Optimal Triangulations of Points and Segments with Steiner Points

Boris Aronov[*], Tetsuo Asano[†], and Stefan Funke[‡]

October 8, 2008

Abstract

Consider a set X of points in the plane and a set E of non-crossing segments with endpoints in X. One can efficiently compute the triangulation of the convex hull of the points, which uses X as the vertex set, respects E, and maximizes the minimum internal angle of a triangle.

In this paper we consider a natural extension of this problem: Given in addition a *Steiner point* p, determine the optimal location of p and a triangulation of $X \cup \{p\}$ respecting E, which is best among all triangulations and placements of p in terms of maximizing the minimum internal angle of a triangle. We present a polynomial-time algorithm for this problem and then extend our solution to handle any constant number of Steiner points.

Keywords: Computational geometry, constrained Delaunay triangulation, polynomial-time algorithm, Steiner point, triangulation, Voronoi diagram, geometric optimization

1 Introduction

The problem of triangulating a set of points X in the plane that respects a set of non-crossing line segments E with endpoints in X is to construct a decomposition of the convex hull of $X \cup E$ into triangles using only points of X as vertices, so that all line segments in E are triangle edges of the decomposition. One important special case of the triangulation of $X \cup E$ which arises in many applications is the triangulation of a simple polygon, i.e., where the

[*]Department of Computer and Information Science, Polytechnic Institute of NYU, Brooklyn, NY, USA. Work by B.A. has been partially supported by a grant from the U.S.-Israel Binational Science Foundation and by NSA MSP Grant H98230-06-1-0016.

[†]School of Information Science, JAIST, Japan. Work by T.A. was partially supported by the Ministry of Education, Science, Sports and Culture, Grant-in-Aid for Scientific Research on Priority Areas and Scientific Research (B). Part of the work of T.A. was performed while he was visiting Max-Planck-Institut für Informatik, Germany, Carleton University, Canada, and New York University, USA, all in 2006.

[‡]Department of Mathematics and Computer Science, Ernst-Moritz-Arndt Universität Greifswald, Germany. Work by S.F. was partially supported by the Max Planck Center for Visual Computing and Communication (MPC-VCC) funded by the German Federal Ministry of Education and Research (FKZ 01IMC01).

edges in E form a simple polygon with vertex set X. Some triangulation of a given simple polygon can even be computed in linear time using Chazelle's algorithm [5]. Optimizing some criterion over all such triangulations is also possible. For example, a popular optimization criterion is to maximize the minimum angle of any triangle (which informally speaking yields a decomposition into triangles which are as 'fat' or 'round' as possible). Such a triangulation is known as a *constrained Delaunay triangulation* and can be found in $O(n \log n)$ time [6] for a set E of n non-crossing line segments, hence in particular for an n-gon. We could also find, in polynomial time, a *minimum-weight triangulation* of a simple polygon that minimizes the total length of chords required for triangulation using dynamic programming (the more general version of this problem, on an arbitrary set X, even with $E = \varnothing$, is NP-hard [12]). Dynamic programming is also powerful enough to find a triangulation in which the worst aspect ratio of the resulting triangles is minimized, where the *aspect ratio* of a triangle is the ratio of length of the longest side to its *width*, i.e., its smallest height.

In this paper we are interested in what happens when we allow one Steiner point in the triangulation of a set X of points respecting a set E of line segments. More precisely, given X and E, we want to find a point $p \notin X$ such that the quality of the optimal triangulation after the insertion of the point p is the best under a given optimization criterion. If maximization of the minimum internal angle is the goal, which is the focus of this paper, we want to find a location of p such that the minimum angle of the optimal triangulation after the insertion of the point p is maximized among all possible placements of point p. As far as the authors know, there is no previous study of this question. Our main concern in this paper is to develop a polynomial-time algorithm; the algorithm we present has worst-case running time $O(n^{4+\varepsilon})$, for any $\varepsilon > 0$, with the implicit constant depending on ε.

A natural extension of this problem is to allow for more Steiner points to be inserted or to use different optimization criteria for the triangulation. The extension to multiple Steiner points is not trivial at all. In fact, it seems no simple algorithm exists for finding an optimal set of Steiner points. We present a fairly involved polynomial-time algorithm that optimally places any constant number k of Steiner points. Even solving the case of two Steiner points seems to require quite complex structures as we will exhibit below. Therefore our presentation will focus on the case of inserting two Steiner points but the presented ideas can be applied to any constant number of Steiner points in a straightforward manner. Using a different optimization criterion is also an interesting direction for future work. Minimization of the largest internal angle, minimization of the largest slope, and minimization of the largest aspect ratio are rather popular criteria [2, 3], but no $O(n \log n)$ algorithm is known for these criteria except for that of maximization of the smallest angle, even in a simple polygon without the introduction of Steiner points.

The question considered in this paper is partially motivated by problems arising in the area of mesh improvement. Given a triangulation of some bounded domain, one interesting goal is to improve the quality of the triangulation by relocating internal vertices (we assume internal vertices can be moved while vertices on the domain boundary are fixed). In the so-called Laplacian method (see, e.g., [8]) an internal vertex is relocated to the barycenter of the polygon defined by its incident triangles. This works well in practice, but the barycenter is not always the best location for a vertex. We want to emphasize that in the Laplacian method the *topology* of the triangulation, i.e., its connectivity structure, is unchanged. Hence the barycenter is just one candidate for a good location when the topology is fixed. It is not known what a good location for the vertex is when the topology is allowed to change. Naturally one might expect better triangulations when topology changes are allowed.

