

Computing Geometric Minimum-Dilation Graphs is NP-Hard

Panos Giannopoulos* Rolf Klein† Christian Knauer‡ Martin Kutz
Dániel Marx *§

Abstract

We prove that computing a geometric minimum-dilation graph on a given set of points in the plane, using not more than a given number of edges, is an NP-hard problem, no matter if edge crossings are allowed or forbidden. We also show that the problem remains NP-hard even when a minimum-dilation tour or path is sought; not even an FPTAS exists in this case.

Keywords: dilation, geometric network, plane graph, tour, path, spanning ratio, stretch factor, NP-hardness

1 Introduction

One of the fundamental problems in many application areas is the following: Given a set of locations, P , construct a network G that provides good connections between them at low cost.

This problem comes in various types, depending on the measures of cost and connection quality. Also, the setting of the problem can be different. In the graph-theoretic case, the locations in P are vertices of a given graph G_0 , and the desired network G must be a subgraph of G_0 . In the geometric case, the locations are points in d -dimensional space, and the required network is a geometric graph $G = (P, E)$ whose edges are straight line segments in \mathbb{R}^d . Formally, the second is a special case of the first, because any geometric graph over P can be considered as subgraph of the complete graph $G_0(P)$ that results from connecting every two points of P by a straight edge. Moreover, the geometric case is potentially easier to solve because properties of Euclidean geometry can be exploited.

In this paper we are studying the geometric case. Our cost model is the number of edges of the network. Connection quality is measured by dilation in the following way. For any two points, $a, b \in P$, let $|ab|$ denote their Euclidean distance, and let $d_G(a, b)$ be the weight or length of a shortest path from a to b in G , where the length of a path is given by the sum of the Euclidean lengths of its edges. Then

$$\delta_G(a, b) := \frac{d_G(a, b)}{|ab|}$$

*Humboldt-Universität zu Berlin, Institut für Informatik, Unter den Linden 6, D-10099 Berlin, Germany, panos@informatik.hu-berlin.de; partially supported by the DFG project no. GR 1492/7-2.

†Institute of Computer Science I, 53117 Bonn, Germany, rolf.klein@uni-bonn.de; partially supported by the German Research Foundation DFG, grant no. Kl 655/14-1/2/3.

‡Institut für Informatik, Freie Universität Berlin, Takustraße 9, D-14195 Berlin, Germany, Christian.Knauer@inf.fu-berlin.de.

§dmarx@informatik.hu-berlin.de

denotes the dilation of a, b in G , and

$$\delta(G) = \max_{a, b \in P, a \neq b} \delta_G(a, b)$$

is the *vertex-to-vertex dilation* or simply *dilation* of G . This value is also known as the stretch factor or the spanning ratio of G (it should not be confused with the geometric dilation that takes all points of the network into account, vertices and interior edge points alike). Disconnected graphs have infinite dilation.

The existence of low cost, low dilation geometric networks is guaranteed by the theory of spanners; see Eppstein [7], Smid [17], or Narasimhan and Smid [16] for surveys. For example, one can construct in time $O(n \log n)$ a network of dilation $\leq 1 + \epsilon$ that connects n points in \mathbb{R}^d using only $O(n)$ edges, if ϵ and dimension d are fixed. However, these spanners need not be optimal with respect to dilation or cost.

In the graph-theoretic case the complexity of finding *optimal* spanners has received a lot of attention. Building on previous work by Cai [3], Brandes and Handke [2] proved the following fact for weighted graphs. For each fixed rational number $\delta \geq 4$, it is an NP-complete problem to decide if a given graph H contains a planar subgraph G , whose weight does not exceed a given bound W , such that for any two vertices v, w of H the relation $d_H(v, w) \leq \delta \cdot d_G(v, w)$ holds, where the length of a path is given by the sum of its edge weights. Cai [3] and Cai and Corneil [4] have studied the problem of finding tree spanners of dilation $\leq \delta$ in weighted graphs. They proved that the decision problem is NP-complete for any $\delta \geq 4$, but polynomially solvable for $\delta = 2$, while the case $\delta = 3$ seems to be open. Fekete and Kremer [10] have considered the same problem for unweighted planar graphs. They proved that it is NP-hard to find the minimum dilation δ for which a tree spanner exists. But, surprisingly, the existence of a tree spanner of dilation 3 can be decided in polynomial time.

None of the results on the graph-theoretic case carries over to the geometric case, leaving wide open the complexity of computing optimal geometric spanning networks, even in dimension 2.

In 2005, Eppstein and Wortman [8] showed how to compute, in expected time $O(n \log n)$, a star of minimum dilation for n points. Then, since 2006, four hardness results were independently obtained and made public, in the following order:

- (1) Gudmundsson and Smid [12] proved that it is NP-hard to find a δ -spanner with $\leq m$ edges in a given geometric graph, for given bounds of δ and m .
- (2) Klein and Kutz [14] showed that the following problem is NP-hard: Given a finite set P of n points in the plane and a dilation bound δ , does there exist a (plane) geometric graph over P with $\lfloor \frac{5920}{5919} \cdot n - \frac{7628}{5919} \rfloor$ many edges and dilation $\leq \delta$?
- (3) Cheong, Haverkort, and Lee [5] proved that it is NP-hard to decide if the minimum (plane) dilation tree of a finite point set in the plane has a dilation $\leq \delta$.
- (4) Giannopoulos, Knauer, and Marx [11] showed that it is NP-hard to decide if there exists a closed tour (or an open path) of dilation at most δ that connects a given finite point set in the plane. They also proved that this problem admits no FPTAS, unless $P=NP$.

None of these problems is known to be in NP.

Results (2) and (3) both imply that it is NP-hard to decide if a given finite point set in the plane admits a geometric spanner of $\leq m$ edges and dilation $\leq \delta$. This, in turn,

implies result (1). Results (2), (3), and (4) are logically independent. Also, the proofs work in different ways. While (1) is shown by reduction from 3SAT, (2) and (3) reduce from PARTITION, but use different constructions. Fact (4) is based on HAMILTONIANCIRCUIT for grid graphs.

Each result is interesting in its own right. Result (4), because tours and paths are rather simple structures. Fact (3) on trees is in sharp contradistinction to the complexity of constructing minimum spanning trees; in addition, it solves an open problem of D. Eppstein's survey chapter [7].

Result (2) shows that hardness remains if a few more edges are allowed. This is interesting because Aronov et al. [1] have shown that a small number of extra edges matter a lot in lowering the dilation. In fact, each tree containing the vertices of a regular n -gon has a dilation of $\Omega(n)$, as was shown by Ebbers-Baumann et al. [6] and in [1], but with $n - 1 + k$ edges, where $0 \leq k < n$, a dilation of only $O(n/(k + 1))$ can be achieved, which is optimal. Results (2) and (3) also hold for plane graphs, that is, for geometric graphs without edge crossings.

In view of the recent result by Mulzer and Rote [15] on the minimum weight triangulation, it would be interesting to know if it is also NP-hard to construct the minimum dilation triangulation of a given point set in the plane. While this is certainly suggested by the above results (1)–(4), it is not implied. It remains to be seen if one the proof techniques of (1)–(4) can be generalized to cover triangulations, too.

This paper presents the extended versions of the results (2)¹ and (4).

More precisely, we assume that we are given a set P of n points in the plane and an upper dilation bound $\delta \geq 1$. By a graph over P we mean a geometric graph $G = (P, E)$ with a set E of straight edges. A Hamiltonian circuit over P is a closed tour that visits each vertex exactly once. A Hamiltonian path is a circuit minus one edge.

We show that the following decision problems are NP-hard. Moreover, none of the problems DILATIONTOUR and DILATIONPATH admits an FPTAS.

1. GRAPHCOMPLETION: Given a geometric graph $G = (P, E)$, is it possible to add $\leq \frac{2}{11}|E| - \frac{4}{11}$ edges such that the dilation of the resulting graph is at most 7?
2. DILATIONGRAPH: Is there a geometric graph $G = (P, E)$ of dilation $\delta(G) \leq 7$ with $|E| \leq \lfloor \frac{5920}{5919} \cdot n - \frac{7628}{5919} \rfloor$?
3. PLANEDILATIONGRAPH: Is there a plane geometric graph with the same properties?
4. DILATIONTOUR : Is there a Hamiltonian circuit G on P of dilation $\delta(G) \leq \delta$?
5. DILATIONPATH : Is there a Hamiltonian path G on P of dilation $\delta(G) \leq \delta$?

