

The Density of Iterated Crossing Points and a Gap Result for Triangulations of Finite Point Sets

Rolf Klein *

University of Bonn
Institute of Computer Science I
53117 Bonn, Germany
rolf.klein@uni-bonn.de

Martin Kutz

Max-Planck-Institut für Informatik
66123 Saarbrücken, Germany
mkutz@mpi-inf.mpg.de

ABSTRACT

Consider a plane graph G , drawn with straight lines. For every pair a, b of vertices of G , we compare the shortest-path distance between a and b in G (with Euclidean edge lengths) to their actual distance in the plane. The worst-case ratio of these two values, for all pairs of points, is called the *dilation* of G . All finite plane graphs of dilation 1 have been classified. They are closely related to the following iterative procedure. For a given point set $P \subseteq \mathbb{R}^2$, we connect every pair of points in P by a line segment and then add to P all those points where two such line segments cross. Repeating this process infinitely often, yields a limit point set $P^\infty \supseteq P$. This limit set P^∞ is finite if and only if P is contained in the vertex set of a triangulation of dilation 1.

The main result of this paper is the following gap theorem: For any finite point set P in the plane for which P^∞ is infinite, there exists a threshold $\lambda > 1$ such that P is not contained in the vertex set of any finite plane graph of dilation at most λ . As a first ingredient to our proof, we show that such an infinite P^∞ must lie dense in a certain region of the plane. In the second, more difficult part, we then construct a concrete point set P_0 such that any planar graph that contains this set amongst its vertices must have a dilation larger than 1.0000047.

Categories and Subject Descriptors

G.2.2 [Discrete Mathematics]: Graph Theory

General Terms

Theory

*This work was partially supported by the German Research Foundation DFG, grant no. Kl 655/14-1, 14-2.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'06, June 5–7, 2006, Sedona, Arizona, USA.

Copyright 2006 ACM 1-59593-340-9/06/0006 ...\$5.00.

Keywords

dilation, geometric network, complete graph, lower bound, plane graph, spanning ratio, stretch factor, triangulation

1. INTRODUCTION

Let us consider the following scenario described by D. Eppstein in [10]. A typical university campus contains facilities like lecture halls, dorms, library, mensa, and supermarkets, which are connected by some path system. Students in a hurry are tempted to walk straight across the lawn, if the shortcut seems worth it. After a while, this causes new paths to appear. Since their intersections are frequented by many people, they attract coffee shops or other new facilities. Now, people will walk across the lawn to get quickly to a coffee shop, and so on. A natural question discussed in [10] is:

- (A) What happens if this process of forming intersections and treading new paths continues?

Eppstein gives a qualitative answer by completely classifying all those finite point sets in the plane for which this process comes to a halt. In this paper we strengthen that result by showing that, otherwise, not only will infinitely many new intersections be created but these intersections will also lie dense in a region of positive measure. So actually, part of the lawn will completely be gone.

In the unfortunate case that the process would not come to a halt for an initial set of sites, we now assume that students could be convinced to stay on the official paths if the resulting detours, as compared to straight-line walks, were sufficiently small. In this paper we address this relaxation of the exact problem above, which has been brought up recently by Ebbbers-Baumann et al. [8]:

- (B) Could the original campus map be designed in such a way that the process never starts because the advantage of walking straight is always negligible?

The main result of this paper is a negative answer to Question B. For a formal statement, a few precise definitions are indispensable.

Definition 1. For a point set $P \subseteq \mathbb{R}^2$ let $\mathcal{L}(P)$ denote the set of all line segments between pairs of points in P . We define the *complete intersection set* \hat{P} of P to be the union of P with the set of crossings points of lines in $\mathcal{L}(P)$; formally,

$\dot{P} = P \cup \bigcup \{\ell_1 \cap \ell_2 \mid \ell_1, \ell_2 \in \mathcal{L}(P) \text{ and } \ell_1, \ell_2 \text{ not parallel}\}$. Then define P^k for all $k \geq 0$ by $P^0 = P$ and $P^{k+1} = (P^k)^\cdot$, and finally let $P^\infty := \bigcup_{k \geq 0} P^k$. We call a finite point set P *stable* if $\dot{P} = P$ holds, and *stabilizing* if P^∞ is finite.