Further closely related results are known in the literature under the heading of *Delaunay refinement*. Here one is given a planar straight line graph (PSLG) represented by a set of vertices and non-intersecting edges, and the goal is to triangulate this PSLG using 'fat' triangles, the latter being important if the obtained subdivision is used, for example, in finite element method calculations. 'Fatness' can be achieved by maximizing the smallest angle in the computed triangulation. Delaunay refinement algorithms repeatedly insert Steiner points until a certain minimum angle is achieved; the major goal here is to bound the number of necessary Steiner points. In some way our paper approaches the same problem from the opposite direction: how much can we improve the 'fatness' of our triangulation with few Steiner points? See [15] for an excellent survey on Delaunay refinement techniques.

Our contribution and outline Apart from formulating the problem of finding an optimal triangulation with Steiner points for the first time, the main contribution of this paper is the precise characterization of possible insertion loci of the Steiner points and exhibiting the connection to (a slight modification of) *angular Voronoi diagrams* [1]. While we provide actual algorithms to solve the problem, we do not consider these algorithms very practical. Algorithms that also work well in practice will be the subject of future studies on this topic.

This paper is organized as follows: Section 2 describes an $O(n^{4+\varepsilon})$ time algorithm for computing the optimal location of a single Steiner point. The case of two or more Steiner points is considered in Section 3. Section 4 includes conclusions and future work.

While we did not have to single out the case of one Steiner point, this is the situation where one can give a fairly explicit description of the algorithm and a bound on its running time. Furthermore, we believe that the main ideas are communicated a lot easier when first restricting to the one-Steiner-point case. Apart from these didactic reasons, this is the version of the problem that would be the first candidate for improvement and optimization. It is also likely to be used as an efficient heuristic for a multiple Steiner point improvement algorithm, by adding Steiner points one at a time.

2 Optimal Triangulation with One Steiner Point

The main problem we address in this section is the following:

Problem 1. Given a set of points X in the plane and a set of non-crossing edges E with endpoints in X, find a location for a Steiner point $p \notin X$ and a triangulation of $X \cup \{p\}$ which respects the edges in E and maximizes the smallest internal angle.

Interestingly, if no Steiner points are allowed, we can triangulate X and E so that the smallest internal angle is maximized by computing the so-called *constrained Delaunay triangulation* of X and E, see [4, 11, 6]. This constrained Delaunay triangulation is also the basis of our solution when allowing one or more Steiner points. Our solution will also immediately apply to the following problem which is an important special case of the above:

Problem 2. Given a simple n-gon P, find a triangulation of the interior of P with one Steiner point maximizing the smallest internal angle.

2.1 Preliminaries

Before proceeding, we need several definitions. Given a set of points X and a set of non-crossing edges E with endpoints in X we say $a \in X$ *sees* $b \in X$ if and only if the segment ab does not cross any edge in E (note that, for a segment $ab \in E$, a sees b). Point a is *visible to a set* Y if a can be seen from some point in Y.

Definition 3. The *constrained Delaunay triangulation (CDT)* of a set of points X and a set of non-crossing edges E with endpoints in X contains exactly those edges (a, b), $a, b \in X$ for which

- $(a, b) \in E$, or
- a sees b and there exists a circle C through a and b such that no other $c \in X \cap C$ is visible from ab.

A constrained Delaunay triangulation $\mathrm{CDT}(X, E)$ for (X, E) can be computed in $O(n \log n)$ time and for non-degenerate (free of cocircular quadruples of points) point sets forms a proper triangulation, i.e., a decomposition of the convex hull of X into triangles. (For clarity of presentation, we will assume hereafter that X is *in general position*, i.e., no three points of X are collinear and no four cocircular.) It maximizes the minimum interior angle of any triangulation of (X, E) that uses only the vertices of X as triangulation vertices; even stronger, the CDT lexicographically maximizes the list of angles sorted from smallest to largest, see [4] for an extensive list of references.

2.2 Topology changes upon insertion of a Steiner point

In the following we are interested in how the constrained Delaunay triangulation changes when some point $p \notin X$ is added to X.[1] For the remainder of this section, put $\mathcal{D} := \mathrm{CDT}(X, E)$ and $\mathcal{D}(p) := \mathrm{CDT}(X \cup \{p\}, E)$. Definition 3 implies that the circumcircle of $\triangle abc \in \mathrm{CDT}(X, E)$ cannot contain a vertex other than a, b, c visible from the interior of $\triangle abc$. Also, the insertion of p is only of interest to triangles in $\mathrm{CDT}(X, E)$ in the circumcircle of which p lies. More precisely, we say an edge $e \notin E$ is *invalidated* by p if and only if p lies in the intersection of the circumcircles of the two adjacent triangles of e and p is visible from the interior of both triangles (for an edge on the convex hull of X consider the artificial triangle with one vertex at infinity and the corresponding 'circumhalfplane').

Lemma 4. *An edge $e = ab$ with $e \in \mathcal{D}$ that is invalidated by a point p cannot be part of $\mathcal{D}(p)$.*

[1] In what follows, we will permit p to lie anywhere in the plane, as long as it does not coincide with an existing point of X or lie on an existing edge of E. It is possible to restrict p to the interior of the convex hull of X, by simply restricting the domain over which the optimization is done below.

