

AN ETH-TIGHT EXACT ALGORITHM FOR EUCLIDEAN TSP*

MARK DE BERG[†], HANS L. BODLAENDER[‡], SÁNDOR KISFALUDI-BAK[§], AND
SUDESHNA KOLAY[¶]

Abstract. We study exact algorithms for METRIC TSP in \mathbb{R}^d . In the early 1990s, algorithms with $n^{O(\sqrt{n})}$ running time were presented for the planar case, and some years later an algorithm with $n^{O(n^{1-1/d})}$ running time was presented for any $d \geq 2$. Despite significant interest in subexponential exact algorithms over the past decade, there has been no progress on METRIC TSP, except for a lower bound stating that the problem admits no $2^{o(n^{1-1/d})}$ algorithm unless ETH fails. In this paper we settle the complexity of METRIC TSP, up to constant factors in the exponent and under ETH, by giving an algorithm with running time $2^{O(n^{1-1/d})}$.

Key words. Euclidean traveling salesman, subexponential algorithm, separator theorem

MSC codes. 68W40, 68Q25

DOI. 10.1137/22M1469122

1. Introduction. The traveling salesman problem, or TSP for short, is one of the most widely studied problems in all of computer science. In (the symmetric version of) the problem we are given a complete undirected graph \mathcal{G} with positive edge weights, and the goal is to compute a minimum-weight cycle visiting every vertex exactly once. In 1972 the problem was shown to be NP-hard by Karp [19]. A brute-force algorithm for TSP runs in $O(n!)$, but the celebrated Held–Karp dynamic-programming algorithm, discovered independently by Held and Karp [15] and Bellman [3], runs in $O(2^n n^2)$ time. Despite extensive efforts and progress on special cases, it is still open whether an exact algorithm for TSP exists with running time $O((2 - \varepsilon)^n \text{poly}(n))$.

In this paper we study the Euclidean version of TSP, where the input is a set P of n points in \mathbb{R}^d and the goal is to find a tour of minimum Euclidean length visiting all the points. EUCLIDEAN TSP has been studied extensively, and it can be considered one of the most important geometric optimization problems. In the mid-1970s, EUCLIDEAN TSP was shown to be NP-hard [12, 26]. Nevertheless, its computational complexity is markedly different from that of the general TSP problem. For instance, EUCLIDEAN TSP admits efficient approximation algorithms. Indeed, the famous algorithm by Christofides [7]—which actually works for the more general METRIC TSP problem—provides a $(3/2)$ -approximation in polynomial time, while

*Received by the editors January 4, 2022; accepted for publication (in revised form) February 6, 2023; published electronically June 5, 2023. An earlier version of this article appeared in the Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, IEEE Computer Society, 2018, pp. 450–461, and the article shares content with the third author’s thesis prepared at Technische Universiteit Eindhoven, Department of Mathematics and Computer Science, 2019.

<https://doi.org/10.1137/22M1469122>

Funding: This work was supported by the NETWORKS project funded by the Netherlands Organization for Scientific Research under grant 024.002.003.

[†]Department of Mathematics and Computer Science, Eindhoven University of Technology, Eindhoven 5600 MB, The Netherlands (m.t.d.berg@tue.nl).

[‡]Department of Computer Science, Utrecht University, Utrecht 3508 TB, The Netherlands (h.l.bodlaender@uu.nl).

[§]Department of Computer Science, Aalto University, Espoo FI-00076, Finland (sandor.kisfaludi-bak@aalto.fi).

[¶]Indian Institute of Technology Kharagpur, Kharagpur, India (skolay@cse.iitkgp.ac.in).

no polynomial-time approximation algorithm exists for the general problem (unless $P = NP$). It was a longstanding open problem whether EUCLIDEAN TSP admits a PTAS. The question was answered affirmatively by Arora [1] who provided a PTAS (a polynomial time approximation scheme) with running time $n(\log n)^{O(\sqrt{d}/\varepsilon)^{d-1}}$. Independently, Mitchell [24] designed a PTAS in \mathbb{R}^2 . The running time was improved to $2^{(1/\varepsilon)^{O(d)}} n + (1/\varepsilon)^{O(d)} n \log n$ by Rao and Smith [27], and then to $2^{(1/\varepsilon)^{O(d)}} n$ by Bartal and Gottlieb [2]. Kisfaludi-Bak, Nederlof, and Wegrzycki [21] recently developed a PTAS with running time $2^{O(1/\varepsilon^{d-1})} n \log n$, and they showed that the dependency on ε is optimal under Gap-ETH. Hence, the computational complexity of approximating EUCLIDEAN TSP is well understood.

Results on exact algorithms for EUCLIDEAN TSP—such algorithms are the topic of our paper—are also quite different from those on the general problem. The best known algorithm for the general case runs, as already noted, in exponential time, and there is no $2^{o(n)}$ algorithm under the exponential time hypothesis (ETH) [17] due to classical reductions for HAMILTONIAN CYCLE [8, Theorem 14.6]. EUCLIDEAN TSP, on the other hand, is solvable in subexponential time. For the planar case, this was shown in the early 1990s by Kann [18] and independently by Hwang, Chang, and Lee [16], who presented an algorithm with an $n^{O(\sqrt{n})}$ running time. Both algorithms use a divide-and-conquer approach that relies on finding a suitable separator. The approach taken by Hwang, Chang, and Lee is based on considering a triangulation of the point set such that all segments of the tour appear in the triangulation, and then observing that the resulting planar graph has a separator of size $O(\sqrt{n})$. Such a separator can be guessed in $n^{O(\sqrt{n})}$ ways, leading to a recursive algorithm with $n^{O(\sqrt{n})}$ running time. It seems hard to extend this approach to higher dimensions. Kann obtains his separator in a more geometric way, using the fact that in an optimal tour, there cannot be too many long edges that are relatively close together—see the packing property we formulate in section 2. This makes it possible to compute a separator that is crossed by $O(\sqrt{n})$ edges of an optimal tour, which can be guessed in $n^{O(\sqrt{n})}$ ways. The geometric flavor of this algorithm makes it more amenable to extensions to higher dimensions. Indeed, some years later Smith and Wormald [28] gave an algorithm for EUCLIDEAN TSP in \mathbb{R}^d , which is based on a geometric separator similar to the kind used by Kann. Their algorithm runs in $n^{O(n^{1-1/d})}$ time. (Here and in what follows, we consider the dimension d to be an absolute constant.)

The main question, also posed by Woeginger in his survey [29] on open problems around exact algorithms, is the following: is an exact algorithm with running time $2^{O(n^{1-1/d})}$ attainable for EUCLIDEAN TSP? Similar results have been obtained for some related problems. In particular, Deineko, Klinz, and Woeginger [10] proved that HAMILTONIAN CYCLE on planar graphs can be solved in $2^{O(\sqrt{n})}$ time, and Dorn et al. [11] proved that TSP on weighted planar graphs can be solved in $2^{O(\sqrt{n})}$ time. Marx and Sidiropoulos [23] have shown that EUCLIDEAN TSP does not admit an algorithm with $2^{O(n^{1-1/d-\varepsilon})}$, unless ETH fails. Recently, this conditional lower bound was strengthened to $2^{\Omega(n^{1-1/d})}$ by De Berg et al. [5]. In the past 20 years, the algorithms for EUCLIDEAN TSP have not been improved, however. Hence, even for the planar case, the complexity of EUCLIDEAN TSP is still unknown.

Our contribution. We settle the complexity of EUCLIDEAN TSP up to constant factors in the exponent: we present an algorithm for EUCLIDEAN TSP in \mathbb{R}^d with running time $2^{O(n^{1-1/d})}$, where $d \geq 2$ is a fixed constant. Our algorithm is asymptotically optimal, as De Berg et al. [5] show that there is no $2^{o(n^{1-1/d})}$ algorithm unless ETH fails.

The global approach for obtaining the upper bound is similar to the approach of Kann [18] and Smith and Wormald [28]: we use a divide-and-conquer algorithm based on a geometric separator. A geometric separator for a given point set P is a simple geometric object—we use a hypercube—such that the number of points inside the separator and the number of points outside the separator are roughly balanced. As mentioned above, Kann [18] and Smith and Wormald [28] use a packing property of the edges in an optimal TSP tour to argue that a separator exists that is crossed by only $O(n^{1-1/d})$ edges from the tour. Since P defines $\binom{n}{2}$ possible edges, the set of crossing edges can be guessed in $n^{O(n^{1-1/d})}$ ways.

The first obstacle we must overcome if we want to beat this running time is the fact that the number of subproblems is already too large at the first step of the recursive algorithm. Unfortunately there is no hope of obtaining a balanced separator that is crossed by $o(n^{1-1/d})$ edges from the tour: there are point sets such that any balanced separator that has a “simple” shape (e.g., ball or hypercube) is crossed $\Omega(n^{1-1/d})$ times by an optimal tour. Thus we proceed differently: we prove that there exists a separator such that, even though it can be crossed by up to $\Theta(n^{1-1/d})$ edges from an optimal tour, the total number of candidate subsets of crossing edges we need to consider is only $2^{O(n^{1-1/d})}$. We obtain such a separator in two steps. First, we prove a *distance-based separator theorem* for point sets. Intuitively, this theorem states that any point set P admits a balanced separator such that the number of points from P within a certain distance from the separator decreases rapidly as the distance decreases. This is useful because an optimal tour cannot have many “long” edges crossing the separator due to the packing property. Hence, limiting the number of points “close” to the separator also limits the number of combinations of edges that form a candidate subset. In the second step we then prove that this separator σ indeed has the required properties, namely (i) σ is crossed by $O(n^{1-1/d})$ edges in an optimal tour, and (ii) the number of candidate sets of crossing edges is $2^{O(n^{1-1/d})}$. In order to prove these properties we use the packing property of the edges in an optimal tour.

There is one other obstacle we need to overcome to obtain a $2^{O(n^{1-1/d})}$ algorithm: after computing a suitable separator σ and guessing a set S of crossing edges, we still need to solve many different subproblems. The reason is that the partial solutions on either side of σ need to fit together into a tour on the whole point set. Thus a partial solution on the outside of σ imposes connectivity constraints on the inside. More precisely, if B is the set of endpoints of the edges in S that lie inside σ , then the subproblem we face inside σ is as follows: compute a set of paths visiting the points inside σ such that the paths realize a given matching on B . The number of matchings on $|B|$ boundary points is $|B|^{\Theta(|B|)}$, which is again too much for our purposes. Fortunately, the rank-based approach [6, 9] developed in recent years can be applied here. By applying this approach in a suitable manner, we then obtain our $2^{O(n^{1-1/d})}$ algorithm.

