TAN, TIOW-SENG

Optimal Triangulation Problems

This paper surveys some recent solutions to triangulation problems in 2D plane and surface. In particular, it focuses on three efficient and practical schemes in computing optimal triangulations useful in engineering and scientific computations, such as finite element analysis and surface interpolation.

The edge-insertion paradigm can compute for a set of n vertices, with or without constraining edges, a min-max angle and a max-min height triangulation in $O(n^2 \log n)$ time and $O(n)$ storage, and a min-max slope and a min-max eccentricity triangulation in $O(n^3)$ time and $O(n^2)$ storage.

The subgraph scheme can compute a min-max length triangulation for a set of n vertices in $O(n^2)$ time and storage. Length refers to edge length and is measured by some normed metric such as the Euclidean or any other ℓ_p metric. Additionally, the scheme provides some insight to the minimum weight triangulation problem.

The wall scheme can compute for a given set of n vertices and m constraining edges, a conforming Delaunay triangulation of $O(m^2 n)$ vertices. Additionally, an extension of the wall scheme can refine a triangulation of size $O(n)$ to a quality triangulation of size $O(n^2)$ that has no angle measuring more than $\frac{11}{15}\pi$.

1. Introduction

Triangulation is a prominent meshing method that decomposes a domain into a collection of triangles for purposes of other computations. It is used in many areas of engineering and scientific applications such as finite element methods, approximation theory, numerical analysis, computer-aided geometric design, etc. Algorithmic problems on computing triangulations that are optimal according to measures on the size and shape of their triangles have been popular research topics since 1970s. Substantial advances on these problems were made recently. In this paper, we survey some recent solutions in computing a number of optimal triangulations mentioned in engineering and scientific computing literatures. Most of these algorithms are the first and currently the only ones that construct the required optimal triangulations in time polynomial to the input size.

The emphasis of this paper is somewhat different from that of the detailed survey on the subject by Bern and Eppstein [2]. We view those algorithms mentioned in the previous paragraph as examples of three general schemes in solving optimal triangulation problems. See [3,8,9,10] for complete technical details on the material discussed in this paper as well as many more relevant references omitted by this paper.

2. Problem Definitions

Let S be a set of n *points* or *vertices* in $\mathbb{R}^2$. An *edge* is a closed line segment connecting two points. Let E be a collection of edges determined by vertices of S. Then $\mathcal{G} = (S, E)$ is a *plane geometric graph* if (i) no edge contains a vertex other than its endpoints, that is, $ab \cap S = \{a, b\}$ for every edge $ab \in E$, and (ii) no two edges *cross*, that is, $ab \cap cd \in \{a, b\}$ for every two edges $ab \neq cd$ in E. One example of a plane geometric graph is a (simple) *polygon* where E forms a single cycle.

A *triangulation* is a plane geometric graph $\mathcal{T} = (S, E)$ so that E is maximal. By maximality, edges in E bound the convex hull of S, i.e. the smallest convex set in $\mathbb{R}^2$ that contains S, and subdivide its interior into disjoint faces bounded by triangles. With reference to a polygon, we talk about its triangulation as restricted to only within the cycle bounding the polygon.

A plane geometric graph $\mathcal{G} = (S, E)$ can be augmented with an edge set E' until it is a triangulation $\mathcal{T} = (S, E \cup E')$, referred to as a *triangulation of* $\mathcal{G}$. In this case, E is the set of *constraining edges* if it is not empty. Besides edges, we can also augment $\mathcal{G} = (S, E)$ with a vertex set S'. A triangulation obtained in this manner is called a *Steiner triangulation* of $\mathcal{G}$, and S' the set of *Steiner vertices*. We call a Steiner triangulation with constraining edges a *conforming triangulation* in which each constraining edge is the union of some edges in the triangulation.

A plane geometric graph $\mathcal{G}$ permits many augmentations, with or without Steiner vertices, to different triangulations. Various shape criteria can be used to classify some as *optimal triangulations*. Many of these criteria are defined as *max-min*, short for maximizes the minimum, or *min-max* when no Steiner vertices are used. The first quantifier is over all triangulations of $\mathcal{G}$ and the second is over all measures of triangles of a triangulation. A *measure* μ is a function that maps a triangle xyz to a real value $\mu(xyz)$. Examples of measures are largest angle and largest edge length. When Steiner vertices are allowed, we refer to such optimal triangulation as *optimal conforming triangulations*. Problems abstracted from the engineering and scientific computing literatures as addressed by this paper are: given a plane geometric graph $\mathcal{G} = (S, E)$, find a specific optimal triangulation or find a specific optimal conforming triangulations. The *edge-insertion paradigm* (Section 3) and the *subgraph scheme* (Section 4) address optimal triangulation problems whereas the *wall scheme* (Section 5) solves some optimal conforming triangulation problems.

