjoint union of the optimum matchings inside the various squares. Hence, we run our approximation algorithm independently in the squares, and output the union of those solutions. Note that in each square the ratio of the size of the boundary to the minimum internode distance is $O(A \log n / (A/n \log^2 n))$, which is $O(n \log^3 n)$. Hence, in the square, the depth of the quadtree is $O(\log(n \log^2 n / \log n)) = O(\log n)$. So we can run our usual dynamic programming in the square to find the optimum (m, r) -light solution in $O(n_1 \cdot m^{O(r)} \cdot \log n)$ time, where n_1 is the number of nodes in the square and $m = O(c \log n)$ and $r = O(c)$. Hence, the overall running time is $O(n \cdot (\log n)^{O(c)})$.

The proof of the Structure Theorem for EMCPM mimics the one for the TSP.

4. Conclusions

Our original goal was to extend the inapproximability results in Arora et al. [1992] to Euclidean TSP. Doing this seemed to call for a reduction that produces

⁶ Alternatively, one may use a simple greedy procedure based upon iterative nearest neighbor searching. It computes a factor n approximation to EMCPM, which is good enough for our purposes.

instances in which every near-optimal tour encodes the complicated code-like objects of Arora et al. [1992]. Then, we realized (upon discovering an early version of Theorem 2) that such an encoding may be impossible. This realization led to an approximation algorithm in a natural way.

Our techniques apply to all Euclidean optimization problems in which the objective function is a sum of edge lengths (so that Lemma 7 can be applied) and which satisfy some version of the Patching Lemma (Lemma 3). We don't have any formal characterization of this class of problems. Note furthermore that recent extensions of our techniques use only the charging argument and not the Patching Lemma.

Two major open problems now remain in connection with TSP: (i) Design an approximation algorithm better than Christofides' for general metric spaces. Note that the results of Arora et al. [1992] mentioned in the introduction seem to preclude a PTAS for general metrics. Recently Arora et al. [1998] extended our ideas and those of Grigni et al. [1995] to design a PTAS for any metric that is the shortest-path metric of a weighted planar graph (the paper [Grigni et al. 1995] designed such PTASs for unweighted planar graphs). (ii) Design a more efficient version of our algorithm. Now we discuss this in more detail.

Our algorithm's running time depends exponentially on c even in \mathcal{R}^2 . Since Euclidean TSP is strongly NP-hard (this follows from the reductions in Papadimitriou [1977], for example), any algorithm that computes $(1 + c)$ -approximations must have a running time with some exponential dependence on c (unless, of course, NP-complete problems can be solved in subexponential time). Likewise, Trevisan [1997] has presented evidence that every approximation scheme for \mathcal{R}^d must have a running time with a doubly exponential dependence on d . He shows for some fixed $c > 1$ that finding a $(1 + 1/c)$ -approximate solution in $\mathcal{R}^{O(\log n)}$ is NP-hard.

Trevisan's work does not preclude the possibility of a running time such as $O(n \log^2 n + 2^{O(c)})$ in \mathcal{R}^2 . Recently Rao and Smith [1998] improved upon our ideas to achieve a running time of $(c\sqrt{d})^{O(d(\sqrt{dc})^{d-1})}n + O(d \log n)$ in \mathcal{R}^d ; thus their running time in \mathcal{R}^2 is $O(c^{O(c)}n + n \log n)$.

Finally, it would be interesting to extend our techniques to other Euclidean problems. Asano et al. [1997] have designed approximation schemes for *Vehicle Routing* problems with certain restrictions on the number of vehicles; they use our Structure Theorem. Arora et al. [1998] have extended our techniques to design approximation schemes for Euclidean k -median and facility location. The Patching Lemma does not hold for these problems, and they use only (a modification of) the charging argument. This may be an indication that the charging argument (together with the idea of a randomly shifted quadtree) is a more general tool than the Patching Lemma. We suspect that other geometric problems including *Minimum Weight Steiner Triangulation* may be amenable to these techniques.

Appendix A

The following well-known fact was used several times in the paper, notably in the proof of the various Patching Lemmas. We outline a proof.

