Computers Math. Applic. Vol. 20, No. 7, pp. 45–49, 1990
Printed in Great Britain.

0097-4943/90 $3.00 + 0.00

GENERALIZED DELAUNAY TRIANGULATIONS OF NON-CONVEX DOMAINS

C. BÖRGERS
Department of Mathematics, University of Michigan, Ann Arbor, MI 48109, U.S.A.

(*Received* 17 *May* 1989; *in revised form* 26 *September* 1989)

Abstract—The Delaunay triangulation associated with a finite set S of points in the plane is a triangulation of the convex hull of S. Delaunay triangulations have been used in a number of computational methods, for example in Lagrangian fluid dynamics. In this paper, we discuss two natural generalizations of the notion of Delaunay triangulation to non-convex domains. We show that the two generalizations are equivalent, and describe a method for the construction of generalized Delaunay triangulations.

1. INTRODUCTION

Consider the following problem. Given a finite set S of points in the plane, determine a triangulation of the convex hull of S whose vertices are exactly the points in S, and whose triangles are as non-degenerate as possible. This task arises, for example, in free Lagrangian methods for the computation of fluid flow [1].

If the requirement that the triangles be as non-degenerate as possible is made precise in a suitable way, described in Theorem 2 below, the solution of this problem is given by the so-called Delaunay triangulation [e.g. 2]. This triangulation can be defined as follows. For each point x in S, define $P(x)$ to be the set of all points in the plane which are at least as close to x as to any other point in S. $P(x)$ is an intersection of closed half planes, therefore a closed, convex, possibly unbounded polygon. The collection of the boundaries of the polygons $P(x)$, $x \in S$, is called the Voronoi diagram associated with S. Connecting two points x, y in S by a straight line exactly if their Voronoi polygons $P(x)$, $P(y)$ share an edge, one obtains the Delaunay triangulation associated with S.

There is a small difficulty with this definition: the Delaunay triangulation associated with S is nearly, but not completely unique. To see this, consider a vertex shared by n Voronoi polygons. We call the vertex simple if $n = 3$, multiple if $n > 3$. We always think of multiple vertices as being resolved into several simple ones, by introducing additional edges of length zero; see Fig. 1. This resolution is always possible, but never unique. Therefore, the Delaunay triangulation associated with S is not unique if there are multiple vertices in the Voronoi diagram.

The following theorem is an immediate consequence of the definition.

Theorem 1

Let τ be a triangulation of the convex hull of S, with the property that S is the set of vertices of τ. Then τ is a Delaunay triangulation iff for every triangle $\triangle \in \tau$, the interior of the circumscribed circle of $\triangle$ contains no point in S.

Theorem 2 gives a well-known equivalent characterization of Delaunay triangulations [e.g. 3]. This characterization makes clear in which sense Delaunay triangulations are as non-degenerate as possible.

Theorem 2

Let τ be a triangulation of the convex hull of S, with the property that S is the set of vertices of τ. Then τ is a Delaunay triangulation iff any pair $(\triangle_1, \triangle_2)$ of adjacent triangles in τ satisfies one of the following two conditions:

(a) $\bar{Q} = \bar{\triangle}_1 \cup \bar{\triangle}_2$ is not convex.
(b) The sum of the angles in $\bar{Q}$ divided by the shared edge of $\bar{\triangle}_1$ and $\bar{\triangle}_2$ is at least as large as the sum of the two other angles in $\bar{Q}$.