3. The Edge-Insertion Paradigm

The edge-insertion paradigm is an iterative improvement method that computes min-max or max-min optimal triangulations. In the case of a min-max criterion, we consider the construction of a triangulation of a plane geometric graph $\mathcal{G} = (S, E)$ whose maximum measure $\mu(xyz)$ over all its triangles xyz is the smallest among all possible triangulations of $\mathcal{G}$. Formally, the *measure* of a triangulation $\mathcal{A}$ is defined as $\mu(\mathcal{A}) = \max\{\mu(xyz) : xyz \text{ a triangle of } \mathcal{A}\}$. If $\mathcal{A}$ and $\mathcal{B}$ are two triangulations of $\mathcal{G}$, then $\mathcal{B}$ is called an *improvement* of $\mathcal{A}$, if $\mu(\mathcal{B}) < \mu(\mathcal{A})$, or $\mu(\mathcal{B}) = \mu(\mathcal{A})$ and the set of triangles xyz in $\mathcal{B}$ with $\mu(xyz) = \mu(\mathcal{B})$ is a proper subset of the set of such triangles in $\mathcal{A}$. A triangulation $\mathcal{A}$ is *optimal* for μ if there is no improvement of $\mathcal{A}$. A max-min criterion can be treated analogously.

The basic idea of the edge-insertion paradigm is to improve iteratively a current triangulation $\mathcal{A}$ by a simple *edge-insertion* step which adds an appropriate new edge say qs to $\mathcal{A}$, deletes edges in $\mathcal{A}$ that cross qs, and retriangulates the resulting polygons to the left and the right of qs. In other words, the method starts by constructing an arbitrary triangulation $\mathcal{A}$ of $\mathcal{G}$, then iteratively applies the edge-insertion step until no further improvement to the current triangulation is found. Obviously, such simple idea does not work for all measures μ as some may lead to sub-optimal solutions. In fact, the paradigm is known to be applicable if the so-called *Cake-Cutting Lemma* that guarantees improvement in each iteration and thus convergent to optimal solution is true; see [1,5,9] for details. The next result is obtained by the edge-insertion paradigm.

Theorem 1. *For a plane geometric graph $\mathcal{G}$ of n vertices,*
(1) a min-max angle triangulation of $\mathcal{G}$ can be computed in time $O(n^2 \log n)$ and storage $O(n)$,
(2) a max-min height triangulation of $\mathcal{G}$ can be computed in time $O(n^2 \log n)$ and storage $O(n)$,
(3) a min-max eccentricity triangulation of $\mathcal{G}$ can be computed in time $O(n^3)$ and storage $O(n^2)$, and
(4) a min-max slope triangulation of $\mathcal{G}$ can be computed in time $O(n^3)$ and storage $O(n^2)$.

Let us define those terms mentioned in the theorem. A *min-max angle triangulation* of $\mathcal{G}$ minimizes the maximum angle of its triangles, over all triangulations of $\mathcal{G}$. The *height* $\eta(xyz)$ of triangle xyz is the minimum distance from a vertex to the opposite edge. We write $\eta(\mathcal{A}) = \min\{\eta(xyz) : xyz \text{ a triangle of } \mathcal{A}\}$ for the measure of a triangulation $\mathcal{A}$ of $\mathcal{G}$. A *max-min height triangulation* of $\mathcal{G}$ maximizes $\eta(\mathcal{A})$ over all triangulations $\mathcal{A}$ of $\mathcal{G}$. The *eccentricity* of triangle xyz, $\epsilon(xyz)$, is the infimum over all distances between the center of the circumcircle of xyz and points in the closure of xyz. Clearly, $\epsilon(xyz) = 0$ iff the center of the circumcircle lies in the closure of xyz. We define $\epsilon(\mathcal{A}) = \max\{\epsilon(xyz) : xyz \text{ a triangle of } \mathcal{A}\}$. A *min-max eccentricity triangulation* of $\mathcal{G}$ minimizes $\epsilon(\mathcal{A})$ over all triangulations $\mathcal{A}$ of $\mathcal{G}$. The definition of slope given in the next paragraph involves surfaces.

