Bitonicity of Euclidean TSP in Narrow Strips*

Henk Alkema[1], Mark de Berg[2], and Sándor Kisfaludi-Bak[3]

1 **Department of Mathematics and Computer Science, TU Eindhoven, the Netherlands**
h.y.alkema@tue.nl
2 **Department of Mathematics and Computer Science, TU Eindhoven, the Netherlands**
m.t.d.berg@tue.nl
3 **Max Planck Institute for Informatics, Germany**
sandor.kisfaludi-bak@mpi-inf.mpg.de

Abstract

We investigate how the complexity of Euclidean TSP for point sets $P \subset (-\infty, +\infty) \times [0, \delta]$ depends on the strip width δ. We prove that if the points have distinct integer x-coordinates, a shortest bitonic tour (which can be computed in $O(n \log^2 n)$ time using an existing algorithm) is guaranteed to be a shortest tour overall when $\delta \leqslant 2\sqrt{2}$, a bound which is best possible.

1 Introduction

In the Traveling Salesman Problem one is given an edge-weighted complete graph and the goal is to compute a tour—a simple cycle visiting all nodes—of minimum total weight. Due to its practical as well as theoretical importance, the Traveling Salesman Problem and its many variants are among the most famous problems in computer science and combinatorial optimization. In this paper we study the Euclidean version of the problem. In Euclidean TSP the input is a set P of n points in $\mathbb{R}^d$, and the goal is to compute a minimum-length tour visiting each point. Euclidean TSP in the plane was proven to be NP-hard in the 1970s [1, 2]. Unlike the general (metric) version, however, it can be solved in *subexponential* time, that is, in time $2^{o(n)}$. In particular, Kann [3] and Hwang *et al.* [4] presented algorithms with $n^{O(\sqrt{n})}$ running time. Smith and Wormald [5] gave a subexponential algorithm that works in any (fixed) dimension; its running time in $\mathbb{R}^d$ is $n^{O(n^{1-1/d})}$. Very recently De Berg *et al.* [6] improved this to $2^{O(n^{1-1/d})}$, which is tight up to constant factors in the exponent, under the Exponential-Time Hypothesis (ETH) [7].

There has also been considerable research on special cases of Euclidean TSP that are polynomial-time solvable. One example is Bitonic TSP, where the goal is to find a shortest *bitonic* tour. (A tour is bitonic if any vertical line crosses it at most twice; here the points from the input set P are assumed to have distinct x-coordinates.) It is a classic exercise [8] to prove that Bitonic TSP can be solved in $O(n^2)$ time by dynamic programming. De Berg *et al.* [9] showed how to speed up the algorithm to $O(n \log^2 n)$.

Our contribution. The computational complexity of Euclidean TSP in $\mathbb{R}^d$ is $2^{\Theta(n^{1-1/d})}$ (for $d \geqslant 2$), assuming ETH. Thus the complexity depends heavily on the dimension d. This is most pronounced when we compare the complexity for $d = 2$ with the trivial case $d = 1$: in the plane Euclidean TSP takes $2^{\Theta(\sqrt{n})}$ time in the worst case, while the 1-dimensional case is trivially solved in $O(n \log n)$ time by just sorting the points. We study the complexity of Euclidean TSP for planar point sets that are "almost 1-dimensional". In particular,

* The work in this paper is supported by the Netherlands Organisation for Scientific Research (NWO) through Gravitation-grant NETWORKS-024.002.003.

36th European Workshop on Computational Geometry, Würzburg, Germany, March 16–18, 2020.
This is an extended abstract of a presentation given at EuroCG'20. It has been made public for the benefit of the community and should be considered a preprint rather than a formally reviewed paper. Thus, this work is expected to appear eventually in more final form at a conference with formal proceedings and/or in a journal.

we assume the point set P is contained in the strip $(-\infty, \infty) \times [0, \delta]$ for some relatively small δ, and that all points have distinct integer x-coordinates. Bitonic TSP can be solved in $O(n \log^2 n)$ time [9]. It is natural to conjecture that for sufficiently small δ, an optimal bitonic tour on P is a shortest tour overall. We give a (partially computer-assisted) proof that this is indeed the case: we prove that when $\delta \leqslant 2\sqrt{2}$ an optimal bitonic tour is optimal overall, and we show that the bound $2\sqrt{2}$ is best possible.