A word on the model of computation. In this paper we are mainly interested in the combinatorial complexity of EUCLIDEAN TSP. The algorithm we describe in sections 2 and 3 therefore works in the real-RAM model of computation, with the capability of taking square roots. In particular, we assume that distances can be added in $O(1)$ time, so that the length of a given tour can be computed exactly in $O(n)$ time. In section 4 we also consider the following “almost Euclidean” version of the problem: we are given a set $P = \{p_1, \dots, p_n\}$ with rational coordinates, together with a distance matrix D such that $D[i, j]$ contains an approximation of $|p_i p_j|$. The property we require is that the ordering of distances is preserved: if $|p_i p_j| < |p_k p_l|$, then $D[i, j] < D[k, l]$. We show that an optimal tour in this setting satisfies the

packing property, which implies that our algorithm can solve the almost Euclidean version of EUCLIDEAN TSP in $2^{O(n^{1-1/d})}$ time.

2. A separator theorem for TSP. In this section we show how to obtain a separator that can be used as the basis of an efficient recursive algorithm to compute an optimal TSP tour. Intuitively, we need a separator that is crossed only a few times by an optimal solution and such that the number of candidate sets of crossing edges is small. We obtain such a separator in two steps: first, we construct a separator σ such that there are only a few points relatively close to σ , and then we show that this implies that σ has all the desired properties.

Notation and terminology. We define a *separator* to be the boundary of an axis-aligned hypercube. A separator σ partitions \mathbb{R}^d into two regions: a region σ_{in} consisting of all points in \mathbb{R}^d inside or on σ , and a region σ_{out} consisting of all points in \mathbb{R}^d strictly outside σ . We define the *size* of a separator σ to be its side length, and we denote it by $\text{size}(\sigma)$. For a separator σ and a scaling factor $t \geq 0$, we define $t\sigma$ to be the separator obtained by scaling σ by a factor t with respect to its center. More precisely, $t\sigma$ is the separator whose center is the same as the center of σ and with $\text{size}(t\sigma) = t \cdot \text{size}(\sigma)$; see Figure 1(i).

Let P be a set of n points in \mathbb{R}^d . A separator σ induces a partition of the given point set P into two subsets, $P \cap \sigma_{\text{in}}$ and $P \cap \sigma_{\text{out}}$. We are interested in δ -balanced separators, which are separators such that $\max(|P \cap \sigma_{\text{in}}|, |P \cap \sigma_{\text{out}}|) \leq \delta n$ for a fixed constant $\delta > 0$. It will be convenient to work with δ -balanced separators for $\delta = 4^d/(4^d + 1)$. From now on we will refer to $(4^d/(4^d + 1))$ -balanced separators simply as *balanced separators*. (There is nothing special about the constant $4^d/(4^d + 1)$, and it could be made smaller by a more careful reasoning and at the cost of some other constants we will encounter later on.)

Distance-based separators for point sets. As mentioned, we first construct a separator σ such that there are only a few points close to it. To this end, we define the *relative distance* from a point p to σ , denoted by $\text{rdist}(p, \sigma)$, as follows:

$$\text{rdist}(p, \sigma) \stackrel{\text{def}}{=} d_\infty(p, \sigma) / \text{size}(\sigma),$$

where $d_\infty(p, \sigma)$ denotes the shortest distance in the ℓ_∞ -metric between p and any point on σ . Note that if t is the scaling factor such that $p \in t\sigma$, then $\text{rdist}(p, \sigma) = |1 - t|/2$. For integers i define

$$P_i(\sigma) \stackrel{\text{def}}{=} \{ p \in P \mid \text{rdist}(p, \sigma) \leq 2^i/n^{1/d} \}.$$




FIG. 1. (i) A separator σ and a point p with $\text{rdist}(p, \sigma) = 0.75$. (ii) Schematic drawing of the weight function $w_p(t)$ of a point p such that $\text{rdist}(p, t\sigma^*) = 0$ for $t = 2.5$.

Note that the smaller i is, the closer to σ the points in $P_i(\sigma)$ are required to be. We now wish to find a separator σ such that the size of the sets $P_i(\sigma)$ decreases rapidly as i decreases.

THEOREM 2.1. *Let P be a set of n points in \mathbb{R}^d . Then there is a balanced separator σ for P such that*

$$|P_i(\sigma)| = \begin{cases} O((3/2)^i n^{1-1/d}) & \text{for all } i < 0, \\ O(4^i n^{1-1/d}) & \text{for all } 0 \leq i. \end{cases}$$

Moreover, such a separator can be found in $O(n^{d+1})$ time.

Proof. Let σ^* be a smallest separator such that $|P \cap \sigma_{\text{in}}^*| \geq n/(4^d + 1)$. We will show that there is a t^* with $1 \leq t^* \leq 3$ such that $t^*\sigma^*$ is a separator with the required properties.

First, we claim that $t\sigma^*$ is balanced for all $1 \leq t \leq 3$. To see this, observe that for $t \geq 1$ we have

$$|P \cap (t\sigma^*)_{\text{out}}| \leq |P \cap \sigma_{\text{out}}^*| = n - |P \cap \sigma_{\text{in}}^*| \leq n - n/(4^d + 1) = (4^d/(4^d + 1))n.$$

Moreover, for $t \leq 3$ we can cover $t\sigma_{\text{in}}^*$ by 4^d hypercubes of size $(3/4) \cdot \text{size}(\sigma^*)$. By definition of σ^* , each of these hypercubes contains less than $n/(4^d + 1)$ points, so $|P \cap (t\sigma^*)_{\text{in}}| < 4^d \cdot (n/(4^d + 1))$, which finishes the proof of the claim. (So far, the proof is similar to earlier separator constructions [5, 13]; the main challenge is to establish the distance properties, which we do next.)

It remains to prove that there is a t^* with $1 \leq t^* \leq 3$ such that $t^*\sigma^*$ satisfies the condition on the sizes of the sets $P_i(t^*\sigma^*)$. To this end, we will define a weight function $w_p : [1, 3] \rightarrow \mathbb{R}$ for each $p \in P$. The idea is that the closer p is to $t\sigma^*$, the higher the value $w_p(t)$. An averaging argument will then show that there must be a t^* such that $\sum_{p \in P} w_p(t^*)$ is sufficiently small, from which it follows that $t^*\sigma^*$ satisfies the condition on the sizes of the sets $P_i(t^*\sigma^*)$. Next, we make this idea precise.

Assume without loss of generality that $\text{size}(\sigma^*) = 1$. For a point $p \in P$, let $i_p(t)$ be the integer such that $2^{i_p(t)-1}/n^{1/d} < \text{rdist}(p, t\sigma^*) \leq 2^{i_p(t)}/n^{1/d}$, where $i_p(t) = -\infty$ if $\text{rdist}(p, t\sigma^*) = 0$. Note that $p \in P_i(t\sigma^*)$ if and only if $i_p(t) \leq i$. We define the weight function $w_p(t)$ as follows (see Figure 1(ii)):

$$w_p(t) \stackrel{\text{def}}{=} \begin{cases} \frac{n^{1/d}}{(3/2)^{i_p(t)}} & \text{if } i_p(t) < 0, \\ \frac{n^{1/d}}{4^{i_p(t)}} & \text{if } i_p(t) \geq 0, \\ \text{undefined} & \text{if } i_p(t) = -\infty. \end{cases}$$

The above definition has different constants in the denominators $((3/2)^{i_p(t)})$ versus $4^{i_p(t)}$ for a technical reason: we will soon need the fact that when multiplied by $2^{i_p(t)}$, the cumulative sums are convergent sequences for both $i_p(t) < 0$ and $i_p(t) \geq 0$. One could instead use any constant from the open interval $(1, 2)$ for the case $i_p(t) < 0$, and any constant from $(2, \infty)$ for the case $i_p(t) \geq 0$, respectively. We now want to bound $\int_1^3 w_p(t) dt$. Note that the function $w_p(t)$ may be undefined for at most one $t \in [1, 3]$, namely when there is a t in this range such that $\text{rdist}(p, t\sigma^*) = 0$. Formally, we should remove such a t from the domain of integration. To avoid cluttering the notation, we ignore this technicality and continue to write $\int_1^3 w_p(t) dt$.

CLAIM 2.2. *For each $p \in P$, we have $\int_1^3 w_p(t) dt = O(1)$.*

Proof. Define $T_p(i) \stackrel{\text{def}}{=} \{t \mid 1 \leq t \leq 3 \text{ and } i_p(t) = i\}$. By the definition of $w_p(t)$, the value $w_p(t)$ is constant over $T_p(i)$. We therefore want to bound $\text{len}(T_p(i))$, the sum of the lengths of the intervals comprising $T_p(i)$. Assume without loss of generality that the center of σ^* lies at the origin of \mathbb{R}^d . Then, depending on the position of p , there is some coordinate x_i such that $\text{rdist}(p, t\sigma^*) = |p_{x_i} - t|/\text{size}(t\sigma^*)$. Assume without loss of generality that $|p_{x_1} - t| = \text{size}(t\sigma^*)$. Since $1 \leq t \leq 3$ and $\text{size}(\sigma^*) = 1$, we then have $\text{rdist}(p, t\sigma^*) \geq |p_{x_1} - t|/3$. Hence, for any $t \in T_p(i)$ we have $|p_{x_1} - t|/3 \leq 2^i/n^{1/d}$. This implies that $\text{len}(T_p(i)) \leq 6 \cdot 2^i/n^{1/d}$, and so

$$\begin{aligned} \int_1^3 w_p(t) dt &= \sum_{i \geq 0} \text{len}(T_p(i)) \cdot \frac{n^{1/d}}{4^i} + \sum_{i < 0} \text{len}(T_p(i)) \cdot \frac{n^{1/d}}{(3/2)^i} \\ &\leq 6 \cdot \sum_{i \geq 0} \left(\frac{1}{2}\right)^i + 6 \cdot \sum_{i < 0} \left(\frac{4}{3}\right)^i \\ &= 30. \end{aligned} \quad \square$$

The above claim implies that $\int_1^3 (\sum_{p \in P} w_p(t)) dt \leq 30n$. Hence there exists a $t^* \in [1, 3]$ such that $\sum_{p \in P} w_p(t^*) \leq 15n$. Now consider a set $P_i(t^*\sigma^*)$ with $i \geq 0$. Each $p \in P_i(t^*\sigma^*)$ has $i_p(t^*) \leq i$, and so

$$|P_i(t^*\sigma^*)| \leq \frac{\sum_{p \in P} w_p(t^*)}{\min_{p \in P_i(t^*\sigma^*)} w_p(t^*)} = \frac{O(n)}{n^{1/d}/4^i} = O(4^i n^{1-1/d}).$$

A similar argument shows that $|P_i(t^*\sigma^*)| = O((3/2)^i n^{1-1/d})$ for all $i < 0$.