Consider a function $f : \mathbb{R}^2 \to \mathbb{R}$ defining a surface $x_3 = f(x_1, x_2)$ in $\mathbb{R}^3$. The *gradient* of f is the vector $\nabla f = (\frac{\partial f}{\partial x_1}, \frac{\partial f}{\partial x_2})$, each component of which is itself a function from $\mathbb{R}^2$ to $\mathbb{R}$. Define $\nabla^2 f = (\frac{\partial f}{\partial x_1})^2 + (\frac{\partial f}{\partial x_2})^2$, and call $\sqrt{\nabla^2 f}$ at a point (x_1, x_2) the *slope* at this point. Let S be a set of n points of $\mathcal{G}$ in $\mathbb{R}^2$ and let $\hat{S}$ be the corresponding set in $\mathbb{R}^3$ where each point of S has a third coordinate called *elevation*. For a point x of S, we write $\hat{x}$ for the "lifted" point, that is, the corresponding point in $\hat{S}$. Analogous to the definitions in $\mathbb{R}^2$, $\hat{x}\hat{y}$ denotes the line segment with endpoints $\hat{x}$ and $\hat{y}$, and $\hat{x}\hat{y}\hat{z}$ denotes the triangle with vertices $\hat{x}, \hat{y}, \hat{z}$. We can think of $\hat{x}\hat{y}\hat{z}$ as a partial function f on $\mathbb{R}^2$, defined within xyz. At each point in the interior of xyz, the gradient is well defined and the same as for any other point in the interior of xyz. We can therefore set $\sigma(xyz)$ equal to the slope at any point in the interior of xyz, and call it the *slope* of xyz. For a triangulation $\mathcal{A}$ of $\mathcal{G}$, we define $\sigma(\mathcal{A}) = \max\{\sigma(xyz) : xyz \text{ a triangle of } \mathcal{A}\}$. A *min-max slope triangulation* of $\mathcal{G}$ minimizes $\sigma(\mathcal{A})$ over all triangulations $\mathcal{A}$ of $\mathcal{G}$.

4. The Subgraph Scheme

The subgraph scheme constructs a desired optimal triangulation by first computing a sub-structure of the optimal triangulation and then complete the computation by solving smaller problems defined by the sub-structure. This scheme works when (i) the sub-structure can be computed efficiently and (ii) the sub-structure can subdivide the problem into smaller problems such as polygons that can be solved efficiently. For instance, the scheme has successfully solved the min-max length triangulation problem [6].

Theorem 2. *A min-max length triangulation of a set of n points in* $\mathbb{R}^2$ *can be constructed in* $O(n^2)$ *time and storage.*

A triangulation that minimizes the length of its longest edge over all possible triangulations of the same point set is called a *min-max length triangulation.* Notice that the theorem is formulated with reference to a set of n points instead of the general plane geometric graph. In fact, the theorem is valid for the latter provided the minimization condition is defined over all edges including the constraining ones; see [9]. In any case, the correctness of the theorem follows from the so-called *Subgraph Theorem* which asserts that every point set S in $\mathbb{R}^2$ has a min-max length triangulation $mlt(S)$ so that $rng(S) \cup ch(S) \subseteq mlt(S)$ where $rng(S)$ is the relative neighborhood graph of S and $ch(S)$ is the set of edges bounding the convex hull of S. The next paragraph provides further details on this.

The relative neighborhood graph of S, $rng(S)$, is a plane geometric graph with vertex set S and edge set containing ab, for a and b in S, iff $\text{dist}(ab) \leq \min_{x \in S-\{a,b\}} \max\{\text{dist}(xa), \text{dist}(xb)\}$ where $\text{dist}(xy)$ denotes the distance between x and y. Since $rng(S)$ and $ch(S)$ can each be computed in $O(n \log n)$ time, and $rng(S) \cup ch(S)$ is a connected graph of S, the problem of computing a $mlt(S)$ can be solved by first computing $rng(S) \cup ch(S)$ and then computing an optimal triangulation within each polygon defined by edges of $rng(S) \cup ch(S)$. The latter was also shown to be solvable in $O(n^2)$ time. Besides Euclidean metric, Theorem 2 can be extended to general normed metrics as stated in the next theorem.

