

Linear-Time Algorithms for Geometric Graphs with Sublinearly Many Crossings

David Eppstein
Dept. of Computer Science
University of California, Irvine
<http://www.ics.uci.edu/~eppstein/>

Michael T. Goodrich
Dept. of Computer Science
University of California, Irvine
<http://www.ics.uci.edu/~goodrich/>

Darren Strash
Dept. of Computer Science
University of California, Irvine
<http://www.ics.uci.edu/~dstrash/>

Abstract

We provide linear-time algorithms for geometric graphs with sublinearly many crossings. That is, we provide algorithms running in $O(n)$ time on connected geometric graphs having n vertices and k crossings, where k is smaller than n by an iterated logarithmic factor. Specific problems we study include Voronoi diagrams and single-source shortest paths. Our algorithms all run in linear time in the standard comparison-based computational model; hence, we make no assumptions about the distribution or bit complexities of edge weights, nor do we utilize unusual bit-level operations on memory words. Instead, our algorithms are based on a *planarization* method that “zeroes in” on edge crossings, together with methods for extending planar separator decompositions to geometric graphs with sublinearly many crossings. Incidentally, our planarization algorithm also solves an open computational geometry problem of Chazelle for triangulating a self-intersecting polygonal chain having n segments and k crossings in linear time, for the case when k is sublinear in n by an iterated logarithmic factor.

1 Introduction

A *geometric graph* [47] is an embedding of a graph $G = (V, E)$ in \mathbf{R}^2 so that each vertex v is associated with a unique point p in \mathbf{R}^2 and each edge is “drawn” as a straight line segment joining the points associated with its end vertices. Moreover, the edges incident on each vertex v are given in angular order around v , so that faces in the embedding of G in \mathbf{R}^2 are well-defined (e.g., using the next-clockwise-edge ordering). Thus, we use the same notation and terminology to refer to G and its embedding. If the edges in G have no crossings, then G is said to be a *plane graph*, while graphs that admit realizations as plane graphs are *planar graphs* [18, 26].

Geometric graphs are natural abstractions of the geometric and connectivity relationships that arise in a number of applications, including road networks, railroad networks, and utility distribution grids, as well as sewer lines and the physical connections defining the Internet¹. An example road network is shown in Figure 1.




Figure 1: A portion of the road network surrounding the location of SODA 2009. This image is from <http://wiki.openstreetmap.org/>, under the Creative Commons attribution-share alike license.

¹The union of these latter two graphs was featured in a recent Google hoax (see <http://www.google.com/tisp/>).




Figure 2: A geometric graph and its planarization.

Although planar graphs and their plane graph realizations have been studied extensively (e.g., see [53]), real-world geometric graphs often contain edge crossings. Recent experimental studies by the first two authors gives empirical evidence that real-world road networks typically have $\Theta(\sqrt{n})$ edge crossings, where n is the number of vertices [23]. Motivated by this real-world example, therefore, we are interested in studying algorithms for connected geometric graphs that have a sublinear number of edge crossings. However, we use a weaker restriction on the number of crossings than the bounds that our evidence suggests for road networks: here we are interested in n -vertex geometric graphs that have at most $O(n/\log^{(c)} n)$ edge crossings, for some constant c , where $\log^{(c)} n$ denotes the c -th iterated logarithm function. We refer to such geometric graphs as *restrained graphs*.

Given an n -vertex geometric graph G , the *planarization*² of G is the graph G' that is defined by the arrangement of the edges in G . That is, as shown in Figure 2, we place a vertex in G' for every vertex and pairwise edge crossing in G , and we create an edge in G' for every maximal edge segment from G that connects exactly two vertices in G' . Likewise, we preserve the (clockwise/counterclockwise) ordering of edges around corresponding vertices in G and G' , and we assume that intersection vertices in G' similarly have their edges given in rotational order. Thus, G' is a plane graph having $n+k$ vertices, where k is the number of pairwise edge crossings among the edges in G . By well-known properties of planar graphs (e.g., see [46, Prop. 2.1.6]), this implies that G' has at most $3n + 3k - 6$ edges, which in

turn implies that G has at most $3n+k-6$ edges. Therefore, by restricting our attention to connected geometric graphs with a sublinear number of edge crossings, we are, by implication, focusing on connected geometric graphs that have $O(n)$ edges in their planarizations.

As mentioned above, a wealth of algorithms are known for planar graphs and plane graphs. Indeed, many of these algorithms, for such problems as single-source shortest paths and minimum spanning trees, run in $O(n)$ time. Much less is known for non-planar geometric graphs, however, which motivates our interest in such graphs in this paper. Specifically, we are interested in the following problems for connected, restrained geometric graphs:

- The *Voronoi diagram* problem, which is also known as the *post office* problem: we are given a set P of k vertices in a geometric graph G and asked to determine for every other vertex v in G the vertex in P that is closest to v according to the graph metric.
- The *single-source shortest path* problem: we are given a vertex s and a geometric graph G and asked to find the shortest paths from s to every other vertex in G .
- The *polygon planarization* problem: given a geometric graph defining a non-simple polygon P having n vertices, compute the arrangement of all the edges of P , including vertices defined by the pairwise crossings of the edges in P .