To find the desired separator, we first compute σ^* . Note that we can always shift σ^* such that it has at least one point on at least d of its $(d - 1)$ -dimensional faces. (Note that an input point on a lower-dimensional face of σ^* is counted towards each incident facet.) Hence, a simple brute-force algorithm can find σ^* in $O(n^{d+1})$ time. Once we have σ^* , we would like to find the value $t^* \in [1, 3]$ minimizing $\sum_{p \in P} w_p(t)$. Recall that each w_p is a step function, and so $\sum_{p \in P} w_p$ is a step function as well. There is one slight issue, however; namely, the number of steps of the functions w_p is unbounded. We deal with this issue by replacing each w_p by a truncated version \bar{w}_p , as explained next.

We define the truncated function \bar{w}_p as follows:

$$\bar{w}_p(t) \stackrel{\text{def}}{=} \begin{cases} 1/n & \text{if } w_p(t) < 1/n, \\ w_p(t) & \text{if } 1/n \leq w_p(t) \leq 15n, \\ 15n & \text{if } w_p(t) > 15n. \end{cases}$$

Each function \bar{w}_p is a step function, and one easily verifies that \bar{w}_p has $O(\log n)$ steps which we can compute in $O(\log n)$ time. Hence, we can find a value \bar{t} that minimizes $\sum_{p \in P} \bar{w}_p(t)$ in $O(n \log n)$ time by scanning over all step functions in parallel and maintaining their sum as we go. Since $\sum_{p \in P} w_p(\bar{t}) = O(n)$ if $\sum_{p \in P} \bar{w}_p(\bar{t}) = O(n)$, the separator $\bar{t}\sigma^*$ has the required properties. \square

Remark 2.3. It is not hard to reduce the time needed to compute the separator by working with an approximation of the smallest hypercube σ^* containing at least $n/(4^d + 1)$ points. We can find an ϵ -approximation to the minimum enclosing ball in linear time [14], whose circumscribed axis-parallel cube is a constant-approximation to the minimum enclosing cube. Note that this would weaken the balance factor of the

separator theorem. Nonetheless, in our application this does not make a difference, and the simple brute-force algorithm to find σ^* suffices.

In the remainder of the paper, we will need a slightly more general version of Theorem 2.1, where we require the separator to be balanced with respect to a given subset $Q \subseteq P$; that is, we require $\max(|Q \cap \sigma_{\text{in}}|, |Q \cap \sigma_{\text{out}}|) \leq \delta |Q|$ for $\delta = 4^d / (4^d + 1)$. Note that the distance condition in the corollary below is still with respect to the points in P . The proof of the corollary is exactly the same as before; we only need to redefine σ^* to be a smallest separator such that $|Q \cap \sigma_{\text{in}}^*| \geq |Q| / (4^d + 1)$.

COROLLARY 2.4. *Let P be a set of n points in \mathbb{R}^d , and let $Q \subseteq P$. Then there is a separator σ that is balanced with respect to Q and such that*

$$|P_i(\sigma)| = \begin{cases} O((3/2)^i n^{1-1/d}) & \text{for all } i < 0, \\ O(4^i n^{1-1/d}) & \text{for all } 0 \leq i. \end{cases}$$

Moreover, such a separator can be found in $O(n^{d+1})$ time.

A separator for TSP. Let P be a set of n points in \mathbb{R}^d , and let $\mathcal{S}(P)$ be the set of segments defined by P , that is, $\mathcal{S}(P) \stackrel{\text{def}}{=} \{pq \mid (p, q) \in P \times P\}$. Now consider a segment $s \in \mathcal{S}(P)$ and a separator σ . We say that s crosses σ if one endpoint of s lies in σ_{in} while the other lies in σ_{out} . Using our distance-based separator for points, we want to find a separator that is crossed only a few times by an optimal TSP tour. Moreover, we want to control the number of ways in which we have to “guess” a set of crossing segments. For this we will need the following crucial property of the segments in an optimal TSP tour.

DEFINITION 2.5. *A set S of segments in \mathbb{R}^d has the packing property if for any separator σ we have*

- (PP1): $|\{s \in S \mid s \text{ crosses } \sigma \text{ and } \text{length}(s) \geq \text{size}(\sigma)\}| = O(1)$,
- (PP2): $|\{s \in S \mid s \subset \sigma_{\text{in}} \text{ and } \text{length}(s) \geq \text{size}(\sigma)/4\}| = O(1)$.

Property (PP2) is actually implied by (PP1), but it will be convenient to explicitly state (PP2) as part of the definition. Note that the constants hidden in the O -notation in Definition 2.5 may (and do) depend on d .

Some variants of the above packing property have been shown to hold for the set of edges of an optimal tour for EUCLIDEAN TSP [18, 28]. For completeness we present a proof in a more general setting, which can be found in section 4.¹ Hence, we can restrict our attention to subsets of $\mathcal{S}(P)$ with the packing property. For a separator σ , we are thus interested in the following collection of sets of segments crossing σ :

$$\mathcal{C}(\sigma, P) \stackrel{\text{def}}{=} \{S \subseteq \mathcal{S}(P) \mid S \text{ has the packing property and all segments in } S \text{ cross } \sigma\}.$$

Our main separator theorem, given as Theorem 2.6, states that we can find a separator σ that is balanced and such that the sets in $\mathcal{C}(\sigma, P)$, as well as the collection $\mathcal{C}(\sigma, P)$ itself, are small. Since the general packing property is hard to test, we will enumerate a slightly larger collection of *candidate sets*, which we denote by $\mathcal{C}'(\sigma, P)$.

THEOREM 2.6. *Let P be a set of n points in \mathbb{R}^d , and let $Q \subseteq P$. Then there is a separator σ and collection $\mathcal{C}'(\sigma, P)$ such that*

¹In this more general setting, we work with a distance matrix D instead of with the exact Euclidean distances. We prove that the packing property holds when the pairwise distances in D have the same ordering as the pairwise Euclidean distances, that is, when $D[i, j] < D[k, \ell]$ if and only if $|p_i p_j| < |p_k p_\ell|$.

- (i) σ is balanced with respect to Q ;
- (ii) each candidate set $S \in \mathcal{C}'(\sigma, P)$ contains $O(n^{1-1/d})$ segments;
- (iii) $\mathcal{C}(\sigma, P) \subseteq \mathcal{C}'(\sigma, P)$ and $|\mathcal{C}(\sigma, P)| \leq |\mathcal{C}'(\sigma, P)| = 2^{O(n^{1-1/d})}$.

Moreover, σ and the collection $\mathcal{C}'(\sigma, P)$ can be computed in $2^{O(n^{1-1/d})}$ time.

Proof. Let σ be the separator obtained by applying Corollary 2.4 to the sets P and Q . Then σ has property (i). It remains to show that there is a set $\mathcal{C}'(\sigma, P)$ that has properties (ii) and (iii) and that can be enumerated in $2^{O(n^{1-1/d})}$ time. We assume without loss of generality that $\text{size}(\sigma) = 1$ and that σ is centered at the origin.

Let $L_{\text{small}} \stackrel{\text{def}}{=} 1/(n^{1/d}n^{(1-1/d)\log_{3/2}2})$. Any set $S \in \mathcal{C}(\sigma, P)$ can be partitioned into the following three subsets:

- $S_{\text{short}} \stackrel{\text{def}}{=} \{s \in S \mid \text{length}(s) \leq L_{\text{small}}\}$,
- $S_{\text{mid}} \stackrel{\text{def}}{=} \{s \in S \mid L_{\text{small}} < \text{length}(s) \leq 1\}$,
- $S_{\text{long}} \stackrel{\text{def}}{=} \{s \in S \mid \text{length}(s) > 1\}$.

We will analyze each of the three subsets separately. This analysis will give an upper bound on the number of candidates for the given type of subset. The argument will be constructive in the sense that it allows us to enumerate the candidates efficiently. By combining the candidates for the three types of subsets, we then get the desired set $\mathcal{C}'(\sigma, P)$. We start by analyzing S_{short} and S_{long} .

CLAIM 2.7. *For any $S \in \mathcal{C}(\sigma, P)$, the set S_{short} consists of $O(1)$ segments, and the number of different subsets S_{short} that can arise over all sets $S \in \mathcal{C}(\sigma, P)$ is $n^{O(1)}$. Similarly, S_{long} consists of $O(1)$ segments, and the number of different subsets S_{long} that can arise over all sets $S \in \mathcal{C}(\sigma, P)$ is $n^{O(1)}$.*

Proof. A segment in S_{short} has both endpoints at distance at most L_{small} from σ , and so both endpoints are in $P_i(\sigma)$ for $i = -\log_{3/2}n^{1-1/d}$. By Corollary 2.4, the number of points in this $P_i(\sigma)$ is $O((3/2)^i n^{1-1/d}) = O(1)$. Hence, $|S_{\text{short}}| = O(1)$, which trivially implies we can choose S_{short} in $n^{O(1)}$ ways. The number of segments in S_{long} is $O(1)$ by packing property (PP1), which again implies that we can choose S_{long} in $n^{O(1)}$ ways. \square

We can also enumerate the candidates for S_{short} and S_{long} in polynomial time. It remains to handle S_{mid} .

CLAIM 2.8. *For any $S \in \mathcal{C}(\sigma, P)$, the set S_{mid} consists of $O(n^{1-1/d})$ segments, and the number of different subsets S_{mid} that can arise over all sets $S \in \mathcal{C}(\sigma, P)$ is $2^{O(n^{1-1/d})}$.*

Proof. Define $S_{\text{mid}}(i) \subseteq S_{\text{mid}}$ to be the following:

$$S_{\text{mid}}(i) \stackrel{\text{def}}{=} \left\{ s \in S_{\text{mid}} \mid 2^{i-1}/n^{1/d} < \text{length}(s) \leq 2^i/n^{1/d} \right\}.$$

Note that

$$S_{\text{mid}} = \bigcup \{ S_{\text{mid}}(i) \mid -\log_{3/2}n^{1-1/d} + 1 \leq i \leq \log n^{1/d} \}.$$

We first analyze $|S_{\text{mid}}(i)|$ and the number of ways in which we can choose $S_{\text{mid}}(i)$ for a fixed i . To this end, we partition each face f of σ into a $(d-1)$ -dimensional grid whose cells have size $2^i/n^{1/d}$. (If $n^{1/d}/2^i$ is not an integer, then we use size $1/\lceil n^{1/d}/2^i \rceil$; all subsequent arguments work in this case as well.) Let G_i be the set of grid points generated over all faces f , and note that

$$|G_i| = O((n^{1/d}/2^i)^{d-1}) = O(n^{1-1/d}/2^{i(d-1)}).$$




FIG. 2. The grid G_i used in the proof of Theorem 2.6. The grid points are shown in black and gray; only the hypercubes H_g of the gray grid points are shown.