We have to exclude intersections of parallel segments in the above definition for the simple reason that such overlaps would give one-dimensional intersections and not discrete points. This choice corresponds well with our campus scenario: four locations on a straight line should not create any new “crossings” between them.

We can now make our strengthened answer to Question A precise.

THEOREM 1. *For every non-stabilizing finite point set $P \subset \mathbb{R}^2$, the limit set P^∞ lies dense in some region of positive measure.*

In order to capture the relaxed version in Question B, we need the concept of the *dilation* of a plane graph.

Definition 2. Let $G = (V, E)$ be a plane graph, i.e., its vertices are mapped to points in the plane and the edges are mapped to connecting curves that pairwise do not cross. The *dilation* of G is defined as

$$\delta(G) := \max_{p \neq q \in V} \frac{d_G(p, q)}{|pq|}, \quad (1)$$

where $d_G(p, q)$ denotes the length of a shortest path in G connecting p and q , and $|pq|$ is the Euclidean distance.

There are obvious connections between stable and stabilizing point sets and the dilation of plane graphs.

LEMMA 1. *If a finite point set V is stable then there exists a unique maximal triangulation T of V and this T is a graph of dilation 1. Converseley, every dilation-1 graph is a triangulation and its vertex set is stable. Exactly the subsets of finite stable sets are stabilizing.*

The following definition provides the formal notion behind Question B.

Definition 3. Let P be a finite set of points in the plane. The *dilation*, $\Delta(P)$, of P is defined as

$$\Delta(P) := \inf \left\{ \delta(T) \mid \begin{array}{l} T = (V, E) \text{ finite triangulation} \\ \text{with } P \subseteq V \end{array} \right\}.$$

We should remark that the dilation of a point set is invariant under scaling because the quotient in (1) is. Our main result now reads:

THEOREM 2. *Every non-stabilizing finite point set in the plane has dilation strictly greater than 1.*

In other words, this gap theorem states that if a point set P is not one of the few nicely behaved ones then there exists a threshold $\lambda > 1$ such that P cannot be contained in any finite plane graph of dilation smaller than λ .

1.1 Related work

The problem of how to connect a given point set by a network of small dilation has been extensively studied in the context of spanners, where the dilation of a graph is also called its stretch factor or spanning ratio. There exists a wealth of results on constructing spanners with a dilation arbitrarily close to 1, that have other nice properties like, e.g., linear size, a weight not much exceeding that of the minimum spanning tree, or bounded degree; see the monographs [9] and [18]. But these spanners can contain many edge crossings.

If crossings are not allowed (since bridges are too expensive) one could ask for the plane graph of lowest possible dilation over the given point set, P . Such a graph can be assumed to be a triangulation of P , because curved edges could be pulled taut and adding non-crossing edges cannot increase the dilation. In his handbook chapter [9], Eppstein has posed the following open problems.

- (Problem 8) Can the minimum dilation triangulation of a given point set be computed in polynomial time?
- (Problem 9) How large a dilation can the minimum dilation triangulation have, in the worst case?

In Problem 9, an upper bound is the dilation of the Delaunay triangulation, which has been shown to be less than 2.42 by Keil and Gutwin [14], while a tight bound of $\pi/2$ is conjectured. A lower bound is $\sqrt{2}$, attained by the vertices of a square. A partial answer to Problem 8 was recently given by Eppstein and Wortman [11] who provided a randomized $O(n \log n)$ algorithm for computing an optimal star center for n points in d -space.

The problem setting we consider in this paper can be seen as the “Steiner point” variant of Problem 9 above: Instead of looking for an optimal triangulation on just the given point set, we allow ourselves the introduction of a finite number of auxiliary “Steiner points” to help reduce the overall dilation. The crucial point here is, of course, that dilation amongst the new points, which we introduce to reduce detours for the given vertices, also needs to be controlled. Theorem 2 shows that essentially every point configuration has an intrinsic irregularity that cannot be smoothed out. Ebberts-Baumann et al. [8] have studied Problem 9 in this model. They provided an upper bound of 1.1247 but left open the question of a non-trivial lower bound for finite point sets. The present paper fills this gap.