Theorem 3. *Let* S *be a set of* n *points in* $\mathbb{R}^2$ *equipped with a normed metric. Given the relative neighborhood graph, a min-max length triangulation of* S *can be constructed in time* $O(n^2)$*.*

Examples of normed metrics are the ℓ_p-metrics, for $p = 1, 2, 3, \ldots$, and the so-called A-metric used in VLSI applications. The above theorem raises the question of how fast the relative neighborhood graph can be constructed for a given normed metric. The trivial algorithm tests all $\binom{n}{2}$ edges, each in time $O(n)$, and therefore takes time $O(n^3)$. Faster algorithms are known for the ℓ_p-metrics where $O(n \log n)$ time suffices.

We note that min-max length is currently the only non-trivial length criterion known to be computable in polynomial time. There are other related problems on length criteria whose complexities remain open. In particular the well-known problem of *minimum weight triangulation*, which is a triangulation that minimizes the sum of edge lengths over all triangulations of the same point set. The recent result of Keil shows that the so-called $\sqrt{2}$-skeleton of S is a subgraph of the minimum weight triangulation [8]. This subgraph, though can be computed efficiently, is not always a connected graph of S and thus does not subdivide the problem into smaller problems that have efficient solutions. Some experimental results in this direction can be found in [4].

5. The Wall Scheme

This scheme addresses conforming triangulation problems of a plane geometric graph $\mathcal{G}$. A common difficulty of adding Steiner vertices is as follows. Let us assume the extreme case that $\mathcal{G}$ is a triangulation for the discussion. We want to refine $\mathcal{G}$ so that it is an optimal conforming triangulation eventually. For the non-optimality of $\mathcal{G}$ at some triangle say pqr in $\mathcal{G}$, the natural attempt is to add a Steiner vertex say t in pqr or on one of its three edges. On the other hand, this new vertex t may in turn affect the optimality of some triangles nearby pqr or adjacent to pqr. In other words, the trouble of non-optimality has been propagated to other triangles and the addition of other Steiner vertices is necessary. For an approach to be successful, it has to terminate this kind of propagations in polynomial number of steps. The wall scheme is such an approach—it first plans sufficient Steiner vertices to build some *logical walls* that can stop arbitrary propagations. We illustrate this idea in the following paragraphs with the conforming Delaunay triangulation problem and the quality conforming triangulation problem. We note that the approach of packing disks in computing non-obtuse polygon triangulation [3] has the similar spirit.

Conforming Delaunay triangulation. A conforming Delaunay triangulation of a plane geometric graph $\mathcal{G} = (S, E)$ is a conforming triangulation $\mathcal{T} = (S \cup S', E')$ of $\mathcal{G}$ such that each edge cd in E' satisfies the so-called *empty disk property* with respect to $S \cup S'$, i.e., there is a circle through c and d so that all other points of $S \cup S'$ lie outside the open disk bounded by the circle. Initially, edges in E may not satisfy the empty disk property with respect to S. So, these edges have to be subdivided with a set S' of Steiner vertices into shorter edges so that each one satisfies the empty disk property with respect to $S \cup S'$. In fact, the main difficulty of constructing a conforming Delaunay triangulation of $\mathcal{G}$ is to find such a set S' so as to apply some well-known method to compute a Delaunay triangulation of $S \cup S'$. The following theorem is obtained by [7] using the wall scheme.

Theorem 4. *Let* $\mathcal{G} = (S, E)$ *be a plane geometric graph with* $|S| = n$ *and* $|E| = m \geq 1$*. A point set* $S \cup S'$ *of size* $O(m^2 n)$ *that admits a conforming Delaunay triangulation of* $\mathcal{G}$ *can be computed in time* $O(m^2 n + n^2)$*.*

The solution is to construct $V = S \cup S'$ in two steps, the *blocking* and the *propagation step*. Initially, V is equal to S. The goal of the blocking step is to find $\mathrm{O}(n)$ pairwise disjoint disks that contain no points of S so that the union of their closures is connected and contains S. Each circle bounding such a disk is called a *blocking circle*. Then, add Steiner vertices to V at the intersections between blocking circles and edges of $\mathcal{G}$, and at locations where blocking circles touch each other. With this, edges of E are subdivided by Steiner vertices into shorter edges of two types. Each *protected* edge is enclosed by a blocking circle; its endpoints lie on the blocking circle. All other edges are *unprotected*. By construction, protected edges satisfy the empty disk property with respect to the current set V. The construction will make sure that no points inside the blocking circles are later added to V so that this property persists with respect to all future sets V. The unprotected edges are further subdivided into shorter edges in the propagation step. Roughly speaking, if ab is unprotected and the circle with ab as a diameter contains a point $c \in V$ then a point c' subdividing ab into ac' and $c'b$ is added to V. This point c' is chosen so that c lies outside both the circle with ac' as a diameter and the circle with $c'b$ as a diameter. Due to the blocking step, a logical wall exists between two adjacent Steiner vertices in each blocking circle, and the collection of all logical walls subdivides the problem into smaller regions where the above subdivisions with c' cannot cross regions and thus bounds the size of V to that stated in the theorem.

