

On the Stretch Factor of the Theta-4 Graph^{*}

Luis Barba^{1,2,**}, Prosenjit Bose¹, Jean-Lou De Carufel¹, André van Renssen¹,
and Sander Verdonschot¹

¹ School of Computer Science, Carleton University, Ottawa, Canada
{jit, andre}@scs.carleton.ca,
{jdecарuf, sander}@cg.scs.carleton.ca

² Département d'Informatique, Université Libre de Bruxelles, Brussels, Belgium
lbarbaf1@ulb.ac.be

Abstract. In this paper we show that the θ -graph with 4 cones has constant stretch factor, i.e., there is a path between any pair of vertices in this graph whose length is at most a constant times the Euclidean distance between that pair of vertices. This is the last θ -graph for which it was not known whether its stretch factor was bounded.

Keywords: computational geometry, geometric spanners, θ -graphs.

1 Introduction

A c -spanner of a weighted graph G is a connected sub-graph H with the property that for all pairs of vertices u and v , the weight of the shortest path between u and v in H is at most c times the weight of the shortest path between u and v in G , for some fixed constant $c \geq 1$. The smallest constant c for which H is a c -spanner of G is referred to as the *stretch factor* or *spanning ratio* of the graph.

The graph G is referred to as the *underlying graph*. In our setting, the underlying graph is the complete graph on a set of n points in the plane and the weight of an edge is the Euclidean distance between its endpoints. A c -spanner of such a graph is called a *geometric c -spanner*. For a comprehensive overview of geometric c -spanners, see the book by Narasimhan and Smid [10].

In this paper, we focus on θ -graphs. Introduced independently by Clarkson [7] and Keil [9], the θ_m -graph is constructed as follows. Given a set P of points in the plane, we consider each point $p \in P$ and partition the plane into m cones (regions in the plane between two rays originating from the same point) with apex p , each defined by two rays at consecutive multiples of $\theta = 2\pi/m$ radians from the negative y -axis. We label the cones $C_0(p)$ through $C_{m-1}(p)$, in counter-clockwise order around p , starting from the negative y -axis; see Fig. 1. In each cone $C_i(p)$, we add an edge between p and p_i , the point in $C_i(p)$ nearest to p . However, instead of using the Euclidean distance, we measure distance in $C_i(p)$ by projecting each point onto the angle bisector of this cone. Formally, p_i is the point in $C_i(p)$ such that for every other point $w \in C_i(p)$, the projection of p_i

* Research supported in part by NSERC and FQRNT.

** Boursier FRiA du FNRS.

onto the angle bisector of $C_i(p)$ lies closer to p than the projection of w . For simplicity, we assume that no two points of P lie on a line parallel to either the boundary or the angle bisector of a cone.

It has been shown that θ_m -graphs are geometric spanners for $m \geq 7$, and their stretch factor approaches 1 as m goes to infinity [4,6,11]. The proofs crucially rely on the fact that, given two points p and q such that $q \in C_i(p)$, the distance between p_i and q is always less than the distance between p and q . This property does not hold for $m \leq 6$ and indeed, the path obtained by starting at p and repeatedly following the edge in the cone that contains q , is not necessarily a spanning path. The main motivation for using spanners is usually to reduce the number of edges in the graph without increasing the length of shortest paths too much. Thus, θ -graphs with fewer cones are more interesting in practice, as they have fewer edges. This raises the following question: “What is the smallest m for which the θ_m -graph is a geometric spanner?” Bonichon *et al.* [1] showed that the θ_6 -graph is a geometric 2-spanner. Recently, Bose *et al.* [5] proved that the θ_5 -graph is a geometric 9.96-spanner. Coming from the other side, El Molla [8] showed that there is no constant c for which the θ_2 - and θ_3 -graphs are geometric c -spanners. This leaves the θ_4 -graph as the only open question. Moreover, its resemblance to graphs like the Yao₄-graph [3] and the L_∞ -Delaunay triangulation [2], both of which are geometric c -spanners, makes this question more tantalizing. In this paper we establish an upper bound of approximately 237 on the stretch factor of the θ_4 -graph. In Section 5, we present a lower bound of 7 that we believe is closer to the true stretch factor of the θ_4 -graph.

2 Existence of a Spanning Path

Let P be a set of points in the plane. In this section, we prove that the θ_4 -graph of P is a spanner. We do this by showing that the θ_4 -graph approximates the L_∞ -Delaunay triangulation. The L_∞ -Delaunay triangulation of P is a geometric graph with vertex set P , and an edge between two points of P whenever there exists an empty axis-aligned square having these two points on its boundary.

Bonichon *et al.* [2] showed that the L_∞ -Delaunay triangulation has a stretch factor of $c^* = \sqrt{4 + 2\sqrt{2}}$, i.e., there is a path between any two vertices whose length is at most c^* times their Euclidean distance. We approximate this path in the L_∞ -Delaunay triangulation by showing the existence of a spanning path in the θ_4 -graph of P joining the endpoints of every edge in the L_∞ -Delaunay triangulation. The main ingredient to obtain this approximation is Lemma 1 whose proof is presented in Section 4. Before stating this result, we need a few more definitions. Given two points s and t , their L_1 distance $d_{L_1}(s, t)$ is the sum of the absolute differences of their x - and y -coordinates.

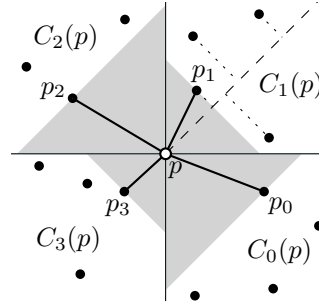


Fig. 1. The neighbors of p in the θ_4 -graph of P . Each edge supports an empty isosceles triangle.