For each $g \in G_i$, let H_g denote the axis-aligned hypercube of size $2^{i+1}/n^{1/d}$ centered at g ; see Figure 2. Let $\mathcal{H}_i \stackrel{\text{def}}{=} \{H_g \mid g \in G_i\}$ be the set of all these hypercubes. Note that for any segment $s \in S_{\text{mid}}(i)$ there is a hypercube $H_g \in \mathcal{H}_i$ that contains s . Furthermore, all points in any H_g have distance at most $2^i/n^{1/d}$ from σ , and so $P \cap H_g \subseteq P_i(\sigma)$ for all g .

Now let n_g denote the number of points from P inside H_g . Since each point $p \in P$ is contained in a constant (depending on d) number of hypercubes H_g , we have $\sum_{g \in G_i} n_g = O(|P_i(\sigma)|)$. Furthermore, by packing property (PP2) we know that a hypercube H_g can contain only $O(1)$ segments from $S_{\text{mid}}(i)$. Thus,

$$\begin{aligned} |S_{\text{mid}}(i)| &= O(\text{number of nonempty hypercubes } H_g) \\ &= O(\min(|G_i|, |P_i(\sigma)|)) \\ &= O(\min(n^{1-1/d}/2^{i(d-1)}, |P_i(\sigma)|)). \end{aligned}$$

For $i < 0$ we have $|P_i(\sigma)| = O((3/2)^i n^{1-1/d})$, which implies

$$\begin{aligned} |S_{\text{mid}}| &= \sum_i |S_{\text{mid}}(i)| \\ &= O\left(\sum_{i < 0} (3/2)^i n^{1-1/d} + \sum_{i \geq 0} n^{1-1/d}/2^{i(d-1)}\right) = O(n^{1-1/d}). \end{aligned}$$

Now consider the number of ways in which we can choose $S_{\text{mid}}(i)$. In a hypercube H_g we have $O(1)$ edges which we can choose in $n_g^{O(1)}$ ways. Hence, if $G_i^* \subseteq G_i$ denotes the collection of grid points g such that $n_g > 0$ (in other words, such that H_g is nonempty), then

$$\text{total number of ways to choose } S_{\text{mid}}(i) = \prod_{g \in G_i^*} n_g^{O(1)} = 2^{O(\sum_{g \in G_i^*} \log n_g)}.$$

We bound $\sum_{g \in G_i^*} \log n_g$ separately for $i \geq 0$ and $i < 0$.

First, consider the case $i \geq 0$. Here we have

$$\sum_{g \in G_i^*} n_g = O(|P_i(\sigma)|) = O(4^i n^{1-1/d}).$$

Moreover, $|G_i^*| \leq |G_i| = O\left(\frac{n^{1-1/d}}{2^{i(d-1)}}\right)$ and $\frac{n^{1-1/d}}{2^{i(d-1)}} \leq 4^i n^{1-1/d}$. Since G_i^* are the nonempty cells, we have $|G_i^*| < c|P_i(\sigma)|$ for all i , where c is the maximum number of cubes in \mathcal{H}_i with a nonempty common intersection. Note that c is a constant depending only on d . Hence,

$$\begin{aligned} \sum_{g \in G_i^*} \log n_g &\leq |G_i^*| \cdot \log \left(\frac{|P_i(\sigma)|}{|G_i^*|} \right) \\ &< |G_i^*| \cdot \log \left(\frac{c \cdot e \cdot |P_i(\sigma)|}{|G_i^*|} \right) \\ &= O \left(\frac{n^{1-1/d}}{2^{i(d-1)}} \cdot \log 2^{i(d+1)} \right) \\ &= O \left(\frac{i(d+1)}{2^{i(d-1)}} \cdot n^{1-1/d} \right). \end{aligned}$$

In this formula, the first step follows from the AM-GM inequality. The third step uses the fact that $x \log(c \cdot e \cdot |P_i(\sigma)|/x)$ is monotone increasing for $x \in (0, c|P_i(\sigma)|)$, and therefore we can replace $|G_i^*|$ with $|G_i|$ (since $|G_i^*| < |G_i|$).

Now consider the case $i < 0$. The number of points in the hypercubes is

$$\sum_{g \in G_i^*} n_g = O(|P_i(\sigma)|) = O((3/2)^i n^{1-1/d}).$$

Because for $i < 0$ we have $n^{1-1/d}/2^{i(d-1)} > (3/2)^i n^{1-1/d}$, the number of points to distribute is smaller than the number of available hypercubes, and so $\sum_{g \in G_i^*} \log n_g$ is maximized when $n_g = 2$ for all g (except for at most one grid point g). Hence,

$$\sum_{g \in G_i^*} \log n_g = O(|P_i(\sigma)|) = O((3/2)^i n^{1-1/d}).$$

Thus the total number of ways in which we can choose S_{mid} is bounded by the following expression, where i ranges from $i_{\min} \stackrel{\text{def}}{=} -\log_{3/2} n^{1-1/d} + 1$ to $i_{\max} \stackrel{\text{def}}{=} \log n^{1/d}$:

$$\begin{aligned} &\prod_i \exp \left(O \left(\sum_{g \in G_i^*} \log n_g \right) \right) \\ &= \prod_{i < 0} \exp \left(O \left((3/2)^i n^{1-1/d} \right) \right) \prod_{i \geq 0} \exp \left(O \left(\frac{i(d+1)}{2^{i(d-1)}} \cdot n^{1-1/d} \right) \right) \\ &= \exp \left(O \left(\sum_{i < 0} (3/2)^i n^{1-1/d} + \sum_{i \geq 0} \frac{i(d+1)}{2^{i(d-1)}} \cdot n^{1-1/d} \right) \right) \\ &= 2^{O(n^{1-1/d})}. \end{aligned}$$

This concludes the proof of Claim 2.8. □

In order to enumerate candidates for S_{mid} , we can combine all segment choices from the hypercubes \mathcal{H}_i for each i that were considered in the above claim. There are $2^{O(n^{1-1/d})}$ possible combinations, and they can be enumerated with polynomial delay.

Combining these sets with the candidates for S_{short} and S_{long} (we already showed how to enumerate these earlier) yields the desired collection $\mathcal{C}'(\sigma, P)$. The size of any combination of three candidate subsets $S_{\text{short}}, S_{\text{mid}}, S_{\text{long}}$ is $O(1) + O(n^{1-1/d}) + O(1)$, and the total number of combinations is $n^{O(1)} \cdot 2^{O(n^{1-1/d})} \cdot n^{O(1)} = 2^{O(n^{1-1/d})}$. The time to enumerate all combinations is asymptotically the same. □

3. An exact algorithm for TSP. In this section we design an exact algorithm for TSP using the separator theorem from the previous section. As a first step, let us take a look at the TSP problem in \mathbb{R}^2 . The separator theorem from the previous section provides us with a separator σ such that the number of segments from an optimal tour that cross σ is $O(n^{1-1/d})$. Moreover, the number of candidate subsets $S \in \mathcal{C}(\sigma, P)$ that we need to try is only $2^{O(n^{1-1/d})}$. We can now obtain a divide-and-conquer algorithm similar to the algorithms of [18, 28] in a relatively standard manner. As we shall see, however, the resulting algorithm would still not run in $2^{O(n^{1-1/d})}$ time. We will therefore need to modify the algorithm and employ the so-called rank-based approach [6] to get our final result. In what follows, we describe an exact algorithm for TSP in \mathbb{R}^d .

A separator-based divide-and-conquer algorithm for EUCLIDEAN TSP works as follows. We first compute a separator using Theorem 2.6 for the given point set. For each candidate subset of edges crossing the separator, we then need to solve a subproblem for the points inside the separator and one for the points outside the separator. In these subproblems, we are no longer searching for a shortest tour but rather a collection of paths of minimum total length that connect the endpoints of the edges crossing the separator in a suitable manner.

To define the subproblems more precisely, let P be a point set, and let M be a perfect matching on a set $B \subseteq P$ of so-called *boundary points*. We say that a collection $\mathcal{P} = \{\pi_1, \dots, \pi_{|B|/2}\}$ of paths *realizes M on P* if (i) for each pair $(p, q) \in M$ there is a path $\pi_i \in \mathcal{P}$ with p and q as endpoints, and (ii) the paths together visit each point $p \in P$ exactly once. We define the length of a path π_i to be the sum of the Euclidean lengths of its edges, and we define the total length of \mathcal{P} to be the sum of the lengths of the paths $\pi_i \in \mathcal{P}$. The subproblems that arise in our divide-and-conquer algorithm can now be defined as follows (see Figure 3 for an illustration):

EUCLIDEAN PATH COVER

Input: A point set $P \subset \mathbb{R}^d$, a set of boundary points $B \subseteq P$, and a perfect matching M on B .

Question: Find a collection of paths covering P of minimum total length that realizes M on P .

Note that the matching in M does *not* specify edges that must be in the final tour, but it specifies which pairs of points form the endpoints of the paths in the cover. Also note that we can solve EUCLIDEAN TSP on a point set P by creating a copy p' of an arbitrary point $p \in P$, and then can solve EUCLIDEAN PATH COVER on $P \cup \{p'\}$ with $B \stackrel{\text{def}}{=} \{p, p'\}$ and $M \stackrel{\text{def}}{=} \{(p, p')\}$.

We will solve this initial instance of EUCLIDEAN PATH COVER recursively, as explained in detail below. Roughly speaking, we will compute a separator σ together




FIG. 3. (i) An instance of EUCLIDEAN PATH COVER with $B = \{p, q, p', q'\}$ and $M = \{(p, p'), (q, q')\}$. (ii) Solution for the instance from (i).

with its candidate sets $\mathcal{C}'(\sigma, P)$ of edges crossing the separator, solve suitable recursive problems inside and outside the separator, and then glue the solutions to the subproblems together to obtain a solution to the overall problem. (Our actual algorithm will be more involved and use the rank-based approach to deal with the matching in an efficient manner.) Note that the “correct” candidate set $S \in \mathcal{C}'(\sigma, P)$ leads to subproblems whose solutions are part of an optimal TSP tour and, hence, satisfy the packing property. “Wrong” candidate sets, however, lead to other subproblems, and it is not clear if an optimal solution to these subproblems satisfies the packing property. Hence, our algorithm may not find an optimal solution to such an instance of EUCLIDEAN PATH COVER. It might even happen that no valid solution satisfying the packing property exists for certain subproblems. Fortunately, this is not a problem, because of the following two properties: (i) we know that there will be a sequence of recursive calls that are consistent with an optimal solution for the original TSP problem, and this sequence will be solved correctly; (ii) the solution (if any) that is reported for any of the subproblems is a valid solution. Thus we are guaranteed of finding an optimal TSP tour and that no shorter but invalid tour can be reported. With this issue out of the way, we now describe our algorithm in detail.