Since edges can always be added without increasing the dilation, Facts 2 and 3 imply NP-hardness of the following problem: Given a finite point set P and upper bounds δ and e , is there a (plane) geometric graph over P of dilation $\leq \delta$ and $\leq e$ many edges?

¹In the conference version [14] of (2) it was also shown that the minimum dilation tree can contain edge crossings, solving an open problem from Eppstein [7]. However, a smaller example was later given in [5], so that we do not include our construction here.

2 Minimum-dilation graphs

In this section we are proving that both DILATIONGRAPH and PLANEDILATIONGRAPH are NP-hard problems.

As a preparation we first provide, in Subsection 2.1, a result on problem GRAPHCOMPLETION. It is NP-hard to decide if the dilation of a given geometric graph $G = (P, E)$ can be decreased below a given bound by inserting at most $\frac{2}{11}|E| - \frac{4}{11}$ new edges. This result complements a positive finding by Farshi et al. [9] that a single edge, whose insertion reduces the dilation as much as possible, can be found in time $O(n^4)$ time, and $(2 + \epsilon)$ -approximated in time $O(nm + n^2(\log n + \epsilon^{-6}))$.

2.1 Adding edges to a geometric graph

First we introduce the main idea of our construction.

Assume we are given the three line segments au and vc of length 4 each, and bd of length 9, as shown in Figure 1. Suppose we were allowed to add two more edges, with four new vertices, in such a way that the resulting graph has the smallest possible dilation. If we used one of the the two extra edges to connect au to vc , and the other for connecting au and to bd , the shortest path distance between c and d would be at least

$$d(c, d) \geq |cv| + |vu| + |ud| \geq 5 + \sqrt{26} > 10,$$

so that a dilation > 10 would result. It is more efficient to connect bd to au and to vc , and the unique best way to do this is by using the vertical edges xx' and yy' depicted in Figure 1. This yields a dilation of 7, attained by each of the vertex pairs (a, b) , (u, v) , and (c, d) . Now




Figure 1: Only by adding edges xx' and yy' can a dilation of 7 be achieved.

let us modify this configuration in the following way. We enlarge the gap between u and v by moving each of these points by a distance of $\eta \leq 1/24$ outwards; see Figure 2. This




Figure 2: In network H one of the two top-down edges can be slanted while the dilation remains ≤ 7 .

modification gives us some freedom for placing the connecting edges. We keep the upper vertices x, y on au and vc fixed. But for each of the lower endpoints we consider two options. In addition to the original vertex x' , we place a new vertex x'' on edge bd , at a distance of 10η to the left of x' . Similarly, we introduce a new vertex y'' at a distance 10η to the right of y' .

Now let's see what happens if we add the edge xx'' , as before, but use yy'' instead of yy' . Clearly, the new path length between c, d equals

$$d(c, d) = 7 - (1 + 10\eta - \sqrt{1 + (10\eta)^2}) \leq 7 - 10\eta + (10\eta)^2. \quad (1)$$

Because of $|cd| = 1$, the same upper bound holds for the dilation $\delta(c, d)$. While the dilation between a, b remains equal to 7, we obtain

$$\delta(u, v) = \frac{7 - 2\eta + 10\eta + \sqrt{1 + (10\eta)^2} - 1}{1 + 2\eta} \quad (2)$$

$$\leq \frac{7 + 10\eta}{1 + 2\eta} \quad (3)$$

$$= \frac{7 + 10\eta}{7 + 14\eta} 7 \quad (4)$$

$$\leq 7; \quad (5)$$

observe that (3) holds because $\eta \leq 1/24$ implies $\sqrt{1 + (10\eta)^2} - 1 \leq 2\eta$.

What would happen if both slanted edges xx'' and yy'' were used? Then the dilation between u, v would be

$$\delta(u, v) = \frac{7 - 2\eta + 2 \cdot 10\eta + 2\sqrt{1 + (10\eta)^2} - 2}{1 + 2\eta} \quad (6)$$

$$> 7, \quad (7)$$

by straightforward calculation. So far, we have shown the following.

Lemma 1 *If the network H depicted in Figure 2 is of dilation ≤ 7 , then only one of its two top-down edges can be slanted. A slanted edge saves $1 + 10\eta - \sqrt{1 + (10\eta)^2}$ in path length between the terminal vertices (a, b resp. c, d) on its respective side.*

Now we are prepared to prove the following preliminary result.

Theorem 2 *Given a geometric graph $G = (P, E)$, it is NP-hard to decide if one can obtain a dilation ≤ 7 by adding to G up to $\leq \frac{2}{11}|E| - \frac{4}{11}$ new edges without introducing new vertices. The same is true for plane geometric graphs, where the new edges must not introduce edge crossings.*

Proof: We use reduction from the PARTITION problem:

Given a set S of s positive integers with $\sum_{r \in S} r = 2 \cdot R$, for some integer R , decide whether there exists a subset $T_1 \subseteq S$ such that $\sum_{r \in T_1} r = R = \sum_{r \in T_2} r$, where $T_2 = S \setminus T_1$.

The idea is to use one network H_i for each integer $r_i \in S$. For the numbers r_i in the prospective subset T_1 of S , the left top-down edge of H_i will be slanted, while the right edge

will be slanted if $r_j \in T_2$. The savings in path length obtained from a slanted edge of H_i should correspond to r_i in size; thus, we let

$$\eta_i := 10^{-(\lambda+1)} \cdot r_i$$

be the parameter used in the construction of H_i . Here, $\lambda \geq 10$ is a global parameter satisfying

$$2 \cdot s \cdot r_{\max}^2 < 10^\lambda \quad (8)$$

for the largest number r_{\max} in S . This choice guarantees in particular that the condition $\eta_i \leq 1/24$ is always fulfilled (if $s \geq 3$).

In order to add up the savings on the right and on the left hand side, respectively, we connect the networks $H_i, 1 \leq i \leq s$, as shown in Figure 3 for $s = 4$. Two adjacent networks




Figure 3: Graph G contains one network H_i for each integer $r_i \in S$, without the dotted edges.

are joined by vertical edges of length 3. From the terminal vertices a_1, c_1 of the topmost network H_1 , and b_s, d_s of the bottommost network H_s , edges of length h extend outwards. Parameter h is defined by

$$h := 9(s-1) + \frac{1}{2} 10^{-\lambda} R - \frac{1}{2} 10^{-2\lambda} s r_{\max}^2 \quad (9)$$

for a reason that will become clear in (14) below. Observe that the negative term is of smaller absolute value than the preceding positive term.

Now graph G is defined to consist of the solid edges shown in Figure 3.²

Assume that instance S of the partition problem has a solution $S = T_1 \cup T_2$. Then the sum over all r_i in T_j equals R , for $j = 1$ and $j = 2$. We add two edges to each H_i in G in the way described before, that is, for $r_i \in T_1$ the left top-down edge of H_i is slanted while the right one is vertical, and vice versa for $r_i \in T_2$.

²This figure is not quite true to scale. The horizontal edges adjacent to the four outermost vertices are more than twice as long as the height of the graph, by definition of h .

By Lemma 1, the path through the resulting graph, G^* , between its leftmost vertices t_1 and t'_1 is of length

$$d_{G^*}(t_1, t'_1) = 2h + 7s + 3(s - 1) - \sum_{r_i \in T_1} (1 + 10\eta_i - \sqrt{1 + (10\eta_i)^2}) \quad (10)$$

$$\leq 2h + 10s - 3 - \sum_{r_i \in T_1} (10\eta_i - (10\eta_i)^2) \quad (11)$$

$$= 2h + 10s - 3 - \sum_{r_i \in T_1} 10^{-\lambda} r_i + \sum_{r_i \in T_1} 10^{-2\lambda} r_i^2 \quad (12)$$

$$\leq 2h + 10s - 3 - 10^{-\lambda} R + 10^{-2\lambda} s r_{\max}^2 \quad (13)$$

$$= 7 \cdot (4s - 3) \quad (14)$$

$$= 7 \cdot |t_1 t'_1|. \quad (15)$$

Here, (14) follows from the definition of h in (9), and $s + 3(s - 1) = 4s - 3$ is the Euclidean distance between t_1 and t'_1 . Consequently, we obtain

$$\delta_{G^*}(t_1, t'_1) \leq 7 \quad \text{and, symmetrically,} \quad \delta_{G^*}(t_2, t'_2) \leq 7.$$

The vertices within each network H_i have dilation ≤ 7 , by the analysis leading to (5). All other pairs of vertices of G^* have dilation far less than 7. Thus, we have shown the following.