Notation and terminology. Let $P := \{p_1, \ldots, p_n\}$ be a set of points with distinct integer x-coordinates in a horizontal strip of width δ—we call such a strip a *δ-strip*—which we assume without loss of generality to be the strip $(-\infty, \infty) \times [0, \delta]$. We denote the x-coordinate of a point $p \in \mathbb{R}^2$ by $x(p)$, and its y-coordinate by $y(p)$. To simplify the notation, we also write x_i for $x(p_i)$, and y_i for $y(p_i)$. We sort the points in P such that $x_i < x_{i+1}$ for all $1 \leqslant i < n$.

For two points $p, q \in \mathbb{R}^2$, we write pq to denote the *directed* edge from p to q. The length of an edge pq is denoted by $|pq|$, and the total length of a set E of edges is denoted by $\|E\|$.

A *separator* is a vertical line not containing any of the points in P that separates P into two non-empty subsets. For our purposes, two separators s, s' that induce the same partitioning of P are equivalent. Therefore, we can define $\mathcal{S} := \{s_1, \ldots, s_{n-1}\}$ as the set of all combinatorially distinct separators, obtained by taking one separator between any two points p_i, p_{i+1}. Let E, F be sets of edges with endpoints in P. The *tonicity of E at a separator s*, written as $\mathrm{ton}(E, s)$, is the number of edges in E crossing s. We say that E has *lower tonicity* than F, denoted by $E \preccurlyeq F$, if $\mathrm{ton}(E, s_i) \leqslant \mathrm{ton}(F, s_i)$ for all $s_i \in \mathcal{S}$. E has *strictly lower tonicity* than F, denoted by $E \prec F$, if there also exists at least one i for which $\mathrm{ton}(E, s_i) < \mathrm{ton}(F, s_i)$. Finally, we call E *bitonic* if $\mathrm{ton}(E, s_i) = 2$ for all $s_i \in \mathcal{S}$.

2 Bitonicity for points with integer x-coordinates

The goal of this section is to prove the following theorem.

▶ **Theorem 1.** *Let P be a set of points with distinct and integer x-coordinates in a δ-strip. When $\delta \leqslant 2\sqrt{2}$, a shortest bitonic tour on P is a shortest tour overall. Moreover, for any $\delta > 2\sqrt{2}$ there is a point set P in a δ-strip such that a shortest bitonic tour on P is not a shortest tour overall.*

The construction for the case $\delta > 2\sqrt{2}$ is shown in Fig. 1. It is easily verified that, up to symmetrical solutions, the tours T_1 and T_2 are the only candidates for the shortest tour. Observe that $\|T_2\| - \|T_1\| = |p_1p_4| - |p_4p_5| = 3 - \sqrt{1+\delta^2}$. Hence, for $\delta > 2\sqrt{2}$ we have $\|T_2\| < \|T_1\|$, which proves the second statement of Theorem 1. The remainder of the section is devoted to proving the first statement.

Let P be a point set in a δ-strip for $\delta = 2\sqrt{2}$, where all points in P have distinct integer x-coordinates. Among all shortest tours on P, let T_{opt} be one that is minimal with respect to the $\preccurlyeq$-relation; T_{opt} exists since the number of different tours on P is finite. We claim that T_{opt} is bitonic, proving the upper bound of Theorem 1.

Suppose for a contradiction that T_{opt} is not bitonic. Let $s^* \in \mathcal{S}$ be the rightmost separator for which $\mathrm{ton}(T_{\mathrm{opt}}, s^*) > 2$. We must have $\mathrm{ton}(T_{\mathrm{opt}}, s^*) = 4$ because otherwise $\mathrm{ton}(T_{\mathrm{opt}}, s) > 2$ for the separator $s \in \mathcal{S}$ immediately to the right of s^*, since there is only one point between s^* and s. Let F be the four edges of T_{opt} crossing s^*, and let E be the remaining edges of T_{opt}. Let Q be the set of endpoints of the edges in F. We will argue that there exists a set F' of edges with endpoints in Q such that $E \cup F'$ is a tour and (i) $\|F'\| < \|F\|$, or (ii) $\|F'\| = \|F\|$ and $F' \prec F$. We will call such an F' *superior to F*. Option (i) contradicts that T_{opt} is a shortest tour, and (ii) contradicts that T_{opt} is a shortest