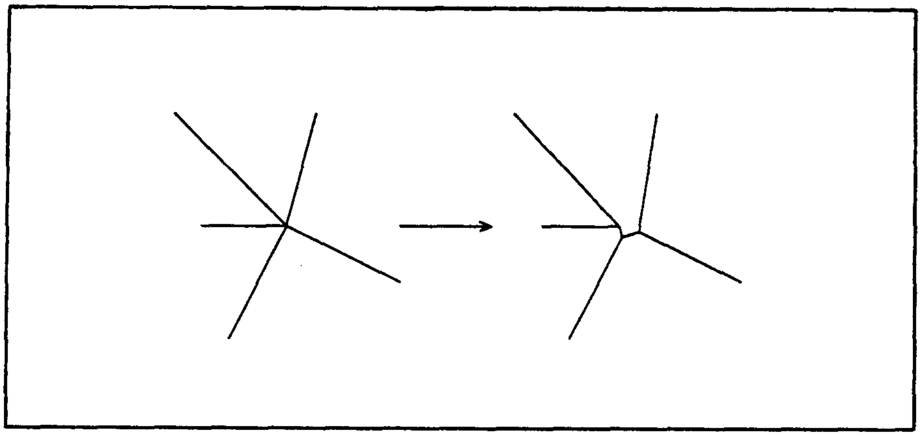

Fig. 1. Resolution of a multiple vertex into a set of simple vertices.

Delaunay triangulations are, by definition, triangulations of convex domains in the plane. Theorem 1 suggests a generalization of the notion of Delaunay triangulation to non-convex domains; see Section 2. Theorem 2 leads to another natural generalization; see Section 3. We prove in Section 3 that these two generalizations are equivalent. In Section 4, we show how the second characterization of generalized Delaunay triangulations leads to a proof of existence of such triangulations, and an algorithm for their construction. Section 5 summarizes this algorithm.

2. A GENERALIZED DEFINITION SUGGESTED BY THEOREM 1

Definition

Let Ω be a bounded open region in the plane with a polygonal boundary. Let S be a finite set of points in the closure $\bar{\Omega}$ of Ω, containing all vertices of $\partial\Omega$, and possibly other points. Let τ be a triangulation of Ω with the property that S is the set of vertices of τ. τ is called a generalized Delaunay triangulation of Ω with respect to S if it has the following property. Let $\triangle abc \in \tau$. By convention, we take $\triangle abc$ to be an open set. Let $\bigcirc abc$ be the interior of the circumscribed circle of $\triangle abc$. Then for any $p \in \triangle abc$, $x \in \bigcirc abc$, $x \in S$, the straight line segment $[x, p]$ intersects $\partial\Omega$. Briefly, $\partial\Omega$ separates x from $\triangle abc$.

Remarks

(1) The convention that $\bigcirc abc$ denotes the *interior* of the circumscribed circle of $\triangle abc$ excludes the possibility that x is one of the vertices, a, b, c.
(2) Ω is allowed to have slits, i.e. segments of the boundary may lie in the interior of the closure $\bar{\Omega}$ of Ω.
(3) This way of generalizing the Delaunay triangulation to non-convex domains was suggested in Ref. [4, p. 257], and also in Ref. [5].

3. AN EQUIVALENT CHARACTERIZATION OF GENERALIZED DELAUNAY TRIANGULATIONS

The following theorem makes clear in which sense generalized Delaunay triangulations are as non-degenerate as possible.

Theorem 3

Let τ be a triangulation of Ω with the property that S is the set of vertices of τ. Then τ is a generalized Delaunay triangulation iff any pair $(\triangle_1, \triangle_2)$ of adjacent triangles in τ satisfies one of the following three conditions:

(a) $\bar{Q} = \bar{\triangle}_1 \cup \bar{\triangle}_2$ is not convex.
(b) The sum of the angles in $\bar{Q}$ divided by the shared edge of $\bar{\triangle}_1$ and $\bar{\triangle}_2$ is at least as large as the sum of the two other angles in $\bar{Q}$.
(c) The edge shared by $\bar{\triangle}_1$ and $\bar{\triangle}_2$ is a part of $\partial\Omega$.

Proof. (1) We shall first show that in a generalized Delaunay triangulation, any pair $(\triangle_1, \triangle_2)$ of triangles satisfies one of the conditions (a)–(c). Suppose that $\triangle_1 = \triangle abc$ and $\triangle_2 = \triangle bcd$ violate conditions (a)–(c). Then the point d lies in the circumscribed circle $\bigcirc abc$, violating the definition of Section 2.