Lemma 3 *If partition instance S is solvable then we can add $2s$ edges to graph $G = (V, E)$ such that the dilation of the resulting graph does not exceed 7.*

Conversely, let us assume that we have obtained a dilation ≤ 7 by adding $2s$ edges to G . Then each H_i must have received two top-down edges—or its dilation would exceed 7. Thus, all $2s$ edges are accounted for. By Lemma 1, only one, of the two edges H_i has received, can be slanted. As before, let T_1 denote the set of all $r_i \in S$ where the left edge of H_i is slanted. We want to prove that T_1 and $T_2 := S \setminus T_1$ are a solution of partition instance S . Let us assume that this is not the case. Then,

$$\sum_{r_i \in T_j} r_i \leq R - 1 \quad (16)$$

must hold for $j = 1$ or $j = 2$; let's suppose it holds for $j = 1$. This implies

$$d_{G^*}(t_1, t'_1) = 2h + 7s + 3(s - 1) - \sum_{r_i \in T_1} (1 + 10\eta_i - \sqrt{1 + (10\eta_i)^2}) \quad (17)$$

$$\geq 2h + 10s - 3 - \sum_{r_i \in T_1} 10\eta_i \quad (18)$$

$$= 2h + 10s - 3 - \sum_{r_i \in T_1} 10^{-\lambda} r_i \quad (19)$$

$$\geq 2h + 10s - 3 - 10^{-\lambda} R + 10^{-\lambda} \quad (20)$$

$$> 7 \cdot (4s - 3) \quad (21)$$

$$= 7 \cdot |t_1 t'_1|. \quad (22)$$

Here, (20) follows from (16), and (21) is due to the definition of h in (9) and the inequality $10^\lambda > s r_{\max}^2$, which follows from (8). Thus, we obtain $\delta_{G^*}(t_1, t'_1) > 7$, a contradiction.

This proves the following converse of Lemma 3.

Lemma 4 *If the dilation of graph G can be reduced to a value ≤ 7 by adding $2s$ edges to G , then partition instance S is solvable.*

It remains to ensure that the graph G depicted in Figure 3 can be constructed by a Turing machine in time polynomial in the bit length of partition instance S , which is in $\Omega(s + \log R)$. Clearly, graph G is of combinatorial complexity $O(s)$. All vertices of G have rational coordinates whose numerators are powers of ten. Due to (8), parameter λ is in $O(\log s + \log R)$. Thus, the bit length of the parameters η_i and h , that occur in the vertex coordinates, lies in $O(\log s + \log R)$, too. Consequently, the vertices of G have numerators and denominators consisting of $O(\log s + \log R)$ many bits.

Finally, we observe that graph G depicted in Figure 3 has $e := 9s + 2(s - 1) + 4 = 11s + 2$ edges, to which $2s = \frac{2}{11}e - \frac{4}{11}$ edges have been added.

This concludes the proof of Theorem 2. \square

2.2 Adding edges to a set of points

Now we prove NP-hardness of the problems DILATIONGRAPH and PLANE DILATIONGRAPH. Again, we are presented with an instance S of the PARTITION problem that involves s positive integers; but this time we have to construct a point set P , rather than a graph, such that there exists a low-dilation graph with few edges over P if, and only if, instance S is solvable.

We are going to employ the same construction as in Subsection 2.1, and we shall use again the numbers η_i and λ as defined in (8). Our point set P is shown in Figure 4. It consists of sampling points taken from the edges of the graph depicted in Figure 3.

The spacing of the sample points is as follows. In general, two neighboring sample points are a distance of 10^{-2} apart. But there are three exceptions to this rule.

- In each substructure $H_i, 1 \leq i \leq s$, the white point on the left hand side is at distance $10\eta_i$ to its right black neighbor. Symmetrically, the white point on the right hand side has distance $10\eta_i$ to its left black neighbor. Still, the two black neighbors of each white point are at distance 10^{-2} from each other.
- The vertical point sequences leave gaps of width 10^{-1} at their upper and lower endpoints.
- Let $[t_1, a_1)$ denote the horizontal outward group of points including t_1 but excluding a_1 . It contains exactly $913(s - 1) + 1$ points, which are a distance of $\psi \approx 10^{-2}$ apart; the precise value of ψ will be defined below. The rightmost point of this group is at distance $1/3$ to its right neighbor, a_1 . Analogous statements hold for $[t'_1, b_4), (c_1, t_2]$, and $(d_4, t'_2]$.

Lemma 5 *For a partition instance S of s integers, point set P consists of $5919s - 4210$ points.*

Proof: There are $913(s - 1) + 1$ points on each of the four horizontal outward groups. Each structure H_i contains 3 segments of total length 17 that are sampled at density 10^{-2} , which results in 1703 points, plus the two extra points painted white in Figure 4. Consecutive H_i are connected by two vertical rows of length $3 - 2 \cdot 10^{-1}$ each, sampled at density 10^{-2} . Thus, we obtain

$$|P| = 4 \cdot 913(s - 1) + 4 + s \cdot 1705 + (s - 1) \cdot 2 \cdot 281 = 5919s - 4210. \quad (23)$$

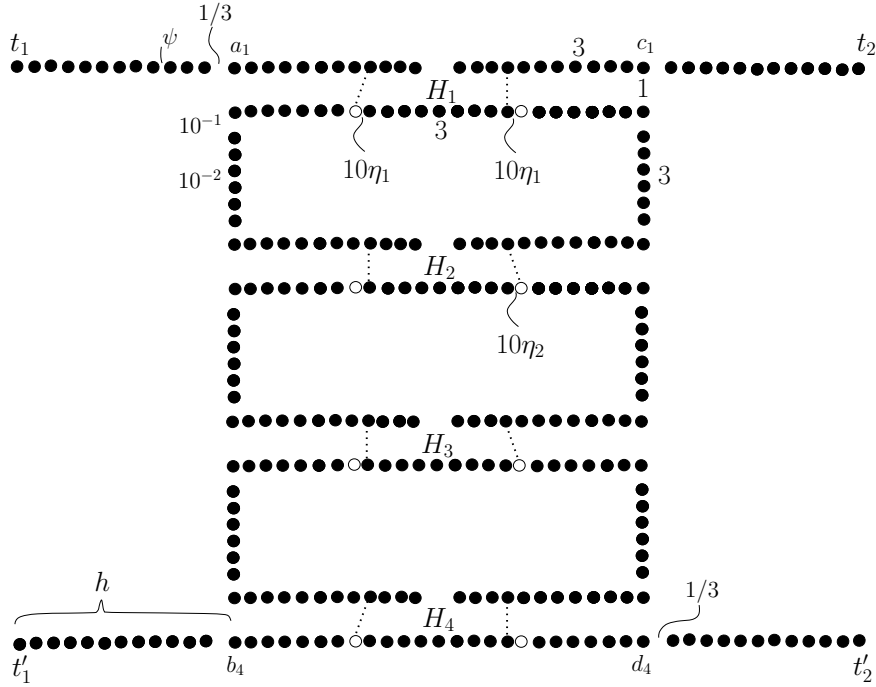


Figure 4: A sample P of the graph shown in Figure 3.

□

Now we proceed by proving the analogue of Lemma 3.

Lemma 6 *If partition instance S is solvable then there exists a plane graph of dilation ≤ 7 over vertex set P that contains $5920s - 4212$ many edges.*

Proof: Suppose that $S = T_1 \cup T_2$ is a solution of the given partition instance. Let graph G over P be defined as follows. First, we introduce

$$4 \cdot 913(s - 1) + s \cdot 1702 + (s - 1) \cdot 2 \cdot 280 = 5914s - 4212$$

edges of length 10^{-2} that connect neighboring points in each of the horizontal and vertical groups of sample points; we shall refer to these edges as being “short”. Then the 4 outward groups are connected, by an edge of length $1/3$ each, to the points a_1, c_1, b_s , and d_s , respectively; see Figure 4. As before, each H_i receives two top-down edges, exactly one of which is slanted. The left edge of H_i is slanted if $r_i \in T_1$ holds, and the right edge, if $r_i \in T_2$. Finally, we employ $(s - 1) \cdot 2 \cdot 2$ edges to connect the vertical point groups to their adjacent horizontal groups of the structures H_i . Altogether, we have used $5920s - 4212$ many edges.