An instance of EUCLIDEAN PATH COVER is solved by a separator-based recursive algorithm as follows. Let (P, B, M) be an instance of EUCLIDEAN PATH COVER, and let σ be a separator for P . We consider each candidate set $S \in \mathcal{C}'(\sigma, P)$ of edges crossing the separator σ . In fact, it is sufficient to consider candidate sets where the number of segments from S incident to any point $p \in P \setminus B$ is at most two, and the number of segments from S incident to any point in B is at most one.

We now wish to define subproblems for σ_{in} and σ_{out} (the regions inside and outside σ , respectively), the combination of which yields a solution for the given problem on P . Let $P_1(S) \subseteq P$ denote the set of endpoints with precisely one incident segment from S , and let $P_2(S) \subseteq P$ be the set of endpoints with precisely two incident segments from S . Note that in a solution to the problem (P, B, M) , the points in B need one incident edge—they must become endpoints of a path—while points in $P \setminus B$ need two incident edges. This means that the points in $B \cap P_1(S)$ and the points in $P_2(S)$ now have the desired number of incident edges, so they can be ignored in the subproblems. Points in $B \Delta P_1(S) \stackrel{\text{def}}{=} (B \setminus P_1(S)) \cup (P_1(S) \setminus B)$ still need one incident edge, while the remaining points in $P \setminus ((B \cap P_1(S)) \cup P_2(S))$ still need two incident edges. Hence, for σ_{in} we obtain subproblems of the form $(P_{\text{in}}, B_{\text{in}}, M_{\text{in}})$ where

$$(3.1) \quad \begin{aligned} P_{\text{in}} &\stackrel{\text{def}}{=} \left(P \setminus ((B \cap P_1(S)) \cup P_2(S)) \right) \cap \sigma_{\text{in}}, \\ B_{\text{in}} &\stackrel{\text{def}}{=} (B \Delta P_1(S)) \cap \sigma_{\text{in}}, \\ M_{\text{in}} &\text{ is a perfect matching on } B_{\text{in}}. \end{aligned}$$

See Figure 4. For σ_{out} we obtain subproblems of the form $(P_{\text{out}}, B_{\text{out}}, M_{\text{out}})$ defined in a similar way. As already noted, we can restrict our attention to candidate sets $S \in \mathcal{C}'(\sigma, P)$ that contain at most one edge incident to any given point in B , and at most two edges incident to any given point in $P \setminus B$. Moreover, S should be such that $|B_{\text{in}}|$ and $|B_{\text{out}}|$ are even. We define $\mathcal{C}^*(\sigma, P)$ to be the family of candidate sets $\mathcal{C}'(\sigma, P)$ restricted in this way. Also note that while $P_{\text{in}}, B_{\text{in}}$ and $P_{\text{out}}, B_{\text{out}}$ are determined once S is fixed, the algorithm has to find the best matchings M_{in} and M_{out} . These matchings together should realize the matching M on P , and, among all such matchings, we want the pair that leads to a minimum-length solution.




FIG. 4. (i) A EUCLIDEAN PATH COVER instance and a separator σ . Blue squares indicate boundary points. (ii) The subproblem generated inside σ consists of the black points and blue squares.

The number of perfect matchings on k points is $k^{\Theta(k)}$. Unfortunately, in the first call of the recursive algorithm, $|B_{\text{in}}|$ and $|B_{\text{out}}|$ already can be as large as $\Theta(n^{1-1/d})$. Hence, recursively checking all matchings will not lead to an algorithm with the desired running time. In \mathbb{R}^2 we can use the fact that an optimal TSP tour is crossing-free, so it is sufficient to look for “crossing-free matchings,” of which there are only $2^{O(k)}$. (This approach would actually require a different setup of the subproblems; see the papers by Deineko, Klinz, and Woeginger [10] and Dorn et al. [11].) However, the crossing-free property has no analogue in higher dimensions, and it does not hold in \mathbb{R}^2 for our “almost-Euclidean” setting either. Hence, we need a different approach to rule out a significant proportion of the available matchings.

Applying the rank-based approach. Next, we describe how we can use the rank-based approach [6, 9] in our setting. We try to avoid the intricate notation introduced in the original papers, but our terminology is mostly compatible with [6]. A standard application of the rank-based approach works on a tree-decomposition of the underlying graph, where the bags represent vertex separators. In our application the underlying graph is a complete graph on the points—all segments are potentially segments of the TSP tour—and we have to use a separator for the edges in the solution.

Let P be a set of points in \mathbb{R}^d , and let $B \subseteq P$ be a set of boundary points such that $|B|$ is even. Let $\mathcal{M}(B)$ denote the set of all perfect matchings on B , and consider a matching $M \in \mathcal{M}(B)$. We can turn M into a weighted matching by assigning to it the minimum total length of any solution realizing M . In other words, $\text{weight}(M)$ is the length of the solution of EUCLIDEAN PATH COVER for input (P, B, M) . Whenever we speak of weighted matchings, we always mean perfect matchings on a set $B \subseteq P$ weighted as above. We use $\mathcal{M}(B, P)$ to denote the set of all such weighted matchings on B . Note that $|\mathcal{M}(B, P)| = |\mathcal{M}(B)| = 2^{O(|B| \log |B|)}$. The key to reducing the number of matchings we have to consider is the concept of representative sets, as explained next.

We say that two matchings $M, M' \in \mathcal{M}(B)$ *fit* if their union is a Hamiltonian cycle on B . Consider a pair P, B . Let \mathcal{R} be a set of weighted matchings on B , and let M be another matching on B . We define

$$\text{opt}(M, \mathcal{R}) \stackrel{\text{def}}{=} \min\{\text{weight}(M') \mid M' \in \mathcal{R} \text{ and } M' \text{ fits } M\},$$

that is, $\text{opt}(M, \mathcal{R})$ is the minimum total length of any collection of paths on P that, together with the matching M , forms a cycle.

A set $\mathcal{R} \subseteq \mathcal{M}(B, P)$ of weighted matchings is defined to be *representative* of another set $\mathcal{R}' \subseteq \mathcal{M}(B, P)$ if for any matching $M \in \mathcal{M}(B)$ we have $\text{opt}(M, \mathcal{R}) = \text{opt}(M, \mathcal{R}')$. Note that our algorithm is not able to compute a representative set of $\mathcal{M}(B, P)$, because it is also restricted by the packing property, while a solution of EUCLIDEAN PATH COVER for a generic P, B, M may not satisfy it. Let $\mathcal{M}_{\text{PP}}(B, P)$ denote the set of weighted matchings in $\mathcal{M}(B, P)$ that have a corresponding EUCLIDEAN PATH COVER solution satisfying the packing property.

The basis of the rank-based method is the following result.

LEMMA 3.1 (Bodlaender et al. [6, Theorem 3.7]). *There exists a set \mathcal{R}^* consisting of $2^{|B|-1}$ weighted matchings that is representative of the set $\mathcal{M}(B, P)$. Moreover, there is an algorithm Reduce that, given a representative set \mathcal{R} of $\mathcal{M}(B, P)$, computes such a set \mathcal{R}^* in $|\mathcal{R}| \cdot 2^{O(|B|)}$ time.*

Lemma 3.1 also holds for our case, where \mathcal{R} is representative of $\mathcal{M}_{\text{PP}}(B, P)$. The result of Bodlaender et al. is actually more general than stated above, as it applies not only to matchings but also to other types of partitions. Moreover, for matchings, the bound has been improved to $2^{|B|/2-1}$ [9]. However, Lemma 3.1 suffices for our purposes.

We now briefly sketch the main ideas behind the rank-based approach. Consider a matrix \mathcal{A} whose rows and columns are indexed by perfect matchings of B . One can think of the row index as the matching realized by the subproblem inside the separator, and the column index as the matching realized by the subproblem outside the separator. An entry in \mathcal{A} is 1 if the pair of corresponding matchings fit, that is, their union is a Hamiltonian cycle on B , and the entry is 0 otherwise.

Consider a set X of rows in \mathcal{A} that is linearly dependent over the field $\text{GF}(2)$. Suppose that M is a matching in X of largest weight. Bodlaender et al. [6] show that $X \setminus \{M\}$ is a representative set for the set X —this is because every matching (column) that can be combined with M can also be combined with a matching (row) in $X \setminus \{M\}$ to a solution (of equal or lower weight). Thus, when we have a set of rows in \mathcal{A} that is linearly dependent over the field $\text{GF}(2)$, the heaviest matching in this set is redundant and can be removed. Consequently, it is sufficient to compute a minimum weight row basis of \mathcal{A} , as the matchings corresponding to the rows in such a basis form a representative set. The basis can be computed using standard linear algebra; this is done by the Reduce algorithm of Lemma 3.1. In order to bound the size of the obtained representative set, it is enough to bound the rank of \mathcal{A} . Bodlaender et al. prove that \mathcal{A} has rank at most $2^{|B|-1}$ by providing a factorization $\mathcal{A} = \hat{A}^T \hat{A}$ where \hat{A} has $2^{|B|-1}$ rows.

Lemma 3.1 bounds the size of the representative set in terms of $|B|$, the number of boundary points. In the first call of our algorithm $|B| = O(n^{1-1/d})$ because of the properties of our separator, but we have to be careful that the size of B stays under control in recursive calls.

Algorithm 3.1 describes how we deal with this. A key step in the algorithm is step 4, where we invoke the balance condition of the separator with respect to B or P depending on the size of B relative to the size of P . (The constant γ will be specified in the analysis of the running time.) Steps 8–12 combine the representative sets \mathcal{R}_{in} and \mathcal{R}_{out} . Next, we explain these steps in more detail.

Consider a set $S \in \mathcal{C}^*(\sigma, P)$, a matching $M_{\text{in}} \in \mathcal{M}(B_{\text{in}})$, and a matching $M_{\text{out}} \in \mathcal{M}(B_{\text{out}})$. Let $G = G_S(M_{\text{in}}, M_{\text{out}})$ be the graph with vertex set $V(G) \stackrel{\text{def}}{=} B \cup P_1(S) \cup P_2(S)$ and edge set $E(G) \stackrel{\text{def}}{=} M_{\text{in}} \cup M_{\text{out}} \cup S$. We say that M_{in} and M_{out} are *compatible*

Algorithm 3.1 *TSP-Repr*(P, B).