Proof. Let C_1, C_2 be the circumcircles of the two triangles adjacent to e in $\mathcal{D}$, m_1, m_2 their respective centers. Assume that C is a circle, centered at m, certifying the presence of e in $\mathcal{D}(p)$. Point m has to lie on the line segment m_1m_2, otherwise C would contain the third corner of one of the two triangles adjacent to e in $\mathcal{D}$ which also would be visible from ab. As $p \in C_1 \cap C_2 \subseteq C$ we obtain a contradiction to e being part of $\mathcal{D}(p)$ since p is visible from ab. □

Lemma 5. *$\mathcal{D}(p)$ can be obtained from $\mathcal{D}$ by deleting all edges $\{e_1, e_2, \ldots e_l\}$ in $\mathrm{CDT}(X, E)$ invalidated by p and re-triangulating the resulting 'hole' H in a star fashion from p.*

Proof. Clearly all edges in $\mathcal{D}$ not invalidated by p are part of $\mathcal{D}(p)$ according to Definition 3. Furthermore, it is not possible that $\mathcal{D}(p)$ contains an edge vw which was not present in $\mathcal{D}$ and $v, w \neq p$, since the insertion of p only decreases the number of admissible edges on the original vertices. Therefore, new edges in $\mathcal{D}(p)$ have to have p as one endpoint. Connecting p to all visible vertices of H is the only way to obtain a triangulation again. □

Consider the arrangement $\mathcal{A}$ defined by the triangles of $\mathcal{D}$ and their circumcircles. (For technical reasons we further refine $\mathcal{A}$ into "trapezoids," by drawing a vertical ray up and down from every vertex of $\mathcal{A}$ until the first intersection with another arrangement edge, or to infinity. A "trapezoidal" cell of the resulting arrangement is delimited by two vertical sides on left and right and by a line segment or circular arc on its top and bottom. Slightly abusing the notation we use $\mathcal{A}$ to refer to the refined arrangement.) We argue that the topology of $\mathcal{D}(p)$ does not change when p is moved within one cell σ of this arrangement.

Lemma 6. *If some point $p \in \sigma$ can be seen from the interior of some triangle $\triangle abc \in \mathcal{D}$ whose circumcircle C contains p, then all points within σ can be seen from the interior of the triangle $\triangle abc$.*

Proof. Since $p \in \sigma$ and $p \in C$, $\sigma \subset C$, by construction of $\mathcal{A}$.

Let x be a point in the interior of $\triangle abc$ which sees p and assume there exists a point $p' \in \sigma$ which cannot be seen from x. Let γ be any curve in σ connecting p to p'. Consider the line segment xq as q moves from p towards p' along γ. If p' cannot be seen from x, at some point xq must hit a segment $e \in E$ obstructing the view, at an endpoint of e. Hence this endpoint must be visible from x and must lie on the segment $xq' \subset C$, contradicting $\triangle abc \in \mathcal{D}$. □

Since all points within a cell σ of $\mathcal{A}$ lie within the same set of circumcircles, we have shown that all points within σ invalidate the same set of edges. It remains to show that all points p within this cell σ behave the same in terms of visibility from vertices of triangles of which at least one edge was invalidated by p.

Lemma 7. *Let $e = bc$ be an edge of $\mathcal{D}$, $\triangle abc$ and $\triangle bcd$ the respective triangles adjacent to e. If e is invalidated by p then p sees a, b, c, and d.*

Proof. Let $x \in \triangle abc$ be visible from p and suppose, without loss of generality, that p cannot see a. Consider the segment yp as y moves along xa towards a. At some point yp must meet a constraining edge $e \in E$ at a vertex f of X lying in the interior of the circumcircle of $\triangle abc$, contradicting the assumption that $\triangle abc \in \mathcal{D}$, since f would be visible from a point in the interior of $\triangle abc$. □

We have shown that all points within a cell σ behave identically in terms of invalidation of edges as well as visibility hence leading to the identical topology of $\mathcal{D}(p)$. Furthermore observe that the complexity of $\mathcal{A}$ is $O(n^2)$ since the arrangement of $O(n)$ circles and line segments has complexity $O(n^2)$; refinement into vertical trapezoids does not increase the complexity, asymptotically. We summarize our findings in the following corollary.

Corollary 8. *The arrangement $\mathcal{A}$ defined by the triangles of $\mathcal{D}$ and their circumcircles characterizes the different topologies of $\mathcal{D}(p) = \mathrm{CDT}(X \cup \{p\}, E)$ after insertion of a Steiner point p in the sense that all placements of p within the same cell of $\mathcal{A}$ lead to the same topology. The combinatorial complexity of $\mathcal{A}$ is $O(n^2)$.*

We note that the arrangement will in general be overrefined in the sense that points in different cells of $\mathcal{A}$ might lead to the same topology of $\mathcal{D}(p)$.

For our Problem 2 of triangulating a polygon P with one Steiner point p we compute the arrangement $\mathcal{A}$ with respect to $\mathrm{CDT}(X, E)$ where X is the set of polygon vertices and E is the set of polygon edges, and only consider the cells of $\mathcal{A}$ inside P.