Let $S_t(s)$ be the smallest axis-aligned square centered on t that contains s . Let ℓ_t^- and ℓ_t^+ be the lines with slope -1 and $+1$ passing through t , respectively.

Throughout this paper, we repeatedly use t to denote a *target* point of P that we want to reach via a path in the θ_4 -graph. Therefore, we typically omit the reference to t and write ℓ^-, ℓ^+ and $S(s)$ when referring to ℓ_t^-, ℓ_t^+ and $S_t(s)$, respectively.

We say that an object is *empty* if its interior contains no point of P . An *s-t-path* is a path with endpoints s and t .

Lemma 1. *Let s and t be two points of P such that t lies in $C_0(s)$. If the top-right quadrant of $S(s)$ is empty and $C_1(s)$ contains no point of P below ℓ^- , then there is an s - t -path in the θ_4 -graph of P of length at most $18 \cdot d_{L_1}(s, t)$.*

For ease of readability, the proof of Lemma 1 is deferred to Section 4.

Given a path φ , let $|\varphi|$ denote the sum of the lengths of the edges in φ . Using Lemma 1, we obtain the following.

Lemma 2. *Let s and t be two points of P . If the smallest axis-aligned square enclosing s and t , that has t as a corner, is empty, then there is an s - t -path in the θ_4 -graph of P of length at most $(\sqrt{2} + 36) \cdot |st|$.*

Proof. Assume without loss of generality that s lies in $C_1(t)$. Then, the top-right quadrant of $S(s)$ is empty as it coincides with the smallest axis-aligned square enclosing s and t that has t as a corner; see Fig. 2(a). Recall that s_3 is the neighbor of s in the θ_4 -graph inside the cone $C_3(s)$. Assume that $s_3 \neq t$ as otherwise the result follows trivially. Consequently, s_3 must lie either in $C_0(t)$ or in $C_2(t)$. Assume without loss of generality that s_3 lies in the top-left quadrant of $S(s)$. As s_3 lies in the interior of $S(s)$, $S(s_3) \subset S(s)$ and hence, the top-right quadrant of $S(s_3)$ is empty. Moreover, s_3 lies above ℓ^- and hence $C_1(s_3)$ contains no point of P below ℓ^- . Therefore, by Lemma 1 there is an s_3 - t -path φ of length at most $18 \cdot d_{L_1}(s_3, t)$. Since s_3 lies inside $S(s)$, $|s_3t| \leq \sqrt{2} \cdot |st|$ and hence $|\varphi| \leq 18 \cdot d_{L_1}(s_3, t) \leq 18\sqrt{2} \cdot |s_3t| \leq 18\sqrt{2}\sqrt{2} \cdot |st| = 36 \cdot |st|$. Moreover, the length of edge ss_3 is at most $d_{L_1}(s, t) \leq \sqrt{2} \cdot |st|$ since s_3 must lie above ℓ^- . Thus, $ss_3 \cup \varphi$ is an s - t -path of length $|ss_3| + |\varphi| \leq (\sqrt{2} + 36) \cdot |st|$. \square

The following observation is depicted in Fig. 2(b).

Observation 1 *Let S be an axis-aligned square. If two points a and b lie on consecutive sides along the boundary of S , then there is a square S_{ab} containing the segment ab such that $S_{ab} \subseteq S$ and either a or b lies on a corner of S_{ab} .*

Lemma 3. *Let ab be an edge of the L_∞ -Delaunay triangulation of P . There is an a - b -path φ_{ab} in the θ_4 -graph of P such that $|\varphi_{ab}| \leq (1 + \sqrt{2}) \cdot (\sqrt{2} + 36) \cdot |ab|$.*

Proof. Let $T = (a, b, c)$ be a triangle in the L_∞ -Delaunay triangulation of P . By definition of this triangulation, there is an empty square S such that every vertex of T lies on the boundary of S . By the general position assumption, a, b and c must lie on different sides of S . If a and b lie on consecutive sides of the boundary

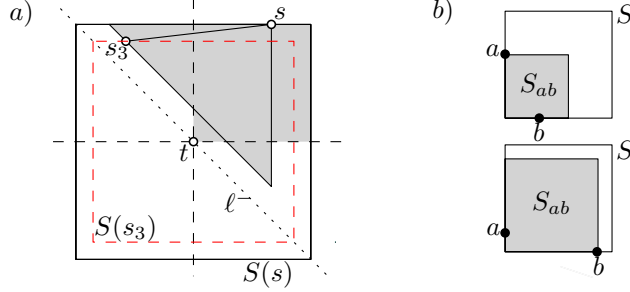


Fig. 2. a) Configuration used in the proof of Lemma 2, grey areas represent empty regions. b) If a and b lie on consecutive sides of a square S , there is a square S_{ab} such that $ab \subset S_{ab} \subseteq S$ and either a or b lies on a corner of S_{ab}

of S , then by Observation 1 and Lemma 2 there is a path φ_{ab} contained in the θ_4 -graph of P such that $|\varphi_{ab}| \leq (\sqrt{2} + 36) \cdot |ab|$.