Input: A set P of points in \mathbb{R}^d and a subset $B \subseteq P$
Output: A set $\mathcal{R} \subset \mathcal{M}(B, P)$ with $|\mathcal{R}| \leq 2^{|B|-1}$ representing $\mathcal{M}_{PP}(B, P)$

- 1: **if** $|P| \leq 1$ **then return** $\{(\emptyset, 0)\}$
- 2: $\mathcal{R} \leftarrow \emptyset$
- 3: Compute σ (by Theorem 2.6) with $Q=P$ if $|B| \leq \gamma|P|^{1-1/d}$ and $Q=B$ otherwise.
- 4: **for all** candidate sets $S \in \mathcal{C}^*(\sigma, P)$ **do**
- 5: $\mathcal{R}_{\text{in}} \leftarrow \text{TSP-Repr}(P_{\text{in}}, B_{\text{in}})$, where P_{in} and B_{in} are defined as in (3.1)
- 6: $\mathcal{R}_{\text{out}} \leftarrow \text{TSP-Repr}(P_{\text{out}}, B_{\text{out}})$, where P_{out} and B_{out} are defined as in (3.1)
- 7: **for all** $M_{\text{in}} \in \mathcal{R}_{\text{in}}$ and $M_{\text{out}} \in \mathcal{R}_{\text{out}}$ **do**
- 8: **if** M_{in} and M_{out} are compatible **then**
- 9: Add $(\text{Join}_S(M_{\text{in}}, M_{\text{out}}), \text{weight}(M_{\text{in}}) + \text{length}(S) + \text{weight}(M_{\text{out}}))$ to \mathcal{R}
- 10: $\mathcal{R} \leftarrow \text{Reduce}(\mathcal{R})$
- 11: **return** \mathcal{R}

if G consists of $|B|/2$ disjoint paths covering $V(G)$ whose endpoints are exactly the points in B . A pair of compatible matchings induces a perfect matching on B , where for each of these $|B|/2$ paths we add a matching edge between its endpoints. We denote this matching by $\text{Join}_S(M_{\text{in}}, M_{\text{out}}) \in \mathcal{M}(B)$. To get a set \mathcal{R} of weighted matchings on B we thus iterate in steps 8–10 through all pairs $M_{\text{in}}, M_{\text{out}}$ where M_{in} and M_{out} are compatible, and for such pairs, we add to \mathcal{R} the matching $\text{Join}_S(M_{\text{in}}, M_{\text{out}})$. The weight of this matching is $\text{weight}(M_{\text{in}}) + \text{length}(S) + \text{weight}(M_{\text{out}})$.

Recall that the initial call to *TSP-Repr*(P, B) is done by creating a copy p' of an arbitrary point $p \in P$, and then solving EUCLIDEAN PATH COVER on $P \cup \{p'\}$ with $B \stackrel{\text{def}}{=} \{p, p'\}$ and $M \stackrel{\text{def}}{=} \{(p, p')\}$. This solves EUCLIDEAN TSP as desired, since the representative set \mathcal{R} consists of a single matching $\{(p, p')\}$ whose weight is the optimum Euclidean TSP tour of P . One can use standard techniques to obtain the optimum tour itself with a slightly modified algorithm. The correctness of our algorithm is implied by the following claim.

CLAIM 3.2. *The set \mathcal{R} created in steps 4–9 of Algorithm 3.1 is representative of $\mathcal{M}_{PP}(B, P)$.*

Proof. The proof is by induction on $|P|$. Clearly, for $|P| \leq 1$ the claim holds. Otherwise, let $S \in \mathcal{C}^*(\sigma, P)$ be fixed. The set S is considered in some iteration of the outer loop. Define $P_{\text{in}}, P_{\text{out}}, B_{\text{in}}, B_{\text{out}}$ as in this iteration, and let $\mathcal{R}_{\text{in}}, \mathcal{R}_{\text{out}}$ be the sets returned by the recursive calls, which are representative sets of $\mathcal{M}_{PP}(B_{\text{in}}, P_{\text{in}})$ and $\mathcal{M}_{PP}(B_{\text{out}}, P_{\text{out}})$, respectively, by induction. Notice that S can be regarded as a EUCLIDEAN PATH COVER solution for $(B_{\text{in}} \cup B_{\text{out}}, B_{\text{in}} \cup B_{\text{out}}, M_S)$, where M_S is the matching realized by S on $B_{\text{in}} \cup B_{\text{out}}$. Let $\text{length}(S)$ be the weight assigned to M_S . Clearly $\{M_S\}$ is representative of $\{M_S\}$. Now our Join_S operation can be regarded as the succession of two join operations as defined by [6], applied first to $\{M_S\}$ and \mathcal{R}_{in} and then to the result and \mathcal{R}_{out} . By Lemma 3.6 in [6], the join operation preserves representation, and therefore the matchings added to \mathcal{R} in this iteration of the outer loop form a representative set of

$$\widehat{\mathcal{M}}_S \stackrel{\text{def}}{=} \{ \text{Join}_S(M_{\text{in}}, M_{\text{out}}) \mid M_{\text{in}} \in \mathcal{M}_{PP}(B_{\text{in}}, P_{\text{in}}), \\ M_{\text{out}} \in \mathcal{M}_{PP}(B_{\text{out}}, P_{\text{out}}) \}.$$

Consequently, the set \mathcal{R} that is created at the end of the outer loop is a representative set of $\widehat{\mathcal{M}} \stackrel{\text{def}}{=} \bigcup_{S \in \mathcal{C}^*(P, \sigma)} \widehat{\mathcal{M}}_S$. The set $\widehat{\mathcal{M}}_S$ contains the subset of $\mathcal{M}_{\text{PP}}(B, P)$ that has a corresponding optimum with the packing property that intersects σ in S , because for any such optimum path cover \mathcal{P} , the subpaths of \mathcal{P} induced by P_{in} also have the packing property, and they form an optimal EUCLIDEAN PATH COVER for the input $(P_{\text{in}}, B_{\text{in}}, M_{\text{in}})$, i.e., there is a corresponding weighted matching $M_{\text{in}} \in \mathcal{M}_{\text{PP}}(B_{\text{in}}, P_{\text{in}})$. (The analogous statement is true for the subpaths induced by P_{out} .) Since $\mathcal{C}^*(\sigma, P)$ contains all sets S that can arise as the set of segments intersecting σ in an optimum EUCLIDEAN PATH COVER solution with the packing property, it follows that $\widehat{\mathcal{M}} \supseteq \mathcal{M}_{\text{PP}}(B, P)$, which concludes the proof of the claim. \square

Notice that steps 5–10 can be implemented using a brute-force algorithm; this takes

$$O(|\mathcal{R}_{\text{in}}| \cdot |\mathcal{R}_{\text{out}}| \cdot \text{poly}(|B| + |S|))$$

time. However, by combining \mathcal{R}_{in} and \mathcal{R}_{out} in this way, the size of \mathcal{R} may be more than $2^{|B|-1}$. Hence, we apply the *Reduce* algorithm [6] to create a representative set of size at most $2^{|B|-1}$ in $|\mathcal{R}| \cdot 2^{O(|B|)}$ time. Since our recursive algorithm ensures that $|\mathcal{R}_{\text{in}}| \leq 2^{|B_{\text{in}}|-1}$ and $|\mathcal{R}_{\text{out}}| \leq 2^{|B_{\text{out}}|-1}$, all of the above steps run in $2^{O(|B|+|S|)} = \exp(O(|B| + |P|^{1-1/d}))$ time.

3.1. Detailed analysis of the running time. In this section, we provide a detailed analysis of the running time of our algorithm.

The running time of *TSP-Repr*(P, B) essentially satisfies the following recurrence, where c_0, c_1, c_2 are positive constants, and we use the notation $n \stackrel{\text{def}}{=} |P|$ and $b \stackrel{\text{def}}{=} |B|$:

$$T(n, b) \leq \begin{cases} c_0 & \text{if } n \leq 1, \\ 2^{c_1(n^{1-1/d}+b)} T(\delta n, b + c_2 n^{1-1/d}) & \text{if } b \leq \gamma n^{1-1/d}, \\ 2^{c_1(n^{1-1/d}+b)} T(n, \delta b + c_2 n^{1-1/d}) & \text{if } b > \gamma n^{1-1/d}. \end{cases}$$

The actual recurrence is a bit more subtle, as explained next.

For each $S \in \mathcal{C}^*(\sigma, P)$, let $n_{S,\text{in}} \stackrel{\text{def}}{=} |P_{\text{in}}|$, let $b_{S,\text{in}} \stackrel{\text{def}}{=} |B_{\text{in}}|$, let $n_{S,\text{out}} \stackrel{\text{def}}{=} |P_{\text{out}}|$, and let $b_{S,\text{out}} \stackrel{\text{def}}{=} |B_{\text{out}}|$. By the discussion in section 3, we can bound the running time of the two inner loops, the *Reduce* algorithm, and the rest of the operations outside the recursive calls by $\exp(c_3(n^{1-1/d} + b))$ for some positive constant c_3 . Therefore, the algorithm *TSP-Repr*(P, B) obeys the following recursion, where σ_P and σ_B are separators balanced with respect to P and B .

$$T(n, b) \leq \begin{cases} c_0 & \text{if } n \leq 1, \\ \sum_{S \in \mathcal{C}^*(\sigma_P, P)} (\exp(c_3(n^{1-1/d} + b)) + T(n_{S,\text{in}}, b_{S,\text{in}}) + T(n_{S,\text{out}}, b_{S,\text{out}})) & \text{if } b \leq \gamma n^{1-1/d}, \\ \sum_{S \in \mathcal{C}^*(\sigma_B, P)} (\exp(c_3(n^{1-1/d} + b)) + T(n_{S,\text{in}}, b_{S,\text{in}}) + T(n_{S,\text{out}}, b_{S,\text{out}})) & \text{if } b > \gamma n^{1-1/d}. \end{cases}$$

Note that the terms in the second and third cases are the same, except that the second case uses a separator σ_P , while the third case uses a separator σ_B . This in turn will influence the bounds on $n_{S,\text{in}}$ and $b_{S,\text{in}}$ (and, similarly, $n_{S,\text{out}}$ and $b_{S,\text{out}}$).