In all these cases, we desire comparison-based algorithms that require no additional assumptions regarding the distribution of edge weights, so that our algorithms can apply to a wide variety of possible edge weights that

²Our use of this term differs from its use in the graph drawing literature (e.g., see [18]), where it refers to the problem of removing a minimal number of edges to make G be planar.

may vary for different users, including combinations of distance, travel time, toll charges, and subjective scores rating safety and scenic interest [21].

1.1 Previous Related Work

In the algorithms community, there has been considerable prior work on shortest path algorithms for Euclidean graphs (e.g., see [29, 35, 37, 49, 50, 54]), which are geometric graphs where edges are weighted by the lengths of the corresponding line segments. This prior work takes a decidedly different approach than we take in this paper, however, in that it focuses on using special properties of the edge weights that do not hold in the comparison model, whereas we study road networks as geometric graphs with a sublinear number of edge crossings and we desire linear-time algorithms that hold in the comparison model.

The specific problems for which we provide linear-time algorithms are well known in the general algorithms and computational geometry literatures. For general graphs with n vertices and m edges, excellent work can be found on efficient algorithms in the comparison model, including single-source shortest paths [16, 31, 48], which can be found in $O(n \log n + m)$ time [27], and Voronoi diagrams [4, 5], whose graph-theoretic version can be constructed in $O(n \log n + m)$ time [25, 41]. None of these algorithms run in linear time, even for planar graphs. Linear-time algorithms for planar graphs are known for single-source shortest paths [34], but these unfortunately do not immediately translate into linear-time algorithms for non-planar geometric graphs. In addition, there are a number of efficient shortest-path algorithms that make assumptions about edge weights [28, 29, 42, 52]; hence, they are not applicable in the comparison model.

Chazelle [8] shows that any simple polygon can be triangulated in $O(n)$ time and that this algorithm can be extended to determine in $O(n)$ time, for any polygonal chain P , whether or not P contains a self-intersection. In addition, Chazelle posed as an open problem whether or not one can compute the arrangement of a non-simple polygon in $O(n+k)$ time, where k is the number of pairwise edge crossings. Clarkson, Cole, and Tarjan [13, 14] answer this question in the affirmative for polygons with a super-linear number of crossings, as they give a randomized algorithm that solves this problem in $O(n \log^* n + k)$ expected time. There is, to our knowledge, no previous algorithm that solves Chazelle’s open problem, however, for non-simple polygons with a sublinear number of edge crossings.

1.2 Our Results

In this paper, we provide the first linear-time algorithm for planarizing a non-planar connected geometric graph having a number of pairwise edge crossings, k , that is sublinear in the number of vertices, n , by an iterated logarithmic factor. Specifically, we provide a randomized algorithm for planarizing geometric graphs in $O(n + k \log^{(c)} n)$ expected time, which is linear for restrained geometric graphs. Given such a planarization, we show how it can be used to help construct an $O(\sqrt{n})$ -separator decomposition of the original graph in $O(n)$ time. Furthermore, we discuss how such separator decompositions can then be used to produce linear-time algorithms for a number of problems, including Voronoi diagrams and single-source shortest paths. We also show how our planarization algorithm can be used to solve Chazelle’s open problem of planarizing non-simple polygons in expected linear time for polygons having a number of pairwise edge crossings that is sublinear in n by an iterated logarithmic factor. Thus, combining this result with the polygon planarization algorithm of Clarkson, Cole, and Tarjan [13, 14] provides a method for planarizing an n -vertex polygon with k edge crossings in optimal $O(n+k)$ expected time, for all values of k except those in the range $[n/\log^{(c)} n, n \log^* n]$. Our result also implies that the convex hull of restrained non-simple polygons can be constructed in $O(n)$ expected time, which, to the best of our knowledge, was also previously open.

Besides planar separator decompositions, which we discuss below, another one of the techniques we use in this paper is a method for constructing a $(1/r)$ -cutting for the edges of a geometric graph, G . This is a proper triangulation³, T , of the interior of the bounding box containing G such that any triangle t in T intersects at most $(1/r)n$ edges of G . Using existing methods (e.g., see [1, 17, 33]), one can construct such a $(1/r)$ -cutting for G in $O(n \log r + (r/n)k)$ time, where n is the number of vertices in G and k is the number of pairwise edge crossings. However, in our application such a bound would be nonlinear, as we require r to be large. We show, in Section 4, that for connected geometric graphs such a cutting can be constructed in the faster expected time bound $O(ns + (r/n)k)$, where $r \leq n/\log^{(s)} n$.

2 Separator Decompositions

One of the main ingredients we use in our algorithms is the existence of small separators in certain graph families (e.g., see [39, 43]). Several of the algorithms

³A *proper triangulation* is a connected planar geometric graph such that every face is a triangle and every triangular face has exactly three vertices on its boundary.




Figure 3: Trapezoidal decomposition of a sampled subset of input graph edges.

in this paper are based on the use of separators: we use them both as part of our algorithm for finding cuttings of geometric graphs, and later, once the graph has been planarized. Hence, we briefly review these tools here.