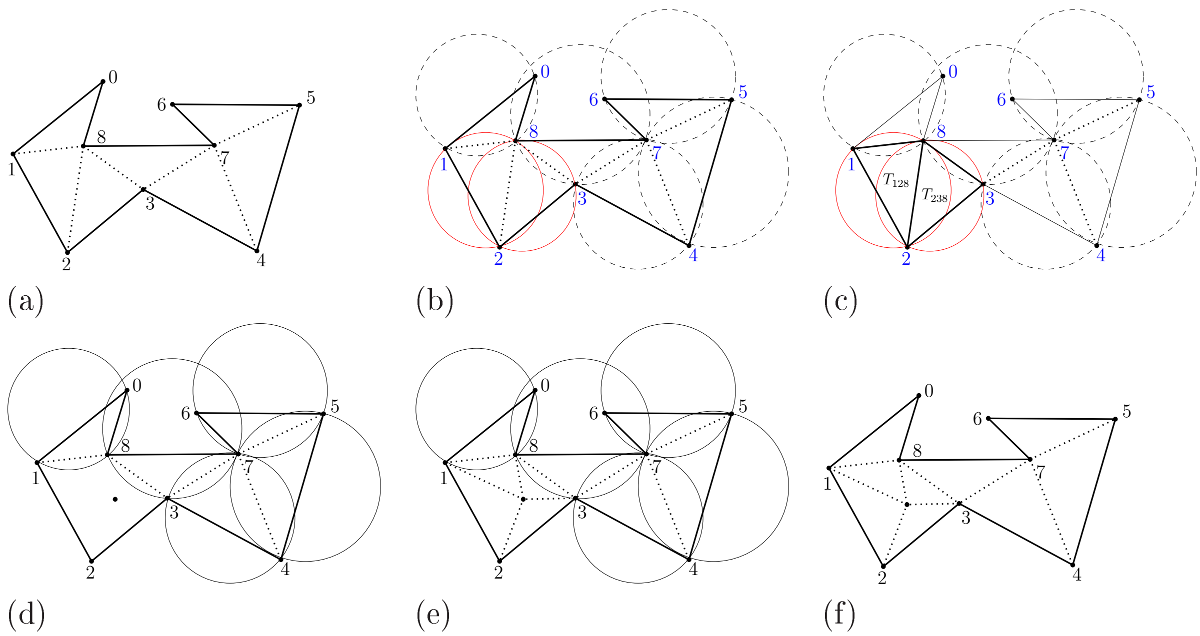

Figure 1: How to obtain a Delaunay triangulation after inserting a new point p. (a) Given polygon with Delaunay edges. (b) Circumcircles of the triangles involved. (c) Two triangles whose circumcircles contain p. (d) Removing invalidated triangular edge. (e) Optimal triangulation of the 'hole'. (f) Final Delaunay triangulation after inserting point p.

Figure 1 gives an example. The interior of the polygon is partitioned into cells by the triangles and their respective circumcircles. (We have used a LEDA [9] function to draw a constrained Delaunay triangulation.) In the figure, the circumcircles of the triangles T_{128} and T_{238} both contain the point p and p is visible from both triangles. So, edge 28 cannot be included in the constrained Delaunay triangulation after insertion of p. On the other hand, p lies outside the circumcircle of T_{378} and thus this triangle is left unchanged after the insertion.

2.3 Optimizing the minimum angle for fixed topology

We have seen that, when a point p is placed somewhere within a cell $\sigma \in \mathcal{A}$, a fixed set of edges is invalidated, producing a star-shaped 'hole' $H = H(\sigma)$ in $\text{CDT}(X, E)$. If we know how to find the best placement of p within a cell σ, we can simply iterate over all cells $\sigma \in \mathcal{A}$ to find the overall best triangulation as follows:

Input: a set X of points and a set of non-crossing edges E with endpoints in X.

1. Compute the constrained Delaunay triangulation $\mathcal{D} := \text{CDT}(X, E)$.
2. Construct the arrangement induced by all triangles of $\mathcal{D}$ and their circumcircles. Refine it by a trapezoidal decomposition to obtain $\mathcal{A}$.
3. For each cell σ of $\mathcal{A}$:
   - Determine the set of edges invalidated by any Steiner point in σ and remove them to form the hole H
   - Compute an angular Voronoi diagram for H, truncated to within σ.
   - For each Voronoi edge in the truncated Voronoi diagram, find a point maximizing the minimum angle along the edge.
   - For each connected component of a boundary edge of the cell σ lying in the same cell of the truncated Voronoi diagram, find a point maximizing the minimum angle along this curve.
4. Return the triangulation yielding the best angle found.

In the following we will describe in detail how to actually treat a cell $\sigma \in \mathcal{A}$, i.e., how to determine the best placement for p within σ, using angular Voronoi diagrams.

Notice that a slight modification of the above algorithm also solves our original Problem 2 of optimizing the triangulation of a simple polygon P using one Steiner point: In step 3 of the algorithm we only need consider cells that lie in the interior of P and when maximizing the minimum angle we also only consider triangles contained in the interior of P.

2.3.1 Optimal placement via angular Voronoi diagrams

It remains to find an optimal point p within a specified cell σ bounded by circles and lines within a given star-shaped polygon $H = H(\sigma)$ that maximizes the smallest interior angle in the star triangulation of H from p.

A related problem Given a star-shaped polygon H, we can find an optimal point p that maximizes the smallest *visual angle* from p to all edges of H, i.e., the angle at which any edge of H is seen from p. Matoušek, Sharir, and Welzl [14] gave an almost linear-time algorithm within the framework of LP-type problems. Asano et al.[1] also gave efficient algorithms for the same problem using parametric search or the so-called *angular Voronoi diagram*. Our question is slightly different. It is not enough to maximize the smallest angle incident to the point p to be inserted: the smallest internal angle may be incident to the boundary of H rather than p. Another difficulty is that we want to find an optimal point p constrained to lie in σ, a cell in the arrangement $\mathcal{A}$ bounded by circular arcs and line segments, rather than anywhere within H. It seems difficult to adapt the aforementioned algorithms based on LP-type problem formulation or parametric search for this purpose, but fortunately the one using the angular Voronoi diagram can be adapted here.

Modified angular Voronoi diagrams The *angular Voronoi diagram* for a star-shaped polygon H is defined as the partition of the plane according to the polygon edge that gives the smallest visual angle [1]. A point p belongs to the Voronoi region of a polygon edge e if the visual angle from p to e is smaller than that to any other polygon edge of H. It is known that it consists of straight-line segments or curves of low degree and has total complexity $O(n^{2+\varepsilon})$, for any $\varepsilon > 0$ and with implied constant depending on ε.