QUALITY CONFORMING TRIANGULATIONS. For the discussion here, a quality conforming triangulation is a conforming triangulation that has angle bounded away from π. For a plane geometric graph $\mathcal{G} = (S, E)$, the algorithm mentioned in Theorem 1 can augment $\mathcal{G}$ to a triangulation $\mathcal{T}$ that minimizes its maximum angle over all possible augmentations. So, the construction is complete if $\mathcal{T}$ has angles measuring at most the targeted angle bound. If not, then $\mathcal{T}$ will be refined as discussed in [10] with six steps. Very briefly, the solution is to first build the so-called *fences* and *dead-ends* in the first two steps. These, as extensions to the notion of walls, can control the propagations of vertices in later steps. Specifically, the six steps successfully bounds the propagation threads to linear in number and each propagation thread to linear in length; the total number of Steiner vertices is, thus, quadratic in the size of the input. The following theorem summarizes the result.

Theorem 5. *Triangulating a plane geometric graph $\mathcal{G} = (S, E)$ of $|S| = n$ vertices using angles no larger than $\frac{11}{15}\pi$ requires $O(n^2)$ storage and $O(n^2 \log n)$ time.*

6. References

1 BERN M., EDELSBRUNNER H., EPPSTEIN D., MITCHELL S., TAN T. S.: Edge insertion for optimal triangulations, *Discrete & Computational Geometry* **10** (1993), 47–65. Also, computer programs are available via anonymous ftp from the directory "/SGI/MinMaxer" at site "ftp.ncsa.uiuc.edu".

2 BERN M., EPPSTEIN D.: Mesh generation and optimal triangulation, in *Computing in Euclidean Geometry*, 2nd Edition, edited by Du D. Z. and Hwang F. K., World Scientific, Singapore, 1995, 47–123.

3 BERN M., MITCHELL S., RUPPERT J.: Linear-size non-obtuse triangulation of polygons, in *Proc. 10th Annual ACM Symposium on Computational Geometry* (1994), 221–230.

4 CHENG S.W., GOLIN M.J., TSANG J.C.F.: Expected case analysis of β-skeletons with applications to the construction of minimum weight triangulations, in Proc. 7th Canadian Conference on Computational Geometry (1995).

5 EDELSBRUNNER H., TAN T. S., WAUPOTITSCH R.: An $\mathrm{O}(n^2 \log n)$ time algorithm for the minmax angle triangulation, *SIAM Journal on Scientific & Statistical Computing* **13** (1992), 994–1008.

6 EDELSBRUNNER H., TAN T. S.: A quadratic time algorithm for the minmax length triangulation, *SIAM Journal on Computing* **22** (1993), 527–551.

7 EDELSBRUNNER H., TAN T. S.: An upper bound for conforming Delaunay triangulations, *Discrete & Computational Geometry* **10** (1993), 197–213.

8 KEIL M.: Computing a subgraph of the minimum weight triangulation, *Computational Geometry: Theory and Applications* **4** (1994), 13–26.

9 TAN T. S.: Optimal Two-Dimensional Triangulations, PhD Thesis, Technical Report UIUCDCS-R-92-1783, Dept. Comput. Sci., University of Illinois, Urbana, 1992. Available as file "/pub/TechReports/UIUCDCS-R-92-1783.ps.Z" through anonymous ftp at site "ftp.cs.uiuc.edu".

10 TAN T. S.: An optimal bound for high-quality conforming quality triangulations, *Discrete & Computational Geometry*, to appear. Available as file "/pub/techreports/1994/TRB3-94.ps.gz" through anonymous ftp at site "ftp.iscs.nus.sg".

Address: TIOW-SENG TAN, National University of Singapore, Department of Information Systems and Computer Science, Lower Kent Ridge Road, Singapore 0511, Republic of Singapore. email: tants@iscs.nus.sg.

Contribution to the International Congress on Industrial & Applied Mathematics, 3–7 July 1995, Germany