Given a graph $G = (V, E)$, a subset W of V is an $f(n)$ -separator if the removal of the vertices in W separates G into two subgraphs G_1 and G_2 , each containing at most δn vertices, for some constant $0 < \delta < 1$. It is well known that planar graphs have $O(\sqrt{n})$ -separators with $\delta = 2/3$, and that such separators can be constructed in $O(n)$ time [39]. Such separators are typically used in divide-and-conquer algorithms, which involve finding a separator, recursively solving the problem in the two separated subgraphs, and then merging the solutions together. If the merge and divide steps can be solved in $o(n)$ time, however, it is useful to have the entire recursive separator decomposition computed in advance; for otherwise there is no way to beat an $O(n \log n)$ time bound. Such a *separator decomposition* defines a binary tree B , such that the root of B is associated with the $f(n)$ -separator for G and the subtrees of this root are defined recursively for the graphs G_1 and G_2 , respectively.

Previous work on separators includes the seminal contribution of Lipton and Tarjan [39], who show that $O(\sqrt{n})$ -sized separators exist for n -vertex planar graphs and these can be computed in $O(n)$ time. Goodrich [30] shows that recursive $O(\sqrt{n})$ -separator decompositions can be constructed for planar graphs in $O(n)$ time. A related concept is that of geometric separators, which use geometric objects to define separators in graphs defined by systems of intersecting disks (e.g., see [3, 44, 45, 51]). Eppstein *et al.* [24] provide a linear-time construction algorithm for geometric separators which translates into an $O(n \log n)$ recursive separator

decomposition algorithm.

Because restrained graphs are not planar, the result of Goodrich does not immediately apply. However, it can be applied once we have planarized the graph, and it can also be applied to planar structures formed from subsets of the graph, such as the one we describe in the next section.

3 Trapezoidal Decomposition of a Sample

Suppose we are given a geometric graph G having n vertices and k pairwise intersections among its edges. In this section, we describe our algorithm for constructing a trapezoidal decomposition of a random sample of the edges of G . That is, given the sample of edges, we construct the arrangement of these edges together with a set of vertical line segments through each edge endpoint and crossing, where each such segment is maximal with respect to the property of not crossing any other sampled edge, as shown below. (See Figure 3.) Our method is parameterized by s where $r \leq n / \log^{(s)} n$, and the sample probability is inversely proportional to $\log^{(s)} n$. We will later show how to refine this sample so that we can produce a cutting and then a planarization of G .

This first step of our algorithm is essentially the same as performing s levels of the Clarkson, Cole, and Tarjan algorithm, except that their method is for polygonal chains, whereas ours is for geometric graphs. Thus, we describe it at a high level.

Our algorithm begins with a trivial trapezoidal decomposition T_0 containing a single trapezoid that encloses all of G . Call this trapezoid t . Let $C(t) = E$ be the *conflict list* for t , that is, the set of edges from G that intersect the interior of t . Then, for $i = 1$ to s , we perform the following computation.

1. Find a random sample S_i of size $n/\log^{(i)} n$, of the edges in G , and for each trapezoid t in T_{i-1} , use the Bentley-Ottmann algorithm [6] to construct the trapezoidal decomposition of the arrangement of the segments in $C(t) \cap S_i$. Once all these trapezoidal decompositions are constructed, merge them together to create a single trapezoidal decomposition, T_i , for the segments in S_i . To be consistent with Clarkson, Cole, and Tarjan, we choose the samples such that $S_1 \subset S_2 \subset \dots \subset S_s$.
2. Perform a depth-first traversal of G , while keeping track of the trapezoids in the trapezoidal decomposition that are intersected during the walk, so as to determine, for each trapezoid t in T_i , the set $C(t)$. Since the geometric graph is connected, we never have to restart the depth-first traversal from a node whose location we do not already know. We can therefore use the arrangement of the sampled line segments to keep track of the intersected trapezoids at each step of the traversal. Thus we eliminate the need for time-consuming point-location data structure lookups.

Let $T = T_s$ be the resulting final trapezoidal decomposition we get from this computation, and let $S = S_s$ be the final random sample. Using the framework established by Clarkson and Shor [15] for randomized divide-and-conquer algorithms, such as this, we can show that

$$(3.1) \quad E(|T|) = O\left(r + \left(\frac{r}{n}\right)^2 k\right)$$

and

$$(3.2) \quad E\left(\sum_{t \in T} |C(t)|\right) = O\left(n + \left(\frac{r}{n}\right) k\right).$$

In particular, Equation (3.1) is from their Lemma 4.1 and Equation (3.2) follows from their Corollary 4.4. The number of steps in the depth-first traversal is proportional to the total size of the conflict lists of the input geometric graph with the trapezoidal decomposition, which as we have seen above is small. A step from one trapezoid to a horizontally adjacent trapezoid may be accomplished in constant time, but a single trapezoid may have a non-constant number of neighbors above and below it, causing steps in those directions to take longer. But as Clarkson, Cole, and Tarjan show, the sum over all trapezoids of the conflict list size of the trapezoid multiplied by its number of neighbors remains linear in expectation, and this sum bounds the time to step vertically from one trapezoid to another using a sequential search along the trapezoid boundary to find the neighboring trapezoid. Therefore, we have the following preliminary result:

LEMMA 3.1. *Given a connected geometric graph G with n edges and k pairwise edge crossings, and a parameter s , we can in expected time $O(ns + (r/n)k)$ find a random sample of $r = O(n/\log^{(s)} n)$ edges from G , the trapezoidal decomposition induced by the sample, and the set of edges of G crossing each trapezoid of the sample.*