We have to modify the definition of the angular Voronoi diagram to take into account the angles associated with polygon edges as well. Given a star-shaped polygon H as a set of its bounding edges $\{e_0, e_1, \ldots, e_n := e_0\}$ and a point p in the plane, for each edge e_i we form the triangle $\mathrm{Tr}(p, e_i)$ by connecting the endpoints of the edge to p. The value $f(p, e_i)$ is defined to be the smallest internal angle in the triangle $\mathrm{Tr}(p, e_i)$. The region $\mathrm{Vor}(e_i)$ of an edge e_i is the locus of points p at which $f(p, e_i)$ is smallest among all edges, that is,

$$\mathrm{Vor}(e_i) := \{p \in \mathbb{R}^2 \mid f(p, e_i) \leq f(p, e_j) \quad \forall e_j \in H\}.^2$$

More concretely, let us consider the geometry of $f(\cdot, ab)$, for a line segment ab. Consider the two circles C_a and C_b centered at a and b, respectively, with radius $|ab|$, and the perpendicular bisector ℓ_{ab} of ab. Refer to Fig. 2. If point p lies outside of both circles, then the smallest internal angle of $\triangle abp$ is $\angle apb$. If p lies in the interior of C_a, on the same side of ℓ_{ab} as a, then $\angle abp$ is smallest. Symmetrically, if p lies in the interior of C_b on the same side of ℓ_{ab} as b, the smallest angle is $\angle bap$. Fig. 2 illustrates three different situations, with points p_1, p_2, p_3 lying outside, on the boundary of, and inside C_a, respectively.

In a similar manner we can derive an equation of a curve along which two polygon edges e_i and e_j give the same smallest angle, i.e., $f(p, e_i) = f(p, e_j)$, which is a curve of degree at most seven; see the appendix. Using this curve for a pair of polygon edges, we can draw Voronoi edges of the modified angular Voronoi diagram.

[2] This is a special case of the 'minimization diagram' construction; see, for example, [13, p. 175]. Given a region of space S and a collection of (partial) functions $\mathcal{F}$ defined on S, each point $x \in S$ is labeled according to the function f (or functions) achieving the smallest value at x, i.e., with $f(x) = \min g(x)$, with the minimum taken over all functions $g \in \mathcal{F}$ defined at x. We now partition S into maximal connected regions with the same label. This partitioning is the *minimization diagram* of $\mathcal{F}$ over S. In our case, the Voronoi diagram is minimization diagram of the functions $f(\cdot, e_i)$.

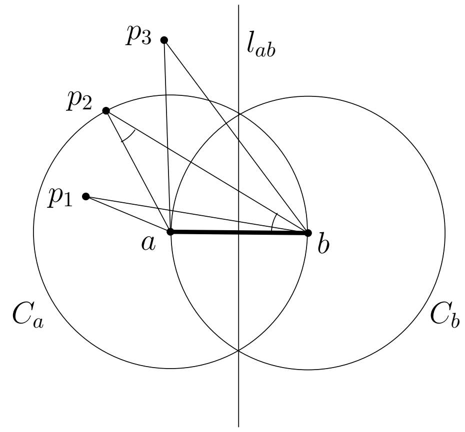

Figure 2: Partition of the plane into regions according to which angle of $\triangle abp$ is smallest.

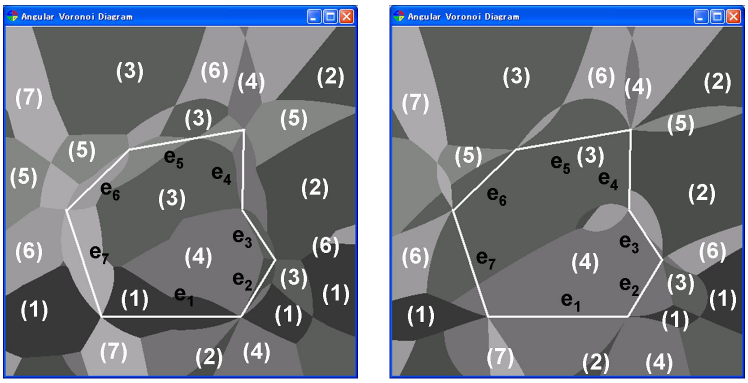

Figure 3: Modified angular Voronoi diagram (left) and original angular Voronoi diagram.

Fig. 3 in which Voronoi regions are painted by colors associated with polygon edges gives an example of such a modified angular Voronoi diagram. In the figure the given polygon is shown by bold white lines.

Optimizing over the angular Voronoi diagram Once we construct the modified angular Voronoi diagram, we can look for an optimal placement for p at Voronoi vertices, along Voronoi edges, or along the boundary of σ just like in the original angular Voronoi diagram [1]; it is easy to see that the maximum does not occur in the interior of a Voronoi region. Recall that we have refined $\mathcal{A}$ so that a cell $\sigma \in \mathcal{A}$ is a 'trapezoid' delimited by four line segments and/or circular arcs. In particular, each function $f(p, e_i)$, viewed as truncated to within σ, is a well-behaved function with a constant-complexity domain (the theorem we invoke applies here, since the functions are algebraic of bounded degree and the domain can be described by a Boolean formula of constant length and constant degree). Theorem 7.7 of [13] implies that the modified angular Voronoi diagram truncated to within σ has at most $O(n^{2+\varepsilon})$ features, i.e., edges, vertices, and faces. The features include connected portions of Voronoi edges within σ and portions of Voronoi cells lying along the boundary of σ; the envelope can be computed and the interesting features extracted in $O(n^{2+\varepsilon})$ time [13, Theorem 7.16]. Therefore in our algorithm the computations can be performed in time $O(n^{2+\varepsilon})$ per cell, for a total of $O(n^{4+\varepsilon})$, since cell processing dominates the runtime of the algorithm.