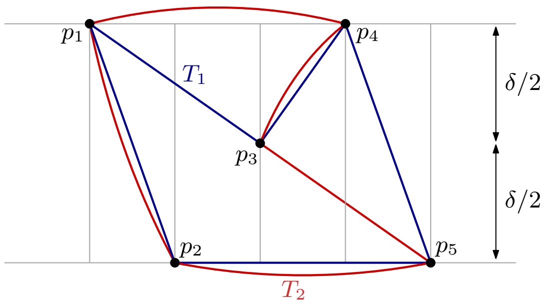

Figure 1 Construction for $\delta > 2\sqrt{2}$ for Theorem 1. The grey vertical segments are at distance 1 from each other. If $\delta > 2\sqrt{2}$ then T_1, the shortest bitonic tour (in blue), is longer than T_2, the shortest non-bitonic tour (in red).

tour, minimal with respect to $\preccurlyeq$ (since $E \cup F' \prec E \cup F$ if and only if $F' \prec F$). Hence, proving a superior set F' exists finishes the proof.

The remainder of the proof proceeds in two steps. In the first step we argue that we can assume without loss of generality that Q uses consecutive integer x-coordinates. In the second step we then give a computer-assisted proof that a superior set F' exists.

Step 1: Reduction to an instance where Q has consecutive x-coordinates. The goal of Step 1 is to move the points in Q to obtain a set $\overline{Q}$ of points with consecutive x-coordinates in such a way that finding a superior set $\overline{F}'$ for $\overline{Q}$ also gives us a superior set F' for Q. Let $\overline{F}$ be the same set as F, but now on the moved point set $\overline{Q}$, and define $\widetilde{E}$, the *connectivity pattern* of E, to be the set of edges obtained by contracting each path in E to a single edge.

▶ **Lemma 2.** *Let $T_{\mathrm{opt}}, s^*, E, \widetilde{E}, F, \overline{F}$ and Q be defined as above. Then there exists a $\overline{Q}$ s.t.:*
1. *$\widetilde{E}, \overline{F}$ and $\overline{Q}$ adhere to one of the six cases in Figure 2.*
2. *If there exists an $\overline{F}'$ superior to $\overline{F}$, there exists an F' superior to F.*

Sketch of proof. We start by taking $\overline{Q} = Q$. Now, property 2 trivially holds. We will now transform $\overline{Q}$, making sure that property 2 keeps holding, until property 1 also holds.
Let p be a point of $\overline{Q}$ which only has one incident edge pq in F. Suppose we move p some distance d along pq towards q. The total length of any candidate $\overline{F}'$ decreases by at most d by doing so, while $\|\overline{F}\|$ decreases by exactly d. Similar reasoning can be made about the tonicity. Therefore, this move does not affect the desired properties. If a point r has two incident edges pr, rq in $\overline{F}$, we can split it into two points r_1, r_2, and add an edge $r_1 r_2$ between them. Then, we can move them towards p and q, respectively. Doing so also does not affect the desired properties. Since all edges of $\overline{F}$ cross the separator s^*, the points can be moved towards each other such that they have consecutive integer x-coordinates. ◀
The complete proof can be found in the full version of this paper.

Step 2: Finding the set F'. The goal of Step 2 of the proof is the following: given a tour $\widetilde{E} \cup \overline{F}$ on a point set $\overline{Q}$ adhering to one of the six cases in Figure 2, show that there exists a set $\overline{F}'$ of edges superior to $\overline{F}$. Lemma 2 then implies that a superior set of edges exists for any Q, E, F, finishing the proof of Theorem 1.