If a and b lie on opposite sides of S , then both ac and cb have their endpoints on consecutive sides along the boundary of S . Let S_{ac} be the square contained in S existing as a consequence of Observation 1 when applied on the edge ac . Thus, either a or c lies on a corner of S_{ac} . Furthermore, as S_{ac} is contained in S , it is also empty. Consequently, by Lemma 2, there is a a - c -path φ_{ac} such that $|\varphi_{ac}| \leq (\sqrt{2} + 36) \cdot |ac|$. Analogously, there is a path φ_{cb} such that $|\varphi_{cb}| \leq (\sqrt{2} + 36) \cdot |cb|$. Using elementary geometry, it can be shown that since a and b lie on opposite sides of S , $|ac| + |cb| \leq (1 + \sqrt{2}) \cdot |ab|$. Therefore, the path $\varphi_{ab} = \varphi_{ac} \cup \varphi_{cb}$ is an a - b -path such that $|\varphi_{ab}| \leq (1 + \sqrt{2}) \cdot (\sqrt{2} + 36) \cdot |ab|$. \square

Theorem 1. *The θ_4 -graph of P is a spanner whose stretch factor is at most*

$$(1 + \sqrt{2}) \cdot (\sqrt{2} + 36) \cdot \sqrt{4 + 2\sqrt{2}} \approx 237$$

Proof. Let ν be the shortest path joining s with t in the L_∞ -Delaunay triangulation of P . Bonichon *et al.* [2] proved that the length of ν is at most $\sqrt{4 + 2\sqrt{2}} \cdot |st|$. By replacing every edge in ν with the path in the θ_4 -graph of P that exists by Lemma 3, we obtain an s - t -path of length at most

$$(1 + \sqrt{2}) \cdot (\sqrt{2} + 36) \cdot |\nu| \leq (1 + \sqrt{2}) \cdot (\sqrt{2} + 36) \cdot \sqrt{4 + 2\sqrt{2}} \cdot |st| \quad \square$$

3 Light Paths

We introduce some tools that will help us prove Lemma 1 in Section 4.

For a point p in the θ_4 -graph, recall that there is an edge between p and p_i , the point in $C_i(p)$ nearest to p . We call the edge pp_i an i -edge. Notice that every i -edge is associated with an empty isosceles right triangle. For a point p , the empty triangle generated by its i -edge is denoted by $\Delta_i(p)$.

Let φ be a path that follows only 0- and 1-edges. A 0-edge pp_0 of φ is *light* if no edge of φ crosses the horizontal ray shooting to the right from p . We say

that φ is a *light* path if all its 0-edges are light. In this section we show how to bound the length of a light path with respect to the Euclidean distance between its endpoints.

Lemma 4. *Given a light path φ , every pair of 0-edges of φ has disjoint orthogonal projection on the line defined by the equation $y = -x$.*

Proof. Let s and t be the endpoints of φ . Let pp_0 be any 0-edge of φ and let ν_{p_0} be the *diagonal line* extending the hypotenuse of $\Delta_0(p)$, i.e., ν_{p_0} is a line with slope $+1$ passing through p_0 . Let γ be the path contained in φ that joins p_0 with t . We claim that every point in γ lies below ν_{p_0} . If this claim is true, the diagonal lines constructed from the empty triangles of every 0-edge in φ split the plane into disjoint slabs, each containing a different 0-edge of φ . Thus, their projection on the line defined by the equation $y = -x$ must be disjoint.

To prove that every point in γ lies below ν_{p_0} , notice that every point in γ must lie to the right of p since φ is x -monotone, and below p since pp_0 is light, i.e., γ is contained in $C_0(p)$. Since $\Delta_0(p)$ is empty, no point of γ lies above ν_{p_0} and inside $C_0(p)$ yielding our claim. \square

Given a point w of P , we say that a point p of P is *w-protected* if $C_1(p)$ contains no point of P below or on ℓ_w^- , recall that ℓ_w^- is the line with slope -1 passing through w . In other words, a point p is *w-protected* if either $C_1(p)$ is empty or p_1 lies above ℓ_w^- . Moreover, every point lying above ℓ_w^- is *w-protected* and no point in $C_3(w)$ is *w-protected*.

Given two points s and t such that s lies to the left of t , we aim to construct a path joining s with t in the θ_4 -graph of P . The role of *t-protected* points will be central in this construction. However, as a first step, we relax our goal and prove instead the existence of a light path $\sigma_{s \rightarrow t}$ going from s towards t that does not necessarily end at t .

To construct $\sigma_{s \rightarrow t}$, start at a point $z = s$ and repeat the following steps until reaching either t or a *t-protected* point w lying to the right of t .

- If z is not *t-protected*, then follow its 1-edge, i.e., let $z = z_1$.
- If z is *t-protected*, then follow its 0-edge, i.e., let $z = z_0$.

The pseudocode of this algorithm can be found in Algorithm 1.

Algorithm 1. Given two points s and t of P such that s lies to the left of t , algorithm to compute path $\sigma_{s \rightarrow t}$

```

1: Let  $z = s$ .
2: Append  $s$  to  $\sigma_{s \rightarrow t}$ .
3: while  $z \neq t$  and  $z$  is not a  $t$ -protected point lying to the right of  $t$  do
4:   if  $z$  is  $t$ -protected then  $z = z_0$  else  $z = z_1$ 
5:   Append  $z$  to  $\sigma_{s \rightarrow t}$ .
6: end while
7: return  $\sigma_{s \rightarrow t}$ 
    
```

Lemma 5. *Let s and t be two points of P such that s lies to the left of t . Algorithm 1 produces a light x -monotone path $\sigma_{s \rightarrow t}$ joining s with a t -protected point w such that either $w = t$ or w lies to the right of t . Moreover, every edge on $\sigma_{s \rightarrow t}$ is contained in $S(s)$.*

Proof. By construction, Algorithm 1 finishes only when reaching either t or a t -protected point lying to the right of t . Since every edge of $\sigma_{s \rightarrow t}$ is either a 0-edge or a 1-edge traversed from left to right, the path $\sigma_{s \rightarrow t}$ is x -monotone. The left endpoint of every 0-edge in $\sigma_{s \rightarrow t}$ lies in $C_2(t)$ as it must be t -protected and no t -protected point lies in $C_3(t)$. Thus, if vv_0 is a 0-edge, then v lies in $C_2(t)$ and hence, v_0 lies inside $S(s)$ and above ℓ^+ . Otherwise t would lie inside $\Delta_0(v)$. Therefore, every 0-edge in $\sigma_{s \rightarrow t}$ is contained in $S(s)$.