The latter connections are delicate. In order to minimize the overall path length between t_1, t'_1 and t_2, t'_2 , we use shortcuts, as shown in Figure 5, instead of straight edges between the points w and b . But the shorter we cut, the larger gets the dilation between w and b ! It is not hard to verify, using, e. g., Maple, that the biggest saving

$$\kappa := \frac{46}{100} - \sqrt{\left(\frac{23}{100}\right)^2 + \left(\frac{23}{100}\right)^2} = \frac{23}{50} - \frac{23}{100}\sqrt{2} \quad (24)$$

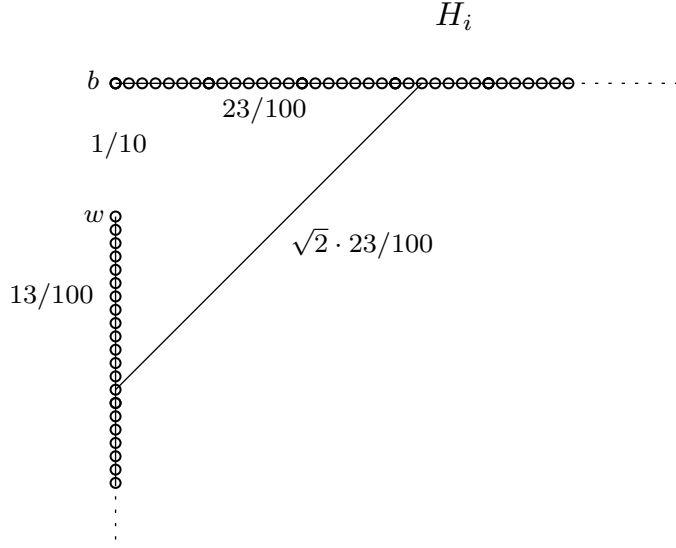


Figure 5: Using shortcuts to connect vertical to horizontal point groups.

respecting $\delta(w, b) \leq 7$ can be obtained by using the diagonal that connects the 13th point below w to the 23rd point to the right of b . Then,

$$\delta(w, b) = \frac{18}{5} + \frac{23}{10}\sqrt{2} \approx 6.8526.$$

Now let q be a rational approximation of κ satisfying

$$|\kappa - q| < 10^{-(2\lambda+1)} \quad (25)$$

We define the distance ψ between neighboring points of the four horizontal outward groups by

$$\psi := \frac{(9 + q)(s - 1) + \frac{1}{2}10^{-\lambda}R - \frac{1}{3} - \frac{1}{4}(10^{-\lambda} + 10^{-2\lambda}sr_{\max}^2)}{913(s - 1)}; \quad (26)$$

observe that $\psi \approx 10^{-2}$ holds since $9 + q \approx 9.13$ and because the other terms in the numerator of ψ are bounded in size. The definition of ψ implies for the length $h := |t_1 a_1|$ of each horizontal outward group

$$h = (9 + q)(s - 1) + \frac{1}{2}10^{-\lambda}R - \frac{1}{4}(10^{-\lambda} + 10^{-2\lambda}sr_{\max}^2) \quad (27)$$

because we have to add the $1/3$ gap. This completes the definition of graph G . Clearly, G is crossing-free.

It remains to show that $\delta(G) \leq 7$ holds. We start with the dilation values that are crucial for our construction, $\delta_G(t_1, t'_1)$ and $\delta_G(t_2, t'_2)$. Taking the shortcuts κ into account we obtain

$$d_G(t_1, t'_1) = 2h + 7s + 3(s-1) - 2\kappa(s-1) - \sum_{r_i \in T_1} (1 + 10\eta_i - \sqrt{1 + (10\eta_i)^2}) \quad (28)$$

$$\leq 2h + (10 - 2\kappa)(s-1) + 7 - 10^{-\lambda}R + 10^{-2\lambda} s r_{\max}^2 \quad (29)$$

$$= (28 + 2(q - \kappa))(s-1) + 7 - \frac{1}{2}10^{-\lambda} + \frac{1}{2}10^{-2\lambda} s r_{\max}^2 \quad (30)$$

$$< 28(s-1) + 7 + 2 \cdot 10^{-(2\lambda+1)}(s-1) - \frac{1}{2}10^{-\lambda} + \frac{1}{2}10^{-2\lambda} s r_{\max}^2 \quad (31)$$

$$< 28(s-1) + 7 + \frac{1}{10}10^{-\lambda} - \frac{1}{2}10^{-\lambda} + \frac{1}{4}10^{-\lambda} \quad (32)$$

$$< 28(s-1) + 7 \quad (33)$$

$$= 7 \cdot (4s - 3) \quad (34)$$

$$= 7 \cdot |t_1 t'_1|. \quad (35)$$

Here, (29) is analogous to (10)–(13), and (30) follows from (27). Formula (31) is implied by (25), and (32) is a consequence of (8), which implies $2(s-1) < 10^\lambda$ and $10^{-2\lambda} s r_{\max}^2 < \frac{1}{2}10^\lambda$. Similarly, $d_G(t_2, t'_2) < 7 \cdot |t_2 t'_2|$ holds.

Now we turn to the other vertex pairs of G . Clearly, two vertices from horizontal outward groups can have a dilation only smaller than $\delta_G(t_1, t'_1)$ resp. $\delta_G(t_2, t'_2)$.

If both vertices are located in the same substructure H_i then $\delta_G(p, q) \leq 7$ holds by construction, because only the vertex pairs (a, b) , (u, v) , and (c, d) shown in Figure 1 are local dilation maxima. If one vertex belongs to H_i and the other to H_{i+k} , where $k \geq 1$, then their Euclidean distance is at least $3k$ whereas the shortest connecting path in G has length at most $10k + 3.5$ (attained by a topmost and a bottommost point in the middle of each structure); see Figure 6. So, the dilation stays far below 7. A similar argument applies if both vertices belong to vertical point groups.




Figure 6: The shortest path in G between p and q is no longer than $10k + 3.5$.

Now assume that one vertex is situated in H_i , whereas the other one lies in a vertical point group. Among the vertices shown in Figure 5, the maximum dilation, attained by (b, w) , is < 6.86 . If we consider, instead of vertex b of H_i , the vertex a of H_i above b (see Figure 1),

we obtain

$$\delta_G(w, a) = \left(\frac{13}{100} + \frac{23}{100} \sqrt{2} + \left(7 - \frac{23}{100} \right) \frac{1}{\left(\frac{1}{10} + 1 \right)} \right) \approx 6.6.$$

All other vertex pairs of this type have a smaller dilation in G . For example, the dilation between vertex w and the bottommost vertex of the next vertical group above H_i is only ≈ 6.32 .

It remains to study the case where one vertex, p , belongs to a horizontal outward group, and the other, q , to a vertical group; see Figure 7. Suppose that p lies at distance y to the left




Figure 7: The dilation between p and q is maximized for $p = r$ and $q = b_1$. Only the left part of graph G is shown.

of the rightmost vertex r of the horizontal group, and that q lies on a vertical group leading downwards from H_k , at distance z . Then,

$$\delta_G(p, q) \leq f(y, z) = \frac{y + \frac{1}{3} + 10k - 3 + z}{\sqrt{\left(y + \frac{1}{3}\right)^2 + (4k - 3 + z)^2}}.$$

In the definition of function f we ignore the shortcuts and assume that the top-down edges in the H_i are vertical—which only increases the path length. Function f takes on its maximum value 6.597 for $k = 1$ at $y = z = 0$, that is, for $p = r$ and $q = b_1$.³

This completes the proof of Lemma 6. \square

³As point r moves from a_1 to the left, the dilation $\delta(r, b_1)$ first increases to a value > 7 , but then decreases again. We left the gap of $1/3$ between a_1 and r just to make sure that $\delta(r, b_1) < 7$ holds.

Now we prove the analogue claim to Lemma 4. Let S denote a partition instance of size s , and let P be the point set depicted in Figure 4.

Lemma 7 *Suppose there exists a geometric graph $G = (P, E)$ over P such that $\delta(G) \leq 7$ and $|E| \leq 5920s - 4212$. Then partition instance S is solvable. Moreover, there exists a crossing-free graph with these properties.*

Proof: Among all graphs $G = (P, E)$ satisfying $\delta(G) \leq 7$ and $|E| \leq 5920s - 4212$, consider those that, (i), pareto-minimize the dilations $\delta_G(t_1, t'_1)$ and $\delta_G(t_2, t'_2)$. In the set of all graphs satisfying (i), let G be one of minimum weight, (ii).