Each of the six cases has several subcases, depending on the left-to-right order of the vertices inside the gray rectangles in the figure. Once we fixed the ordering, we can still vary the y-coordinates in the range $[0, \delta]$, which may lead to scenarios where different sets $\overline{F}'$ are required. We handle this potentially huge amount of cases in a computer-assisted

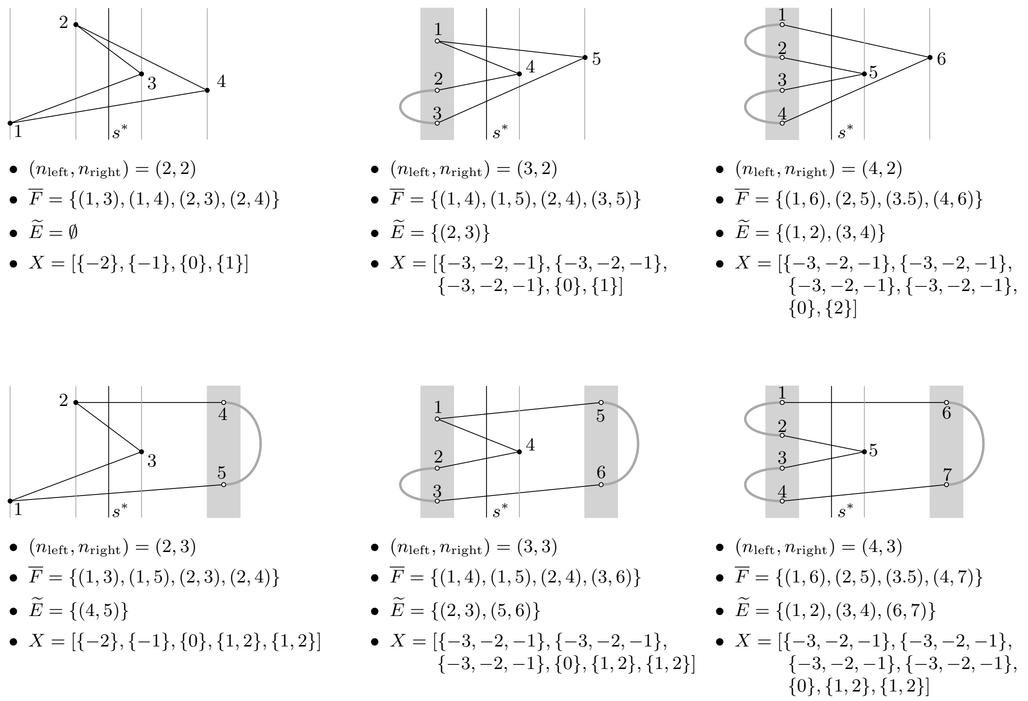

Figure 2 The six different cases that result after applying Step 1 of the proof. Points indicated by filled disks have a fixed x-coordinate. The left-to-right order of points drawn inside a grey rectangle, on the other hand, is not known yet. The vertical order of the edges is also not fixed, as the points can have any y-coordinate in the range $[0, 2\sqrt{2}]$.

manner, using an automated prover *FindShorterTour*$(n_{\text{left}}, n_{\text{right}}, \overline{F}, \widetilde{E}, X, \delta, \varepsilon)$. The input parameter X is an array where $X[i]$ specifies the set from which the x-coordinate of the i-th point in the given scenario may be chosen, where we assume w.l.o.g. that $x(s^*) = -1/2$; see Fig. 2. The role of the parameter ε will be explained below.

The output of *FindShorterTour* is a list of *scenarios* and an *outcome* for each scenario. A scenario contains for each point q an x-coordinate $x(q)$ from the set of allowed x-coordinates for q, and a range y-range$(q) \subseteq [0, 2\sqrt{2}]$ for its y-coordinate, where the y-range is an interval of length at most ε. The outcome is either Success or Fail. Success means that a set $\overline{F}'$ has been found with the desired properties: $\widetilde{E} \cup \overline{F}'$ is a tour, and for all possible instantiations of the scenario—that is, all choices of y-coordinates from the y-ranges in the scenario—we have $\|\overline{F}'\| < \|\overline{F}\|$. Fail means that such an $\overline{F'}$ has not been found, but it does not guarantee that such an $\overline{F'}$ does not exist for this scenario. The list of scenarios is complete in the sense that for any instantiation of the input case there is a scenario that covers it.

