Optimal Triangulation with Steiner Points

Boris Aronov[1], Tetsuo Asano[2], and Stefan Funke[3]

[1] Polytechnic University, Brooklyn, NY, USA
[2] School of Information Science, JAIST, Japan
[3] Max-Planck-Institut für Informatik, Saarbrücken, Germany

Abstract. There are many ways to triangulate a simple n-gon; for certain optimization criteria such as maximization of the smallest internal angle it is known how to efficiently compute the best triangulation with respect to this criterion. In this paper we consider a natural extension of this problem: Given a simple polygon P and one Steiner point p in its interior, determine the optimal location of p and a triangulation of P and p which is best amongst all triangulations and placements of p. We present a polynomial-time algorithm for this problem when the optimization criterion is maximization of the minimum angle. Furthermore, we also provide a more general polynomial-time algorithm for finding the optimal placement of a constant number of Steiner points under the same optimization criterion.

1 Introduction

Triangulations of simple polygons arise in many applications. Some triangulation of a given simple polygon can even be computed in linear time using Chazelle's algorithm [6]. Optimizing some criterion over all triangulations is also possible. For example, a popular optimization criterion is to maximize the minimum angle of any triangle. Such a triangulation is known as a *constrained Delaunay triangulation*; it can be obtained in $O(n \log n)$ time for an n-gon [7]. We could also find a *minimum-weight triangulation* that minimizes the total length of chords required for triangulation using dynamic programming. Dynamic programming is also powerful enough to find a triangulation in which the worst aspect ratio of 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. More precisely, given a simple polygon P, we want to find a point p in the interior of P such that the quality of the optimal triangulation of $P + \{p\}$ is optimized under a given optimization criterion. If maximization of the minimum internal angle is the goal, we want to find a location of an interior point p such that the minimum angle of the optimal triangulation of $P + \{p\}$ is maximized among all possible interior points p. As far as the authors know, there is no previous study of the question. Our main concern in this paper is to develop a polynomial-time algorithm.

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. Using a different optimization criterion is also interesting. Minimization of the largest internal angle, minimization of the largest slope, and minimization of the largest aspect ratio are rather popular criteria [3, 4], but no $O(n \log n)$ algorithm is known for these criteria except for that of maximization of the smallest angle (if adding Steiner points is not permitted). So, although it is challenging to extend our ideas to other criteria, in this paper we shall only consider the maxmin angle criterion for which we can design polynomial-time algorithms for the case of one or a constant number of Steiner points.

This problem is closely related to mesh improvement. Given a triangulation of some bounded domain, we sometimes want 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., [9]) an internal vertex is relocated to the barycenter of the polygon defined by its incident triangles. It 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, that is its underlying graph, is unchanged. Hence the barycenter is just a 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 [14] for an excellent survey on Delaunay refinement techniques.

This paper is organized as follows: Section 2 describes a polynomial-time algorithm for computing the optimal location for a single Steiner point. The more general case of a constant number of Steiner points is considered in Section 3. Section 4 includes conclusions and future work.

2 Triangulation using One Steiner Point

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

Problem 1. Given a simple n-gon P, find a triangulation of the interior of P with one Steiner point maximizing the smallest internal angle.

To that end we will consider the following more general problem:

Problem 2. Given a set of points X and a set of non-crossing edges E with endpoints in X. Find a triangulation of X and one Steiner point which respects the edges in E and maximizes the smallest internal angle.

As we will see below, our solution to Problem 2 also provides us with a solution to Problem 1.

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* $b \in X$ iff the line segment ab does not cross any edge in E (note that, for a line 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 1. *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 either $(a, b) \in E$, or (1) a sees b and (2) there exists a circle through a and b such that no $c \in X$ contained in the interior of the circle 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. It maximizes the minimum interior angle of any triangulation of (X, E) that uses only the points of X as triangulation vertices; in fact, the CDT lexicographically maximizes the list of angles from smallest to largest, see [5] for an extensive list of references.

In the following we are interested in how the constrained Delaunay triangulation changes when some point p is added to X. Definition 1 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 $\triangle abc$. Hence the insertion of p can only affect 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 iff 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 1. *$\mathcal{D}(p) := \mathrm{CDT}(X \cup \{p\}, E)$ can be obtained from $\mathcal{D} := \mathrm{CDT}(X, E)$ by deleting all edges in $\mathrm{CDT}(X, E)$ invalidated by p and retriangulating 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 1. Furthermore, it is not possible that $\mathcal{D}(p)$ contains an edge vw which was not present in $\mathcal{D}$, with $v, w \neq p$, 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 a technical reason explained below, we further refine $\mathcal{A}$ into constant-size "trapezoids," by replacing it with its trapezoidal decomposition [2, p. 124].) We argue that the topology of $\mathcal{D}(p)$ does not change when p is moved within one cell σ of this arrangement.

