

Delaunay Triangulations of Imprecise Points in Linear Time after Preprocessing

Maarten Löffler

Department of Information
and Computing Sciences
Utrecht University
Utrecht, the Netherlands
loffler@cs.uu.nl

Jack Snoeyink

Department of Computer Science
University of North Carolina
Chapel Hill, NC, USA
snoeyink@cs.unc.edu

ABSTRACT

An assumption of nearly all algorithms in computational geometry is that the input points are given precisely, so it is interesting to ask what is the value of imprecise information about points. We show how to preprocess a set of n disjoint unit disks in the plane in $O(n \log n)$ time so that if one point per disk is specified with precise coordinates, the Delaunay triangulation can be computed in linear time. From the Delaunay, one can obtain the Gabriel graph and a Euclidean minimum spanning tree; it is interesting to note the roles that these two structures play in our algorithm to quickly compute the Delaunay.

Categories and Subject Descriptors

F.2.2 [ANALYSIS OF ALGORITHMS AND PROBLEM COMPLEXITY] Nonnumerical Algorithms and Problems

General Terms

Theory, Algorithms

Keywords

Data imprecision, Delaunay triangulation, Gabriel graph, Minimum spanning tree

1. INTRODUCTION

A fundamental assumption of most algorithms in computational geometry is that the input data given is exact. There are actually two good justifications for this assumption: First, by carefully studying the predicates to perform exact computation on the data given, computational geometers can compute a result that is guaranteed to terminate, be self-consistent, and correct on the given input, which is at least close to the input desired. Second, we geometers don't really know what else to do when someone gives us numbers

or coordinates but to believe them. Somehow, these justifications are not reassuring to the application practitioners who know that their data is inexact before they throw it over the wall into the geometer's realm.

In this paper we wanted to explore the question, "What is the value of imprecise information given to an algorithm?" To give a particular direction to our query, we answer a question posed by Marc van Kreveld: Suppose that we are given a set of n disjoint unit disks, which represent imprecise information about the coordinates of corresponding points. Can we preprocess these disks so that if we are given m point sets $P_1 \dots P_m$, with each P_i consisting of exactly one point from each disk, then we can compute their m Delaunay triangulations in $o(mn \log n)$ total time? We show that after $O(n \log n)$ time processing the disks using $O(n)$ memory, one can compute each Delaunay triangulation in $O(n)$ time. And once the Delaunay triangulation is obtained, one can compute other structures from it, including the convex hull, the Gabriel graph, or a Euclidean minimum spanning tree.

Our solution actually uses Gabriel graphs and Euclidean minimum spanning trees for the disk centers to allow us to compute, in linear time, a connected subset of Delaunay edges for the specified points, from which the Delaunay computation can be completed by the algorithm of Chin and Wang [10]. Unfortunately for our algorithm's practicality, this last step involves rather heavy machinery. Some of our worst-case constants are over 100, meaning that our result is primarily theoretical, but it does demonstrate that an algorithm can benefit from imprecise information about the location of points.

2. RELATED WORK

Problems of exact computation with imprecise geometric operations or data are being attacked from several directions in computational geometry, often with notable success. To set our work in context, we briefly survey imprecise geometric computations, and remind the reader of several standard graphs and computations involving points and disks. We focus on geometric computation, even though ad-hoc and wireless networking applications have stimulated renewed interest in graphs defined by disks, as models of broadcast reachability or of uncertainty [6].

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

SCG'08, June 9–11, 2008, College Park, Maryland, USA.
Copyright 2008 ACM 978-1-60558-071-5/08/06 ...\$5.00.

2.1 Geometric computation on imprecise points

Two different but related issues have dominated the research in robust geometric algorithms [31]: closing the gap between between the precise mathematics of Euclidean geometry and the inexact primitives offered by the computer, and handling degenerate cases.

Both are important to the practical application of the theory of computational geometry: advances in exact geometric computing (including interval arithmetic, floating point filters, lazy evaluation, root bounds, and real number types) make correct practical code possible for software like Triangle [25] and libraries like LEDA [23] and CGAL [7]; techniques like simulation of simplicity [11] make correct handling of special cases easier. Both, however, assume that the input is exact—even when it is acknowledged that input coordinates represented in floating point are imprecise, it is assumed that the result of predicate on that input is desired, and that degeneracies can be correctly detected and need only be consistently handled.

One way to think about imprecise input is to say that the predicates may return incorrect responses. An early model of inexact predicates is *epsilon geometry* [17, 18]: a predicate on a tuple of points would return *true* or *false* if every tuple within a specified ε either was true or was false; it would return *unknown* if both true and false values could be found within ε of the given tuple. This models the uncertainty of a computation by saying that each point lies in an imprecision region—a disk of radius ε centered at the input coordinates.