First, we argue that G equals the graph we have constructed in the proof of Lemma 6—up to the fine positions of the two top-down edges in each H_i . To this end, the following observation is helpful. For two points p, p' , let $F(p, p')$ denote the ellipse with foci p and p' whose boundary points z satisfy $|pz| + |zp'| = 7|pp'|$.

Lemma 8 *Let p, p' denote two points of P . Then the shortest path connecting p, p' in G is contained in $F(p, p')$. Moreover, if $F(p, p') \cap P$ consists only of points situated on the line through p, p' , but not between p and p' , then (p, p') is an edge of E .*

Proof: Assume that $(p, p') \notin E$. Each vertex v of the shortest path π , that connects p and p' in G , must be contained in $F(p, p')$, because of

$$7 \geq \delta(G) \geq \delta_G(p, p') = \frac{|\pi|}{|pp'|} \geq \frac{|pv| + |vp'|}{|pp'|}.$$

Now assume that all of these vertices are situated on the line L through p, p' in—say—left-to-right order (v_1, v_2, \dots, v_r) , in such a way that no v_i lies between p and p' . If we replace the edges of π in G with the segments $v_i v_{i+1}$, $1 \leq i \leq r - 1$, the resulting graph over P has the same number of edges, but a smaller weight than G . Moreover, the dilation of any vertex pair is at most as large as in G , because the long edges of π can be obtained as concatenation of shorter, co-linear edges $v_i v_{i+1}$. This implies that the dilation values $\delta(t_1, t'_1)$ and $\delta(t_2, t'_2)$, that were minimal before, remain unchanged, so that the resulting graph still satisfies (i). But its weight has been reduced—a contradiction to (ii). \square

For two neighboring points p, p' of a horizontal or vertical point group of P , the assumptions of Lemma 8 are fulfilled because the diameter of $F(p, p')$ equals $7|pp'| \leq 7 \cdot 10^{-2} < 10^{-1}$, so that no point of another group is contained in the ellipse. Consequently, all neighboring pairs are connected by an edge of G ; as calculated in the proof of Lemma 6, as many as $5914s - 4212$ edges of G are now accounted for.

The first statement of Lemma 8 also implies the following facts.

- HO) Each of the four horizontal outward groups must be linked, by (at least) one edge each, to the chain of structures H_i .
- VE) Each vertical group must be linked, by (at least) two edges each, to its adjacent structures H_i, H_{i+1} .

This leaves us with (at most) $2s$ edges. We claim that each structure H_i must receive two of them, and that they must be positioned as depicted in Figure 2. Figure 8 shows the




Figure 8: In a network H of dilation ≤ 7 , the two top-down edges must have their upper endpoints at x, y , and their lower endpoints at x', x'' and y', y'' , correspondingly.

discretized version of a structure H . With the solid edges xx' and yy'' we would have dilations $\delta(a, b) = 7$, $\delta(c, d) < 7$, and $\delta(u, v) < 7$. Instead of edge xx' , let us now consider an edge e whose upper endpoint lies k points to the left of x . Since the length of e exceeds the length of the old edge xx' by at least

$$\sqrt{1 + 10^{-4}} - 1 \geq \frac{2}{5} 10^{-4}, \quad (36)$$

the lower endpoint of e can be at most $k - 1$ points to the right of x' , or $\delta(a, b)$ would become bigger than 7. So, the path length between u and v increases by at least 10^{-2} plus $4 \cdot 10^{-5}$. Since η is by some powers of ten smaller than these numbers, the dilation between u and v would be pushed above 7, unless edge yy' is also repositioned in such a way that $d(u, v)$ shrinks by at least $2 \cdot 10^{-2}$. But then $d(c, d)$ grows by the same amount, causing $\delta(c, d)$ to exceed 7. Similar arguments show that the upper endpoint of edge xx' cannot be moved to the right. The same holds on the right hand side of H . Consequently, each H_i must receive two top-down edges whose upper endpoints are in the intended positions. It is easy to see that there are exactly two choices for their lower endpoints, x', x'' and y', y'' . Now, Lemma 1 implies that only one top-down edge can be slanted.

Having established where the edges of G are located, we take a closer look to the way the four outward groups and the $2(s - 1)$ vertical groups are connected. By statement HO from above, exactly one edge connects the horizontal group of point r shown in Figure 7 to the rest of the graph. If this edge is co-linear with the horizontal group containing point a_1 then it must be the edge ra_1 , because G is of minimum weight; here the same argument as in the proof of Lemma 8 applies.

If not, the shortest path from r to a_1 would pass through x_1 , causing $\delta(r, a_1)$ to be larger than $(3 + 3)/(1/3) = 18$.

Now let us consider a one-edge connection between a vertical group and the adjacent structure H_i ; refer to Figure 5. A shortcut saving more than κ on path length $d_G(t_1, t'_1)$ would cause $\delta_G(w, b)$ to be greater than 7. A shortcut saving less could be adjusted to save exactly κ . This would decrease $d_G(t_1, t'_1)$ without increasing $\delta(G)$, which is impossible by the minimality property (i) of G . So, G contains exactly the shortcuts depicted in Figure 5.

At this point, we have seen that G equals the crossing-free graph we have constructed in the proof of Lemma 6—up to the fine positioning of the two top-down edges in each H_i . We claim that the positions of these edge correspond to a solution of partition instance S . As in (16), let T_1 denote the set of all numbers $r_i \in S$ where the left edge of H_i is slanted. For the sake of a contradiction, let us assume that

$$\sum_{r_i \in T_1} r_i \leq R - 1 \quad (37)$$

holds. Then,

$$d_G(t_1, t'_1) = 2h + 7s + 3(s-1) - 2\kappa(s-1) - \sum_{r_i \in T_1} (1 + 10\eta_i - \sqrt{1 + (10\eta_i)^2}) \quad (38)$$

$$\geq 2h + (10 - 2\kappa)(s-1) + 7 - 10^{-\lambda}R + 10^{-\lambda} \quad (39)$$

$$= (28 + 2(q - \kappa))(s-1) + 7 + \frac{1}{2}10^{-\lambda} - \frac{1}{2}10^{-2\lambda} s r_{\max}^2 \quad (40)$$

$$\geq 28(s-1) + 7 - 2 \cdot 10^{-(2\lambda+1)}(s-1) + \frac{1}{2}10^{-\lambda} - \frac{1}{2}10^{-2\lambda} s r_{\max}^2 \quad (41)$$

$$\geq 28(s-1) + 7 - \frac{1}{10}10^{-\lambda} + \frac{1}{2}10^{-\lambda} - \frac{1}{4}10^{-\lambda} \quad (42)$$

$$> 28(s-1) + 7 \quad (43)$$

$$= 7 \cdot (4s - 3) \quad (44)$$

$$= 7 \cdot |t_1 t'_1|, \quad (45)$$

which gives the desired contradiction. We observe that (39) follows from (37), as in (17)–(20). Estimate (40) is a consequence of the definition of h in (27), and (41) is implied by (25). Finally, (42) is obtained in the same way as (32).

This completes the proof of Lemma 7. \square

Now we can prove our main result. As in the introduction, let

$$e(n) := \left\lfloor \frac{5920}{5919} \cdot n - \frac{7628}{5919} \right\rfloor.$$

Theorem 9 *The decision problems DILATIONGRAPH and PLANEDILATIONGRAPH are NP-hard. More precisely, given a set P of n points in the plane, it is NP-hard to decide if there exists a (plane) geometric graph $G = (P, E)$ such that $\delta(G) \leq 7$ and $|E| = e(n)$.*

Proof: We start with a PARTITION instance S of s numbers, and construct the set P of $n := 5919s - 4210$ points addressed in Lemma 5. This can be done by a Turing machine in time polynomial in the input size of S ; in addition to the arguments given at the end of Subsection 2.1 we point out that the rational approximation q of κ , the only irrational number involved, can be constructed in time $O(\lambda)$, by Newton's method; see (24) and (25).