One may legitimately wonder if the aforementioned discussion provides a gross overestimate of the worst-case performance of the algorithm. While we cannot prove that the analysis is tight, it is easy to construct examples illustrating that no trivial improvement of the above bounds is immediately possible. For example, starting with a set of segments forming the boundary of a nearly-regular n-gon P, we see that the size of the arrangement $\mathcal{A}$ can reach $\Omega(n^2)$. Moreover, when the Steiner point p is inserted into cell σ, many edges of P will appear on the boundary of the hole $H(\sigma)$ formed by inserting p, for many choices of $\sigma \in \mathcal{A}$. Hence, for a typical cell σ, we will need to construct a modified angular Voronoi of a linear number of edges. Since we only have a quadratic bound on the complexity of such a diagram, our upper bound on the running time of the argument must be nearly tight, unless one can argue that a 'typical' modified angular Voronoi has much lower complexity. Such an improvement may indeed be possible, especially when one recalls that each diagram must be viewed as truncated to within the corresponding cell σ. However, we have no proof of this assertion.

3 Triangulation with Two or More Steiner Points

We now turn our attention to the situation where two Steiner points, p and q, are to be added to X. We start with the triangulation $\mathcal{D} := \mathrm{CDT}(X, E)$. We consider the space $Q := \mathbb{R}^2 \times \mathbb{R}^2$ of all possible placements of the two points.[3] We aim to identify the best placement of p, q in order to maximize the smallest angle in the resulting constrained Delaunay triangulation

[3] Technically, we must require $p \neq q$ and $p, q \notin X$. In addition, we should not distinguish (p, q) and (q, p) in Q, for $p \neq q$. Moreover, one may want to restrict p and q to lie in the interior of the convex hull of X. In the case of E forming the boundary of a simple polygon P, one may want to restrict p, q to lie in the interior of P. All of the above constraints amount to appropriately restricting the search space for our algorithm,

$\mathcal{D}(p,q) := \mathrm{CDT}(X \cup \{p,q\}, E)$. As in the previous section, we partition Q according to the topology of $\mathcal{D}(p,q)$, then use the analog of the modified angular Voronoi diagram from the previous section to determine which angle is smallest in the triangulation for every choice of $(p,q) \in Q$ and search the resulting diagrams for the placement maximizing the minimum angle. This plan is complicated by the need to explicitly identify all possible triangulation topologies. Instead, we will arrive at this partition indirectly, as detailed below.

In this section, we focus on constructing a polynomial-time algorithm, without any attempt at optimizing the running time. Such an optimization might be a good topic for further research towards a practical algorithm for this problem, especially when coupled with a heuristic to eliminate infeasible placements of p and q in order to reduce the search space.

We first recall a standard fact, the analogue of Definition 3 [6].

Fact 9. *Given a set of points Y and a set of segments with endpoints in Y, a triangle $\triangle abc$, for $a,b,c \in Y$ is present in $\mathrm{CDT}(Y,E)$ if and only if a,b,c are pairwise visible and no other vertex of Y visible from any point in $\triangle abc$ lies in the circumcircle of $\triangle abc$.*

Consider $\mathcal{D}(p,q)$ as defined above and consider a potential triangle $\triangle abc$ in it. Let $f(a,b,c;p,q)$ be a partial function defined as follows: it is defined for $(p,q) \in Q$ if and only if $\triangle abc$ is present in $\mathcal{D}(p,q)$ and the value of f is the measure of $\angle abc$. Then clearly the smallest angle in $\mathcal{D}(p,q)$ is

$$m(p,q) := \min_{(a,b,c)} f(a,b,c;p,q),$$

where the minimum is taken over all ordered triples of distinct vertices in $X \cup \{p,q\}$. The desired triangulation maximizing the minimum angle is just the one determined by (p,q) maximizing the function $m(p,q)$ over all of Q.

In order to bound the complexity of $m(\cdot,\cdot)$, we would like to apply Theorem 7.17 from [13]. The theorem states that the complexity of the *lower envelope* (i.e., the total number of vertices, edges, and faces of all dimensions on the graph of the pointwise minimum $m(\cdot,\cdot)$) of a set of N d-variate functions is $O(N^{d+\varepsilon})$, for all positive values of ε, *provided* the functions are well-behaved (e.g., algebraic of bounded degree) *and* defined over domains of constant description complexity (say, by a constant-length Boolean formula constructed of polynomial inequalities of bounded degree). In our case, the form of functions f is sufficiently simple to apply the theorem (if one uses, for example, an appropriate trigonometric function of the measure of $\angle abc$ to compare angles, to avoid transcendental functions, this is a low-degree algebraic function of the coordinates of the points).

However, the domains of definition of functions f are in general not of constant complexity and hence the theorem is not applicable directly. We instead decompose Q into a polynomial number of constant-size cells σ in such a manner that each function is either defined everywhere in σ or nowhere in σ (thus the combinatorial structure of $\mathcal{D}(p,q)$ must be constant over all $p,q \in \sigma$); this is the analogue of arrangement $\mathcal{A}$ in section 2. Then Theorem 7.17 of [13] can be applied to each σ separately, to conclude that the minimization

which can be done easily without slowing it down asymptotically. We will proceed with the description assuming $Q = \mathbb{R}^4$, for simplicity of presentation.

diagram of the functions f has polynomial complexity. Moreover, the diagram can be computed in polynomial time (even though efficient methods analogous to Theorem 7.16 of [13] are not known for functions of four or more variables that arise here) and therefore one can identify the placement of p and q that gives rise to the largest minimum angle in each σ and thus overall, in polynomial time.