LEMMA 3.3. *For a suitable choice of γ in step 4 of *TSP-Repr* we have $T(n, 2) = 2^{O(n^{1-1/d})}$.*

Proof. We prove by induction that $T(n, b) \leq \exp(d_1 n^{1-1/d} + d_2 b)$ for some constants d_1 and d_2 and for all $1 \leq b \leq n$. This clearly holds for $b, n \leq 1$, when, as we will ensure, we have $d_1, d_2 \geq c_0$. Hence, by induction, for each S we have

$$\begin{aligned} & \exp(c_3(n^{1-1/d} + b)) + T(n_{S,\text{in}}, b_{S,\text{in}}) + T(n_{S,\text{out}}, b_{S,\text{out}}) \\ & \leq \exp\left(c_3(n^{1-1/d} + b)\right) + \exp\left(d_1 n_{S,\text{in}}^{1-1/d} + d_2 b_{S,\text{in}}\right) + \exp\left(d_1 n_{S,\text{out}}^{1-1/d} + d_2 b_{S,\text{out}}\right) \\ & \leq \exp\left(c_3(n^{1-1/d} + b)\right) + 2 \exp\left(d_1 n_{S,\text{max}}^{1-1/d} + d_2 b_{S,\text{max}}\right), \end{aligned}$$

where the additions were replaced by multiplications. Note that for $x, y \geq 2$ we have $\exp(x) + 2\exp(y) \leq \exp(x + y)$. Therefore, since $c_3 > 2$ and we may assume that $d_1, d_2 > 2$, we can conclude that

$$\begin{aligned} & \exp(c_3(n^{1-1/d} + b)) + T(n_{S,\text{in}}, b_{S,\text{in}}) + T(n_{S,\text{out}}, b_{S,\text{out}}) \\ & \leq \exp\left(c_3(n^{1-1/d} + b) + d_1 n_{S,\text{max}}^{1-1/d} + d_2 b_{S,\text{max}}\right). \end{aligned}$$

Let c_2 and c_4 be the constants from parts (ii) and (iii) of Theorem 2.6. For any $S \in \mathcal{C}^*(\sigma_B, P)$, we have $b_{S,\text{in}} \leq \delta b + c_2 n^{1-1/d}$ and $b_{S,\text{out}} \leq \delta b + c_2 n^{1-1/d}$; similarly, for any $S \in \mathcal{C}^*(\sigma_P, P)$, we have $n_{S,\text{in}} \leq \delta n$ and $n_{S,\text{out}} \leq \delta n$. Note that we always have the trivial bounds $b_{S,\text{max}} \leq b + c_2 n^{1-1/d}$ and $n_{S,\text{max}} \leq n$ as well. Since $|\mathcal{C}^*(\sigma, P)| \leq \exp(c_4 n^{1-1/d})$, we get the following:

$$T(n, b) \leq \begin{cases} c_0 & \text{if } n \leq 1, \\ \exp\left(c_1(n^{1-1/d} + b) + d_1(\delta n)^{1-1/d} + d_2(b + c_2 n^{1-1/d})\right) & \text{if } b \leq \gamma n^{1-1/d}, \\ \exp\left(c_1(n^{1-1/d} + b) + d_1 n^{1-1/d} + d_2(\delta b + c_2 n^{1-1/d})\right) & \text{if } b > \gamma n^{1-1/d}, \end{cases}$$

where $c_1 = c_3 + c_4$. We set $c \stackrel{\text{def}}{=} \max(c_1, c_2)$ and let $\gamma \stackrel{\text{def}}{=} \frac{2c}{1-\delta}$. (Notice that the definition of γ here is valid: it is independent of d_1 and d_2 .)

If $b \leq \gamma n^{1-1/d} = \frac{2c}{1-\delta} n^{1-1/d}$, we have the following:

$$\begin{aligned} T(n, b) & \leq \exp\left(c n^{1-1/d} + c b + d_1(\delta n)^{1-1/d} + d_2 b + d_2 c n^{1-1/d}\right) \\ & \leq \exp\left(\left(c + \frac{2c}{1-\delta} + d_1 \delta^{1-1/d} + d_2 c\right) n^{1-1/d} + d_2 b\right) \\ & \leq \exp(d_1 n^{1-1/d} + d_2 b), \end{aligned}$$

where the second inequality uses $b \leq \frac{2c}{1-\delta} n^{1-1/d}$, and the third uses

$$c + \frac{2c}{1-\delta} + d_1 \delta^{1-1/d} + d_2 c \leq d_1.$$

This can be ensured by setting

$$d_1 \stackrel{\text{def}}{=} \max\left(c_0, \frac{c + 2c/(1-\delta) + d_2 c}{1 - \delta^{1-1/d}}\right).$$

Note that $0 < \delta < 1$ and $c > 0$ are fixed constants. Moreover, d_2 will be a positive constant as well; see below. Hence, d_1 is a positive constant.

Finally, if $b > \gamma n^{1-1/d} = \frac{2c}{1-\delta} n^{1-1/d}$, we have the following:

$$\begin{aligned} T(n, b) &\leq \exp(cn^{1-1/d} + cb + d_1 n^{1-1/d} + d_2 \delta b + d_2 c n^{1-1/d}) \\ &< \exp\left(d_1 n^{1-1/d} + \left(\frac{1-\delta}{2} + c + \delta d_2 + \frac{1-\delta}{2} d_2\right) b\right) \\ &\leq \exp(d_1 n^{1-1/d} + d_2 b), \end{aligned}$$

where the strict inequality uses $cn^{1-1/d} < \frac{1-\delta}{2} b$, and the final inequality uses

$$\frac{1-\delta}{2} + c + \frac{1+\delta}{2} d_2 \leq d_2.$$

We can ensure this by setting

$$d_2 \stackrel{\text{def}}{=} \max\left(c_0, \frac{(1-\delta)/2 + c}{1 - (1+\delta)/2}\right),$$

and therefore d_2 is a positive constant. Since there exists positive constants d_1 and d_2 satisfying the above inequalities, we have that $T(n, b) \leq \exp(d_1 n^{1-1/d} + d_2 b)$, and, in particular, for the initial call we have $T(n, 2) = 2^{O(n^{1-1/d})}$. \square

4. Almost Euclidean TSP. So far, we considered EUCLIDEAN TSP in the real-RAM model of computation. We now consider a slightly more general scenario in the Word-RAM model. Here we assume that the input is a set P of n points in \mathbb{R}^d , specified by rational coordinates, as well as a distance matrix D . The basic assumption we make is that the distances in D approximate the real Euclidean distances well. More precisely, we require that the ordering of pairwise distances on the given point set $P \stackrel{\text{def}}{=} \{p_1, \dots, p_n\}$ be preserved: if $|p_i p_j| < |p_k p_\ell|$, then $D[i, j] < D[k, \ell]$. We remark that the entries of D need not satisfy the triangle inequality. We call this the *almost Euclidean* version of TSP.

In order to show that our algorithms work in this setting, we only need to show that an optimal tour in this setting satisfies the packing property. Note that the packing property for the almost Euclidean version immediately implies that the packing property also holds for the Euclidean version (where, as noted in section 2, similar properties were already known). Recall that a set S of segments in \mathbb{R}^d has the *packing property* if for any separator σ we have

- (PP1): $|\{s \in S \mid s \text{ crosses } \sigma \text{ and } \text{length}(s) \geq \text{size}(\sigma)\}| = O(1)$,
- (PP2): $|\{s \in S \mid s \subset \sigma_{\text{in}} \text{ and } \text{length}(s) \geq \text{size}(\sigma)/4\}| = O(1)$.

THEOREM 4.1. *Let $P \stackrel{\text{def}}{=} \{p_1, \dots, p_n\}$ be a point set in \mathbb{R}^d , and let D be a distance matrix for P such that we have $|p_i p_j| < |p_k p_\ell|$ if and only if $D[i, j] < D[k, \ell]$. Let T be a tour on P that is optimal for the distances given by D . Then the set of edges of T has the packing property.*

Proof. We first prove packing property (PP1) and then argue that (PP2) follows from (PP1).

Let σ be a hypercube, and suppose without loss of generality that $\text{size}(\sigma) = 1$. Suppose for a contradiction that there are more than c tour edges of length at least 1 that cross σ , where c is a suitably large constant (which depends on d). By the pigeonhole principle, we can then find three edges in T such that (i) the pairwise Euclidean distance between the endpoints of these edges that lie inside σ_{in} is at most $1/10$, and (ii) the pairwise angle between these edges is at most $\pi/30$. Here the angle between two edges is measured as the smaller angle between two lines going through the origin and parallel to the given edges.




FIG. 5. Since $|p_i p_j| > |p_j p_\ell|$ and $|p_k p_\ell| > |p_i p_k|$, we can exchange $p_i p_j$ and $p_k p_\ell$ for $p_i p_k$ and $p_j p_\ell$ in the tour and get a shorter tour. Note that for $d > 2$ the points do not all have to lie in the same plane.

Now fix an orientation on the tour T such that at least two of the three edges cross σ from inside to outside, and orient these edges accordingly. Let $p_i p_j$ and $p_k p_\ell$ denote these two oriented edges; see Figure 5. Thus $p_i, p_k \in \sigma_{\text{in}}$ and $|p_i p_k| \leq 1/10$. Assume without loss of generality that $|p_i p_j| \geq |p_k p_\ell|$.

Our goal now is to show that $|p_k p_\ell| > |p_i p_k|$ and $|p_i p_j| > |p_j p_\ell|$; this will help us find a way to reroute the current tour into a shorter one, leading to a contradiction. The first inequality is easy: $|p_k p_\ell| > |p_i p_k|$ since $|p_k p_\ell| \geq 1$ and $|p_i p_k| \leq 1/10$. In order to show $|p_i p_j| > |p_j p_\ell|$, we proceed as follows.

In the triangle $p_i p_k p_\ell$, we have $|p_k p_\ell| \geq 1$ and $|p_i p_k| \leq 1/10$, so by the law of sines, $\angle(p_i p_\ell p_k) \leq \arcsin(1/10) < \pi/30$. Let $\bar{p}_\ell = p_i + (p_\ell - p_k)$. Here and in what follows we use the fact that a point can also be thought of as a vector, so we can add and subtract points to get new points. Then we have that $\angle(p_\ell p_i \bar{p}_\ell) = \angle(p_i p_\ell p_k) \leq \pi/30$. Due to our choice of $p_i p_j$ and $p_k p_\ell$, their angle is at most $\pi/30$, so we have

$$\angle(\bar{p}_\ell p_i p_j) = \angle((p_j - p_i), (p_\ell - p_k)) \leq \pi/30.$$

Therefore, $\angle(p_\ell p_i p_j) \leq \pi/15$. Note that this is also true if the points do not all lie in the same plane. Now observe that

$$|p_i p_\ell| \leq |p_k p_\ell| + |p_i p_k| \leq |p_i p_j| + 1/10,$$

that is, $p_i p_\ell$ cannot be much longer than $p_i p_j$. Thus, if we look at the triangle $p_i p_\ell p_j$, then we have $|p_i p_j| \geq 1$ and $\angle(p_\ell p_i p_j) \leq \pi/15$, and $|p_i p_\ell| \leq |p_i p_j| + 1/10 \leq \frac{11}{10} |p_i p_j|$.

By the law of cosines, we have

$$\begin{aligned} |p_\ell p_j|^2 &= |p_i p_j|^2 + |p_i p_\ell|^2 - 2|p_i p_j||p_i p_\ell| \cos(\angle(p_\ell p_i p_j)) \\ &= |p_i p_j|^2 + |p_i p_\ell|(|p_i p_\ell| - 2|p_i p_j| \cos(\angle(p_\ell p_i p_j))) \\ &< |p_i p_j|^2 + |p_i p_\ell| \left(\frac{11}{10} |p_i p_j| - 2 \cos\left(\frac{\pi}{15}\right) |p_i p_j| \right) \\ &< |p_i p_j|^2, \end{aligned}$$

where the first inequality uses our earlier bounds on $|p_i p_\ell|$ and $\angle(p_\ell p_i p_j)$, and the second uses that $\frac{11}{10} - 2 \cos(\frac{\pi}{15}) < 0$. This shows that $|p_i p_j| > |p_j p_\ell|$ as claimed earlier.