Lemma 2. *If a point $p \in \sigma$ can be seen from the interior of a triangle $\triangle abc \in \mathcal{D}$ whose circumcircle contains p, then all points within σ can be seen from the interior of $\triangle abc$.*

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

Since all points within a cell σ also lie within the same set of circumcircles, we have shown that all points within σ of the above arrangement 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 3. *Let $e = (b, c)$ be an edge of* $\mathrm{CDT}(X, E)$, *$\triangle abc$ and $\triangle bcd$ the respective adjacent triangles 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 x cannot see a. Consider the line segment $\overline{yp}$ as y moves along $\overline{xa}$ towards a. At some point $\overline{xp}$ must meet a constraining edge $e \in E$ at a vertex of X lying in the interior of the circumcircle of $\triangle abc$, contradicting the assumption that $\triangle abc \in \mathcal{D}$. □

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 has complexity $O(n^2)$ and the additional $O(n)$ line segments only add $O(n^2)$ intersection points. We summarize our findings in the following corollary.

Corollary 1. *The arrangement $\mathcal{A}$ defined by the triangles of $\mathcal{D}$ and their circumcircles characterizes the different topologies of* $\mathrm{CDT}(X \cup \{p\}, E)$ *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 size of $\mathcal{A}$ is $O(n^2)$.*

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

For our original Problem 1 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

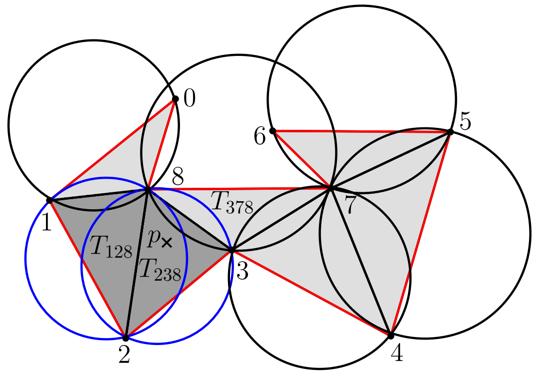

Fig. 1. A point p in the interior of a simple polygon is contained in the circumcircles of the two Delaunay triangles T_{128} and T_{238}, but not in that of T_{378}.

are the polygon vertices and E the polygon edges, and only consider the cells of $\mathcal{A}$ inside P.

Fig. 1 shows an example. The interior of the polygon (shaded) is partitioned into cells by all triangles and their respective circumcircles. The circumcircles of the triangles T_{128} and T_{238} (shaded darker) contain the point p and p is visible from both triangles. So, edge $\overline{28}$ 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.

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 $\mathrm{CDT}(X, E)$. We then optimize the minimum angle in the triangulation, over all possible placements of $p \in \sigma$, only focusing on the interior angles in the star triangulation of H, as the rest of the triangulation is unaffected by the insertion of p.

Algorithm for finding an optimal location of a Steiner point
Input: a set X of points and a set of non-crossing edges E with endpoints in X

1. Compute the constrained Delaunay triangulation $\mathcal{D} := \mathrm{CDT}(X, E)$.
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.

The next section will describe in detail how to actually treat a cell $\sigma \in \mathcal{A}$ using angular Voronoi diagrams. Note that the above algorithm also solves our original Problem 1 of optimizing the triangulation of a simple polygon P using one Steiner point: In step 3 of the algorithm we only need to consider cells that lie in the interior of P and when maximizing the minimum angle we also only consider triangles that lie in the interior of P.

Finding an optimal triangulation with fixed topology 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.

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 [12] 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 visual angle around 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 straightline segments, rather than ranging over all of H. It seems to be hard 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.

The *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 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 $Tr(p, e_i)$ by connecting the endpoints to p. The value $f(p, e_i)$ is defined to be the smallest internal angle in the triangle $Tr(p, e_i)$. The region $\mathrm{Vor}(e_i)$ of an edge e_i is defined by a set of points p at which $f(p, e_i)$ is smallest among all edges, that is,

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

Given a line segment ab, we partition the plane into four 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. Refer to Fig. 2. If the point p lies to the left of the line l_{ab} (more precisely, the halfplane defined by the line l_{ab} that contains the point a) and in the exterior of C_b, then the smallest internal angle of $\triangle abp$ is $\angle apb$. If p lies in the same halfplane but in the interior of C_b, then $\angle abp$ is smallest. The smallest angle in the right halfplane is similarly defined. 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.