4 Cuttings

At this stage we take a detour from the Clarkson, Cole, and Tarjan algorithm. For each trapezoid g in T , let $\alpha_g = |C(g)|r/n$. That is, α_g is the degree of excess that the conflict list for g has beyond what we would like for a $(1/r)$ -cutting. For each trapezoid t with $\alpha_t > 1$, we form a random sample, R_t , of $C(t)$ of size $2b\alpha_t \log \alpha_t$, where b is the constant K_{\max} from Corollary 4.4 of Clarkson-Shor [15]. We then form the trapezoidal decomposition, T_t of the arrangement of the segments in R_t using any quadratic-time line segment arrangement algorithm [2, 7, 10, 19]. Thus, by Corollary 4.4 from Clarkson-Shor [15], the maximum size of any conflict list of a trapezoid in T_t is expected to be less than

$$\begin{aligned} \left(\frac{|C(t)|}{|R_t|}\right) \log |R_t| &= \binom{n}{r} \left(\frac{1}{\log \alpha_t^2}\right) \log(2\alpha_t \log \alpha_t) \\ &\leq \frac{n}{r}, \end{aligned}$$

for $\alpha_t \geq 4$. Thus, we can repeat the above algorithm an expected constant number of times until we have this condition satisfied, which gives us one of the crucial properties of a $(1/r)$ -cutting: namely, that each cell intersects at most (n/r) edges of G .

In addition, the number of new trapezoids created inside t , as well as the running time for creating the trapezoidal diagram T_t , is certainly at most $O(|R_t|^2)$, which is $O(\alpha_t^2 \log^2 \alpha_t)$. More importantly, we have the following:

LEMMA 4.1. *Given the above construction applied to each trapezoid t in T , then*

$$E\left(\sum_{t \in T} \alpha_t^2 \log^2 \alpha_t\right) = O\left(r + \left(\frac{r}{n}\right)^2 k\right).$$

Proof. Our proof is based on an application of Theorem 3.6 from the Clarkson-Shor framework. To apply this theorem, we bound

$$E\left(\sum_{t \in T} \alpha_t^2 \log^2 \alpha_t\right)$$

by bounding the term, $\alpha_t^2 \log^2 \alpha_t$, by

$$W\left(\binom{|C(t)|}{c}\right),$$

where W is a positive concave function on \mathbf{R}^+ and c is a constant. Here, for the sake of an upper bound, we take $c = 3$ and we define

$$W(x) = \left(\frac{x^{1/3}}{N}\right)^2 \log^2 \frac{x^{1/3} + N}{N},$$

where $N = n/r$. Finally, to apply Theorem 3.6 from [15], we need to observe that the number of trapezoids in T that have a conflict list size at most c is proportional to the number of trapezoids in T that have a conflict list size at least 0, which is $|T|$. To see this, note that we can extend the vertical edges of any trapezoid in T in at most $O(1)$ ways until it hits $i = 1, 2, 3$ other edges of the random sample, S , at which point we can extend this trapezoid horizontally in $O(1)$ ways until we hit 3 segments in total. Therefore, by Theorem 3.6 from [15],

$$E\left(\sum_{t \in T} \alpha_t^2 \log^2 \alpha_t\right)$$

is

$$O\left(r + \left(\frac{r}{n}\right)^2 k\right).$$

□

Thus, our refined trapezoidal decomposition, T' , will have size proportional to $|T|$. It is still not quite a $(1/r)$ -cutting, however, as it is not a proper triangulation. Indeed, some trapezoids may have many more than 4 vertices on their boundaries (see Figure 4).




Figure 4: Many trapezoids may be adjacent to another trapezoid along its top or bottom edges.

To refine T' into a proper triangulation, we borrow an idea from the fractional cascading framework of Chazelle and Guibas [?] to first refine T' into a trapezoidal decomposition such that each trapezoid has $O(1)$ vertices on its boundary, while keeping the total number of trapezoids to be $O(|T'|)$, which is expected to be

$$O\left(r + \left(\frac{r}{n}\right)^2 k\right).$$

By triangulating the interior of each such trapezoid, we will get a $(1/r)$ -cutting whose size is still $O(|T'|)$.

Construct the graph-theoretic planar dual U to T' , and note that we can direct the edges of U so as to define four directed-acyclic graphs, which respectively define the partial orders “below,” “above,” “left-of,” and “right-of” among the trapezoids. Without loss of generality, let us direct U according to the “below” relation, perform a topological sort, and process the trapezoids of T' from top to bottom according to this ordering. When processing a trapezoid, t , we assume inductively that we have determined the ordered list of vertices $V_t = (v_1, v_2, \dots, v_j)$ on t 's upper edge, which are bottom vertices of trapezoids above t . To process t we choose every other vertex, v_{2i} , in V_t and extend a vertical segment from v_{2i} to the bottom of t to split t in two for each such v_{2i} . Doing this for every other vertex in V_t , therefore, splits t and increases the number of trapezoids by $\lfloor |V_t|/2 \rfloor$. We then repeat this computation by considering the new set of trapezoids according to the “above” relation, from bottom to top. Next, we do a similar computation for the “left-of” and “right-of” relations (except that now we extend segments parallel to the top or bottom edges of our trapezoid in a way that partitions its interior into non-crossing trapezoids). When we have completed this last scan of the trapezoids, we will have created a trapezoidal decomposition such that each trapezoid has $O(1)$ vertices on its edges. More importantly, we also have the following:

LEMMA 4.2. *The total number of trapezoids created by the above refinement process is $O(|T|)$, which has expected value $O(r + (r/n)^2 k)$.*

Proof. We have already established that $E(|T|)$ is $O(r + (r/n)^2 k)$ and that $E(|T'|)$ is $O(E(|T|))$. So we have yet to show that the number of new trapezoids created during any of our splitting processes is $O(|T'|)$. We do this by an accounting argument. Without loss of generality, consider the processing according to the “below” relation. Assume, for the sake of our analysis, that, at the beginning of our computation, we give each vertical edge in our trapezoidal decomposition \$2 and we require every vertical edge at the end of the process to have at least \$1. When we extend a vertical ray from an even numbered vertex v_{2i} at the top of a trapezoid, t we can assume inductively that the vertical edge above v_{2i} has \$2, as does the vertical edge directly to the left of this edge (which hits t at vertex v_{2i-1}). Let us take \$1 from this vertical edge and from the one that hits t at v_{2i} , which leaves \$1 at each of those edges, and use the \$2 to pay for the new vertical edge that we then extend through t . Therefore, since the

two vertical edges we just took money from will not be processed again, we can process each trapezoid and pay for every action, while keeping \$1 for each trapezoid in our refined trapezoidal decomposition. Repeating this accounting argument for the “above,” “left-of,” and “right-of” relations completes the proof. \square

Given a trapezoidal diagram having $O(1)$ vertices on the boundary of each trapezoid, and each trapezoid intersecting at most (n/r) edges of our geometric graph G we can easily triangulate each trapezoidal face in this diagram to turn it into a $(1/r)$ -cutting with a number of triangles that is proportional to the number of trapezoids. Thus, putting all the pieces together, we get the following.

THEOREM 4.1. *Given a connected geometric graph G having n vertices and k pairwise edge crossings, one can construct a $(1/r)$ -cutting for the edges of G of expected size $O(r + (r/n)^2 k)$ in expected time $O(ns + (r/n)k)$, for $r \leq n/\log^{(s)} n$.*

Taking s as a constant gives us such a $(1/r)$ -cutting of expected size $O(r + (r/n)^2 k)$ in expected time $O(n + (r/n)k)$, and taking $s = \log^* n$ gives us a $(1/r)$ -cutting of the same expected size (but with a potentially larger r) in expected time $O(n \log^* n + (r/n)k)$, for any $r \leq n$. Since, in our applications involving restrained geometric graphs, k is sublinear in n by an iterated logarithmic factor, we will be taking s to be a constant.

5 Planarization

In this section, we describe how to planarize a connected geometric graph G having n vertices and k edge crossings. We begin by using the method of Theorem 4.1 to construct a $(1/r)$ -cutting, C , of the edges of G of expected size $O(r + (r/n)^2 k)$ in expected time $O(n + (r/n)k)$, where $r = n/\log^{(c+1)} n$, for a fixed constant $c \geq 1$. We then do a depth-first search of G , keeping track of the triangles we cross in C as we go, to compute, for each triangle t in C , the set, $C(t)$, of at most (n/r) edges of G that intersect t . This takes $O(|C|n/r)$ time, which has expectation $O(n + (r/n)k)$.

We then apply Goodrich’s separator decomposition algorithm [30] to construct an $O(\sqrt{|D|})$ -separator decomposition of the graph-theoretic dual, D , to C . Rather than taking this decomposition all the way to the point where we would have subgraphs of D of constant size, however, we stop when subgraphs have size $O(\log^2(n/r))$; hence, have separators of size $O(\log(n/r))$. Since C is a triangulation, D has degree 3; hence, any vertex separator for D of size g also gives us an edge separator for D of size at most $3g$. Moreover, each edge of D corresponds to a triangle edge in

C , which in turn crosses at most (n/r) edges of G . For each separator H in our decomposition, therefore, we can sort the edges of G that cross each boundary of a triangle in the separator in time $O((n/r)\log(n/r))$ time. There are $O(|D|/\log^2(n/r))$ nodes at this level of the separator decomposition tree; hence, there are $O(|D|/\log^2(n/r)) \times O(\log(n/r)) = O(|D|/\log(n/r))$ triangles involved. Thus, the total time for all these sorts is $O(|D|/\log(n/r)) = O(n + k)$.

After performing all these sorts of edges on the boundaries of triangles in our separators, we can imagine that we have used these boundaries to cut G into $O(|D|/\log(n/r))$ regions (including each triangle in one of our separators), such that the edges of G intersecting each region boundary are given in sorted order. The total size of each subgraph is $O((n/r)\log^2(n/r))$. Moreover, the boundaries of these regions form a planar subdivision. Thus, we have just subdivided our geometric graph G into $O(|D|/\log(n/r))$ disjoint geometric graphs. In other words, all k edge crossings in G have been isolated into these small subgraphs.

For each subgraph G_i , use Chazelle’s algorithm [8] to test if all the faces of G_i are simple in $O(|G_i|)$ time. If all the faces of G_i are in fact simple, then G_i clearly contains no edge crossings. Thus, we can identify each small subgraph in this partition that contains an intersection in time $O(|C|(n/r) + |G|)$, which has expectation $O(n + k)$.