Thanks to Lemma 6 and Lemma 7, S admits a partition if, and only if, there exists a (plane) graph over P with $5920s - 4212$ edges, whose dilation is at most 7. Now we observe that the number of edges can also be written as

$$5920s - 4212 = n + s - 2 = n + \frac{n + 4210}{5919} - 2 = \frac{5920}{5919}n - \frac{7628}{5919},$$

which completes the proof. \square

3 Minimum-dilation tour (and path)

In this section we prove that both DILATIONTOUR and DILATIONPATH are NP-hard.




Figure 9: A grid graph G , its smallest enclosing rectangle R , and the point-sets (‘handles’) S and T .

3.1 Reduction

For a point $a \in \mathbb{R}^2$, we denote by $a(1)$ and $a(2)$ its x - and y -coordinate, respectively.

Let G^∞ be the infinite graph whose vertex set contains all points of the plane with integer coordinates and in which two vertices are connected if and only if the Euclidean distance between them is equal to 1. A *grid graph* is a finite, node-induced subgraph of G^∞ . Note that a grid graph is completely specified by its vertex set. Let HAMILTONIANCIRCUIT be the problem of deciding whether a given grid graph has a Hamiltonian circuit. It is well-known that HAMILTONIANCIRCUIT is NP-hard [13]. We reduce HAMILTONIANCIRCUIT to DILATIONTOUR. As an initial step, we adjust the input instances of HAMILTONIANCIRCUIT as follows.

Let G be a grid graph with vertex set V and $|V| = n$. Using V , we construct a point set W that will be used later in the reduction. We assume that G has no 0- or 1-degree vertices and that it is connected since otherwise there is no Hamiltonian circuit in G . Both properties can be checked in polynomial time. Consider the smallest enclosing rectangle R of G , see Fig. 9. Since G is finite and connected, and $|V| = n$, rectangle R has finite dimensions and its height is at most n . Let $v \in V$ be the vertex that is closest to the lower-left corner of R and lies on the left vertical edge of R . Then v must be a degree-2 vertex and have a neighbor on the same edge of R ; let u be this vertex. Clearly, vertex u is above v . We append two point-sets S and T , called *handles*, to G as shown in Fig. 9. Handle S consists of a point s_0 at horizontal distance 1 from u and a vertical sequence of $n + 1$ points s_1, \dots, s_{n+1} with a gap of distance 1 between consecutive points. Set T is defined similarly. We have

$$S = \{s_0 = (u(1) - 1, u(2))\} \cup \{s_i = (u(1) - 2, u(2) + i - 1) \mid i = 1, \dots, n + 1\}$$

and

$$T = \{t_0 = (v(1) - 1, v(2))\} \cup \{t_i = (v(1) - 2, v(2) - i + 1) \mid i = 1, \dots, n + 1\}.$$

Let $W = V \cup S \cup T$. We have that $|W| = 3n + 4$. Copies of W will be included later in

P . Consider the points $s, t \in W$ with $s = s_{n+1}$ and $t = t_{n+1}$. We have that $|st| = 2n + 1$. We start with a simple lemma.

Lemma 10 *There exists a Hamiltonian s - t path on W with length $3n + 3$ if and only if there exists a Hamiltonian circuit in G .*

Proof: Assume that there exists a Hamiltonian s - t path on W with length $3n + 3$. Since W contains $3n + 4$ points, any such path must contain only edges with length 1. Every point s_i with $i = 2, \dots, n$ is at distance one only from two points, namely, s_{i+1} and s_{i-1} . Hence, the s - t path must contain the edges $s_{i+1}s_i$ and $s_i s_{i-1}$. Similarly, the path must contain the edges $t_{i+1}t_i$ and $t_i t_{i-1}$ for $i = 2, \dots, n$. The edge $s_1 t_1$ cannot be in the path, since, otherwise, the path cannot visit all points in W . Thus, s_1 and t_1 have to connect to s_0 and t_0 respectively. Similarly, $s_0 t_0$ cannot be in the path, and so, the edges $s_0 u$ and $t_0 v$ must be in the path. The remaining of the s - t path must have a length of $3n + 3 - 2(n + 2) = n - 1$ and visit the remaining $n - 2$ vertices of V starting from u and ending at v . This implies that there is a u - v Hamiltonian path H_G in G . Since u and v are neighbors in G , edge uv and the u - v path H_G form a Hamiltonian circuit in G .

Conversely, assume that there is a Hamiltonian circuit in G . Since v has degree two, any such circuit contains uv . Thus, there is a u - v Hamiltonian path on V with length $n - 1$. We append to the latter path the edges $s_0 u$, $s_1 s_0$, $t_0 v$, $t_1 t_0$, and $s_{i+1} s_i$, $s_i s_{i-1}$, $t_{i+1} t_i$, $t_i t_{i-1}$ for $i = 2, \dots, n$. This forms a Hamiltonian s - t path on W with length $3n + 3$. \square


Next, using W we construct a point set P such that, for some δ to be defined later, a Hamiltonian circuit on P with dilation at most δ exists if and only if G has a Hamiltonian circuit.

First, we choose points on a rectangle R' of width α and height β , with

$$\alpha = (2n^2 + 1)n^6 + 2n^2 n^3 \quad \text{and} \quad \beta = 2n^6 + 3n^3.$$

Let a, b, c , and d be the upper-right, upper-left, bottom-left, and bottom-right corner points of R' respectively. Consider a straight-line segment of length n^6 . We choose a set B of points on the segment at regular intervals such that the distance between any two consecutive points is $n/2$. We have that $|B| = 2n^5 + 1$. We use B as a building block: starting from a and going on the rectangle in anti-clockwise direction, we place copies of B , simply referred to as *blocks*, at regular intervals such that the distance between two blocks is n^3 ; see Fig. 10 (to avoid cluttering, the edges of the rectangle are not shown). Let K, L, M , and N be the sets of points on the right, upper, left, and lower sides of the rectangle respectively (we define the corner points to be in L and N). Sets K and M are unions of two vertical blocks each, while L and N are unions of $2n^2 + 1$ horizontal ones. The right-most and left-most point of a horizontal block are called the right and left *end-points* of the block. Similarly, the lower and upper-most point of a vertical block are called the lower and upper end-points of the block. Let $K = K_1 \cup K_2$, where K_1, K_2 is the upper and lower block respectively, as shown in Fig. 10. Also, let e be the lower end-point of K_1 and f be the upper end-point of K_2 . In the gap between K_1 and K_2 , we place point set W such that s and t lie on the right side of the R' and $V \subset W$ is to the right of R' . Additionally, we require that

$$|es| = |ft| = (n^3 - |st|)/2 = (n^3 - 2n - 1)/2.$$



◦ ... ◦ block

◦ block end-point

Figure 10: Constructing point set P .

Since the height of the minimum enclosing rectangle R of V is at most n , the distance between any point of a block and any point of W is at least $(n^3 - 2n - 1)/2$ as well. A reflected copy of W , denoted by W' , is placed between the two blocks (subsets) of M in a similar way.

Let

$$P = K \cup L \cup M \cup N \cup W \cup W'.$$

We have that $|P| = (2n^2 + 1)(2n^5 + 1) + 2(2n^5 + 1) = O(n^7)$.

Let p and q be the ‘middle’ points of L and N respectively, that is,

$$p = (a + b)/2 \quad \text{and} \quad q = (c + d)/2.$$

Also, let

$$\delta = \frac{\alpha + \beta - (2n + 1) + 3n + 3}{\beta} = \frac{\alpha + \beta + n + 2}{\beta} = 1 + \frac{\alpha}{\beta} + \frac{n + 2}{(2n^3 + 3)n^3} \quad (46)$$

$$\begin{aligned} &= 1 + \frac{(2n^2 + 1)n^6 + 2n^2n^3}{2n^6 + 3n^3} + h(n) \\ &= 1 + n^2 + \frac{n^3 - n^2}{2n^3 + 3} + h(n) \\ &= 1 + n^2 + g(n) + h(n), \end{aligned} \quad (47)$$

with

$$g(n) = \frac{n^3 - n^2}{2n^3 + 3} \quad \text{and} \quad h(n) = \frac{n + 2}{(2n^3 + 3)n^3}.$$

Note that $g(n), h(n) < 1$ for every $n \geq 1$.




Figure 11: The Hamiltonian circuit H_P and example positions of points p' and q' .