Every 1-edge in $\sigma_{s \rightarrow t}$ has its two endpoints lying below ℓ^- ; otherwise, we followed the 1-edge of a t -protected point which is not allowed by Step 4 of Algorithm 1. Thus, every 1-edge in $\sigma_{s \rightarrow t}$ lies below ℓ^- and to the right of s . As 1-edges are traversed from bottom to top and the 0-edges of $\sigma_{s \rightarrow t}$ are enclosed by $S(s)$, every 1-edge in $\sigma_{s \rightarrow t}$ is contained in $S(s)$.

Let vv_0 be any 0-edge of $\sigma_{s \rightarrow t}$. Since we followed the 0-edge of v , we know that v is t -protected and hence no point of P lies in $C_1(v)$ and below ℓ^- . As every 1-edge has its two endpoints lying below ℓ^- and $\sigma_{s \rightarrow t}$ is x -monotone, no 1-edge in $\sigma_{s \rightarrow t}$ can have an endpoint in $C_1(v)$. In addition, every 0-edge of $\sigma_{s \rightarrow t}$ joins its left endpoint with a point below it. Thus, no 0-edge of $\sigma_{s \rightarrow t}$ can cross the ray shooting to the right from v . Consequently, vv_0 is light and hence $\sigma_{s \rightarrow t}$ is a light path; see Fig 3. \square

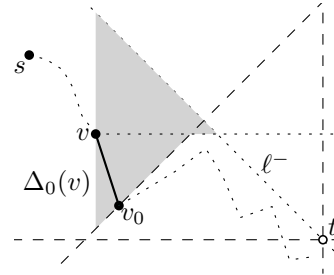


Fig. 3. If v is a t -protected point, then edge vv_0 is light in any path $\sigma_{s \rightarrow t}$ that contains it

Given two points p and q , let $|pq|_x$ and $|pq|_y$ be the absolute differences between their x - and y -coordinates, respectively, i.e., $d_{L_1}(p, q) = |pq|_x + |pq|_y$.

Lemma 6. *Let s and t be two points of P such that s lies to the left of t . If s is t -protected, then $|\sigma_{s \rightarrow t}| \leq 3 \cdot d_{L_1}(s, t)$.*

Proof. To bound the length of $\sigma_{s \rightarrow t}$, we bound the length of its 0-edges and the length of its 1-edges separately. Let Z be the set of all 0-edges in $\sigma_{s \rightarrow t}$ and consider their orthogonal projection on ℓ^- . By Lemma 4, all these projections are disjoint. Moreover, the length of every 0-edge in Z is at most $\sqrt{2}$ times the length of its projection. Let s_\perp be the orthogonal projection of s on ℓ^- and let δ be the segment joining s_\perp with t . Since s is t -protected and $\sigma_{s \rightarrow t}$ is x -monotone, the orthogonal projection of every 0-edge of Z on ℓ^- is contained in δ and hence $\sum_{e \in Z} |e| \leq \sqrt{2} \cdot |\delta|$. Since $|\delta| = d_{L_1}(s, t) / \sqrt{2}$ as depicted in Fig. 4(a), we conclude that $\sum_{e \in Z} |e| \leq d_{L_1}(s, t)$.

Let O be the set of all 1-edges in $\sigma_{s \rightarrow t}$ and let η be the horizontal line passing through t . Since $\sigma_{s \rightarrow t}$ is x -monotone, the orthogonal projections of all edges in

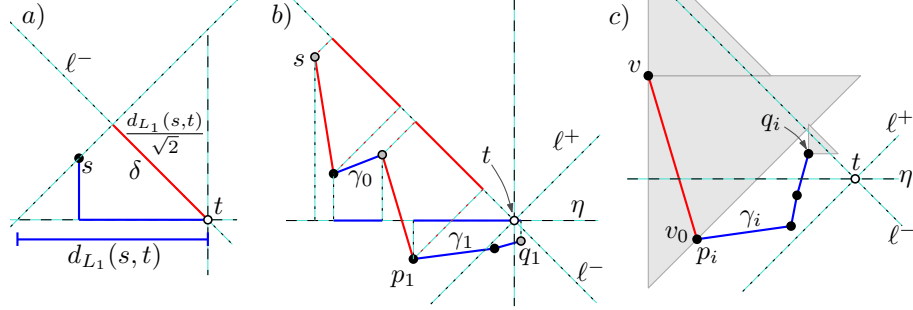


Fig. 4. *a)* The segment δ having length $d_{L_1}(s,t)/\sqrt{2}$. *b)* The 0-edges of $\sigma_{s \rightarrow t}$ have disjoint projections on ℓ^- and the 1-edges have disjoint projections on the horizontal line passing through t . The slope between the endpoints of the maximal paths γ_0 and γ_1 is less than 1. *c)* The slope between p^i with q^i is smaller than 1.