(2) We shall now show that a triangulation τ with the property that any pair $(\triangle_1, \triangle_2)$ of triangles satisfies one of the conditions (a)–(b) is a generalized Delaunay triangulation. Assume that τ violates the definition of Section 2. We show that there is a pair of adjacent triangles violating conditions (a)–(c). Since τ violates the definition of Section 2, there is a triangle $\triangle abc$ in τ, a point $p \in \triangle abc$ and a point $x \in \bigcirc abc$, $x \in S$, such that $\partial\Omega$ does not separate x from p. $x \notin \overline{\triangle abc}$, otherwise τ would not even be a triangulation. Without loss of generality, assume that $[b, c]$ is the edge of $\triangle abc$ separating x and p from each other. $[b, c]$ is not part of $\partial\Omega$, otherwise $\partial\Omega$ would separate x from p. Thus, there is a triangle $\triangle bcy$ in τ, where y lies on the same side of the edge $[b, c]$ as x. There are two possibilities: $y \in \bigcirc abc$ (possibly $y = x$), or $y \notin \bigcirc abc$. If $y \in \bigcirc abc$, then the triangles $\triangle abc$ and $\triangle bcy$ violate conditions (a)–(c). Thus, we only have to consider the case $y \notin \bigcirc abc$. Clearly, $x \notin \overline{\triangle bcy}$, otherwise τ would not even be a triangulation. It is then clear that x is separated from $\triangle abc$ either by the edge $[b, y]$, or by the edge $[c, y]$. Without loss of generality, assume that $[b, y]$ separates $\triangle abc$ from x. Then $[b, y]$ is not part of $\partial\Omega$, since $\partial\Omega$ does not separate x from p. From $x \in \bigcirc abc$ and $y \notin \bigcirc abc$ follows by a simple argument that $x \in \bigcirc bcy$. Therefore, we may repeat our construction, with $\bigcirc bcy$ playing the role which $\triangle abc$ played before. In this way, we obtain a sequence of triangles, all of which contain x in the interiors of their circumscribed circles, and all of which contain interior points which are not separated from x by $\partial\Omega$. To avoid a violation of conditions (a)–(c), this sequence would have to be infinite. Therefore, repetitions of triangles would have to occur. But the angle at x is strictly increasing. That is, the angle $\angle bxc$ is strictly smaller than the angle $\angle bxy$. Therefore repetitions are impossible. □

4. PROOF OF EXISTENCE FOR GENERALIZED DELAUNAY TRIANGULATIONS

In the proof of Theorem 4, we shall show how generalized Delaunay triangulations can be obtained in an iterative way. The iteration starts with an arbitrary triangulation of Ω whose set of vertices is S. Lemma 2 states that there is such a triangulation, and the proof of Lemma 2 shows a way of constructing one.

Lemma 1

Let Ω be a bounded open region in the plane with a polygonal boundary. Then there is a triangle of positive area contained in Ω whose vertices are vertices of $\partial\Omega$.

Proof. Consider an arbitrary edge $[a, b]$ of $\partial\Omega$. Let x be a point moving in the following way. Initially, x is the midpoint of the edge $[a, b]$. x then moves into Ω along the perpendicular bisector of $[a, b]$. Initially, $\triangle abx$ will be contained in Ω, but eventually one of the edges $[a, x]$ or $[b, x]$ will touch $\partial\Omega$. At the first instant when this happens, there are two possibilities. Either there is a vertex c of $\partial\Omega$ different from a, b which lies on one of the edges $[a, x]$ or $[b, x]$. In this case, $\triangle abc$ is a triangle of the desired kind. Or one of the edges $[a, x]$ or $[b, x]$ is entirely contained in an edge of $\partial\Omega$. We now let x move along that boundary edge, in the direction away from a, until the edge $[b, x]$ touches $\partial\Omega$. At the first instant when this happens, there is a vertex c of $\partial\Omega$ different from b which lies on the edge $[b, x]$, and $\triangle abc$ is a triangle of the desired kind. □

Corollary 1

Let Ω be a bounded open region in the plane with a polygonal boundary. Then there is a partition of Ω into triangles Ω_i, $i = 1, \ldots, k$ such that the union of the sets of vertices of the $\partial\Omega_i$ equals the set of vertices of $\partial\Omega$.