FindShorterTour works brute-force, by checking all possible combinations of x-coordinates and subdividing the y-coordinate ranges until a suitable $\overline{F}'$ can be found or until the y-ranges have length at most ε. The implementation details of the procedure can be found in the full version of this paper.

Note that case $(n_{\text{left}}, n_{\text{right}}) = (2,3)$ in Fig. 2 is a subcase of case $(n_{\text{left}}, n_{\text{right}}) = (3,2)$, if we exchange the roles of the points lying to the left and to the right of s^*. Hence, we ignore this subcase and run our automated prover on the remaining five cases, where we set $\varepsilon := 0.001$. It successfully proves the existence of a suitable set $\overline{F}'$ in four cases; the

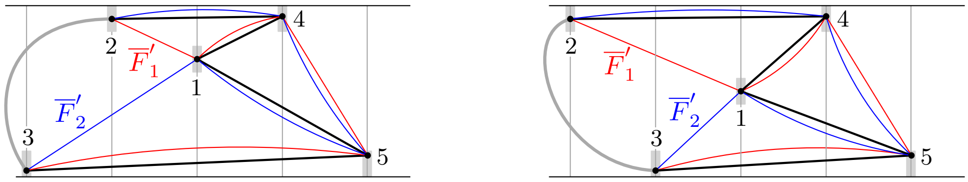

- $X = [\{-1\}, \{-2\}, \{-3, \}, \{0\}, \{1\}]$
- $Y = [[1.61, 1.62], [2.82, 2\sqrt{2}], [0, 0.01], [2.82, 2\sqrt{2}], [0, 0.01]]$

- $X = [\{-1\}, \{-3\}, \{-2, \}, \{0\}, \{1\}]$
- $Y = [[1.41, 1.42], [2.82, 2\sqrt{2}], [0, 0.01], [2.82, 2\sqrt{2}], [0, 0.01]]$

Figure 3 Two scenarios covering all subscenarios where the automated prover fails. Each point has a fixed x-coordinate and a y-range specified by the array Y; the resulting possible locations are shown as small grey rectangles (drawn larger than they actually are for visibility). For all subscenarios, at least one of $\overline{F}'_1$ (in red) and $\overline{F}'_2$ (in blue) is at most as long as $\overline{F}$ (in black).

case where the prover fails is the case $(n_{\text{left}}, n_{\text{right}}) = (3, 2)$. For this case it fails for the two scenarios depicted in Fig. 3; all other scenarios for these cases are handled successfully (up to symmetries). For both scenarios we consider two alternatives for the set $\overline{F}'$: the set $\overline{F}'_1$ shown in red in Fig. 3, and the set $\overline{F}'_2$ shown in blue in Fig. 3. We will show that in any instantiation of both scenarios, either $\overline{F}'_1$ or $\overline{F}'_2$ is at least as short as $\overline{F}$; since both alternatives are bitonic this finishes the proof.

For $1 \leqslant i \leqslant 5$, let q_i be the point labeled i in Fig. 3. We first argue that (for both scenarios) we can assume without loss of generality that $y(q_2) = y(q_4) = 2\sqrt{2}$ and $y(q_3) = y(q_5) = 0$. To this end, consider arbitrary instantiations of these scenarios, and imagine moving q_2 and q_4 up to the line $y = 2\sqrt{2}$, and moving q_3 and q_5 down to the line $y = 0$. It suffices to show, for $i \in \{1, 2\}$, that if we have $\|\overline{F}'_i\| \leqslant \|\overline{F}\|$ after the move, then we also have $\|\overline{F}_i\| \leqslant \|\overline{F}\|$ before the move. This can easily be proven by repeatedly applying the following observation.