It remains to construct such a decomposition. Observe that the boundaries of the domain of definition of a function $f(a, b, c; p, q)$ are given precisely by Fact 9. Namely, if we view p and q as moving, a triangle $\triangle abc$ formed by three points from among $X \cup \{p, q\}$ can cease to belong to $\mathcal{D}(p, q)$ only when some visibility constraint is violated (either a vertex of X becomes collinear with pq, or one of p, q becomes collinear with a line defined by two vertices of X) or when a cocircularity is created or destroyed (again, either p or q becomes cocircular with three vertices of X, or p and q become cocircular with two vertices of X). Each of the four possibilities corresponds to a low-degree hypersurface in $\mathbb{R}^4$. (For example, the cocircularity condition corresponds to the coordinates of the four points making the determinant of the INCIRCLE test vanish. It is a degree-4 polynomial equation in the coordinates of the points involved.) The number of possible constraints that we should check is clearly polynomial, since there are $n + 2$ vertices in all (conservatively, one can just check the cocircularity condition for every quadruple of points and the collinearity condition for every triple of points involving at least one of p, q, for a total of $\Theta(n^3)$ constraint surfaces).

To summarize, leaving the domain of one of the functions $f(a, b, c; p, q)$ requires crossing one of the following surfaces in Q:

- A visibility constraint, when p (or q) crosses a line through two vertices of X, or when p and q move in such a way that one of the vertices of X crosses the line pq.
- An INCIRCLE constraint, where the points move in such a way that the circle defined by three points is crossed by the fourth; one or two of the points involved is p and/or q and remaining ones are points of X.

We collect all these hypersurfaces and construct their arrangement (if $Q \neq \mathbb{R}^4$, we also add the boundary of Q to the arrangement and truncate it to within Q). We then refine the resulting partition of Q to contain only constant-size cells (e.g., via a cylindrical algebraic decomposition or a vertical decomposition [7, 10]); the resulting decomposition $\mathcal{A}$ is still of polynomial size. Recall that, now that we computed a decomposition of Q into polynomially many cells σ of simple shape, for each σ, we apply Theorem 7.17 from [13] to those functions that are defined over σ, to conclude that the lower envelope or the minimization diagram of these functions over σ, which is the higher-dimensional analog of the modified angular Voronoi diagram has polynomial number of features. Moreover, as outlined about, standard tools can be used to enumerate all the features and thus identify the placement of p, q maximizing the smallest angles in $\mathcal{D}(p, q)$.

This completes the proof of the following theorem.

Theorem 10. *Given a polygon P with n vertices, or more generally a set of n points X and a collection E of non-crossing edges connecting them, we can find two points p and q such that the minimum angle of the constrained Delaunay triangulation* $\mathrm{CDT}(X \cup \{p, q\}, E)$ *(and* *thus of any other triangulation with vertices $X \cup \{p, q\}$ respecting E) is maximized, in time polynomial in n.*

The same method applies almost verbatim to any constant number k of Steiner points. Q becomes $\mathbb{R}^{2k}$, with appropriate deletions. With more Steiner points, several additional types of constraints may arise: three of the Steiner points being collinear, or four points being cocircular, with three or four of them Steiner points. The number of constraints and corresponding surfaces is still polynomial in n, as long as k is a constant and the rest of the argument goes through essentially unmodified.

We now roughly estimate the running time of the algorithm, for a constant number $k > 1$ of Steiner points. The space Q of possible placements is $2k$-dimensional. We decompose it using $O(n^3)$ constraint surfaces. A vertical decomposition of the resulting arrangement consists of $O((n^3)^{2 \cdot 2k-4+\varepsilon}) = O(n^{12k-12+\varepsilon})$ cells [10] and can be computed in that amount of time. In each cell, the topology of the constrained Delaunay triangulation is fixed: There are at most $O(n)$ competing triangles, each defining a well-behaved function that is completely defined over the cell, so using the same reasoning as before, the complexity of the minimization diagram is at most $O(n^{2k+\varepsilon})$ per cell. However, it is not known how to compute the minimization diagram in higher dimensions efficiently. It is sufficient (though probably an overkill—computing the arrangement of the surfaces might be enough) to compute the vertical decomposition of the arrangement of the graphs of these functions, whose complexity is $O(n^{2 \cdot 2k-4+\varepsilon}) = O(n^{4k-4+\varepsilon})$ per cell and which can be computed in this time. Summing over all cells, we obtain an estimate of $O(n^{16k-16+\varepsilon})$, for $k > 1$, for the running time of the algorithm for computing optimal placement of k Steiner point, for any $\varepsilon > 0$.

4 Conclusions and Future Work

In this paper we have presented polynomial-time algorithms for determining the optimal position(s) of one or a constant number of Steiner points to be added to a given set of points and non-crossing segments with the goal of constructing a triangulation of minimum internal angle. Adding one Steiner point takes $O(n^{4+\varepsilon})$ time, while inserting $k > 1$ points requires polynomial time, with the exponent of the polynomial depending linearly on k.

It would be interesting to develop a faster and more practical algorithm for optimally inserting a single Steiner point, to improve the dependence of the runtime of the multiple-Steiner-point algorithm on the number of Steiner points, and to extend our analysis to other optimization criteria. Is it NP-complete to determine, given $k > 1$ and α, if there exists a triangulation with k Steiner points in which the smallest interior angle is at least α?