Lemma 11 *If there is a Hamiltonian s - t path on W with length $3n + 3$, then there is a Hamiltonian circuit on P with dilation at most δ .*

Proof: Let H_W be a Hamiltonian s - t path on W with length $3n + 3$ (Lemma 10). We construct a Hamiltonian circuit H_P on P by simply connecting the points in K, L, M, N in the ‘canonical’ way along the sides of rectangle R' , as shown in Fig. 11. First, every two consecutive points in each block are connected by an edge. Second, in L, N , the left end-point of each block is connected to the right end-point of its immediate neighbor block. Finally, the upper end-point of K_1 and the lower end-point of K_2 are connected to points a and d respectively, while e connects to s and f connects to t ; the blocks of M are connected to b, c , and the point set W' in a similar way. We prove that $\delta(H_P) \leq \delta$. Note that any path from p to q in H_P must go through either W or W' . By the symmetry of the construction of H_P , we have that

$$\begin{aligned}
 d_{H_P}(p, q) &= |pa| + |ad| + |dq| - |ef| + |es| + |ft| + d_{H_W}(s, t) \\
 &= \alpha + \beta - |st| + d_{H_W}(s, t) \\
 &= \alpha + \beta - (2n + 1) + 3n + 3 = \alpha + \beta + n + 2.
 \end{aligned} \tag{48}$$

Note that the total length of H_P is equal to $2d_{H_P}(p, q)$.

Using (48) and from the derivation of δ in (46) we have that

$$\delta_{H_P}(p, q) = \frac{d_{H_P}(p, q)}{|pq|} = \frac{\alpha + \beta + n + 2}{\beta} = \delta.$$

We now prove that for any other pair of points $p', q' \in P$, $\delta_{H_P}(p', q') \leq \delta$. We distinguish the following cases, see Fig. 11:

(i) p', q' lie on opposite sides of R' , or $p' \in W$ and $q' \in M$ (symmetrically, $p' \in K$ and $q' \in W'$), or $p' \in W'$ and $q' \in W$. In this case we have that $|p'q'| \geq |pq|$. Since $d_{H_P}(p, q)$ is exactly half of the total length of H_P , this is the maximum distance between any two points, that is $d_{H_P}(p', q') \leq d_{H_P}(p, q)$. Thus, $\delta_{H_P}(p', q') = d_{H_P}(p', q')/|p'q'| \leq d_{H_P}(p, q)/|pq| = \delta$.

(ii) Either $p', q' \in W$ or $p', q' \in W'$. In this case, we have that $d_{H_P}(p', q') \leq d_{H_W}(s, t) = 3n + 3$ and $|p'q'| \geq 1$, hence $\delta_{H_P}(p', q') \leq 3n + 3 \leq \delta$, for any $n \geq 4$.

(iii) None of the previous two cases apply. We consider two subcases. First, none of p', q' is in $W \cup W'$. Then the shortest path from p' to q' in H_P goes on the boundary of R' with a possible detour either in W or in W' . Assume that such a detour goes through W . (The case where it goes through W' is symmetric.) Since we are not in case (i), the path touches at most one horizontal and at most one vertical side of R' . Thus,

$$d_{H_P}(p', q') < |p'(1) - q'(1)| + |p'(2) - q'(2)| + d_{H_W}(s, t) < 2|p'q'| + 3n + 3.$$

Second, one of p', q' is in $W \cup W'$. Without loss of generality, assume that $p' \in W$. Then, either $q' \in L \cup K_1$ or $q' \in N \cup K_2$. Consider again the shortest path from p' to q' in H_P : it goes from p' to either s or t and from there to q' . Assume it goes through s , i.e. s is the last point of W on this path. Then, by construction, s is closer to q' than p' is, i.e. $|sq'| \leq |p'q'|$. As before, the path from s to q' touches at most one horizontal and at most one vertical side of R' . Thus,

$$d_{H_P}(s, q') < |s(1) - q'(1)| + |s(2) - q'(2)| < 2|s'q'| \leq 2|p'q'|.$$

In total we have

$$d_{H_P}(p', q') = d_{H_P}(p', s) + d_{H_P}(s, q') < d_{H_W}(s, t) + 2|p'q'| = 3n + 3 + 2|p'q'|.$$

Since we are not in case (ii), $|p'q'| \geq n/2$. Therefore, for both subcases we have

$$\delta_{H_P}(p', q') = \frac{d_{H_P}(p', q')}{|p'q'|} < \frac{2|p'q'| + 3n + 3}{|p'q'|} \leq 2 + \frac{3n + 3}{n/2} < n^2 < \delta,$$

for any $n \geq 4$. □

Conversely, we now prove the following.

Lemma 12 *If there is a Hamiltonian circuit on P with dilation at most δ , then there is a Hamiltonian s - t path on W with length $3n + 3$.*

Proof: Let H_P be a Hamiltonian circuit on P with $\delta(H_P) \leq \delta$. We prove that H_P must contain a path from p to q that is ‘locally optimal’ in the sense that firstly, it connects p to s and t to q in the ‘canonical’ way on the sides of R' (as it was described in the proof of Lemma 11), and secondly, it connects s to t via a Hamiltonian path on W with length $3n + 3$. In particular, we show that δ is small enough to ensure that the following requirements are met:

(i) Once inside a block, H_P visits all the points of the block before leaving it. To see this, consider a block B such that H_P has more than one visits in B , i.e., B induces more than one




Figure 12: Examples of case (iii) in the proof of Lemma 12.

components of H_p . In that case there has to be a vertex p_k of B such that p_k and its neighbor p_{k+1} belong to different visits. This means that the path on H_P from p_k to p_{k+1} contains at least one edge leaving B and at least one edge entering B . Recall that the distance between any two blocks is at least n^3 and that the distance between any block and W or W' is at least $|es| = (n^3 - 2n - 1)/2$. We have that $d_{H_P}(p_{k+1}, p_k) \geq 2 \cdot (n^3 - 2n - 1)/2$ and

$$\begin{aligned} \delta_{H_P}(p_{k+1}, p_k) &= \frac{d_{H_P}(p_{k+1}, p_k)}{|p_{k+1}p_k|} \geq \frac{(n^3 - 2n - 1)}{n/2} = 2n^2 - 4 - \frac{2}{n} \\ &> 2n^2 - 6 > n^2 + 2 > \delta, \end{aligned}$$

for any $n \geq 3$.

(ii) Once inside W (or W'), H_P visits all the points of W (or W') before leaving it. This can be seen by using arguments similar to the ones in case (i). We show that if the distance between two points p', q' in W is 1, then p' and q' belong to the same visit of W . Clearly, this implies that there is only one visit. If p' and q' belong to different visits, then (arguing as in the previous case) $d_{H_P}(p', q') > 2|es| = 2(n^3 - 2n - 1)$ and

$$\delta_{H_P}(p', q') = \frac{d_{H_P}(p', q')}{|p'q'|} > 2(n^3 - 2n - 1) > n^2 + 2 > \delta,$$

for any $n \geq 3$.

(iii) Any two blocks that are consecutive along the sides of R' must be ‘connected’ by an edge in H_P , as long as W or W' does not lie between the two blocks. To see this, consider a block B and a neighbor of it, B' , and let p' and q' be the endpoints of B and B' respectively, with $|p'q'| = n^3$. Assume that H_P contains no edge connecting a point of B to a point of B' ; see Fig 12. Then, any path from p' to q' in H_P must visit some other block B'' different from B and B' , or visit one of W and W' . In the first case, we have $d_{H_P}(p', q') > n^6$: if two blocks are not consecutive, then their distance is greater than n^6 (this is true even if the two blocks are on different sides of R'), hence B'' is at distance greater than n^6 from at least one of B and B' . For the second case, observe that W (W') is at distance greater than n^6 from at least one of B and B' (since we have excluded the case when the two blocks are K_1 and K_2 , or M_1 and M_2). Thus, $\delta_{H_P}(p', q') > n^6/n^3 = n^3 > n^2 + 2 > \delta$, for any $n \geq 2$.

(iv) Blocks K_1, K_2 must be connected to W by an edge in H_P ; this holds also for the blocks in M and W' . Similarly to the case (iii), assume, for example, that H_P contains no edge ‘connecting’ K_1 and W . Then, any path from e to s in H_P must visit some other block

as well. By case (i), once in the latter block, H_P has to visit all its points before exiting. Hence, $d_{H_P}(e, s) > n^6$ and

$$\delta_{H_P}(e, s) = \frac{d_{H_P}(e, s)}{|es|} > \frac{n^6}{(n^3 - 2n - 1)} > n^3 > \delta,$$

for any $n \geq 2$.