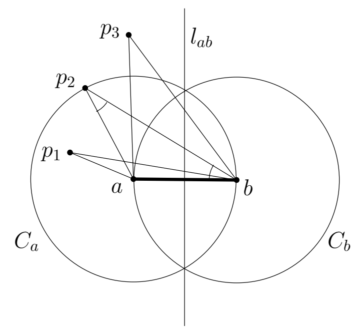

Fig. 2. Partition of the plane into regions according to which angle of $\triangle abp$ is smallest.

Fig. 3 in which Voronoi regions are painted by colors associated with polygon edges gives an example of such a modified angular Voronoi diagram. The given polygon H is indicated by solid white lines.

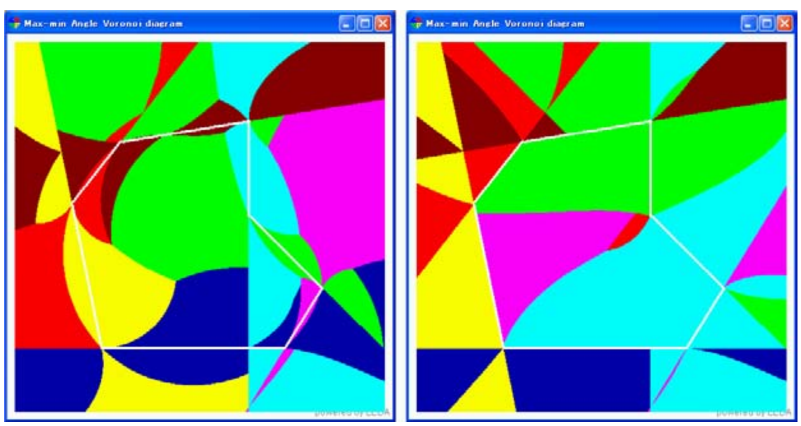

Fig. 3. Modified angular Voronoi diagram (left) and original angular Voronoi diagram.

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 Voronoi regions. Recall that we have refined $\mathcal{A}$ so that a cell $\sigma \in \mathcal{A}$ is a constant-complexity region. In particular, each function $f(p, e_i)$, viewed as truncated to within σ, is a well-behaved function with a constant-complexity domain. Hence standard results of envelope theory [13] imply that the modified angular Voronoi diagram truncated to within σ has at most $O(n^{2+\varepsilon})$ edges, including connected portions of Voronoi edges within σ and portions of Voronoi cells lying along the boundary of σ. 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.

3 Triangulation using Several Steiner Points

We now turn our attention to the situation when two Steiner points, p and q, are permitted to be placed in a simple n-gon P. We start with the triangulation $\mathcal{D} := \mathrm{CDT}(X, E)$. We consider the space $P^2 := P \times P$ of all possible placements of the two points. We aim to identify the best placement of p, q in order to maximize the smallest angle in the resulting constrained Delaunay triangulation $\mathcal{D}(p, q) := \mathrm{CDT}(X \cup \{p, q\}, E)$ (where as before X is the set of vertices of P and E is the set of its edges). As in the previous section, we partition P^2 according to the topology of $\mathcal{D}(p, q)$, then use an analog of the modified angular Voronoi diagram from previous section to determine which angle is smallest in the triangulation for every choice of $(p, q) \in P^2$ 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, especially when coupled with some heuristics to eliminate infeasible placements of p and q in order to reduce the search space, which we have developed but have been unable to include in this version due to space limitations.

We first recall a standard fact, the analogue of Definition 1 [7].

Fact 1 *A triangle $\triangle abc$, for $a, b, c \in X$ is present in* $\mathrm{CDT}(X, E)$ *if and only if a, b, c are pairwise visible and no other vertex of E 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 $(p, q) \in P^2$ 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 triples of distinct elements in $X \cup \{p, q\}$, for which the function $f(a, b, c; p, q)$ is defined at (p, q). The desired triangulation maximizing the minimum angle is just the maximum of function $m(p, q)$ over all of P^2.

We would like to apply envelope theory to compute $m(p, q)$ for all p, q. In a typical lower envelope argument [13], however, functions are well-behaved (e.g., algebraic of bounded degree) *and* defined over domains of constant description complexity (say, by a constant-length semialgebraic condition of bounded degree). In our case, the form of functions f is simple enough (if one uses, for example, $\cos^2 \angle abc$ to compare angles, to avoid transcendental functions, this is a low-degree rational function of the coordinates of the points).