Clearly, there are at most k such subgraphs that contain edge crossings. We complete our planarization algorithm, therefore, by running the Bentley-Ottmann algorithm [6] for each subgraph of G that is identified as having at least one edge crossing. The time for each such invocation of the Bentley-Ottmann algorithm is $O((n/r)\log^3(n/r) + k' \log(n/r))$, where $k' \geq 1$ is the number of edge crossings found. Summing this over k regions implies that the total time needed to complete the planarization of G is $O(k(n/r)\log^3(n/r))$. Substituting for r , we see that this time is $O(k \log^{(c+1)} n \log^3 \log^{(c+1)} n)$, which is $O(k \log^{(c)} n)$. Therefore, we have the following:

THEOREM 5.1. *Suppose one is given a connected geometric graph G with n vertices and k edge crossings, together with a $(1/r)$ -cutting of the edges of G of size $O(r + (r/n)^2 k)$, for $r = n/\log^{(c+1)} n$. Then one can construct a planarization of G (and the trapezoidal decomposition of the arrangement of G ’s edges), in time $O(n + k \log^{(c)} n)$.*

Combining this result with Theorem 4.1, we get the following corollary.

COROLLARY 5.1. *Given a connected geometric graph G having n vertices and k pairwise edge crossings, one can construct a planarization of G in expected time $O(n + k \log^{(c)} n)$.*

6 Applications

In this section, we provide a number of applications of the above algorithms.

6.1 Separator Decompositions of Restrained Geometric Graphs

The algorithms in this section are based on the use of separators. As mentioned above, the separator-decomposition algorithm of Goodrich [30] applies only to planar graphs. Nevertheless, given the tool of geometric graph planarization, we can adapt Goodrich’s result to restrained geometric graphs in a fairly straightforward manner. Given a restrained geometric graph G , we planarize it using the algorithm above, creating the planar graph G' . As observed above, G' has total size $O(n)$. Thus, we can use the result of Goodrich [30] to compute a recursive $O(\sqrt{n})$ -separator decomposition of G' in $O(n)$ time. We convert this separator decomposition into a $O(\sqrt{n})$ -separator for G by the following transformation. For each node v in a separator W of G' at a node w in the separator decomposition tree B , we do the following:

- If v is also a vertex in G , then we add v to the separator for G corresponding to w , provided v is not already a member of a separator associated with an ancestor of w .
- If v is an intersection point in G' , between edges (a, b) and (c, d) in G , then we add each of a , b , c , and d to the separator for G corresponding to w , provided it is not already a member of a separator associated with an ancestor of w .

This gives us the following:

THEOREM 6.1. *Suppose we are given an n -vertex geometric graph G and its planarization, G' , which is of size $O(n)$. Then we can construct a recursive $O(\sqrt{n})$ -separator decomposition of G in $O(n)$ time, for $\delta = 2/3$.*

6.2 Single-Source Shortest Paths and Voronoi Diagrams

Given an n -vertex bounded-degree graph G and a recursive $O(\sqrt{n})$ -separator decomposition for G , Henzinger *et al.* [34] show that one can compute shortest paths from a single source s in G to all other vertices in G in $O(n)$ time. Using the separator decomposition algorithms presented above, then, we can show that their

algorithm applies to restrained geometric graphs, even ones that do not have bounded degree, by a simple transformation that replaces high-degree vertices with bounded-degree trees of zero-weight edges.

Suppose we are given K distinguished vertices in an n -vertex restrained geometric graph G and we wish to construct the *Voronoi diagram* of G , which is a labeling of each vertex v of G with the name of the distinguished vertex closest to v . As before, by replacing high degree vertices with bounded-degree trees of zero-weight edges we can assume without loss of generality that G has constant degree. In this case, we construct a recursive $O(\sqrt{n})$ -separator decomposition of G using one of the algorithms of the previous section. Let B be the recursion tree and let us label each vertex v in G with the internal node w in B where v is added to the separator or with the leaf w in B corresponding to a set containing v where we stopped the recursion (because the set’s size was below our stopping threshold). Given this labeling, we can trace out the subtree B' of B that consists of the union of paths from the root of B to the distinguished nodes in G in $O(n)$ time. Let us now assign each edge in B' to have weight 0 and let us add B' to G to create a larger graph G' . Note that if we add each internal node v in B' to the separator associated with node v in B , then we get a recursive $O(\sqrt{n})$ -separator decomposition for G' , for each separator in the original decomposition increases by at most one vertex. Thus, we can apply the algorithm of Henzinger *et al.* [34] to compute the shortest paths in G' from the root of B' to every other vertex in G' in $O(n)$ time. Moreover, since the edges of G' corresponding to edges of B' have weight 0, this shortest path computation will give us the Voronoi diagram for G . Therefore, we have the following:

THEOREM 6.2. *Given a connected n -vertex restrained graph G , together with its planarization, one can compute shortest paths from any vertex s or the Voronoi diagram defined by any set of K vertices in G in $O(n)$ time.*

Incidentally, the above approach also implies a linear-time Voronoi diagram construction algorithm for planar graphs, which was not previously known.