Proof. Let Ω_1 be the triangle whose existence is asserted in Lemma 1. Then apply Lemma 1 to the domain $\Omega - \bar{\Omega}_1$. The conclusion is that there exists a triangle Ω_2 of positive area contained in $\Omega - \bar{\Omega}_1$ whose vertices are vertices of $\partial(\Omega - \bar{\Omega}_1)$, and thus of $\partial\Omega$. Continuing this process, we obtain a partition of the desired kind. □

Lemma 2

Let Ω and S be as described in the definition of Section 2. Then there is a triangulation τ_0 of Ω with the property that S is the set of vertices of τ_0.

Proof. Consider a partition of Ω into convex polygons $\Omega_1, \ldots, \Omega_k$ such that the union of the sets of vertices of the $\partial\Omega_i$ equals the set of vertices of $\partial\Omega$. Such a partition exists by Corollary 1. Let

$$S_j := S \cap \bar{\Omega}_j.$$

We remark that $S_j \cap S_i$ may be non-empty for $i \neq j$. The Delaunay triangulation associated with S_j is a triangulation of Ω_j. The union of all Delaunay triangulations associated with the S_j is a triangulation of Ω whose set of vertices is S. □

Theorem 4

Let Ω and S be as described in the definition of Section 2. Then there is a generalized Delaunay triangulation of Ω with respect to S.

Proof. We use an algorithm which has been proposed by a number of authors for the construction of ordinary Delaunay triangulations [e.g. 6].

Let τ_0 be as in Lemma 2. Look for a pair of adjacent triangles $\triangle_1$, $\triangle_2$ in τ_0 violating conditions (a)–(c) in Theorem 3. If there is no such pair, a generalized Delaunay triangulation has been found.

Otherwise, let $\bar{Q} := \bar{\triangle}_1 \cup \bar{\triangle}_2$. Re-divide $\bar{Q}$ into two triangles, along the diagonal which is not the common edge of $\bar{\triangle}_1$ and $\bar{\triangle}_2$.

This process can be repeated. One obtains a sequence of triangulations which ends in a generalized Delaunay triangulation, provided that it ends at all. We shall now show that the sequence must end. In Ref. [3], it was proved that whenever a quadrilateral $\bar{Q}$ is redivided, the circumscribed circles of both new triangles have strictly smaller radii than the circumscribed circles of both old triangles. Therefore, the sum of the radii of the circumscribed circles strictly decreases each time a diagonal is exchanged, and thus repetitions of triangulations are impossible. This completes the proof, since there are only finitely many triangulations which can possibly occur in the sequence. □

5. CONSTRUCTION ALGORITHM FOR GENERALIZED DELAUNAY TRIANGULATIONS

The arguments of Section 4 constitute an algorithm for the construction of generalized Delaunay triangulations. In this section, we state this algorithm in a more explicit form:

Step 1. Set $k := 0$ and $\tilde{\Omega} := \Omega$.

Step 2. If $\tilde{\Omega}$ is empty, go to Step 6.

Step 3. Choose an edge $[a, b]$ of $\partial\tilde{\Omega}$.

Step 4. For every vertex c of $\partial\tilde{\Omega}$ which does not equal a or b, test whether $\triangle abc$ is contained in $\tilde{\Omega}$, until a vertex c with this property is found.

Step 5. Set $k := k + 1, \Omega_k := \triangle abc$ and $\tilde{\Omega} := \tilde{\Omega} - \bar{\Omega}_k$. Go to Step 2.

Step 6. Compute $S_j := S \cap \bar{\Omega}_j$, for $j = 1, \ldots, k$.

Step 7. Compute the Delaunay triangulations associated with the S_j, $j = 1, \ldots, k$, and denote their union, a triangulation of Ω whose set of vertices is S, by τ.