We have already emphasized that a point set is stabilizing if and only if it is contained in the vertex set of a triangulation of dilation 1. Figure 1 shows the complete classification of such triangulations from [10]. There are two infinite families and one exceptional case.

Quite recently, a related measure called *geometric dilation* has been introduced, see [1, 3, 4, 5, 6, 7], where *all* points of the graph, vertices and interior edge points alike, are considered in computing the dilation. The small difference in definition leads to rather different results. For example, plane graphs of minimum geometric dilation tend to have curved edges. Yet, it is interesting to observe that in this model non-trivial upper and lower bounds could be established. Each finite point set can be embedded into a plane graph of geometric dilation ≤ 1.678 , and there are point sets for which each plane graph containing them has




Figure 1: The triangulations of dilation 1.

a geometric dilation $\geq (1 + 10^{-11})\pi/2$; see [3, 7]. These results are based on methods from convex geometry that do not apply here.

For a point set P in general position we can expect a bi-quadratic number of crossings in the complete graph $C(P)$, due to the Crossing-Number Theorem; see, e. g., Matoušek [16]; but this result does not imply where the crossing points are located. Bezdek and Pach [2] have studied iterated intersections of unit circles. Ismailescu and Radoičić [12] have recently shown that the iterated construction of the intersections of all *lines* through the point pairs of a given set P is dense in the whole plane, except for two special cases. Since their methods heavily use intersection points that lie outside the connecting line segments, we cannot adapt them to our problem. In fact, we shall see that our Theorem 1 is actually a stronger statement and implies the result of [12].

1.2 Overview

The rest of this paper is organized as follows. In Section 2 we prove that for each set P of five points in convex position, their limit set P^∞ lies dense in the convex hull of P^1 ; see Figure 2. This result has independently been shown by Kamali [13] by different methods. As a consequence, we obtain the density result stated in Theorem 1. This also yields, as a byproduct, a new simple proof for the density result on iterated line intersections from [12].

In Section 3 we consider triangulations whose dilation is almost equal to 1, and show that they contain good approximations of the complete intersection sets P^k , for any finite subset P of vertices, and for each fixed order k . Those results form the bridge between Theorem 1 about dense iterated intersections and Theorem 2 about point sets with dilation greater than 1. The concept of “approximations” of point sets will be required frequently throughout this paper, so let us introduce a formal notion.

Definition 4. We say that a set $Q \subseteq \mathbb{R}^2$ is an ϵ -cover of a set $P \subseteq \mathbb{R}^2$ if the (closed) ϵ -neighborhood of every point of P contains a point of Q .

In other words, Q is an ϵ -cover for P if the (directed) Hausdorff distance from P to Q is at most ϵ .

In Section 4 we provide a lower bound to the Steiner-point version of Problem 9 above by constructing a point set of dilation larger than 1.0000047. Precisely, the set of 200 points evenly placed on the boundary B of the unit square turns out to have this property (Theorem 5 and Corollary 1).

Our construction uses the following novel technique. We show that if some set V is an ϵ -cover for B (for some sufficiently small ϵ) then the point set V^2 approximates a scaled copy B' of B with a tolerance smaller than the scaling factor times ϵ . If, in addition, V is the vertex set of a triangulation of sufficiently small dilation, then, by the result obtained in Section 3, the set V^2 , and therefore B' , are closely approximated by vertices in V . By infinite descent, this argument shows that V cannot be finite.

Finally, we combine the above results to prove that for each point set P that is not contained in the vertex set of one of the triangulations in Figure 1, there exists a bound $\lambda > 1$ such that each finite plane graph containing P in its vertex set has dilation at least λ . We conclude with some open problems and suggestions for further research.

2. THE DENSITY OF THE SET OF ITERATED CROSSING POINTS

We begin with the density results for dilation 1. Any set of four points in the plane is easily seen to be stabilizing. So let us consider five points in convex position and show that their limit is dense. The generalization to arbitrary non-stabilizing point sets will be easy.