PROPOSITION A1. *The length of the optimum salesman tour on k nodes in \mathcal{R}^d that lie inside a rectangle of size L is at most $O(k^{1-1/d}L)$. This is also true when*

distance is measured using any Minkowski norm (except the constant in the $O(\cdot)$ depends on the norm).

PROOF. We give a proof for the Euclidean case in \mathcal{R}^2 ; the proof for general d uses a simple induction (see Chapter 6 of Lawler et al. [1985]).

Partition the rectangle into horizontal “strips” of width L/\sqrt{k} . Then find a “strip tour,” which first visits all the nodes in the first strip, then all the nodes in the second strip, and so on. If k_i is the number of nodes in the i th strip, then the cost of the strip tour is at most

$$L + \sum_i \left(L + k_i \times O\left(\frac{L}{\sqrt{k}}\right) \right) = L + k \times O\left(\frac{L}{\sqrt{k}}\right) + \sqrt{k}L,$$

where the extra terms of L appear because of the need to traverse the strip to reach the starting point in the next strip. The above expression is $O(\sqrt{k}L)$. \square

A1. How Practical is our PTAS?

Many years of research and rapid increase in computer power have led to approaches—the best of which currently is branch-and-cut [Applegate 1995]—that in less than an hour finds *provably* optimal tours on instances with a few hundred nodes. Finding optimal tours on a few thousand nodes takes longer, however. Furthermore, the convergence rate of branch-and-cut is variable. In particular, there exist unsolved instances with less than a thousand nodes.

We see some hope that our techniques, in combination with recent branch-and-cut implementations, could be used to find *provably* near-optimal tours on 5,000 to 50,000 nodes in reasonable amount of time. The main idea would be to use our divide and conquer strategy for a few levels, say 7. Heuristically speaking, this would partition 25,000 nodes into 2^7 instances with $25,000/2^7 \approx 200$ nodes each, which one could hope to solve optimally using existing code. Our analysis would need to be redone (with a careful attention to constants) for a 7-level divide-and-conquer strategy.

Our techniques could also prove useful for structured TSP instances, such as those arising in circuit design. On random instances—a frequent testbed for experimental work on the TSP—our algorithm will probably not do any better than Karp’s algorithm.

Our techniques may prove especially useful for Minimum Steiner trees or k -MST, since current heuristics for those problems do not give good results in practice.

A2. The Local-Exchange View

Local-exchange algorithms for the TSP work by identifying possible edge exchanges in the current tour that lower the cost. For example each step of K -OPT identifies a set of K tour edges $\{(u_1, v_1), (u_2, v_2), \dots, (u_K, v_K)\}$, deletes them, and finds a way to re-link the $K - 1$ paths thus created so as to reduce the tour cost. If the algorithm cannot reduce the cost using any set of K edges, it stops. (We will not discuss the numerous heuristics used to identify possible edge-exchanges.)

Our dynamic programming algorithm can be restated as a slightly more inefficient divide-and-conquer with a running time $n^{O(c \log n)}$. The algorithm starts by making the instance well-rounded, then computes a trivial tour and a shifted quadtree with random shifts. Let $r = O(c)$ be the same constant as in the Structure Theorem. By working on one quadtree region at a time, the algorithm modifies the tour until it finds the best salesman tour that crosses the boundary of each region at most $4r$ times. This process is the same as the dynamic programming described earlier, except the algorithm does not maintain a lookup table. Thus, it resembles K -OPT for $K = O(c)$, except that cost-increasing exchanges have to be allowed in order to undo bad guesses. Maybe it is closer in spirit to more ad-hoc heuristics such as genetic algorithms, which do allow cost-increasing exchanges.

ACKNOWLEDGMENTS. This paper contains work that evolved over a couple of years and benefitted from comments from many colleagues. I particularly thank Mic Grigni, David Johnson, Ravi Kannan, Haim Kaplan, Sanjeev Khanna, Tom Leighton, Dick Lipton, Rajeev Motwani, Umesh Vazirani, and Mihalis Yannakakis. I also thank Pankaj Agarwal, Moses Charikar, David Eppstein, and Jeff Erickson for their help with improving the manuscript. Finally, I am indebted to Warren Smith and Michel Goemans for suggesting the use of Steiner points or “portals” in the description of the PTAS for TSP; this greatly simplified the exposition and the proof.