For some computations, epsilon geometry can bound the accuracy of the output as a function of the ε bound on the input. To obtain such results, it is usually necessary to impose a restriction that the points are separated by 2ε (i.e. the regions are disjoint) or that at most a constant number of imprecise regions contain any point in the plane, because otherwise the information in the input may not constrain the order type of points chosen in the imprecise regions.

Various regions have been used to bound the imprecision for some questions in pattern matching [15], but in these cases the output is a simple pairing of points, and less geometric. There are many examples of using hierarchical structures (quadtrees or octrees, for example) that approximate objects to calculate simple combinatorial or metric properties, such as intersection or distance [5]. Van Kreveld and Löffler [22, 29] consider a variety of problems such as determining the largest and smallest convex hulls possible given regions that contain the imprecise points—note that convex hull of the regions is typically larger than the largest hull that can be obtained by selecting one point from each region.

For the *Voronoi diagram*, which is the decomposition of the plane by finite number of *sites* induced by labeling each point in the plane by its set of closest sites, and its dual *Delaunay triangulation*, Fortune analyzed the numerical precision of the predicates [13], and pointed out that *geometric rounding*—rounding the output back down to the precision of the input—is an important step in geometric algorithms that is often not explicitly considered. Sugihara and Iri [28, 27] advocated designing algorithms to guarantee topological properties even if the primitives are faulty.

Abellanas *et al.* [1] and Weller [30] have considered the smallest perturbation of sites that can change the combinatorial structure of Delaunay or Voronoi diagrams. Bandy-

opadhyay and Snoeyink [4, 3] compute the set of “almost-Delaunay simplices,” which are the tuples of points that could define a Delaunay simplex if the entire point set is perturbed by at most $\varepsilon > 0$, is again a simple, local output structure, albeit one that is fitting in a protein analysis application that depends upon identifying potential neighboring atoms as coordinates are perturbed. The algorithms to identify almost-Delaunay simplices were relatively brute-force.

Ely and Leclerc [12] and Khanban and Edalat [20] consider the epsilon geometry versions of the In-Circle predicate for Delaunay triangulation with imprecise points modeled as disks or rectangles, respectively. Khanban and co-authors [19, 21] developed a theory for returning partial Delaunay or Voronoi diagrams, consisting of the portion of the diagram that is certain.

Van Kreveld’s question was motivated by the desire to statistically sample the possible triangulations given n regions that model imprecision. Our aim in solving this question is not to compute a partial Delaunay diagram, but to compute enough structure that we can recover a connected set of Delaunay edges for a given sample, then complete the Delaunay triangulation in linear time. To explain further, we need some more definitions.

2.2 Disks, graphs, and algorithms

We remind the reader of some standard graphs defined by finite sets of points and disks, and the geometric algorithms to compute them. We also define some notation and observe properties that we will use in subsequent sections. We assume general position for all point sets in this abstract; namely, we assume that no four points are co-circular. This assumption can be removed by symbolic perturbation if desired [11].

We follow the idea of epsilon geometry, and model input points as unit-radius disks: Let \mathcal{R} be a set of n disjoint open unit disks in the plane, and let $P = \{p_1, p_2, \dots, p_n\}$ be their center points. An exact sample for \mathcal{R} is a set of points $\hat{P} = \{\hat{p}_1, \hat{p}_2, \dots, \hat{p}_n\}$ drawn one from each disk: i.e., for all $1 \leq i \leq n$, the length $|p_i \hat{p}_i| < 1$.

The Delaunay triangulation of P can be defined directly as the graph in which an edge joins two sites $p, q \in P$ if and only if there exists a circumcircle for edge pq that has all other sites of P outside.

In general, computing the Delaunay triangulation of n points requires $\Theta(n \log n)$ time. This lower bound implies that we cannot completely eliminate the disjointness condition and allow all disks to have a common intersection. The lower bound holds even if the points are sorted along x and y coordinated directions [24]; therefore we also cannot hope to do anything for general convex regions, since for a set of vertical lines we would not know anything more than the sorted order.