THEOREM 3. *For any set P of five points in convex position, the limit P^∞ lies dense in the region $A = \text{conv}(P^1 \setminus P)$. (See Figure 2.) By compactness, this means that for every $\epsilon > 0$, there exists an index k such that the ϵ -neighborhood of every point in A contains a point from P^k , i.e., P^k is an ϵ -cover for A .*




Figure 2: Five points in convex position with their dense limit shaded.

PROOF. We show that some P^k is an ϵ -cover for the boundary of the region A . Then the claim for the interior of A immediately follows.

Therefore consider the configuration in Figure 3. We assume that the two points x and y there lie in some P^j and so \bar{x} and \bar{y} , their projections onto the boundary line ab , lie in P^{j+1} . Starting from these four points, we now perform a sequence of line intersections to create a new point \bar{z} on the boundary ab , between \bar{x} and \bar{y} , that shall split the segment $\bar{x}\bar{y}$ approximately in half.

As shown in the figure, we form pairs of intersection points x_t, y_t along the lines $x\bar{x}$ and $y\bar{y}$ by zig-zagging between v and w . This way we can obtain a pair x_t, y_t that lies arbitrarily close to the line ab so that the line segments $\bar{x}\bar{y}$ and $x_t y_t$ are almost parallel and also, the distances $|x_t \bar{x}|$ and $|y_t \bar{y}|$ are much smaller than $|\bar{x}\bar{y}|$ and $|x_t y_t|$.




Figure 3: Zig-zagging towards the boundary segment ab .

For the resulting elongated quadrangle $\bar{x}\bar{y}y_t x_t$ we form the intersection z of the diagonals $x_t\bar{y}$ and $y_t\bar{x}$, as shown in Figure 4. The projection of z (along the segment zu) onto the line ab yields the desired point \bar{z} between \bar{x} and \bar{y} on the boundary ab .




Figure 4: Constructing the splitting point \bar{z} between \bar{x} and \bar{y} .

Since the pair x_t, y_t could be pressed arbitrarily close towards the line ab , the point \bar{z} can be placed arbitrarily close to the exact center of the segment $\bar{x}\bar{y}$. The clue here is that the actual slopes of the segments ux, uy , and uz become irrelevant for the position of \bar{z} when we make the quadrangle thin enough.

In order to cover the whole segment ab , we first apply the above construction with the corner points $\bar{x} = a$ and $\bar{y} = b$, using the P^1 -points c and d as x and y , respectively. This yields a pair \bar{z}, z collinear with u such that \bar{z} splits the segment ab almost exactly in half. Then we repeat the construction on the two segments $a\bar{z}$ and $\bar{z}b$ and recurse on subsegments. Note that each step of this process needs the inner points x and y in order to create the subdivision \bar{z} of the segment $\bar{x}\bar{y}$. But since the point \bar{z} is created together with its interior partner z , those points are always available.

Eventually, after having subdivided the whole boundary of A to a sufficiently fine resolution, we easily cover the interior of A by intersections of connections between boundary points. \square

Now the general density result of Theorem 1 follows from the easily verified fact that for any non-stabilizing point set P , the set P^4 must contain five points in convex position.

2.1 A related concept: exterior intersections

If the definition of the intersection set \dot{P} was modified to include all intersections of infinite lines through pairs

of points (instead of connecting segments, only), we would be able to create new points outside the convex hull of the initial set P . This setting was investigated by Ismailescu and Radoičić in [12]. They proved that except for two exceptional cases, all point sets grow densely into the whole plane.

THEOREM 4 (ISMAILESCU AND RADOIČIĆ, 2004). *The limit set of line intersections of a point set that contains four points, one of them in the interior of the convex hull of the others, is everywhere dense in the whole plane.*

We remark that our Theorem 3 directly implies this result, simply because the dense region A in the pentagon can be used to shoot at any point in the plane with arbitrary precision. (The initial five-point configuration required for our result is easily obtained from the four points of Theorem 4 in a few steps.) This gives a very simple new proof for Theorem 4, which in [12] requires almost ten pages of geometric construction and algebraic argument.

3. APPROXIMATING EXACT INTERSECTIONS

The above density results rely crucially on knowing the exact intersection points as they are guaranteed by dilation 1. Next we show that if we allow for only slightly larger dilation, we can still get arbitrarily good approximations of the exact results.