REFERENCES

- AGARWAL, P., EDELSBRUNNER, H., SCHWARTZKOPF, O., AND WELZL, E. 1991. Euclidean MST and bichromatic closest pairs. *Disc. Comput. Geom.* 6, 407–422.
- APPLEGATE, D., BIXBY, R., CHVATAL, V., AND COOK, W. 1995. Finding cuts in the TSP (a preliminary report). Report 95-05. DIMACS, Rutgers Univ., New Brunswick, N.J.
- ARORA, S. 1996. Polynomial-time approximation schemes for Euclidean TSP and other geometric problem. In *Proceedings of the 37th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 2–12.
- ARORA, S., GRIGNI, M., KARGER, D., KLEIN, P., AND WOLOSZYN, A. 1998. A polynomial-time approximation scheme for weighted planar graph TSP. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, pp. 33–41.
- ARORA, S., LUND, C., MOTWANI, R., SUDAN, M., AND SZEGEDY, M. 1992. Proof verification and interactability of approximation problems. In *Proceedings of the 33rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 13–22.
- ARORA, S., RAGHAVAN, P., AND RAO, S. 1998. Approximation schemes for the Euclidean k -medians and related problems. In *Proceedings of the 30th Annual ACM Symposium on Theory of Computing* (Dallas, Tex., May 23–26). ACM, New York, pp. 106–113.
- ARORA, S., AND SAFRA, S. 1992. Probabilistic checking of proofs: A new characterization of NP. In *Proceedings of the 33rd IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 2–12.
- ASANO, T., KATO, N., TAMAKI, H., AND TOKUYAMA, T. 1997. Covering points in the plane by k -tours: Towards a polynomial time approximation scheme for general k . In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (El Paso, Tex., May 4–6). ACM, New York, pp. 275–283.
- BEARDWOOD, J., HALTON, J. H., AND HAMMERSLEY, J. M. 1959. The shortest path through many points. *Proc. Cambridge Philos. Soc.* 55, 299–327.
- BENTLEY, J. 1992. Fast algorithms for geometric traveling salesman problem. *ORSA J. Comput.* 4, 387–411.
- BERN, M., AND EPPSTEIN, D. 1996. Approximation algorithms for geometric problems. In *Approximation Algorithms for NP-Hard Problems*. PWS Publishing, Boston, Mass.