Aggarwal *et al.* [2] gave a clever linear-time algorithm to compute the Voronoi diagram or Delaunay triangulation of points in convex position in the plane. Chin and Wang extended this to compute the constrained Delaunay triangulation of a simple polygon [10]. (See [9] for an exposition of similar ideas applied to compute the medial axis.) Rather than define the constrained Delaunay triangulation here, we simply note that if all edges of the simple polygon satisfy the Delaunay empty circle criterion, then the constrained Delaunay is the Delaunay. Chin and Wang’s algorithm does re-




Figure 1: (a) Expanded Gabriel circle C_{pq}^+ contains centers of disks that can prevent $\hat{p}\hat{q}$ from being Delaunay in an exact sample \hat{P} . (b) Empty lune for EMST edge.

quire that the polygon is decomposed into trapezoids, which can theoretically be done in linear time by Chazelle’s algorithm [8].

Gabriel and Sokal [14] defined the Gabriel graph for points sites in a similar manner. First, for two sites p and q , let C_{pq} denote the circle with diameter pq . Sites p and q are joined by edge pq if and only if the circumscribing circle C_{pq} has all other sites outside. It is well known, and obvious from this definition, that the Gabriel graph is a subgraph of the Delaunay triangulation.

Any Euclidean minimum spanning tree (EMST) for P is a subgraph of the Gabriel graph of P . This fact is also well-known, and easy to observe: removing a tree edge uv partitions the EMST into two connected components; no vertex in the component of u can lie strictly inside the circle of radius $|uv|$ around v , and vice versa, so the interior of the lune that is the intersection of both circles is empty. This lune contains C_{uv} except for u and v , so all other points must be outside C_{uv} . One consequence is that any EMST has maximum vertex degree 6.

3. EXPANDED GABRIEL CIRCLES AND EMST EDGES

Define the expanded Gabriel circle, C_{pq}^+ , as the circle with center $(p+q)/2$ and radius $|pq|/2+2$. The expanded Gabriel circle contains the centers of disks that could, in an exact sample, prevent $\hat{p}\hat{q}$ from being a Delaunay edge.

OBSERVATION 1. *For disk centers $p, q \in P$, if no point $r \in P$ lies in the expanded Gabriel circle C_{pq}^+ , then in any exact sample \hat{P} , the edge $\hat{p}\hat{q}$ is Delaunay in \hat{P} .*

PROOF. Consider the smallest disk D enclosing the unit disks centered at p and q ; specifically, the disk D centered at the midpoint $(p+q)/2$ with radius $|pq|/2+1$, whose boundary is drawn dotted in Figure 1(a). There is a circle inside D that has the samples \hat{p} and \hat{q} on its boundary: shrink D about its center until the first point, say \hat{p} , lies on the boundary, then continue to shrink about \hat{p} until \hat{q} is also on the boundary. An exact sample \hat{r} can lie inside D only if the corresponding unit disk center satisfies $r \in C_{pq}^+$. \square

Let $T = (P, E)$ be the Euclidean minimum spanning tree (EMST) of P .

We will now show that each point in the plane (and therefore also the sites of P) can lie in at most a constant number

of the expanded Gabriel circles defined by the edges in E . We use this in later sections to bound the amount of repair work necessary to find a spanning tree of Delaunay edges for a particular sample from the unit disks centered at P .

We do an initial partitioning of spanning tree edges into *long* and *short*, depending on whether an edge’s length is greater than, or at most, $L = 2 + 2\sqrt{3} \approx 5.464$. This threshold value is chosen so that we can identify the connected components of the EMST when a long edge is removed.

LEMMA 2. *Let uv be a long edge of the EMST of P . Any point $w \in P \cap C_{uv}^+$ for which $|uw| \leq |vw|$ satisfies $|uw| < |uw| \leq |vw|$, and, if uv is removed from the EMST, w remains connected to the closer endpoint, u .*

PROOF. Recall that when uv is an edge of the Euclidean minimum spanning tree, the lune that is the intersection of the circles of radius $|uv|$ centered at u and at v has no sites in its interior. When $|uv| > L$, this lune pokes outside the expanded Gabriel circle C_{uv}^+ . So, no portion of the perpendicular bisector of uv lies outside the lune but inside C_{uv}^+ ; see Figure 1(b). Thus, we can partition $P \cap C_{uv}^+$ into the sets U and V , closer to u and v , with no ambiguity.

The distance from u to $w \in U$ is maximized if w is at the intersection of the lune with C_{uv}^+ . If we let $\ell = |uv|/2$, then because $\ell > L/2$ we know that $\ell + 2 < \ell\sqrt{3}$, and the triangle uvw cannot be equilateral, but must have $|uw| < |uv|$. Now, uv was chosen as a spanning tree edge, which implies that edge uw was unavailable, so w must have been in the component of u . \square

For a given point p in the plane (possibly a site from P), let E_p denote the set of edges of the EMST that have P in their expanded Gabriel circle, that is, $E_p = \{uv \in E \mid p \in C_{uv}^+\}$. We partition E_p into two groups: the *near* edges, for which both endpoints are at most $L + 2$ away from p , and the *far* edges, for which at least one endpoint is $L + 2$ or more away from p . Note that the far edges must necessarily all be long, and the near edges can be either short or long. We separately bound the sets of near edges and far edges for p .