O on η are disjoint. Let $\gamma_0, \dots, \gamma_k$ be the connected components induced by O , i.e., the set of maximal connected paths that can be formed by the 1-edges in O ; see Fig. 4(b). We claim that the slope of the line joining the two endpoints p^i, q^i of every γ_i is smaller than 1. If this claim is true, the length of every γ_i is bounded by $|p^i q^i|_x + |p^i q^i|_y \leq 2 \cdot |p^i q^i|_x$ as each γ_i is x - and y -monotone.

To prove that the slope between p^i and q^i is smaller than 1, let vv_0 be the 0-edge of $\sigma_{s \rightarrow t}$ such that $v_0 = p^i$. Since vv_0 is in $\sigma_{s \rightarrow t}$, v is t -protected by Step 4 of Algorithm 1 and hence, as $\Delta_0(v)$ is empty, q^i must lie below the line with slope +1 passing through p^i yielding our claim; see Fig. 4(c) for an illustration.

Let ω be the segment obtained by shooting a ray from t to the left until hitting the boundary of $S(s)$. We bound the length of all edges in O using the length of ω . Notice that the orthogonal projection of every γ_i on η is contained in ω , except maybe for γ_k whose right endpoint q^k could lie below and to the right of t . Two cases arise: If the projection of γ_k on η is contained in ω , then $\sum_{i=0}^k |\gamma_i| \leq \sum_{i=0}^k 2 \cdot |p^i q^i|_x \leq 2 \cdot |\omega|$. Otherwise, since q^k is t -protected, q^k lies below ℓ^- and hence $d_{L_1}(p^k, q^k) \leq d_{L_1}(p^k, t)$. Moreover, p^k must lie above ℓ^+ as p^k is reached by a 0-edge coming from above η , i.e., $|p^k t|_y < |p^k t|_x$. Therefore,

$$|\gamma_k| \leq d_{L_1}(p^k, q^k) \leq d_{L_1}(p^k, t) = |p^k t|_x + |p^k t|_y \leq 2 \cdot |p^k t|_x$$

Consequently, $\sum_{i=0}^k |\gamma_i| \leq 2 \cdot |p^k t|_x + \sum_{i=0}^{k-1} 2 \cdot |p^i q^i|_x \leq 2 \cdot |\omega|$. Since $|\omega| \leq d_{L_1}(s, t)$, we get that $\sum_{e \in O} |e| = \sum_{i=0}^k |\gamma_i| \leq 2 \cdot d_{L_1}(s, t)$. Thus, $\sigma_{s \rightarrow t}$ is a light path of length at most $\sum_{e \in O} |e| + \sum_{e \in Z} |e| \leq 3 \cdot d_{L_1}(s, t)$. \square

By the construction of the light path in Algorithm 1, we observe the following.

Lemma 7. *Let s and t be two points of P such that s lies to the left of t . If the right endpoint w of $\sigma_{s \rightarrow t}$ is not equal to t , then w lies either above ℓ^+ if $w \in C_1(t)$, or below ℓ^- if $w \in C_0(t)$.*

Proof. If w lies in $C_1(t)$, then by Step 4 of Algorithm 1, w was reached by a 0-edge pw such that p is a t -protected point lying above and to the left of t . As $\Delta_0(p)$ is empty, t lies below the hypotenuse of $\Delta_0(p)$ and hence w lies above ℓ^+ .

Assume that w lies in $C_0(t)$. Notice that w is the only t -protected point of $\sigma_{s \rightarrow t}$ that lies to the right of t ; otherwise, Algorithm 1 finishes before reaching w . By Step 4 of Algorithm 1, every 0-edge of $\sigma_{s \rightarrow t}$ needs to have a t -protected left endpoint. Moreover, every t -protected point of $\sigma_{s \rightarrow t}$, other than w , lies above and to the left of t . Therefore, w is not reached by a 0-edge of $\sigma_{s \rightarrow t}$, i.e., w must be the right endpoint of a 1-edge pw of $\sigma_{s \rightarrow t}$. Notice that w cannot lie above ℓ^- since otherwise p is t -protected and hence Algorithm 1 finishes before reaching w yielding a contradiction. Thus, w lies below ℓ^- . \square

4 One Empty Quadrant

Before stepping into the proof of Lemma 1, we need one last definition. Given a point p of P , the \max_1 -path of p is the longest x -monotone path having p as an endpoint that consists only of 1-edges and contains edge pp_1 . We restate Lemma 1 using the notions of t -protected and s - t -path.

Lemma 1. *Let s and t be two points of P such that t lies in $C_0(s)$. If the top-right quadrant of $S(s)$ is empty and s is t -protected, then there is an s - t -path in the θ_4 -graph of P of length at most $18 \cdot d_{L_1}(s, t)$.*

Proof. Since s is t -protected, no point of P lies above s , to the right of s and below ℓ^- ; see the dark-shaded region in Fig. 5. Let R be the smallest axis-aligned rectangle enclosing s and t and let k be the number of t -protected points inside R , by the general position assumption, these points are strictly contained in R . We prove the lemma by induction on k .

Base Case: Assume that R contains no t -protected point, i.e., $k = 0$. We claim that R must be empty and we prove it by contradiction. Let q be a point in R and note that q cannot lie above ℓ^- as it would be t -protected yielding a contradiction. If q lies below ℓ^- , we can follow the \max_1 -path from q until reaching a t -protected point p lying below ℓ^- . Since s is t -protected, p must lie inside R which is also a contradiction. Thus, R must be empty.