This fact is based on the following observation. Let $T = (V, E)$ denote a triangulation of dilation δ . By definition, each pair of vertices $p, q \in V$ is connected by a path $\pi(p, q)$ in T whose length does not exceed $\delta \cdot |pq|$. In particular, for each vertex t on π , we know that

$$\frac{|pt| + |tq|}{|pq|} \leq \delta$$

holds. Thus, path π is contained in the ellipse of diameter $\delta|pq|$, whose foci are p and q ; see Figure 5.¹

The width of this ellipse is equal to $w = |pq| \sqrt{\delta^2 - 1}$; it tends to zero as the dilation tends to 1. This has a useful consequence.

Suppose that vertices p, w, q, v of T are in convex position, and that dilation δ is sufficiently small. Then two shortest paths $\pi(p, q)$ and $\pi(v, w)$ in T must cross at some vertex t that is contained in the intersection of the ellipses. Clearly, the smaller dilation δ is, the closer must vertex t be to the crossing point, z , of the segments pq and vw .

LEMMA 2. *For every finite point set $P \subseteq \mathbb{R}^2$, every natural number k , and any $\epsilon > 0$, there exists a bound $\delta > 1$ such that the following holds. If $T = (V, E)$ is a triangulation of dilation $\delta(T) \leq \delta$ such that V contains P , then V is an ϵ -cover for P^k .*

PROOF. Otherwise, there would be a sequence of triangulations $T_n = (V_n, E_n)$ of dilation $\delta(T_n) \leq 1 + \frac{1}{n}$ such that each vertex set V_n contains P but none of them is an ϵ -cover for P^k . For each crossing point $a_i, 1 \leq i \leq t$, of P^1 we can find a sequence $(a_{i,n})_n$, where $a_{i,n} \in V_n$, such that $\lim_{n \rightarrow \infty} a_{i,n} = a_i$, by the ellipse argument from above.

¹The diameter of an ellipse with focal points p, q denotes the constant sum $|pb| + |bq|$, where b is any point on the boundary.




Figure 5: In a triangulation of dilation δ , the shortest path between vertices p and q is contained in an ellipse of width $|pq| \sqrt{\delta^2 - 1}$.

Now, let b_j be a crossing point in P^2 ; we can find in each complete intersection set of the set $\{a_{i,n}; 1 \leq i \leq t\} \subseteq V_n$ an element $b'_{j,n}$ such that $\lim_{n \rightarrow \infty} b'_{j,n} = b_j$. By the same argument as before, each set V_n contains some vertex $b_{j,n}$ near $b'_{j,n}$ such that $\lim_{n \rightarrow \infty} b_{j,n} = b_j$ holds, too. Since the total number of points in $\bigcup_{i=1}^k P^i = P^k$ is finite, this argument, repeated k times, shows that some set V_n , of large enough index n , will contain a subset that is an ϵ -cover for P^k —a contradiction. \square

4. CONSTRUCTING POINT SETS OF DILATION LARGER THAN 1

Let us quickly recall the central definitions from the preceding sections. The *dilation*, $\Delta(P)$, of a finite point set P is the infimum of the dilations of all finite triangulations containing P amongst their vertices. We say that a set Q is an ϵ -cover for some point set P if each point of P has a point of Q in its ϵ -neighborhood. Finally, P^k denotes the set of all crossing points that result from iterating the construction of the complete set of intersections k times, starting with point set P .

The goal of this section is to provide concrete point sets Q whose dilation $\Delta(Q)$ can be bounded away from 1. In principle, it could be quite conceivable that by adding an astronomical number of points to a given point set P , we could embed P in a triangulation with dilation arbitrarily close to 1. Adding more and more points could help to reduce the dilation arbitrarily.

Consider, for example, the 10-point triangulation for the corners of a regular 5-gon in Figure 4. This graph, due to Lorenz [15], is the best known embedding for five points in regular position and yields a dilation of 1.02046. Could it be possible to produce arbitrarily better constructions by using arbitrarily many additional points?