An easy packing argument bounds the set of near edges for p , which includes all short EMST edges.

LEMMA 3. *For any point p in the plane E_p contains at most 70 near edges; i.e., p is in at most 70 expanded Gabriel circles for edges of the EMST of P that have both endpoints entirely within $L + 2$ of p .*

PROOF. If a center from P is within $L + 2$ of p , the corresponding disk from \mathcal{R} is within $L + 3$. At most $\lfloor (L + 3)^2 \rfloor = 71$ unit disks from \mathcal{R} can fit into this area, inducing at most 70 edges of the minimum spanning tree. \square

The constant of 70 is rather pessimistic. The best penny packing known for a circle of radius $L + 3$ has only 57 disks [16, 26], and even then it seems hard to draw many spanning tree edges between them that actually have p in their expanded Gabriel circle.

An angle packing argument in the next lemma shows that an input point $p \in P$ has few far edges.

LEMMA 4. *For any point $p \in P$, E_p contains at most 8 far edges.*

PROOF. We consider far edges $F \subset E_p$ in order of decreasing length, removing them from the EMST of P , and




Figure 2: (a) EMST edge uv is longer than pu' . (b) The region inside circles A and B , and outside C , is free from other far vertices of F .

keeping track of the connected component containing p . We assume, without loss of generality, that each far edge is labeled so that the first endpoint is the closer to p ; e.g., for uv , we have $|pu| \leq |pv|$.

Let T be the current EMST component, which partitions into $\{T_u, uv, T_v\}$ by removing uv , the longest edge of $F \cap T$. By Lemma 2, we know that p remains in the component of u , namely T_u , and that $|pu| < |uv| \leq |pv|$. We claim that all other edges of F in T are found in T_u : consider another edge $u'v' \in F \cap T$, as illustrated in Figure 2(a). Since $u'v'$ is long, Lemma 2 gives $|pu'| < |u'v'| \leq |pv'|$, and ordering by length gives $|u'v'| \leq |uv|$. But uv was chosen as the EMST edge joining T_u and T_v , and the shorter edge pu' was not; therefore u' must be in T_u with p .

Next, we show that v indicates a sector of the plane as seen from p that contains no other second endpoints of edges of F – no other far vertices of F . By definition, we know that the circle A of radius $L + 2$ around p contains no such vertices, and by the previous paragraph we know that the circle B of radius $|uv|$ around v contains no such vertices. Now consider the farthest vertex v' among all vertices in F , so all remaining vertices are inside a circle C of radius $|pv'|$ around p . This vertex must be part of an edge $u'v'$ of length at least $|pv'| - 2$, otherwise it would not be in E_p . Therefore, also $|uv| \geq |u'v'| \geq |pv'| - 2$. Now, all remaining far vertices of F must be in the region $C \setminus (A \cup B)$, see Figure 2(b).

To define the free sector, consider now the angle that pv makes with the intersections between A and B , and the angle it makes with the intersections between B and C . The smaller of those two angles bounds the sector.

Thus, we consider triangles of side lengths $L + 2$, $|pv|$, and $|uv|$ and of $|pv|$, $|pv|$, and $|uv|$. We know that $|pv| < |uv| + 2$ and $|uv| > L$. The angle at p is minimized as $|pv|$ approaches $L + 2$, which would give, in both cases, the isosceles triangle with angle

$$2 \arcsin\left(\frac{L/2}{L+2}\right) > 0.7494 > 2\pi/9.$$

Thus, inside the empty circle around v we find two sectors of angle $> \pi/4$ on either side of \overline{pv} that contain no far points closer to p than v . At most two empty sectors can overlap—one from the clockwise (CW) and one from the counter-clockwise (CCW) direction around p , which implies that there are at most 8 far edges. \square

We can summarize:




Figure 3: (a) The point in the plane p and the furthest endpoint v ensure that all endpoints of F_u lie within the shaded area. (b) The point u can be closer to p , but there is a circle with radius $|uv|$ around it that contains no endpoints of F_v .