7 Conclusions and Future Work

We have provided linear-time algorithms for a number of problems on connected restrained geometric graphs, which includes real-world road networks. Our results allow for linear-time trapezoidalization, triangulation, and planarization of geometric graphs except for the very narrow range of the number of crossings for which neither our algorithm nor the previous $O(n \log^* n + k)$ al-

gorithm is linear. In addition, our methods imply linear-time algorithms for other problems on such graphs as well. For example, one can use our algorithm to planarize a restrained non-simple polygon and then construct its convex hull in linear time by computing the convex hull of the outer face of our planarization (e.g., by an algorithm from [32, 38]). There are a number of interesting open problems and future research directions raised by this paper, including:

- Can one close the $\log^{(c)} n$ gap on values of k that admit optimal solutions to Chazelle’s open problem of computing a trapezoidal decomposition of an n -vertex non-simple polygon in $O(n+k)$ time, where k is the number of its edge crossings?
- Can we planarize restrained geometric graphs deterministically in linear time? Such a result would allow us to apply separator-based divide and conquer techniques for minimum spanning trees [22] to construct them in linear time for this family of graphs. Known linear-time minimum spanning tree algorithms for arbitrary graphs require randomization [36], and known deterministic algorithms for this problem are superlinear [9], although deterministic linear-time algorithms are known for planar graphs and minor-closed graph families [12, 20, 40].

Acknowledgment. We would like to thank Bernard Chazelle for several helpful discussions regarding possible approaches to solving his open problem involving non-simple polygons. This research was supported in part by the National Science Foundation, under grants 0724806, 0713046, and 0830403.

References

- [1] P. K. Agarwal. Geometric partitioning and its applications. In J. E. Goodman, R. Pollack, and W. Steiger, editors, *Computational Geometry: Papers from the DIMACS Special Year*. American Mathematical Society, 1991.
- [2] R. J. Anderson and G. L. Miller. An optimal algorithm for intersecting line segments in the plane. *Algorithmica*, 6:859–868, 1991.
- [3] D. Armon and J. Reif. A dynamic separator algorithm. In *Proc. 3rd Workshop Algorithms Data Struct.*, volume 709 of *Lecture Notes Comput. Sci.*, pages 107–118. Springer-Verlag, 1993.
- [4] F. Aurenhammer. Voronoi diagrams: A survey of a fundamental geometric data structure. *ACM Comput. Surv.*, 23(3):345–405, Sept. 1991.
- [5] F. Aurenhammer and R. Klein. Voronoi diagrams. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 201–290. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.
- [6] J. L. Bentley and T. A. Ottmann. Algorithms for reporting and counting geometric intersections. *IEEE Trans. Comput.*, C-28(9):643–647, Sept. 1979.
- [7] J.-D. Boissonnat, O. Devillers, R. Schott, M. Teillaud, and M. Yvinec. Applications of random sampling to on-line algorithms in computational geometry. *Discrete Comput. Geom.*, 8:51–71, 1992.
- [8] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6(5):485–524, 1991.
- [9] B. Chazelle. A minimum spanning tree algorithm with inverse-ackermann type complexity. *J. ACM*, 47(6):1028–1047, 2000.
- [10] B. Chazelle and H. Edelsbrunner. An optimal algorithm for intersecting line segments in the plane. *J. ACM*, 39(1):1–54, 1992.
- [11] B. Chazelle and L. J. Guibas. Fractional cascading. In *Proc. Internat. Colloq. Automata Lang. Program.*, 1985. To appear.
- [12] D. Cheriton and R. E. Tarjan. Finding minimum spanning trees. *SIAM J. Comput.*, 5:724–742, 1976.
- [13] K. L. Clarkson, R. Cole, and R. E. Tarjan. Erratum: Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2(3):341–343, 1992.
- [14] K. L. Clarkson, R. Cole, and R. E. Tarjan. Randomized parallel algorithms for trapezoidal diagrams. *Internat. J. Comput. Geom. Appl.*, 2(2):117–133, 1992.
- [15] K. L. Clarkson and P. W. Shor. Applications of random sampling in computational geometry, II. *Discrete Comput. Geom.*, 4:387–421, 1989.
- [16] T. H. Cormen, C. E. Leiserson, R. L. Rivest, and C. Stein. *Introduction to Algorithms*. MIT Press, Cambridge, MA, 2nd edition, 2001.
- [17] M. de Berg and O. Schwarzkopf. Cuttings and applications. *Internat. J. Comput. Geom. Appl.*, 5:343–355, 1995.
- [18] G. Di Battista, P. Eades, R. Tamassia, and I. G. Tollis. *Graph Drawing*. Prentice Hall, Upper Saddle River, NJ, 1999.
- [19] H. Edelsbrunner, J. W. Jaromczyk, J. W. Penny, and G. Świątek. Primal canoes: optimal arrangements of segments. In *Proc. 3rd Canad. Conf. Comput. Geom.*, pages 224–227, Aug. 1991.
- [20] D. Eppstein. Spanning trees and spanners. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, chapter 9, pages 425–461. Elsevier, 2000.
- [21] D. Eppstein. Setting parameters by example. *SIAM J. Computing*, 32(3):643–653, 2003.
- [22] D. Eppstein, Z. Galil, G. F. Italiano, and T. H. Spencer. Separator based sparsification i: planarity testing and minimum spanning trees. *J. Computer and Systems Sciences*, 52(1):3–27, 1996.
- [23] D. Eppstein and M. T. Goodrich. Studying (non-planar) road networks through an algorithmic lens. Submitted, 2008.
- [24] D. Eppstein, G. L. Miller, and S.-H. Teng. A deterministic linear time algorithm for geometric separators