Figure 6: An embedding of five points in regular position into a triangulation of dilation 1.02046, due to Lorenz [15].

Think of the construction of such a low-dilation a triangulation $T = (V, E)$ with $V \supseteq P$ for a given point set P as an iterative process, similar to the construction of intersection sets P^k . In order to satisfy a given upper bound on the dilation of T , we would have to introduce a new “intersection point” x for every quadruple of the point that are already vertices of T . However, in contrast to the dilation-1 case, we have some freedom in choosing the exact position of this point x , as characterized by the ellipse construction from the preceding section. It could be possible that a careful choice of all those “near intersection points” eventually, after finitely many steps, produces a stable set of points.

Our approach to showing that such refinements are *not* possible, is to provide a set B that exerts a certain stiffness with respect to the above process; in the sense that if the vertex set of a low-dilation triangulation covers B sufficiently closely, then it must also cover a smaller copy of B .

LEMMA 3. *There are values $0 < a < A$, $0 < \epsilon < a/2$, and $1 < \delta$, such that the following holds. Let V be a point set such that*

1. V is an ϵ -cover for the boundary B of a square of side length $A + a$ and
2. V is vertex set of some triangulation T of dilation $\leq \delta$.

Then V is an $\frac{A-a}{A+a} \cdot \epsilon$ -cover for the boundary B' of a smaller square of side length $A - a$, centered in B .

The crucial point about Lemma 3 is that, relatively to the sidelength of the respective box, the precision at which the smaller box B' is covered by V , is no worse than the precision at which the big box B is covered. This will allow us to apply the lemma recursively to smaller and smaller boxes, producing infinitely many points in V .

PROOF OF LEMMA 3 (SKETCH). Here we present the ideas behind the proof; the technical details are given in the Appendix.

Let v be a point on the lower horizontal edge, e , of the smaller square B' , and let w denote the point on B below v ; see Figure 7. Since V is an ϵ -cover for B , the ϵ -neighborhood of w contains some point p of V . As we move a point t from left to right along the upper edge of B , there is a position where the closure of t 's ϵ -neighborhood contains two points, q and r , of V such that segment pq crosses edge e to the left

7. REFERENCES

- [1] P. Agarwal, R. Klein, Ch. Knauer, S. Langerman, P. Morin, M. Sharir, and M. Soss. Computing the Detour and Spanning Ratio of Paths, Trees and Cycles in 2D and 3D. To appear in *Discrete Comput. Geom.*, 2006.
- [2] K. Bezdek and J. Pach. A Point Set Everywhere Dense in the Plane. *Elemente der Mathematik* 40(4): 81–84, 1985.
- [3] A. Dumitrescu, A. Ebbers-Baumann, A. Grüne, R. Klein, and G. Rote, On the Geometric Dilation of Closed Curves, Graphs, and Point Sets. To appear in *Computational Geometry: Theory and Applications*, 2006.
- [4] A. Dumitrescu, A. Ebbers-Baumann, A. Grüne, R. Klein, and G. Rote. On Geometric Dilation and Halving Chords. In F. Dehne, A. López-Ortiz, J.-R. Sack (eds.) Proceedings WADS'05, pp.244–255, LNCS 3608, Springer-Verlag, 2005.
- [5] A. Ebbers-Baumann, R. Klein, E. Langetepe, and A. Lingas. A Fast Algorithm for Approximating the Detour of a Polygonal Chain. *Computational Geometry: Theory and Applications*, 27:123–134, 2004.
- [6] A. Ebbers-Baumann, A. Grüne, and R. Klein. Geometric Dilation of Closed Planar Curves: New Lower Bounds. To appear in *Computational Geometry: Theory and Applications*, 2006.
- [7] A. Ebbers-Baumann, A. Grüne, and R. Klein. On the Geometric Dilation of Finite Point Sets. *Algorithmica* 44(2), pp. 137–149, 2006.
- [8] A. Ebbers-Baumann, A. Grüne, M. Karpinski, R. Klein, Ch. Knauer, and A. Lingas. Embedding Point Sets into Plane Graphs of Small Dilation. In X. Deng und D. Du (eds.) Proceedings ISAAC 2005, pp. 5–16, LNCS 3827, Springer-Verlag, 2005.
- [9] D. Eppstein. Spanning trees and spanners. In Handbook of Computational Geometry, J.-R. Sack and J. Urrutia (eds.), pp. 425-461. Elsevier, 1999.
- [10] D. Eppstein. The Geometry Junkyard. <http://www.ics.uci.edu/~eppstein/junkyard/dilation-free/>.
- [11] D. Eppstein and A. Wortman. Minimum Dilation Stars. Proceedings 21th Annual ACM Symposium on Computational Geomtry, pp. 321-326, 2005.
- [12] D. Ismailescu and R. Radoičić. A Dense Planar Point Set from Iterated Line Intersections. *Computational Geometry: Theory and Applications*, 27(3), pp. 257–267, 2004.
- [13] S. Kamali. On the Density of Iterated Segment Intersections. Diploma Thesis. Bonn, 2005.
- [14] J.M. Keil and C.A. Gutwin. The Delaunay Triangulation Closely Approximates the Complete Euclidean Graph. *Discrete Comput. Geom.* 7, 1992, pp. 13-28.
- [15] D. Lorenz. On the Dilation of Finite Point Sets. Diploma Thesis. Bonn, 2005.
- [16] J. Matoušek. Lectures on Discrete Geometry. Graduate Texts in Mathematics. Springer-Verlag, 2000.
- [17] G. Narasimhan and M. Smid. Approximating the Stretch Factor of Euclidean Graphs. *SIAM J. Comput.* 30(3), 2000, pp. 978–989.
- [18] G. Narasimhan and M. Smid. Geometric Spanner Networks. Cambridge University Press, to appear.