▶ **Observation 3.** *Let a, b, c be three points. Let ℓ be the vertical line through c, and let us move c downwards along ℓ. Let α be the smaller angle between ac and ℓ if $y(c) < y(a)$, and the larger angle otherwise, and let β be the smaller angle between bc and ℓ if $y(c) < y(b)$, and the larger angle otherwise, and suppose $\alpha < \beta$ throughout the move. Then the move increases $|ac|$ more than it increases $|bc|$.*

So now assume $y(q_2) = y(q_4) = 2\sqrt{2}$ and $y(q_3) = y(q_5) = 0$. Consider the left scenario in Fig. 3, and let $y := y(q_3)$. If $y \geqslant (8\sqrt{2})/7$ then

$$|q_2q_1| + |q_4q_5| = \sqrt{1 + (2\sqrt{2} - y)^2} + 3 \leqslant 2 + \sqrt{4 + y^2} = |q_2q_4| + |q_1q_5|,$$

so $\|\overline{F}'_1\| \leqslant \|\overline{F}\|$. On the other hand, If $y \leqslant (8\sqrt{2})/7$ then

$$|q_3q_1| + |q_4q_5| = \sqrt{4 + y^2} + 3 \leqslant \sqrt{1 + (2\sqrt{2} - y)^2} + 4 = |q_1q_4| + |q_3q_5|,$$

so $\|\overline{F}'_2\| \leqslant \|\overline{F}\|$. So either $\overline{F}'_1$ or $\overline{F}'_2$ is at least as short as $\overline{F}$, finishing the proof for the left scenario in Fig. 3. The proof for the right scenario in Fig. 3 is analogous, with cases $y \geqslant \sqrt{2}$ and $y \leqslant \sqrt{2}$. This finishes the proof for the right scenario and, hence, for Theorem 1.

3 Concluding remarks

In our paper, we proved that for points with integer x-coordinates in a strip of width δ, an optimal bitonic tour is optimal overall when $\delta \leqslant 2\sqrt{2}$. The proof of this bound, which is tight in the worst case, is partially automated to reduce the potentially very large number of cases to two worst-case scenarios. It would be interesting to see if a direct proof can be given for this fundamental result. Finally, we note that the proof of Theorem 1 can easily be adapted to point sets of which the x-coordinates of the points need not be integer, as long as the difference between x-coordinates of any two consecutive points is at least 1.

In the full version of this paper, we also investigate the case $\delta > 2\sqrt{2}$. We present a fixed-parameter tractable algorithm with respect to δ. More precisely, our algorithm has running time $2^{O(\sqrt{\delta})}n^2$ for point sets where each $1 \times \delta$ rectangle inside the strip contains $O(1)$ points. For point sets where the points are chosen uniformly at random from the rectangle $[0, n] \times [0, \delta]$, it has an expected running time $2^{O(\sqrt{\delta})}n^2 + O(n^3)$.

References

1. M.R. Garey, R.L. Graham, and D.S. Johnson. Some NP-complete geometric problems. In *Proc. 8th ACM Symp. Theory Comp. (STOC)*, pages 10–22, 1976.
2. C.H. Papadimitriou. The Euclidean traveling salesman problem is NP-complete. *Theoret. Comput. Sci.* 4(3): 237–244 (1977).
3. V. Kann. *On the approximability of NP-complete optimization problems.* Ph.D. Dissertation, Royal Institute of Technology, Stockholm, 1992.
4. 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(4): 398–423 (1993).
5. W.D. Smith and N.C. Wormald. Geometric separator theorems and applications. In *Proc. 39th IEEE Symp. Found. Comput. Sci. (FOCS)*, pages 232–243, 1998.
6. M. de Berg, H.L. Bodlaender, S. Kisfaludi-Bak, and S. Kolay. An ETH-tight exact algorithm for Euclidean TSP. In *Proc. 59th IEEE Symp. Found. Comput. Sci. (FOCS)*, pages 450–461, 2018.
7. R. Impagliazzo and R. Paturi. On the complexity of k-SAT. *J. Comput. Syst. Sci.* 62(2): 367–375 (2001).
8. T.H. Cormen, C.E. Leiserson, R.L. Rivest, and C. Stein. *Introduction to Algorithms (3rd edition).* MIT Press, 2009.
9. M. de Berg, K. Buchin, B.M.P. Jansen, and G. Woeginger. Fine-grained complexity analysis of two classic TSP variants. In *Proc. 43rd Int. Conf. Automata Lang. Prog. (ICALP)*, pages 5:1–5:14, 2016.