Because the ordering of the pairwise distances in the matrix D is the same as for the Euclidean distances, we can conclude that $D[i, j] > D[j, \ell]$ and $D[k, \ell] > D[i, k]$. But then we can exchange $p_i p_j$ and $p_k p_\ell$ for $p_i p_k$ and $p_j p_\ell$ in the tour T —because both edges are oriented from inside σ to outside σ this gives a valid tour—and get a shorter tour. This contradicts the minimality of the tour, concluding the proof of (PP1).

Property (PP2) is a direct consequence of (PP1). Indeed, if we cover σ_{in} by $O(1)$ hypercubes of diameter size $(\sigma_{\text{in}})/5$, then any segment of length at least $\text{size}(\sigma_{\text{in}})/4$ inside σ_{in} crosses at least one such hypercube, and by (PP1) each hypercube is crossed by $O(1)$ edges of length at least $\text{size}(\sigma_{\text{in}})/4$. \square

Notice that the above properties hold for squared Euclidean distances or, more generally, for any distance matrix D where $D[i, j] = f(|p_i p_j|)$, where f is a monotone increasing function. We get the following corollary.

COROLLARY 4.2. *Let d be a fixed integer, and let $f : \mathbb{R} \rightarrow \mathbb{R}$ be a monotone increasing function. Then for a set of n points $P = \{p_1, \dots, p_n\}$ in \mathbb{R}^d and distances $D[i, j] = f(|p_i p_j|)$, the shortest tour of P with respect to D (that is, the minimum-length Hamiltonian cycle) can be computed in $2^{O(n^{1-1/d})}$ time.*

Remark 4.3. It would be useful for applications if the algorithm could work with a distance matrix that is a constant distortion of the Euclidean distances, that is, a matrix D such that $(1/\alpha) \cdot |p_i p_j| \leq D[i, j] \leq \alpha |p_i p_j|$ for some constant $\alpha \geq 1$. Unfortunately, while (PP2) holds also in this scenario, (PP1) does not.

5. Concluding remarks. In this paper we described a new geometric separation technique, which resulted in a faster exact algorithm for EUCLIDEAN TSP. Together with the lower bound in [5], this settles the complexity of EUCLIDEAN TSP assuming ETH and up to the constant in the exponent.

We believe that our separation technique can be useful for other problems in Euclidean geometry as well and, in particular, for problems where one wishes to compute a minimum-length geometric structure that satisfies the packing property. An example of such a problem is RECTILINEAR STEINER TREE. An additional issue to overcome here is that the number of potential Steiner points is $O(n^d)$, which means that a direct application of our techniques does not work. Another challenging problem is finding the minimum-weight triangulation for a set of n points given in \mathbb{R}^2 , which was proven NP-hard by Mulzer and Rote [25] and for which an $n^{O(\sqrt{n})}$ algorithm is known [22]. A minimum-weight triangulation does not have the packing property, because of clusters of points that are far from one another, but finding an optimal triangulation among such clusters can perhaps be handled separately.

REFERENCES

- [1] S. ARORA, *Polynomial time approximation schemes for Euclidean traveling salesman and other geometric problems*, J. ACM, 45 (1998), pp. 753–782.
- [2] Y. BARTAL AND L. GOTTLIEB, *A linear time approximation scheme for Euclidean TSP*, in Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, IEEE Computer Society, 2013, pp. 698–706, <https://doi.org/10.1109/FOCS.2013.80>.
- [3] R. BELLMAN, *Dynamic programming treatment of the travelling salesman problem*, J ACM, 9 (1962), pp. 61–63, <https://doi.org/10.1145/321105.321111>.
- [4] M. DE BERG, H. L. BODLAENDER, S. KISFALUDI-BAK, AND S. KOLAY, *An ETH-tight exact algorithm for Euclidean TSP*, in Proceedings of the 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, IEEE Computer Society, 2018, pp. 450–461, <https://doi.org/10.1109/FOCS.2018.00050>.
- [5] M. DE BERG, H. L. BODLAENDER, S. KISFALUDI-BAK, D. MARX, AND T. C. VAN DER ZANDEN, *A framework for exponential-time-hypothesis-tight algorithms and lower bounds in geometric intersection graphs*, SIAM J. Comput., 49 (2020), pp. 1291–1331, <https://doi.org/10.1137/20M1320870>.
- [6] H. L. BODLAENDER, M. CYGAN, S. KRATSCH, AND J. NEDERLOF, *Deterministic single exponential time algorithms for connectivity problems parameterized by treewidth*, Inform. Comput., 243 (2015), pp. 86–111, <https://doi.org/10.1016/j.ic.2014.12.008>.

- [7] N. CHRISTOFIDES, *Worst-Case Analysis of a New Heuristic for the Travelling Salesman Problem*, Technical report, Graduate School of Industrial Administration, Carnegie Mellon University, 1976.
- [8] M. CYGAN, F. V. FOMIN, L. KOWALIK, D. LOKSHTANOV, D. MARX, M. PILIPCZUK, M. PILIPCZUK, AND S. SAURABH, *Parameterized Algorithms*, Springer, 2015.
- [9] M. CYGAN, S. KRATSCH, AND J. NEDERLOF, *Fast Hamiltonicity checking via bases of perfect matchings*, J. ACM, 65 (2018), 12, <https://doi.org/10.1145/3148227>.
- [10] V. G. DEINEKO, B. KLINZ, AND G. J. WOEGINGER, *Exact algorithms for the Hamiltonian cycle problem in planar graphs*, Oper. Res. Lett., 34 (2006), pp. 269–274, <https://doi.org/10.1016/j.orl.2005.04.013>.
- [11] F. DORN, E. PENNINKX, H. L. BODLAENDER, AND F. V. FOMIN, *Efficient exact algorithms on planar graphs: Exploiting sphere cut decompositions*, Algorithmica, 58 (2010), pp. 790–810, <https://doi.org/10.1007/s00453-009-9296-1>.
- [12] M. R. GAREY, R. L. GRAHAM, AND D. S. JOHNSON, *Some NP-complete geometric problems*, in Proceedings of the 8th Annual ACM Symposium on Theory of Computing, STOC 1976, ACM, 1976, pp. 10–22, <https://doi.org/10.1145/800113.803626>.
- [13] S. HAR-PELED AND K. QUANRUD, *Approximation algorithms for polynomial-expansion and low-density graphs*, SIAM J. Comput., 46 (2017), pp. 1712–1744, <https://doi.org/10.1137/16M1079336>.
- [14] S. HAR-PELED AND B. RAICHEL, *Net and prune: A linear time algorithm for Euclidean distance problems*, J. ACM, 62 (2015), 44, <https://doi.org/10.1145/2831230>.
- [15] M. HELD AND R. M. KARP, *A dynamic programming approach to sequencing problems*, in Proceedings of the 1961 16th ACM National Meeting, ACM '61, ACM, 1961, pp. 71.201–71.204.
- [16] R. Z. HWANG, R. C. CHANG, AND R. C. T. LEE, *The searching over separators strategy to solve some NP-hard problems in subexponential time*, Algorithmica, 9 (1993), pp. 398–423, <https://doi.org/10.1007/BF01228511>.
- [17] R. IMPAGLIAZZO AND R. Paturi, *On the complexity of k -SAT*, J. Comput. System Sci., 62 (2001), pp. 367–375, <https://doi.org/10.1006/jcss.2000.1727>.
- [18] V. KANN, *On the Approximability of NP-Complete Optimization Problems*, Ph.D. thesis, Royal Institute of Technology Stockholm, 1992.
- [19] R. M. KARP, *Reducibility among combinatorial problems*, in 50 Years of Integer Programming, Springer, 2010, pp. 219–241.
- [20] S. KISFALUDI-BAK, *ETH-Tight Algorithms for Geometric Network Problems*, Ph.D. thesis, Department of Mathematics and Computer Science, Technische Universiteit Eindhoven, 2019, <https://research.tue.nl/en/publications/eth-tight-algorithms-for-geometric-network-problems>.
- [21] S. KISFALUDI-BAK, J. NEDERLOF, AND K. WEGRZYCKI, *A gap-ETH-tight approximation scheme for Euclidean TSP*, in Proceedings of the 62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, IEEE, 2021, pp. 351–362, <https://doi.org/10.1109/FOCS52979.2021.00043>.
- [22] A. LINGAS, *Subexponential-time algorithms for minimum weight triangulations and related problems*, in Proceedings of the 10th Canadian Conference on Computational Geometry (CCCG), 1998, 43.
- [23] D. MARX AND A. SIDIROPOULOS, *The limited blessing of low dimensionality: When $1-1/d$ is the best possible exponent for d -dimensional geometric problems*, in Proceedings of the 30th Annual Symposium on Computational Geometry, SoCG 2014, ACM, (2014), pp. 67–76, <https://doi.org/10.1145/2582112.2582124>.
- [24] J. S. B. MITCHELL, *Guillotine subdivisions approximate polygonal subdivisions: A simple polynomial-time approximation scheme for geometric TSP, k -MST, and related problems*, SIAM J. Comput., 28 (1999), pp. 1298–1309, <https://doi.org/10.1137/S0097539796309764>.
- [25] W. MULZER AND G. ROTE, *Minimum-weight triangulation is NP-hard*, J. ACM, 55 (2008), 11, <https://doi.org/10.1145/1346330.1346336>.
- [26] C. H. PAPADIMITRIOU, *The Euclidean traveling salesman problem is NP-complete*, Theoret. Comput. Sci., 4 (1977), pp. 237–244, [https://doi.org/10.1016/0304-3975\(77\)90012-3](https://doi.org/10.1016/0304-3975(77)90012-3).
- [27] S. RAO AND W. D. SMITH, *Approximating geometrical graphs via "spanners" and "banyans"*, in Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, STOC 1998, ACM, 1998, pp. 540–550, <https://doi.org/10.1145/276698.276868>.
- [28] W. D. SMITH AND N. C. WORMALD, *Geometric separator theorems & applications*, in Proceedings of the 39th Annual Symposium on Foundations of Computer Science, FOCS 1998, IEEE Computer Society, 1998, pp. 232–243, <https://doi.org/10.1109/SFCS.1998.743449>.
- [29] G. J. WOEGINGER, *Open problems around exact algorithms*, Discrete Appl. Math., 156 (2008), pp. 397–405, <https://doi.org/10.1016/j.dam.2007.03.023>.