- BERN, M., EPPSTEIN, D., AND TENG, S.-H. 1993. Parallel construction of quadtree and quality triangulations. In *Proceedings of the 3rd WADS*, Lecture Notes in Computer Science, vol. 709. Springer-Verlag, New York, pp. 188–199.
- BERN, M., AND PLASSMANN, P. 1989. The Steiner problem with edge lengths 1 and 2. *Inf. Proc. Lett.* 32, 171–176.
- BLUM, A., CHALASANI, P., AND VEMPALA, S. 1995. A constant-factor approximation for the k -MST problem in the plane. In *Proceedings of the 27th Annual ACM Symposium on Theory of Computing* (Las Vegas, Nev., May 29–June 1). ACM, New York, pp. 294–302.
- CHANDRA, B., KARLOFF, H., AND TOVEY, C. 1994. New results for the old k -OPT algorithm for the TSP. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms*. ACM, New York, pp. 150–159.
- CHRISTOFIDES, N. 1976. Worst-case analysis of a new heuristic for the traveling salesman problem. In *Symposium on New Directions and Recent Results in Algorithms and Complexity*, J. F. Traub, ed. Academic Press, Orlando, Fla., p. 441.
- DU, D. Z., AND HWANG, F. K. 1992. A proof of Gilbert–Pollack’s conjecture on the Steiner ratio. *Algorithmica* 7, 121–135.
- EPPSTEIN, D. 1998. Faster geometric k -point MST approximation. *CGTA*, to appear.
- EPPSTEIN, D., AND ERICKSON, J. 1994. Iterated nearest neighbors and finding minimal polytopes. *Disc. Comput. Geom.* 11, 321–350.
- FEIGE, U., GOLDWASSER, S., LOVÁSZ, L., SAFRA, S., AND SZEGEDY, M. 1991. Approximating clique is almost NP-complete. In *Proceedings of the 32nd IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 2–12.
- FERNANDEZ DE LA VEGA, W., AND LUEKER, G. S. 1981. Bin packing can be solved within $1 + \epsilon$ in linear time. *Combinatorica* 1, 4, 349–355.
- FISCHETTI, M., HAMACHER, H. W., JORNSTEN, K., AND MAFFIOLI, F. 1994. Weighted k -cardinality trees: Complexity and polyhedral structure. *Networks* 32, 11–21.
- GAREY, M. R., GRAHAM, R. L., AND JOHNSON, D. S. 1976. Some NP-complete geometric problems. In *Proceedings of the 8th Annual ACM Symposium on Theory of Computing* (Hershey, Pa., May 3–5). ACM, New York, pp. 10–22.
- GARG, N. 1996. A 3-approximation for the minimum tree spanning k vertices. In *Proceedings of the 37th Annual IEEE Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 302–310.
- GILBERT, E. N., AND POLLAK, R. O. 1968. Steiner minimal trees. *SIAM J. Appl. Math.* 16, 1–29.
- GOEMANS, M. 1995. Worst-case comparison of valid inequalities for the TSP. *Math. Prog.* 69, 336–349.
- GRAHAM, R. L. 1966. Bounds for certain multiprocessing anomalies. *Bell Syst. Tech. J.* 45, 1563–1581.
- GRIGNI, M., KOUTSOUPAS, E., AND PAPDIMITRIOU, C. H. 1995. An approximation scheme for planar graph TSP. In *Proceedings of the 36th Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 640–645.
- HOCHBAUM, D., ED. 1996. *Approximation Algorithms for NP-Hard Problems*. PWS Publishing, Boston, Mass.
- IBARRA, O. H., AND KIM, C. E. 1975. Fast approximation algorithms for the knapsack and sum of subsets problems. *J. ACM* 22, 4 (Oct.) 463–468.
- JOHNSON, D. S. 1974. Approximation algorithms for combinatorial problems. *J. Comput. Syst. Sci.* 9, 256–278.
- JOHNSON, D. S., AND MCGEOCH, L. A. 1997. The traveling salesman problem: A case study in local optimization. In *Local Search in Combinatorial Optimization*, E. H. L. Aarts and J. K. Lenstra, eds. Wiley, New York.
- JOHNSON, D. S., PAPADIMITRIOU, C., AND YANNAKAKIS, M. 1988. How easy is local search. *J. Comput. Syst. Sci.* 37, 79–100.
- JOHNSON, W. B., AND LINDENSTRAUSS, J. 1984. Extensions of Lipschitz mappings into Hilbert space. *Contemp. Math.* 26, 189–206.
- KARMARKAR, N., AND KARP, R. M. 1982. An efficient approximation scheme for the one-dimensional bin-packing problem. In *Proceedings of the 23rd Annual IEEE Symposium on Foundations of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 312–320.