Assume that $s_0 \neq t$ since otherwise the result is trivial. As R is empty and $s_0 \neq t$, s_0 lies below t and above ℓ^+ . Moreover, no point of P lies above t , below ℓ^- and inside $S(s_0)$ since s is t -protected. Thus, if we think of the set of points P rotated 90 degrees clockwise around t , Lemma 6 and Lemma 7 guarantee the existence of an s_0 - t -path γ of length at most $3 \cdot d_{L_1}(s_0, t)$. Since s_0 lies above ℓ^+ , $d_{L_1}(s, s_0) \leq d_{L_1}(s, t)$. Furthermore, $d_{L_1}(s_0, t) \leq 2 \cdot d_{L_1}(s, t)$ as s_0

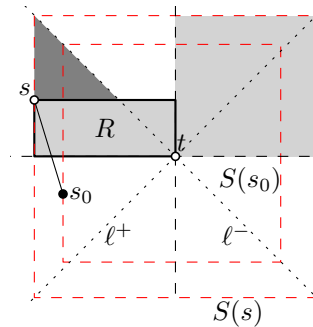


Fig. 5. Base case

lies inside $S(s)$. Thus, by joining ss_0 with γ , we obtain an s - t -path of length at most $7 \cdot d_{L_1}(s, t)$.

Inductive Step: We aim to show the existence of a path γ joining s with a t -protected point $w \in R$ such that the length of γ is at most $18 \cdot d_{L_1}(s, w)$. If this is true, we can merge γ with the w - t -path φ existing by the induction hypothesis to obtain the desired s - t -path with length at most $18 \cdot d_{L_1}(s, t)$. We analyze two cases depending on the position of s_0 with respect to R .

Case 1. Assume that s_0 lies inside R . If s_0 lies above ℓ^- , then s_0 is t -protected and hence we are done after applying our induction hypothesis on s_0 . If s_0 lies below ℓ^- , then we can follow its \max_1 -path to reach a t -protected point w that must lie inside R as s is t -protected. By running Algorithm 1 on s and w , we obtain a path $\sigma_{s \rightarrow w}$ that goes through the edge ss_0 and then follows the \max_1 -path of s_0 until reaching w ; see Fig. 6.

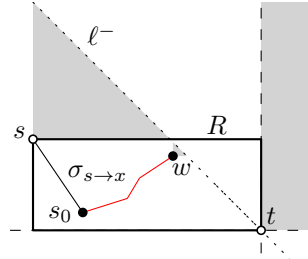


Fig. 6. Case 1

Since s is t -protected and w lies below ℓ^- , s is also w -protected. Therefore, Lemma 6 guarantees that $|\sigma_{s \rightarrow w}| \leq 3 \cdot d_{L_1}(s, w)$. By induction hypothesis on w , there is a w - t -path φ such that $|\varphi| \leq 18 \cdot d_{L_1}(w, t)$. As w lies in R , by joining $\sigma_{s \rightarrow w}$ with φ we obtain the desired s - t -path of length at most $18 \cdot d_{L_1}(s, t)$.

Case 2. Assume that s_0 does not lie in R . This implies that s_0 lies below t . Assume also that $\sigma_{s \rightarrow t}$ does not reach t ; otherwise we are done since $|\sigma_{s \rightarrow t}| \leq 3 \cdot d_{L_1}(s, t)$. Thus, as the top-right quadrant of $S(s)$ is empty, $\sigma_{s \rightarrow t}$ ends at a t -protected point z lying in the bottom-right quadrant of $S(s)$. We consider two sub-cases depending on whether $\sigma_{s \rightarrow t}$ contains a point inside R or not.

Case 2.1. If $\sigma_{s \rightarrow t}$ contains a point inside R , let w be the first t -protected point of $\sigma_{s \rightarrow t}$ after s and note that w also lies inside R since s is t -protected. Notice that the part of $\sigma_{s \rightarrow t}$ going from s to w is in fact equal to $\sigma_{s \rightarrow w}$ since w lies above t and only 1-edges were followed after s_0 by Step 4 of Algorithm 1; see Fig. 7. Thus, as s is also w -protected, the length of $\sigma_{s \rightarrow w}$ is bounded by $3 \cdot d_{L_1}(s, w)$ by Lemma 6. Hence, we can apply the induction hypothesis on w as before and obtain the desired s - t -path.

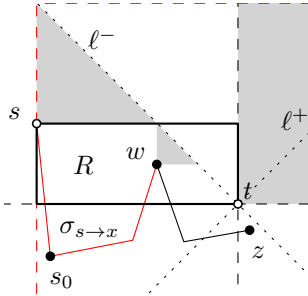


Fig. 7. Case 2.1

Case 2.2. If $\sigma_{s \rightarrow t}$ does not contain a point inside R , then $\sigma_{s \rightarrow t}$ follows only 1-edges from s_0 until reaching z in the bottom-right quadrant of $S(s)$; see Fig. 8(a) for an illustration of this case.

Let P^* be the set of points obtained by reflecting P on line ℓ^+ . Since this reflection preserves 1-edges, it preserves t -protected points. Therefore, if z^* is the reflection of z , then z^* lies in $C_2(t)$ and is also t -protected. Hence, we can use Algorithm 1 to produce a path $\sigma_{z^* \rightarrow t}$ in the θ_4 -graph of P^* . Let $\gamma_{z \rightarrow t}$ be the path in the θ_4 -graph of P obtained by reflecting $\sigma_{z^* \rightarrow t}$ on ℓ^+ . Note that $\gamma_{z \rightarrow t}$ is

a path that uses only 1- and 2-edges. Because the top-right quadrant of $S(s)$ is empty, $\gamma_{z \rightarrow t}$ ends at a point w such that w is either equal to t or w lies in the top-left quadrant of $S(s)$; see Fig. 8(a). Since z lies inside $S(s)$, $d_{L_1}(z, t) \leq 2 \cdot d_{L_1}(s, t)$. Hence, by Lemma 6, the length of $\sigma_{s \rightarrow t} \cup \gamma_{z \rightarrow t}$ is at most $|\sigma_{s \rightarrow t}| + |\gamma_{z \rightarrow t}| \leq 3 \cdot d_{L_1}(s, t) + 3 \cdot d_{L_1}(z, t) \leq 3 \cdot d_{L_1}(s, t) + 6 \cdot d_{L_1}(s, t) = 9 \cdot d_{L_1}(s, t)$.