APPENDIX

Here we provide the technical details of the proof of Lemma 3. They are necessary for computing concrete values of the parameters a, A, ϵ, δ for which this lemma holds and thus, for obtaining Corollary 1.

PROOF OF LEMMA 3. Let v be a point on the lower horizontal edge, e , of B' , and let w be the point on B below v , as shown in Figure 7. Since V is an ϵ -cover for B , there is a point $p \in V$ in the ϵ -neighborhood of w . Let $t \in V$ be a point on the upper horizontal edge of B whose (closed) ϵ -neighborhood contains two points q, r of V such that segment pq passes to the left of v , and pr passes to the right (as we move t from left to right along the upper edge of B' there must be a position where the switch occurs).

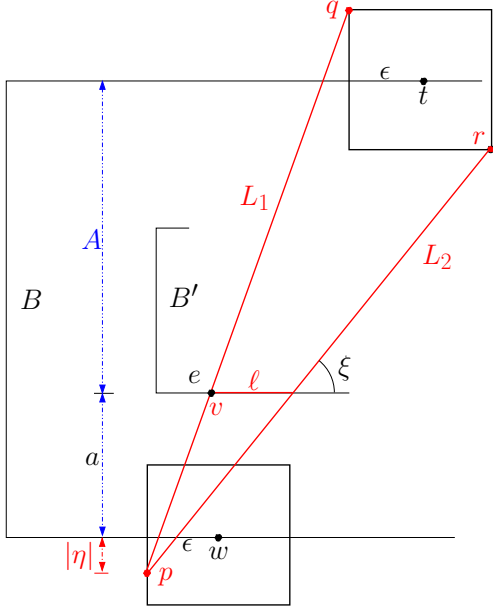


Figure 8: Considering bounding boxes.

We are interested in the maximum distance, ℓ , between the intersection points i, j of pq and pr with edge e , and in the smallest possible angle, ξ , between pq, pr and the X -axis. To simplify computations, we include the ϵ -neighborhoods in their bounding boxes. Also, we move p to the left edge of the lower box, we place q, r in extreme positions in the upper box, and we move the upper box to the right until pq passes through v ; see Figure 8. This can only increase ℓ and decrease ξ .