THEOREM 5. *Let $T = (P, E)$ be the Euclidean minimum spanning tree on the points P . The total number of these points in the expanded circles for all edges is linear in n . That is,*

$$\sum_{uv \in E} |C_{uv}^+ \cap P| = O(n).$$

We can extend the proof of Lemma 4 to bound the number of far edges for an arbitrary point p in the plane, albeit with a large (and overly-pessimistic) constant factor. Since this bound is used only to shorten the description of preprocessing, and not for the algorithm itself, we have not tried to minimize the constant. This lemma implies that the arrangement of all expanded Gabriel circles has linear complexity.

LEMMA 6. *For any point p in the plane $|E_p|$ is constant.*

PROOF. The disk packing argument in Lemma 3 shows that there are at most 71 disk centers within distance $L + 2$ of any point p . As these are vertices in a Euclidean minimum spanning tree (EMST), for which each vertex has degree at most 6, at most 426 edges of E_p can have a vertex within $L + 2$ of p .

We therefore consider only the subset $F \in E_p$ of far edges for p that have both endpoints farther than $L + 2$ from p . We show that the edges of F can be organized into a binary tree whose maximum depth is 8 by the angle packing argument used in Lemma 4. Since such a binary tree has at most $2^9 - 1 = 511$ nodes, F has at most $511 + 426 = 937$ edges.

We build this tree from the root, at depth 0, on down. Each node ν is associated with a subset of edges, $F_\nu \subset F$, as well as an edge of F_ν . The root is associated with F , and an edge uv having one endpoint v farthest from p . Removing uv from the EMST partitions the remaining edges of F into two groups, F_u and F_v , where the first remain connected to u and the second connected to v by Lemma 2. These are the edge sets associated with the children of the root. (In determining connectedness, we include EMST edges and vertices within $L + 2$ of p , even though they are not in $F \setminus \{uv\}$.)

In general, at node ν , the edges F_ν are edges of a connected component of the EMST minus the edges associated with the ancestors of ν , the associated edge $uv \in F_\nu$ is chosen so that the endpoint v is farthest from p , and removing uv partitions the edges of the EMST component into F_u and F_v .




Figure 4: (a) A set of imprecise points. (b) The minimum spanning tree with the expanded Gabriel circles and the disks they intersect associated to its edges.

For the edges in F_u , we know that no endpoint can lie within a circle of radius $|uv|$ centered at v . We also know that all endpoints lie within a circle of radius $|pv|$ centered at p , and that none lie within $L + 2$ of p . These constraints on F_u are depicted in Figure 3(a). As in the proof of Lemma 4, there is a sector with angle greater than $2\pi/9$, as seen from p , that contains no endpoints from F_u .

Even though the point u can lie closer to p , as in Figure 3(b), the same constraints apply, showing that a sector of the same angle is empty of endpoints from F_v . This implies that the tree has depth at most 8, and completes the proof of the lemma. \square

4. DELAUNAY COMPUTATION

Let \mathcal{R} be a set of n disjoint unit-radius disks in the plane that represent the imprecise regions for P , which are the disk center points. Subsection 4.1 details how to preprocess \mathcal{R} in $O(n \log n)$ time into a linear-size data structure $H(\mathcal{R})$. Subsection 4.2 shows that given an exact sample \hat{P} consisting of a point inside each disk of \mathcal{R} , we can compute the Delaunay triangulation of \hat{P} in linear time using $H(\mathcal{R})$.

4.1 Preprocessing

Let P be the set of center points of the n disjoint unit disks of \mathcal{R} . For $H(\mathcal{R})$, we compute a Euclidean minimum spanning tree of P , a list of its edges sorted by length, and for each edge uv the list of points of P that fall inside the expanded Gabriel circle C_{uv}^+ . Figure 4 shows an example.

By Theorem 5 we know that each point of P can fall into at most a constant number of expanded Gabriel circles. Thus, the total size of $H(\mathcal{R})$ is linear.

A minimum spanning tree is easy to compute in $O(n \log n)$ time, since the Delaunay triangulation is a linear-size set of edges that contains all candidates. The sorted list of EMST edges is even easier. Finally, a simple sweep of the arrangement of the unit circles of \mathcal{R} and the points M can locate all points in their circles; because Lemma 6 says that this arrangement has linear size, the sweep can be carried out in $O(n \log n)$ time.