- and its applications. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 99–108, 1993.
- [25] M. Erwig. The graph Voronoi diagram with applications. *Networks*, 36(3):156–163, 2000.
- [26] J. W. Essam and M. E. Fisher. Some basic definitions in graph theory. *Review of Modern Physics*, 42(2):271–288, 1970.
- [27] M. L. Fredman and R. E. Tarjan. Fibonacci heaps and their uses in improved network optimization algorithms. *J. ACM*, 34:596–615, 1987.
- [28] A. V. Goldberg. Scaling algorithms for the shortest paths problem. In *SODA '93: Proceedings of the fourth annual ACM-SIAM Symposium on Discrete algorithms*, pages 222–231, Philadelphia, PA, USA, 1993. Society for Industrial and Applied Mathematics.
- [29] A. V. Goldberg and C. Harrelson. Computing the shortest path: A* search meets graph theory. In *SODA '05: Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 156–165, Philadelphia, PA, USA, 2005. Society for Industrial and Applied Mathematics.
- [30] M. T. Goodrich. Planar separators and parallel polygon triangulation. *J. Comput. Syst. Sci.*, 51(3):374–389, 1995.
- [31] M. T. Goodrich and R. Tamassia. *Algorithm Design: Foundations, Analysis, and Internet Examples*. John Wiley & Sons, New York, NY, 2002.
- [32] R. L. Graham and F. F. Yao. Finding the convex hull of a simple polygon. *J. Algorithms*, 4:324–331, 1983.
- [33] S. Har-Peled. Constructing planar cuttings in theory and practice. *SIAM J. Comput.*, 29:2016–2039, 2000.
- [34] M. R. Henzinger, P. Klein, S. Rao, and S. Subramanian. Faster shortest-path algorithms for planar graphs. *J. Comput. Syst. Sci.*, 55(1):3–23, 1997.
- [35] M. Holzer, F. Schulz, D. Wagner, and T. Willhalm. Combining speed-up techniques for shortest-path computations. *J. Exp. Algorithmics*, 10:2.5, 2005.
- [36] D. R. Karger, P. N. Klein, and R. E. Tarjan. A randomized linear-time algorithm to find minimum spanning trees. *J. ACM*, 42(2):321–328, 1995.
- [37] G. A. Klunder and H. N. Post. The shortest path problem on large-scale real-road networks. *Networks*, 48(4):182–194, 2006.
- [38] D. T. Lee. On finding the convex hull of a simple polygon. *Internat. J. Comput. Inform. Sci.*, 12:87–98, 1983.
- [39] R. J. Lipton and R. E. Tarjan. A separator theorem for planar graphs. *SIAM J. Appl. Math.*, 36:177–189, 1979.
- [40] M. Mareš. Two linear time algorithms for MST on minor closed graph classes. *Archivum Mathematicum (BRNO)*, 40(3):315–320, 2004.
- [41] K. Mehlhorn. A faster approximation algorithm for the Steiner problem in graphs. *Information Processing Letters*, 27:125–128, 1988.
- [42] U. Meyer. Single-source shortest-paths on arbitrary directed graphs in linear average-case time. In *SODA '01: Proceedings of the twelfth annual ACM-SIAM symposium on Discrete algorithms*, pages 797–806, Philadelphia, PA, USA, 2001. Society for Industrial and Applied Mathematics.
- [43] G. L. Miller. Finding small simple cycle separators for 2-connected planar graphs. *J. Comput. Syst. Sci.*, 32(3):265–279, 1986.
- [44] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Geometric separators for finite element meshes. *SIAM J. Sci. Comput.*, ??, 1995. Submitted.
- [45] G. L. Miller, S.-H. Teng, W. Thurston, and S. A. Vavasis. Separators for sphere-packings and nearest neighbor graphs. *J. ACM*, 44:1–29, 1997.
- [46] B. Mohar and C. Thomassen. *Graphs on Surfaces*. Johns Hopkins University Press, 2001.
- [47] J. Pach. *Towards a Theory of Geometric Graphs*, volume 342 of *Contemporary Mathematics*. American Mathematical Society, 2004.
- [48] R. Raman. Recent results on the single-source shortest paths problem. *SIGACT News*, 28(2):81–87, 1997.
- [49] P. Sanders and D. Schultes. Highway hierarchies hasten exact shortest path queries. In *Proceedings 17th European Symposium on Algorithms (ESA)*, volume 3669 of *Springer LNCS*, pages 568–579. Springer, 2005.
- [50] R. Sedgwick and J. S. Vitter. Shortest paths in Euclidean graphs. *Algorithmica*, 1:31–48, 1986.
- [51] D. A. Spielman and S.-H. Teng. Disk packings and planar separators. In *Proc. 12th Annu. ACM Sympos. Comput. Geom.*, pages 349–358, 1996.
- [52] M. Thorup. Undirected single-source shortest paths with positive integer weights in linear time. *J. ACM*, 46(3):362–394, 1999.
- [53] W. T. Trotter. *Planar Graphs*, volume 9 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*. American Mathematical Society, 1993.
- [54] F. B. Zhan and C. E. Noon. Shortest path algorithms: An evaluation using real road networks. *Transportation Science*, 32(1):65–73, 1998.