- KARP, R. M. 1972. Reducibility among combinatorial problems. In *Complexity of Computer Computations*, R. E. Miller and J. W. Thatcher, eds. Advances in Computer Research, Plenum Press, New York, pp. 85–103.
- KARP, R. M. 1977. Probabilistic analysis of partitioning algorithms for the TSP in the plane. *Math. Oper. Res.* 2, 209–224.
- KHULLER, S., RAGHAVACHARI, B., AND YOUNG, N. 1996. Low degree spanning tree of small weight. *SIAM J. Comput.* 25, 355–368.
- KRENTEL, M. W. 1989. Structure in locally optimal solutions. In *Proceedings of the 30th Annual IEEE Symposium on Foundation of Computer Science*. IEEE Computer Society Press, Los Alamitos, Calif., pp. 216–221.
- LAWLER, E. L., LENSTRA, J. K., RINNOOY KAN, A. H. G., AND SHMOYS, D. B. 1985. The traveling salesman problem. Wiley, New York.
- LIN, S. 1965. Computer solutions for the traveling salesman problem. *Bell Syst. Tech. J.* 44, 2245–2269.
- LIN, S., AND KERNIGHAN, B. W. 1973. An effective heuristic algorithm for the traveling salesman problem. *Oper. Res.* 21, 498–516.
- MATA, C. S., AND MITCHELL, J. B. 1995. Approximation algorithms for geometric tour and network design problems. In *Proceedings of the 11th Annual ACM Symposium on Computational Geometry* (Vancouver, B.C., Canada, June 5–7). ACM, New York, pp. 360–369.
- MELZAK, Z. A. 1961. On the problem of Steiner. *Canad. Math Bull.* 4, 143–148.
- MITCHELL, J. S. B. 1996. Guillotine subdivisions approximate polygonal subdivisions: A simple new method for the geometric k -MST problem. In *Proceedings of the 6th Annual ACM-SIAM Symposium on Discrete Algorithms* (Atlanta, Ga., Jan. 28–30). ACM, New York, pp. 402–408.
- MITCHELL, J. S. B. 1998. Guillotine subdivisions approximate polygonal subdivisions: Part II—A simple PTAS for geometric k -MST, TSP, and related problems. *SIAM J. Comput.*, to appear.
- PAPADIMITRIOU, C. H. 1977. Euclidean TSP is NP-complete. *Theoret. Comput. Sci.* 4, 237–244.
- PAPADIMITRIOU, C. H. 1992. The complexity of the Lin–Kernighan heuristic for the traveling salesman problem. *SIAM J. Comput.* 21, 3, 450–465.
- PAPADIMITRIOU, C. H., AND YANNAKAKIS, M. 1984. The complexity of facets (and some facets of complexity). *J. Comput. Syst. Sci.* 28, 244–259.
- PAPADIMITRIOU, C. H., AND YANNAKAKIS, M. 1991. Optimization, approximation and complexity classes. *J. Comput. Syst. Sci.* 43, 425–440.
- RAO, S. B., AND SMITH, W. D. 1988. Approximating geometric graphs via “spanners” and “banyans”. In *Proceedings of the 30th ACM Symposium on the Theory of Computing* (Dallas, Tex., May 23–26). ACM, New York, pp. 540–550.
- REINELT, G. 1991. TSPLIB—A traveling salesman problem library. *ORSA J. Comput.* 3, 376–384.
- SAHNI, S., AND GONZALES, T. 1976. P-complete approximation problems. *J. ACM* 23, 555–565.
- SHMOYS, D. B., AND WILLIAMSON, D. P. 1990. Analyzing the Held–Karp TSP bound: A monotonicity property with application. *Inf. Proc. Lett.* 35, 281–285.
- SMITH, W. D. 1988. Finding the optimum N -city traveling salesman tour in the Euclidean plan in subexponential time and polynomial space. Manuscript.
- TREVISAN, L. 1997. When Hamming meets Euclid: The approximability of geometric TSP and MST. In *Proceedings of the 29th Annual ACM Symposium on Theory of Computing* (El Paso, Tex., Mar. 4–6). ACM, New York, pp. 21–39.
- VAIDYA, P. 1988. Geometry helps in matching. *SIAM J. Comput.* 18, 1201–1225.
- VAIDYA, P. 1989. Approximate minimum weight matching on k -dimensional spaces. *Algorithmica* 4, 569–583.
- WANG, K. 1996. Is the m parameter in Arora’s TSP algorithm the best possible? Junior Project, Princeton University, Princeton, N.J.
- WOLSEY, L. A. 1980. Heuristic analysis, linear programming and branch and bound. *Math. Prog. Study* 13, 121–134.
- ZELIKOVSKY, A. Z. 1993. An 11/6-approximation algorithm for the network Steiner Problem. *Algorithmica* 9, 463–470.
- ZELIKOVSKY, A. Z. 1996. Better approximation bounds for the network and Euclidean Steiner tree problems. Tech. Rep. CS-96-06. Univ. Virginia.

RECEIVED JULY 1997; REVISED JANUARY 1998; ACCEPTED JUNE 1998