LEMMA 7. *Preprocessing the n disjoint unit disks \mathcal{R} produces a linear size data structure $H(\mathcal{R})$ in $O(n \log n)$ time.*

Denote the sorted list of EMST edges by e_1, \dots, e_{n-1} . We define notation for the connected components of the graph consisting of the first k edges of this list: Let \mathcal{I}_k be the partition of the index set $\{1, \dots, n\}$ induced by the connected

components of these first k edges: that is, $i, j \in I$ for some $I \in \mathcal{I}_k$ if and only if m_i and m_j can be joined by edges from $\{e_1, \dots, e_k\}$. We can associate these connected components with $H(\mathcal{R})$ (conceptually, not computationally, as they are needed only for a proof), because our algorithm creates the components (or supersets of them) for points $\hat{P} = \{\hat{p}_1, \dots, \hat{p}_n\}$ drawn from each disk in \mathcal{R} .

4.2 Computing the Delaunay triangulation

Now, given an exact sample $\hat{P} = \{\hat{p}_1, \dots, \hat{p}_n\}$ of \mathcal{R} , and the data structure $H(\mathcal{R})$, we show how to compute in linear time a connected subgraph of the Delaunay triangulation of \hat{P} . Chin and Wang's algorithm [10] then completes the Delaunay triangulation of \hat{P} in linear time.

We form our connected subgraph by finding paths in the Delaunay that make the connected components that are formed as we add edges to the Euclidean minimum spanning tree. (Recall that EMST edges are ordered by increasing length.) We begin by making an observation, illustrated in Figure 5(a), on the portion of a Delaunay triangulation bounded by a circle.

LEMMA 8. *Let P be a set of points in general position in the plane, C be a circle whose interior contains a subset $Q = P \cap \text{int}(C)$, and E be the set of Delaunay edges of P that have empty circles contained inside $\text{int}(C)$. The graph (Q, E) is connected.*

PROOF. Let c be the point of Q closest to the center of C ; we show that any vertex $p \in Q$ is connected to c . Initially, let $a = p$, and, as depicted in Figure 5(b), grow a circle from a towards the center of C , keeping a pinned on the boundary; stop when the circle hits any point $b \in Q$. The edge ab is discovered to be a Delaunay edge in E , and the point b is closer to the center of C than a was. Since P is finite, by setting $a = b$ and repeating this procedure, we eventually construct a path from p to c in the graph (Q, E) . \square

Suppose now that EMST edge e_k joins $u, v \in P$, and consider the expanded Gabriel circle C_{uv}^+ . Lemma 8 says there exists a path of Delaunay edges certified inside C_{uv}^+ that joins the corresponding exact samples $\hat{u}, \hat{v} \in \hat{P}$; our task is to compute one efficiently, or at least to compute a subgraph of the Delaunay triangulation of \hat{P} that contains one or more paths.

THEOREM 9. *Given n disjoint unit disks \mathcal{R} , and the structure $H(\mathcal{R})$ from Lemma 7, the Delaunay triangulation of an*




Figure 5: (a) The Delaunay edges certified by (dotted) empty circles within a bigger circle form a connected graph. (b) Growing a circle from p towards the center. (c) The closest point to the center can be connected to at least one of the points in the other group.

exact sample \hat{P} chosen from these disks can be computed in $O(n)$ time.

PROOF. To reconstruct the Delaunay triangulation, we first want to build up the components of the EMST by adding edges in order; the essential task is to find a path of Delaunay edges joining the exact samples $\hat{u}, \hat{v} \in \hat{P}$ for two centers $u, v \in P$ that form an edge e_k in the EMST. We will do this in C_{uv}^+ , although we could do it in the smaller $C_{\hat{u}\hat{v}}$ with a slightly longer description of the procedure.

Let $Q \subset \hat{P}$ denote the points inside circle C_{uv}^+ and $\mathcal{K} \subset \mathcal{R}$ denote the unit disks with centers inside. When e_k is short, penny packing says there are at most a constant number of disks in \mathcal{K} , so we can process e_k by computing the Delaunay triangulation of the points Q and discarding edges that are not certified by an empty circle inside C_{uv}^+ .

When e_k is long, Lemma 2 says that there are two components that are separated by the perpendicular bisector of uv . Let Q_u and Q_v be the partition of Q by this bisector; it suffices to find a Delaunay edge of \hat{P} from $Q_u \times Q_v$ since the points within Q_u (and Q_v) have already been connected earlier in the algorithm.

Let $c_u \in Q_u$ and $c_v \in Q_v$ be the closest points to the center of C_{uv}^+ , as illustrated in Figure 5(c), and assume that the distance to c_v is greater, meaning the circle concentric with C_{uv}^+ through c_v contains at least one point of Q_u . Shrink this circle with c_v on the boundary by moving its center toward c_v until the last point of Q_u leaves its interior—this point defines the desired Delaunay edge with c_v . Both steps can be carried out in time proportional to $|Q|$.