Two cases arise: If $\gamma_{z \rightarrow t}$ reaches t ($w = t$), then we are done since $\sigma_{s \rightarrow t} \cup \gamma_{z \rightarrow t}$ joins s with t through z and its length is at most $18 \cdot d_{L_1}(s, t)$.

If $\gamma_{z \rightarrow t}$ does not reach t ($w \neq t$), then w lies below ℓ^- by Lemma 7 applied on path $\sigma_{z^* \rightarrow t}$. Moreover, as s is t -protected, no point in $C_1(s)$ can be reached by $\gamma_{z \rightarrow t}$ and hence w must lie inside R . We claim that $d_{L_1}(s, t) \leq 2 \cdot d_{L_1}(s, w)$. If this claim is true, $|\sigma_{s \rightarrow t} \cup \gamma_{z \rightarrow t}| \leq 9 \cdot d_{L_1}(s, t) \leq 18 \cdot d_{L_1}(s, w)$. Furthermore, by the induction hypothesis, there is a path φ joining w with t of length at most $18 \cdot d_{L_1}(w, t)$. Consequently, by joining $\sigma_{s \rightarrow t}$, $\gamma_{z \rightarrow t}$ and φ , we obtain an s - t -path of length at most $18 \cdot d_{L_1}(s, w) + 18 \cdot d_{L_1}(w, t) = 18 \cdot d_{L_1}(s, t)$.

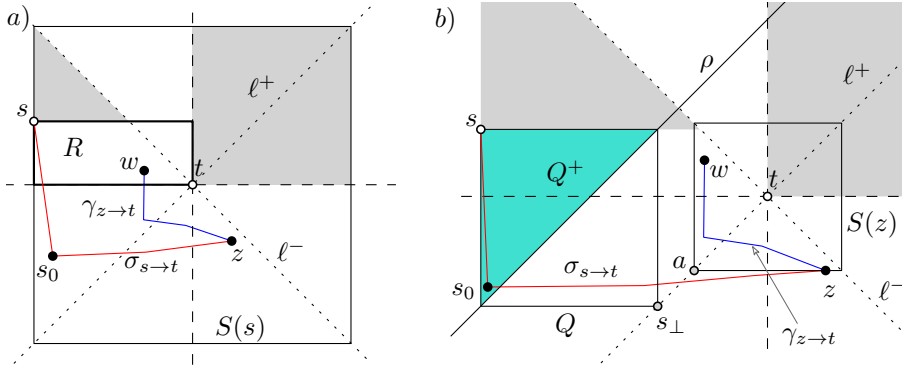


Fig. 8. a) Case 2.2 in the proof of Lemma 1, path $\sigma_{s \rightarrow t}$ has no point inside R and reaches a point z lying in the bottom-right quadrant of $S(s)$. b) The inductive argument proving that the point w , reached after taking the path $\gamma_{z \rightarrow t}$, lies outside of the triangle Q^+ containing all the points above ρ and below s . As s is t -protected, the region above s and below ρ is empty.

To prove that $d_{L_1}(s, t) \leq 2 \cdot d_{L_1}(s, w)$, let s_\perp be the orthogonal projection of s on ℓ^+ . Let ρ be the perpendicular bisector of the segment ss_\perp and notice that for every point y in $C_0(s)$, $d_{L_1}(s, t) \leq 2 \cdot d_{L_1}(s, y)$ if and only if y lies below ρ .

Let Q be the minimum axis-aligned square containing s and s_\perp . Note that ρ splits Q into two equal triangles Q^+ and Q^- as one diagonal of Q is contained in ρ . Assume that Q^+ is the triangle that lies above ρ . Notice that all points lying in $C_0(s)$ and above ρ are contained in Q^+ ; see Fig. 8(b). We prove that w lies outside of Q^+ and hence, that w must lie below ρ .

If s_0 lies below ρ , then the empty triangle $\Delta_0(s)$ contains Q^+ forcing w to lie below ρ . Assume that s_0 lies above ρ . In this case, z lies above s_0 as we only

followed 1-edges to reach z in the construction of $\sigma_{s \rightarrow t}$ by Step 4 of Algorithm 1. Let a be the intersection of ℓ^+ and the ray shooting to the left from z . Notice that w must lie to the right of a as the path $\gamma_{z \rightarrow t}$ is contained in the square $S(z)$ and a is one of its corners. As z lies above s_0 and s_0 lies above s_\perp , we conclude that a is above s_\perp and both lie on ℓ^+ . Therefore, a lies to the right of s_\perp , implying that w lies to the right of s_\perp and hence outside of Q^+ . As we proved that w lies below ρ , we conclude that $d_{L_1}(s, t) \leq 2 \cdot d_{L_1}(s, w)$. \square

5 Lower Bound

We show how to construct a lower bound of 7 for the θ_4 -graph. We start with two points u and w such that w lies in $C_2(u)$ and the difference of their x -coordinates is arbitrarily small. To construct the lower bound, we repeatedly replace a single edge of the shortest u - w -path by placing points in the corners of the empty triangle(s) associated with that edge. The final graph is shown in Fig. 9.