The above requirements assert that H_P does not contain ‘long’ edges (jumps) between any two points of P that belong to different blocks or between a point of a block and a point of W or W' . Note that, since all points of a block lie on the same straight-line segment, the minimum-length Hamiltonian path on a block traverses the points of the block in the ‘canonical’ way, i.e. it has length n^6 ; any detour increases this length by at least n . Also, as already noted before, s and t are the points of W that are closest to e and f respectively. Requirement (ii) also asserts that the part of H_P inside W (W') forms a Hamiltonian s - t path on W (W'); let ℓ be its length. Consider now the pair p, q . By combining all the above, we have that

$$d_{H_P}(p, q) \geq |pa| + |ae| + |es| + \ell + |tf| + |fd| + |dq| = \alpha + \beta - (2n + 1) + \ell. \quad (49)$$

Since $\delta(H_P) \leq \delta$, we have that $\delta_{H_P}(p, q) \leq \delta$ as well. From the derivation of δ in (46), and since $|pq| = \beta$, this implies that

$$d_{H_P}(p, q) \leq \alpha + \beta + n + 2. \quad (50)$$

From (49) and (50), we have that $\ell \leq 3n + 3$. However, any Hamiltonian path on W has length at least $3n + 3$, and the lemma follows. \square

Note that all points in P have rational coordinates with numerators and denominators bounded by a polynomial in n . Also, the construction of P takes $O(|P|) = O(n^7)$ time. By combining Lemmata 10, 11, and 12, we have the following:

Theorem 13 *The decision problem DILATIONTOUR is NP-hard.*

As it is easy to see, this theorem holds also for the decision problem DILATIONPATH by considering a point set P that contains the ‘middle’ points p and q , the points that lie on R' to the right side of p and q , and the points in W . The block containing p now has length $n^6/2$ and only one neighboring block (along the side of R'), which has length n^6 ; the same is true for the block containing q . In this case, the analogue of Lemma 11 can be proved similarly by connecting the points of P with the canonical path from p to q on the sides of R' and through W . As for the analogue of Lemma 12, note that a Hamiltonian path on P with dilation at most δ meets the requirements (i)–(iv): the arguments for this are the same as before since the minimum distance between any two blocks or between a block and W or between two non-consecutive blocks does not change.

Corollary 14 *The decision problem DILATIONPATH is NP-hard.*

3.2 Non-approximability results

The proof of Theorem 13 is based on a reduction from the HAMILTONIAN CIRCUIT problem in a grid graph G . As already mentioned above, any Hamiltonian path on W has length at least $3n + 3$ and, from Lemma 10, this value is achieved if and only if there is a Hamiltonian circuit in G . On the other hand, if there is no Hamiltonian circuit in G , then any Hamiltonian path on W has length at least $(3n + 3) - 1 + \sqrt{2}$: in this case there are two non-adjacent vertices visited consecutively by the path, and non-adjacent vertices have distance at least $\sqrt{2}$. This observation implies that the problem of computing a minimum-dilation tour (or path) admits no FPTAS.

Theorem 15 *The minimum-dilation tour (and path) problem does not admit an FPTAS.*

Proof: Assume that there is an FPTAS, i.e., there is an algorithm that, for every $\epsilon > 0$, computes in time $O((1/\epsilon \cdot n)^{O(1)})$ a Hamiltonian circuit H_P with dilation δ_{apx} satisfying $\delta_{\text{apx}} \leq (1 + \epsilon)\delta_{\text{opt}}$, where δ_{opt} is the dilation of an optimal tour. Consider any $\epsilon < (\sqrt{2} - 1)/(\beta\delta)$, where β, δ are as in the proof of Theorem 13.

Note that $\delta_{\text{apx}} \leq (1 + \epsilon)\delta$ implies that

$$\delta_{\text{apx}} < \delta + \frac{\sqrt{2} - 1}{\beta} = \frac{\alpha + \beta - (2n + 1) + 3n + 3 + (\sqrt{2} - 1)}{\beta}.$$

Consider the ‘middle’ points p, q . Since $|pq| = \beta$, we have that $d_{H_P}(p, q) < \alpha + \beta - (2n + 1) + 3n + 3 + (\sqrt{2} - 1)$. It is easy to check that term $\sqrt{2} - 1$ is small enough to leave no other alternative to H_P but the form imposed by the cases (i), (ii), (iii), and (iv) in the proof of Theorem 13. Moreover H_P must visit W through its the points s and t : if any other point in W is visited first instead, $d_{H_P}(p, q)$ will increase by at least 2. Hence, H_P contains a Hamiltonian path on W of length $3n + 3$.

On the other hand, if $\delta_{\text{apx}} > (1 + \epsilon)\delta$, then $\delta_{\text{opt}} > \delta$. But, for any Hamiltonian circuit H on P we have that $\delta(H) \geq \delta_{\text{opt}} > \delta$, which, by Lemma 11, implies that there is no Hamiltonian path of length $3n + 3$ on W .

For the minimum-dilation path problem, we can use the same observation as for Corollary 14. \square

4 Concluding remarks

We have proved that computing a geometric minimum-dilation graph on a given set of points in the plane, using not more than a given number of edges, is an NP-hard problem, no matter if edge crossings are allowed or forbidden. We have also proved that computing a minimum-dilation (Euclidean) tour or path on a given set of points in the plane is NP-hard as well, and that both problems do not admit an FPTAS. Does any of these problems have a PTAS? Note that no constant-factor, polynomial-time approximation algorithms are even known.

In view of the recent result by Mulzer and Rote [15] on the minimum weight triangulation, is it also NP-hard to construct the minimum dilation triangulation of a given point set?

Acknowledgements

The authors would like to thank two anonymous referees for their valuable comments and suggestions.

References

- [1] B. Aronov, M. de Berg, O. Cheong, J. Gudmundsson, H. Haverkort, and A. Vigneron. Sparse geometric graphs with small dilation. *Computational Geometry: Theory and Applications*, 40(3):207–219, 2008.
- [2] U. Brandes and D. Handke. NP-completeness results for minimum planar spanners. *Discrete Mathematics and Theoretical Computer Science*, 3:1–10, 1998.
- [3] L. Cai. NP-completeness of minimum spanner problems. *Discrete Applied Mathematics*, 48(2):187–194, 1994.
- [4] L. Cai and D. Corneil. Tree spanners. *SIAM J. of Discrete Mathematics*, 8(3):359–387, 1995.
- [5] O. Cheong, H. Haverkort, and M. Lee. Computing a minimum-dilation spanning tree is NP-hard. *Computational Geometry: Theory and Applications*, 41(3):188–205, 2008.
- [6] A. Ebberts-Baumann, A. Grüne, and R. Klein. On the geometric dilation of finite point sets. *Algorithmica*, 44(2):137–149, 2006.
- [7] D. Eppstein. Spanning trees and spanners. In J.R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 425–461. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
- [8] D. Eppstein and K.A. Wortman. Minimum dilation stars. *Computational Geometry: Theory and Applications*, 37(1):27–37, 2007.
- [9] M. Farshi, P. Giannopoulos, and J. Gudmundsson. Improving the stretch factor of a geometric graph by edge augmentation. *SIAM Journal on Computing*, 38(1):226–240, 2008.
- [10] S. Fekete and J. Kremer. Tree spanners in planar graphs. *Discrete Applied Mathematics*, 108(1):85–103, 2001.
- [11] P. Giannopoulos, Ch. Knauer, and D. Marx. Minimum-dilation tour (and path) is NP-hard. In *Abstracts of the 23rd European Workshop Computational Geometry*, pages 18–21, 2007. Full version appears in Technical Report B 07-05, Freie Universitt Berlin, February 2007.
- [12] J. Gudmundsson and M. Smid. On spanners of geometric graphs. *International Journal of Foundations of Computer Science*, 20(1):135–149, 2009.
- [13] A. Itai, C.H. Papadimitriou, and J.L. Szwarcfiter. Hamilton paths in grid graphs. *SIAM J. Computing*, 11(4):676–686, 1982.

- [14] R. Klein and M. Kutz. Computing geometric minimum-dilation graphs is NP-hard. In *Proc. of the 14th International Symposium on Graph Drawing*, LNCS 4372, pages 196–207, 2006.
- [15] W. Mulzer and G. Rote. Minimum-weight triangulation is NP-hard. *J. ACM*, 55(2):1–29, 2008.
- [16] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- [17] M. Smid. Closest point problems in computational geometry. In J.R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 877–935. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.