We spend constant time with each short edge and, by Lemma 5, a total of linear time with the long edges. For each edge we find a path of Delaunay edges of \hat{P} that joins the vertices \hat{u} and \hat{v} , so the connected components induced by the sequence of edges found will be supersets of the components of \mathcal{I}_k of the first k edges of the EMST of \hat{P} . Thus, we obtain a connected graph after processing all EMST edges, and can invoke Chin and Wang [10] to complete the Delaunay triangulation. \square

5. EXTENSIONS

Our algorithm works for a very specific class of imprecise regions: disjoint disks of equal radius. In practice, this may be a rather strong assumption. In this section, we show how to extend the result to less restricted regions.

5.1 Overlapping disks

If we allow the regions to be arbitrarily overlapping disks, then there is little we can hope to prove. In the worst case, all disks could coincide, allowing the constructions that establish the $\Omega(n \log n)$ lower bounds for general Delaunay triangulation [24]. If we limit the depth of overlap, however, our result still holds with the algorithm unchanged.

We say a set of disks is k -overlapping if no point in the plane is contained in more than k disks. In this case, the number of short edges that can contain a point p increases. Clearly, there cannot be more than $k(r+2)^2$ disks touching a circle of radius r . This means the constant grows linearly in k . The arguments involving long edges do not depend on the disjointness of the disks.

5.2 Other extensions

If we allow the disks to have different radii, then in general the problem is open. However, when there is a constant fraction $c = \frac{R}{r}$ between the largest radius R and the smallest radius r , then we can just increase radii until all disks have radius R . Since we know that the sample points lie inside the input disks, they certainly also lie in the grown disks. Of course the disks start overlapping, but not too much: at most $(c+1)^2$ grown disks contain any given point in the plane.

If the input regions are not disks but squares, then we can grow them to the smallest disks containing them, which are 3-overlapping. If the regions are *fat* in the sense that they contain circles of radius r but are contained in circles of radius R (the same radii for all regions), with $c = \frac{R}{r}$, then we can again replace them by disks of radius R that are at most $(c+1)^2$ -overlapping.

Finally, we can also handle combinations of the above (partially overlapping fat regions of restricted different radii) at the expense of an increased constant in the time bound.

6. CONCLUSION AND OPEN PROBLEMS

Our result proves that imprecise information about point coordinates has value: after $O(n \log n)$ time spent preprocessing the n regions of imprecision for the points, we can obtain a Delaunay triangulation of exact points sampled from the imprecise regions in linear time.

We would like to see if we can use the special properties of our points to avoid the several levels of decomposition employed by Chin and Wang’s constrained Delaunay trian-

gulation algorithm, and make this result more practical. We have considered only a couple of the many possible extensions to other models of imprecise points.

Acknowledgments

We thank Marc van Kreveld for the question and discussions and for arranging that the first author could visit UNC Chapel Hill. This research was partially funded by grants from NSF and NGA/Darpa, and by the Netherlands Organization for Scientific Research (NWO) through the project GOGO.