Step 8. For each edge separating two adjacent triangles $\triangle_1$ and $\triangle_2$ in τ, test whether the pair $(\triangle_1, \triangle_1)$ violates conditions (a)–(c) of Theorem 3. If it does, replace the edge by the other diagonal of the convex quadrilateral $\bar{\triangle}_1 \cup \bar{\triangle}_2$. Repeat this procedure until no violation of conditions (a)–(c) occurs. Then τ is a generalized Delaunay triangulation of Ω.

We conclude with some remarks on the operation count for this algorithm. Let N denote the number of points in S, and n the number of vertices of $\partial\Omega$. Using Euler's formula, it is easy to see that $k < 2n - 4$. Let us assume that for any three vertices a, b, c of Ω, the determination whether $\triangle abc$ is contained in Ω can be made in $O(n)$ operations. Then it follows that $\Omega_1, \ldots, \Omega_k$ can be determined in at most $O(n^3)$ operations. Of course, in practice it may often be much less expensive to partition Ω into convex polygons Ω_i with the property that all vertices of the $\partial\Omega_i$ are vertices of $\partial\Omega$. Steps 1–5 of the algorithm only serve the purpose of finding such a partition, and can be replaced by any other procedure leading to such a partition.

Once $\Omega_1, \ldots, \Omega_k$ have been found, the sets S_j must be determined (Step 6). Let us denote the number of vertices of $\partial\Omega_j$ by n_j, and assume that the determination whether a given point $x \in S$ lies in $\bar{\Omega}_j$ can be made in $O(n_j)$ operations. Then the total number of operations needed to determine the sets S_j does not exceed

$$O\left(\sum_{j=1}^{k} n_j\right) \cdot N = O(nN).$$

Let N_j denote the number of points in S_j. The number of operations needed to compute the Delaunay triangulation associated with S_j does not exceed $O(N_j \log N_j)$; see Ref. [2] or Ref. [7, p. 214, Theorem 5.15]. In some applications, heuristic algorithms may have an operation count of $O(N_j)$ [e.g. 8]. Therefore the total number of operations needed for Step 7 is at most $O(N \log N)$, and may be $O(N)$ in practice.

I have not been able to give a realistic general estimate for the complexity of Step 8 of the algorithm. To my knowledge, no such estimate is known even for the case of a convex domain Ω, although numerical experiments for the convex case indicate that the algorithm is quite efficient.

Acknowledgements—I am grateful to Robert Krasny and to a referee for helpful comments. This work was supported in part by NSF Grant DMS-8801991. A portion of the work was carried out while the author was a Visiting Scientist at the IBM Thomas J. Watson Research Center.

REFERENCES

1. M. J. Fritts, W. P. Crowley and H. Trease (Eds), *The Free-Lagrange Method. Lecture Notes in Physics* **238.** Springer, New York (1985).
2. F. P. Preparata and M. I. Shamos, *Computational Geometry*. Springer, New York (1985).
3. R. Sibson, Locally equiangular triangulations. *Computer J.* **21,** 243 (1978).
4. S. H. Lamson and D. G. Holmes, Computational aspects of finite volume calculations on an unstructured triangular mesh. In *Advances in Computer Methods for Partial Differential Equations*, Vol. VI (Eds R. Vichnevetsky and R. S. Stepleman), pp. 255–259 (1987).
5. A. K. Cline and R. J. Renka, A generalized Delaunay triangulation and the solution of closest node problems in the presence of barriers. Preprint (1987).
6. M. J. Fritts, Two-dimensional Lagrangian fluid dynamics using triangular grids. In *Finite-difference Techniques for Vectorized Fluid Dynamics Calculations* (Ed. D. L. Book), p. 98. Springer, Berlin (1981).
7. M. I. Shamos and D. Hoey, Closest point problems. In *Proc. 16th A. IEEE Symp. on Foundations of Computer Science*, pp. 151–162 (1975).
8. C. Börgers and C. S. Peskin, A Lagrangian fractional step method for the incompressible Navier–Stokes equations on a periodic domain. *J. comput. Phys.* **70,** 397 (1987).