References

[1] T. Asano, N. Katoh, H. Tamaki, and T. Tokuyama: "Angular Voronoi Diagram with Applications," *Proc. International Symposium on Voronoi Diagram in Science and Engineering*, Banff, Canada, 32–39 (2006).

[2] M. Bern: “Triangulations and Mesh Generation,” Chapter 25, in *Handbook of Discrete and Computational Geometry*, 2nd Edition, edited by J.E. Goodman and J. O’Rourke, CRC Press LLC, Bocal Raton, FL, 563-582 (2004).

[3] M. Bern, H. Edelsbrunner, D. Eppstein, S. Mitchell, and T.S. Tan: ”Edge Insertion for Optimal Triangulation,” *Discrete and Computational Geometry*, 10 (1): 47-65 (1993).

[4] M. Bern and D. Eppstein: “Mesh Generation and Optimal Triangulation,” *Computing in Euclidean Geometry*, World Scientific (1992).

[5] B. Chazelle: “Triangulating a simple polygon in linear time,” *Discrete & Computational Geometry*, 6(5): 485-524 (1991).

[6] P. Chew: “Constrained Delaunay triangulation,” *Proc. 3rd ACM Symp. on Computational Geometry*, 215–223 (1987).

[7] G. E. Collins: “Quantifier Elimination for Real Closed Fields by Cylindrical Algebraic Decomposition,” *Proc. 2nd GI Conf. Automata Theory Formal Lang.*, Lecture Notes Comput. Sci. 33, Springer-Verlag, 134–183 (1975).

[8] D.A. Field: “Laplacian Smoothing and Delaunay Triangulations,” *Communications in Applied Numerical Methods*, Wiley, 4: 709–712 (1988).

[9] “LEDA: A library for Efficient Data Types and Algorithms,” Algorithmic Solutions Software GmbH, Germany.

[10] V. Koltun: “Almost tight upper bounds for vertical decompositions in four dimensions,” *J. ACM*, 51(5): 699–730 (2004).

[11] D.T. Lee and A.K. Lin: “Generalized Delaunay triangulation for planar graphs,” *Discrete & Computational Geometry*, 1: 201–217 (1986).

[12] W. Mulzer and G. Rote: “Minimum weight triangulation is NP-hard,” *J. ACM*, 55(2): 1–29 (2008).

[13] M. Sharir and P. K. Agarwal, *Davenport-Schinzel Sequences and Their Geometric Applications*. Cambridge University Press, Cambridge, UK (1995).

[14] J. Matoušek, M. Sharir and E. Welzl: “A subexponential bound for linear programming,” *Algorithmica*, 16: 498–516 (1996).

[15] J.R. Shewchuk: “Delaunay Refinement Algorithms for Triangular Mesh Generation,” *Computational Geometry: Theory and Applications* 22(1-3): 21–74 (2002).

[16] R. Seidel: “Constrained Delaunay triangulations and Voronoi diagrams with obstacles,” in “1978-1988 Ten Years IIG,” 178–191 (1988).

Appendix: Equations for Modified Angular Voronoi Edges

This appendix includes equations for a modified angular Voronoi diagram.

The *modified angular Voronoi diagram* for a star-shaped polygon H is defined as a partition of the plane according to the polygon edge that gives the smallest internal angle. A point p belongs to the Voronoi region of a polygon edge e, denoted by $\mathrm{Vor}(e)$, if the smallest internal angle $f(p, e)$ of the triangle $\mathrm{Tr}(p, e)$ defined by connecting two endpoints of the edge e to p is smaller than that of any other polygon edge of H.

Given a line segment ab, we partition the plane into three regions by two circles C_a and C_b centered at a and b, respectively, with the radius $|ab|$ and the perpendicular bisector ℓ_{ab} of ab. If a point p lies in the region R_a that is the intersection of the interior of the circle C_b and the halfplane H_b defined by ℓ_{ab} that contains b, then the angle at a is smallest among three internal angles of the triangle $\triangle pab$. Similarly, the angle at b is smallest if p is in the region R_b given as the intersection between C_a and H_a. Otherwise, that is, if p lies outside the two circles then the angle at p is smallest.

Edges of a modified angular Voronoi diagram for a star-shaped polygon H are defined as curves at which two edges of H give the same smallest internal angle. The cosine value of such an internal angle is given by

$$\frac{|pa|^2 + |pb|^2 - |ab|^2}{2|pa| \cdot |pb|}, \quad \frac{|pb|^2 + |ab|^2 - |pa|^2}{2|pb| \cdot |ab|}, \quad \text{or} \quad \frac{|ab|^2 + |pa|^2 - |pb|^2}{2|ab| \cdot |pa|}.$$

Let (x, y) be coordinates of a point p. Noting that coordinates of the points a and b are constants, the numerators are linear or quadratic expressions in x and y while the denominators are either square roots of quadratic expressions or of a product of two such expressions. Thus, given two edges ab and cd, a set of points p giving the same smallest angle, that is, those points satisfying $f(p, ab) = f(p, cd)$ form a planar curve of degree at most 7.

Fig. 4 shows two different settings of two line segments in the plane. Three regions are defined for each of the line segments by two circles and perpendicular bisector of the line segment. So, given two line segments, the plane is divided into at most 9 regions. Each of the figures shows which angle among $\angle pab$, $\angle abp$, $\angle bpa$, $\angle pcd$, $\angle cdp$, and $\angle dpc$ is smallest using six different colors. The boundaries of two different regions are curves of degree at most 7.

Figure 4: Which angle is smallest? Colors in the figure indicate which angle is smallest. Two different modified angular Voronoi diagram for two line segments.