7. REFERENCES

- [1] M. Abellanas, F. Hurtado, and P. A. Ramos. Structural tolerance and Delaunay triangulation. *Inf. Process. Lett.*, 71(5-6):221–227, 1999.
- [2] A. Aggarwal, L. J. Guibas, J. Saxe, and P. W. Shor. A linear-time algorithm for computing the Voronoi diagram of a convex polygon. *Discrete Comput. Geom.*, 4(6):591–604, 1989.
- [3] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: Nearest neighbor relations for imprecise points. In *ACM-SIAM Symp on Discrete Algorithms*, pages 403–412, 2004.
- [4] D. Bandyopadhyay and J. Snoeyink. Almost-Delaunay simplices: Nearest neighbor relations for imprecise 3D points using CGAL. *Computational Geometry: Theory and Applications*, 38(1-2):4–15, 2007.
- [5] K. Bühler, E. Dyllong, and W. Luther. Reliable distance and intersection computation using finite precision geometry. In *Numerical Software with Result Verification*, number 2991 in LNCS, pages 160–190. Springer Verlag, 2004.
- [6] J. Cartigny, F. Ingelrest, D. Simplot-Ryl, and I. Stojmenovic. Localized LMST and RNG based minimum-energy broadcast protocols in ad hoc networks. *Ad Hoc Networks*, 3(1):1–16, 2005.
- [7] CGAL, Computational Geometry Algorithms Library. <http://www.cgal.org>.
- [8] B. Chazelle. Triangulating a simple polygon in linear time. *Discrete Comput. Geom.*, 6(5):485–524, 1991.
- [9] F. Chin, J. Snoeyink, and C. A. Wang. Finding the medial axis of a simple polygon in linear time. *Discrete Comput. Geom.*, 21(3):405–420, 1999.
- [10] F. Y. L. Chin and C. A. Wang. Finding the constrained Delaunay triangulation and constrained Voronoi diagram of a simple polygon in linear time. *SIAM J. Comput.*, 28(2):471–486, 1998.
- [11] H. Edelsbrunner and E. P. Mücke. Simulation of simplicity: A technique to cope with degenerate cases in geometric algorithms. *ACM Trans. Graph.*, 9(1):66–104, 1990.
- [12] J. S. Ely and A. P. Leclerc. Correct Delaunay triangulation in the presence of inexact inputs and arithmetic. *Reliable Computing*, 6:23–38, 2000.
- [13] S. Fortune. Numerical stability of algorithms for 2-d Delaunay triangulations. *Internat. J. Comput. Geom. Appl.*, 5(1-2):193–213, 1995.
- [14] K. R. Gabriel and R. R. Sokal. A new statistical approach to geographic variation analysis. *Systematic Zoology*, 18:259–278, 1969.
- [15] M. T. Goodrich, J. S. B. Mitchell, and M. W. Orletsky. Practical methods for approximate geometric pattern matching under rigid motion. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 103–112, 1994.
- [16] R. Graham, B. Lubachevsky, K. Nurmela, and P. Östergård. Dense packings of congruent circles in a circle. *Disc. Math.*, 181:139–154, 1998.
- [17] L. J. Guibas, D. Salesin, and J. Stolfi. Epsilon geometry: building robust algorithms from imprecise computations. In *Proc. 5th Annu. ACM Sympos. Comput. Geom.*, pages 208–217, 1989.
- [18] L. J. Guibas, D. Salesin, and J. Stolfi. Constructing strongly convex approximate hulls with inaccurate primitives. *Algorithmica*, 9:534–560, 1993.
- [19] A. A. Khanban. *Basic Algorithms of Computational Geometry with Imprecise Input*. PhD thesis, Imperial College, London, 2005.
- [20] A. A. Khanban and A. Edalat. Computing Delaunay triangulation with imprecise input data. In *Proc. 15th Canad. Conf. Comput. Geom.*, pages 94–97, 2003.
- [21] A. A. Khanban, A. Edalat, and A. Lieutier. Computability of partial Delaunay triangulation and Voronoi diagram. In V. Brattka, M. Schröder, and K. Weihrauch, editors, *Electronic Notes in Theoretical Computer Science*, volume 66. Elsevier, 2002.
- [22] M. Löffler and M. van Kreveld. Largest and smallest convex hulls for imprecise points. *Algorithmica*, 2008 (to appear).
- [23] K. Mehlhorn and S. Näher. *LEDA: A Platform for Combinatorial and Geometric Computing*. Cambridge University Press, Cambridge, UK, 2000.
- [24] R. Seidel. A method for proving lower bounds for certain geometric problems. In G. T. Toussaint, editor, *Computational Geometry*, pages 319–334. North-Holland, Amsterdam, Netherlands, 1985.
- [25] J. R. Shewchuk. Delaunay refinement algorithms for triangular mesh generation. *Computational Geometry: Theory and Applications*, 22(1-3):21–74, 2002.
- [26] E. Specht. The best known packings of equal circles in the unit circle (up to $n = 500$). <http://hydra.nat.uni-magdeburg.de/packing/cci/cci.html>, July 2007.
- [27] K. Sugihara and M. Iri. Two design principles of geometric algorithms in finite-precision arithmetic. *Appl. Math. Lett.*, 2(2):203–206, 1989.
- [28] K. Sugihara and M. Iri. Construction of the Voronoi diagram for ‘one million’ generators in single-precision arithmetic. *Proc. IEEE*, 80(9):1471–1484, Sept. 1992.
- [29] M. van Kreveld and M. Löffler. Largest bounding box, smallest diameter, and related problems on imprecise points. In *Proc. 10th Workshop on Algorithms and Data Structures*, LNCS 4619, pages 447–458, 2007.
- [30] F. Weller. Stability of Voronoi neighborhood under perturbations of the sites. In *Proc. 9th Canad. Conf. Comput. Geom.*, pages 251–256, 1997.
- [31] C. Yap. Robust geometric computation. In J. E. Goodman and J. O’Rourke, editors, *Handbook of Discrete and Computational Geometry, Second Edition*, pages 927–952. CRC Press, 2004.