The difficulty is in their domains of definition—they are in general not of constant complexity and hence the envelope analysis is not applicable directly. We instead decompose P^2 into constant-size cells in such a manner that each function is either total, or totally undefined on every cell c of the decomposition.

In order to construct such a decomposition, we observe that the boundaries of the domain of definition of a function $f(a, b, c; p, q)$ are given precisely by Fact 1. Namely, if we view p and q as moving, a triangle $\triangle abc$ formed by three points from among the vertices of P and/or p, q can cease to belong to $\mathcal{D}(p, q)$ only when some visibility constraint is violated (two possibilities: a vertex becomes collinear with pq, or p or q becomes collinear with a line defined by two vertices) or when a cocircularity is created or destroyed (again, two possibilities: p or q becomes cocircular with three vertices, or both p and q become cocircular with two vertices). All four possibilities correspond to a low-degree hypersurface in $\mathbb{R}^4$, and the number of possibilities is clearly polynomial, since there are n vertices in all.

We collect all these hypersurfaces, add boundaries of P^2 to the arrangement, and truncate it to within P^2. We then refine the resulting partition of P^2 to contain only constant-size cells (e.g., via a cylindrical algebraic decomposition or a vertical decomposition [8, 10]); the resulting decomposition $\mathcal{A}$ is still of polynomial size.

$\mathcal{A}$ is a subdivision of P^2 into polynomial number of constant-complexity cells σ with the property that each function $f(a, b, c; \cdot, \cdot)$ is either defined on the entire cell σ or undefined on all of σ. Functions are algebraic of bounded degree. There is a polynomial number of functions. Hence standard envelope theory [13] concludes that the minimization diagram (i.e., the decomposition of space into maximal connected portions over each of which a fixed function or set of functions achieves the pointwise minimum) of this collection of functions can be computed and, if necessary, further refined to constant-complexity cells, in polynomial time. Over each cell of the minimization diagram, a single function appears on the lower envelope, so we can determine in constant time its largest value. Taking the maximum over all cells, we obtain the placement of $p, q \in P$ that maximizes the minimum angle of any triangulation of P with two Steiner points p, q, in polynomial time, as promised.

We summarize our findings in the following theorem.

Theorem 2. *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* $\mathrm{CDT}(X \cup \{p, q\})$ *(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 of Steiner points. We omit the details due to space limitations.

4 Conclusions and Future Work

In this paper we have presented polynomial-time algorithms for finding optimal placement of one, or a constant number of, Steiner points to be inserted in a simple n-gon to maximize the minimum internal angle of triangulation. It would be interesting to improve the dependence of the latter algorithm on the number of Steiner points, to construct practical algorithms for solving the problems, and to extend our analysis to other optimization criteria.

Acknowledgments

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. 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 Institute für Informatik, Germany, Carleton University, Canada, and New York University, USA, all in 2006. 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).

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. de Berg, M. van Kreveld, M. Overmars, O. Schwarzkopf, Computational Geometry: Algorithms and Applications, Second, revised edition, Springer, Berlin (2000).
3. 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).
4. 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).
5. M. Bern and D. Eppstein: “Mesh Generation and Optimal Triangulation,” Computing in Euclidean Geometry, World Scientific (1992).

6. B. Chazelle: "Triangulating a simple polygon in linear time," Discrete & Computational Geometry, 6(5): 485-524 (1991).
7. P. Chew: "Constrained Delaunay triangulation," Proc. 3rd ACM Symp. on Computational Geometry, 215–223 (1987).
8. 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).
9. D.A. Field: "Laplacian Smoothing and Delaunay Triangulations," Communications in Applied Numerical Methods, Wiley, 4: 709–712 (1988).
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. J. Matoušek, M. Sharir and E. Welzl: "A subexponential bound for linear programming," Algorithmica, 16: 498–516 (1996).
13. M. Sharir and P. K. Agarwal, Davenport-Schinzel Sequences and Their Geometric Applications. Cambridge University Press, Cambridge, UK (1995).
14. J.R. Shewchuk: "Delaunay Refinement Algorithms for Triangular Mesh Generation," Computational Geometry: Theory and Applications 22(1-3): 21–74 (2002).
15. R. Seidel: "Constrained Delaunay triangulations and Voronoi diagrams with obstacles," in "1978-1988 Ten Years IIG," 178–191 (1988).