We start out by removing the edge between u and w by placing two points, one inside $\Delta_2(u)$ and one inside $\Delta_0(w)$, both arbitrarily close to the corner that does not contain u nor w . Let v_1 be the point placed in $\Delta_2(u)$. Placing v_1 and the other point in $\Delta_0(w)$ removed edge uw , but created two new shortest paths, uv_1w being one of them. Hence, our next step is to extend this path.

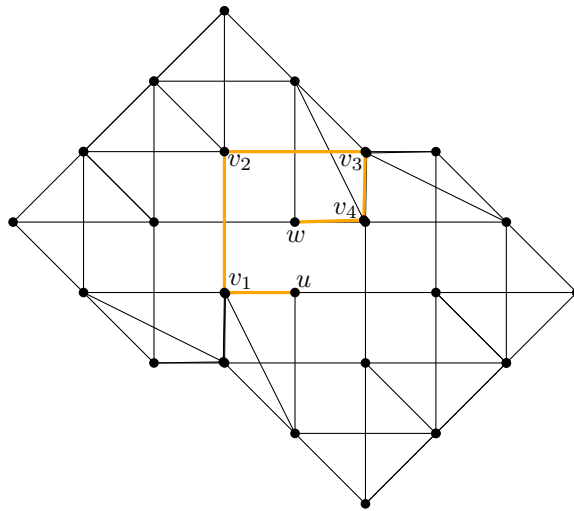


Fig. 9. A lower bound for the θ_4 -graph. One of the shortest paths from u to w goes via v_1, v_2, v_3 , and v_4 .

We remove edge v_1w (and its equivalent in the other path) by placing a point arbitrarily close to the corner of $\Delta_1(v_1)$ and $\Delta_3(w)$ that is farthest from u . Let v_2 be the point placed inside $\Delta_1(v_1)$. Hence, edge v_1w is replaced by the path v_1v_2w . Next, we extend the path again by removing edge v_2w (and its equivalent

edge in the other paths). Like before, we place a point arbitrarily close to the corner of $\Delta_0(v_2)$ and $\Delta_2(w)$ that is farthest from u . Let v_3 be the point placed in $\Delta_0(v_2)$. Hence, edge v_2w is replaced by v_2v_3w .

Finally, we replace edge v_3w (and its equivalent edge in the other paths). For all paths for which this edge lies on the outer face, we place a point in the corner of the two empty triangles defining that edge. However, for edge v_3w which does not lie on the outer face, we place a single point v_4 in the intersection of $\Delta_3(v_3)$ and $\Delta_1(w)$. In this way, edge v_3w is replaced by v_3v_4w . When placing v_4 , we need to ensure that no edge uv_4 is added as this would create a shortcut. This is easily achieved by placing v_4 such that it is closer to v_3 than to w . The resulting graph is shown in Fig. 9.

Lemma 8. *The stretch factor of the θ_4 -graph is at least 7.*

Proof. We look at path $uv_1v_2v_3v_4w$ from Fig. 9. Edges uv_1 , v_3v_4 , and v_4w have length $|uw| - \varepsilon$ and edges v_1v_2 and v_2v_3 have length $2 \cdot |uw| - \varepsilon$, where ε is positive and arbitrarily close to 0. Hence, the stretch factor of this path is arbitrarily close to 7. \square

References

1. Bonichon, N., Gavoille, C., Hanusse, N., Ilcinkas, D.: Connections between Theta-Graphs, Delaunay Triangulations, and Orthogonal Surfaces. In: Thilikos, D.M. (ed.) WG 2010. LNCS, vol. 6410, pp. 266–278. Springer, Heidelberg (2010)
2. Bonichon, N., Gavoille, C., Hanusse, N., Perković, L.: The stretch factor of L_1 - and L_∞ -Delaunay triangulations. In: Epstein, L., Ferragina, P. (eds.) ESA 2012. LNCS, vol. 7501, pp. 205–216. Springer, Heidelberg (2012)
3. Bose, P., Damian, M., Douïeb, K., O’Rourke, J., Seamone, B., Smid, M., Wührer, S.: $\pi/2$ -angle Yao graphs are spanners. *International Journal of Computational Geometry & Applications* 22(1), 61–82 (2012)
4. Bose, P., De Carufel, J.-L., Morin, P., van Renssen, A., Verdonschot, S.: Optimal bounds on theta-graphs: More is not always better. In: *Proceedings of CCCG*, pp. 305–310 (2012)
5. Bose, P., Morin, P., van Renssen, A., Verdonschot, S.: The θ_5 -graph is a spanner. To appear in the proceedings of WG 2013 (2013)
6. Bose, P., van Renssen, A., Verdonschot, S.: On the spanning ratio of theta-graphs. In: Dehne, F., Solis-Oba, R., Sack, J.-R. (eds.) WADS 2013. LNCS, vol. 8037, pp. 182–194. Springer, Heidelberg (2013)
7. Clarkson, K.: Approximation algorithms for shortest path motion planning. In: *Proceedings of STOC*, pp. 56–65 (1987)
8. El Molla, N.M.: Yao spanners for wireless ad hoc networks. Master’s thesis, Villanova University (2009)
9. Keil, J.: Approximating the complete Euclidean graph. In: Karlsson, R., Lingas, A. (eds.) SWAT 1988. LNCS, vol. 318, pp. 208–213. Springer, Heidelberg (1988)
10. Narasimhan, G., Smid, M.: *Geometric Spanner Networks*. Cambridge University Press (2007)
11. Ruppert, J., Seidel, R.: Approximating the d -dimensional complete Euclidean graph. In: *Proceedings of CCCG*, pp. 207–210 (1991)