Let v be the origin. Then the lines L_1, L_2 through pq, pr are given by the equations

$$\begin{aligned} L_1(X) &= \frac{a - \eta}{\epsilon} X \\ L_2(X) &= \frac{(a - \eta)(A + a - \epsilon - \eta)}{\epsilon(A + 3a + \epsilon - 3\eta)} X \\ &\quad - \frac{2(a - \eta)(a - \epsilon - \eta)}{A + 3a + \epsilon - 3\eta}, \end{aligned}$$

where $\eta \in [-\epsilon, \epsilon]$ parametrizes the height of p in its box. It turns out that the size of ℓ is maximized for $\eta = -\epsilon$, which

leads to

$$\ell = \frac{2\epsilon(a + 2\epsilon)}{A + a},$$

whereas the angle ξ is minimized for p in the upper left corner of its box, which results in

$$\tan \xi = \frac{(a - \epsilon)(A + a - 2\epsilon)}{\epsilon(A + 3a - 2\epsilon)}.$$

Next, we consider the pair of points p, r of V . By precondition 2 of Lemma 3, there is a path of length $\leq \delta \cdot |pr|$ in triangulation T that connects p and r . This path must be completely contained in the ellipse of diameter $\delta|pr|$ with foci p and r . Since p, r are at distance at most $(A + a + 2\epsilon)/\sin \xi$, the ellipse is of width at most

$$2 \cdot b = \frac{A + a + 2\epsilon}{\sin \xi} \sqrt{\delta^2 - 1};$$

see Figure 9.




Figure 9: The ellipse with foci p and r and diameter $\delta|pr|$.

Now let v denote the lower left corner of B' and consider the vertical wedge including v . Arguing very generously, we know that within distance ℓ of v , an edge connecting two points of V (namely p, r) of angle at least ξ is passing through the lower edge of B' . Symmetrically, we can consider the horizontal wedge containing v ; see Figure 10. If we surround each of these two segments by a strip of width $2b$, we can be sure that the closest point to v in V cannot be further away from v than the furthest point, m , in the intersection of the strips. Applying the law of sines to the triangle with vertices v, c, m depicted in Figure 10, we obtain for $D = |vm|$:

$$\begin{aligned} \frac{D}{\sin \xi} &= \frac{\ell + \frac{b}{\sin \xi}}{\sin(\xi - \frac{\pi}{4})}, \text{ hence} \\ D &= \left(\ell + \frac{b}{\sin \xi}\right) \frac{\sqrt{2}}{1 - \cot \xi}. \end{aligned}$$

Thus, we have obtained the following intermediate result.




Figure 10: The closest point to v in V is at most distance $D = |vm|$ away from v .

LEMMA 4. V is a D -cover for the four corners of the small square with boundary B' , where D is given by the above expression.

Finally, we approximate an arbitrary point v on the lower edge of B' in the following way. By Lemma 4 there exist points $d, d' \in V$ that approximate the lower corner of B' to an error of D . Clearly, $|dd'| \leq A - a + 2D$, so that the ellipse with foci d, d' of diameter $\delta|dd'|$ is of width at most

$$2b' = (A - a + 2D)\sqrt{\delta^2 - 1}.$$

We consider the intersection of the segment dd' with the least steepest edge, pr , of the vertical wedge depicted in Figure 8. Again, a point of V must be contained in the intersection of the box of width $2b'$ around segment dd' with the box of width $2b$ centered at pr ; see Figure 11.




Figure 11: The distance from v to the closest point in V is at most $|vy|$.

Let y denote the furthest point to v in the strip intersection. Its vertical height over edge e equals $H := D + b'$; hence, $|zy| = H / \sin \xi$, and from the law of cosine we obtain

$$F := |vy| \leq \sqrt{\left(\ell + \frac{b}{\sin \xi}\right)^2 + \left(\frac{H}{\sin \xi}\right)^2 + 2\left(\ell + \frac{b}{\sin \xi}\right)\left(\frac{H}{\sin \xi}\right)\cos \xi}.$$

Using the previous equations, the values of F can now be numerically evaluated, and shown to be less than $\epsilon' := \frac{A-a}{A+a}\epsilon$ for suitably chosen parameters. For example, for $a = 1, A = 15, \epsilon = 0.16$ and $\delta = 1.0000047$ we obtain $F = 0.139904$, which is less than $\epsilon' = 0.14$. This completes the proof of Lemma 3. \square