

WHICH TRIANGULATIONS APPROXIMATE THE COMPLETE GRAPH?

(Conference Abstract)¹

GAUTAM DAS – University of Wisconsin

DEBORAH JOSEPH – University of Wisconsin

Abstract. Chew and Dobkin et. al. have shown that the *Delaunay triangulation* and its variants are sparse approximations of the complete graph, in that the shortest distance between two sites within the triangulation is bounded by a constant multiple of their Euclidean separation. In this paper, we show that other classical triangulation algorithms, such as the *greedy triangulation*, and more notably, the *minimum weight triangulation*, also approximate the complete graph in this sense. We also design an algorithm for constructing extremely sparse (nontriangular) planar graphs that approximate the complete graph.

We define a sufficiency condition and show that *any* Euclidean planar graph constructing algorithm which satisfies this condition always produces good approximations of the complete graph. This condition is quite general because it is satisfied by all the triangulation algorithms mentioned above, and probably by many other graph algorithms as well. We thus partially answer the question posed by the title.

From a theoretical standpoint, our results are interesting because we prove non-trivial properties of minimum weight triangulations, of which little is currently known. From a practical standpoint, the graph algorithms we consider are good alternatives to the Delaunay triangulation, particularly when designing a sparse network under severe constraints on the total edge length. Finally, our general approach may help in identifying or designing other algorithms for constructing sparse networks.

1. INTRODUCTION

The *complete graph* represents an ideal communication network between n sites. However, to conserve resources, sparse networks are often designed which approximate the complete graph in some sense. This problem has also been studied from a geometric context, where the graphs are Euclidean. In particular, recent research [C, DFS] has shown that the *Delaunay triangulation* and its variants are good approximations

¹ This work was supported in part by the National Science Foundation under grant DCR-8402375. The authors' address is: Computer Sciences Department, University of Wisconsin, 1210 West Dayton St., Madison, WI 53706, U.S.A.

of the complete graph, in that the shortest distance between any two sites within the triangulation is bounded by a constant multiple of their Euclidean separation.

In this paper, we show that other classical triangulation algorithms, such as the *greedy triangulation* [MZ], and more notably, the *minimum weight triangulation* [MZ], also produce good approximations of the complete graph in the above sense. We define a sufficiency condition, and show that *any* algorithm for constructing Euclidean planar graphs which satisfies this condition always produces good approximations of the complete graph. This condition is quite general; it is satisfied by all the three triangulation algorithms mentioned above, and by many other graph algorithms as well. We thus partially answer the question posed by the title. Finally, we apply this condition constructively to design an algorithm which generates extremely sparse (nontriangular) planar graphs. The length of the graph produced by this algorithm is proportional to the length of the *minimum spanning tree* between the sites.

From a theoretical standpoint, our results are interesting because we prove some nontrivial properties of minimum weight triangulations, which has always been an enigmatic construct in computational geometry. Other than the basic definition, very few properties are known about it, and its computational complexity is still unknown [L, K]. From a practical standpoint, the graph algorithms we consider are suitable alternatives to the Delaunay triangulation, particularly when designing a sparse network under severe constraints on the total edge length. For instance, an example in [MZ] shows that the ratio of the total edge length between a Delaunay and a greedy triangulation could be as large as $\Omega(n/\log n)$. Finally, our techniques may be useful in identifying or designing other algorithms that construct similar sparse graphs.

In this paper we will restrict our attention to Euclidean planar graphs, unless otherwise mentioned. Throughout we adopt the following definitions. If e is a line segment, $|e|$ denotes its length. If E is a set of line segments, $|E|$ denotes the total length of all segments, rather than the usual notion of cardinality. S is a collection of n points in the plane (called *sites*). Let $K(S)$ be the complete graph, containing all $n(n-1)/2$ straight edges that join the sites in S . A *triangulation* is a maximal nonintersecting subset of $K(S)$. A *graph algorithm* G takes input S , and produces an Euclidean planar graph $G(S)$. A *Delaunay triangulation*, $DT(S)$, is the straight edged Voronoi dual of S [DFS]. A *greedy triangulation*, $GT(S)$, is computed by a greedy algorithm which iteratively inserts the next shortest nonintersecting edge into a partially constructed triangulation [MZ]. A *minimum weight triangulation*, $MWT(S)$, is one with minimum total edge length [MZ]. $SP_{G(S)}(A, B)$ denotes the shortest path between sites A and B , composed of edges of $G(S)$. Let $measure(G(S)) = \max_{A, B \in S} \frac{|SP_{G(S)}(A, B)|}{|AB|}$. Intuitively, this quantity measures how badly a graph approximates the complete graph.

The algorithm G produces *good approximations* if there exists a constant c such that, for all S , $measure(G(S)) \leq c$.

In [DFS] it was shown that, for all S , $measure(DT(S)) \leq \pi(1 + \sqrt{5})/2$. Our methods are more general because we isolate a condition satisfied by many graph algorithms, including DT , GT , and MWT . (This is described in Section 3) First we show, in Section 2, that any graph algorithm satisfying this condition always produces good approximations of the complete graph. Finally in Section 4, we apply this condition in designing an algorithm for constructing general (nontriangular) graphs.

The reader should note that in this version of the paper, the various constants are roughly computed, and do not represent tight upper bounds. We feel that to establish tighter bounds, it is necessary to exploit properties unique to a particular graph algorithm. Our general approach thus helps in easily identifying algorithms that deserve further study.

2. GRAPHS WHICH APPROXIMATE THE COMPLETE GRAPH

We begin this section by defining the graph property that will be sufficient for our first result.

Definition. Let α be an angle less than $\pi/2$, d a constant, and G a graph algorithm. The condition $P_{\alpha,d}(G)$ is defined as follows.

$P_{\alpha,d}(G) = true$, if for every S , the following are true.

- i) *The diamond property.* For every edge e of $G(S)$, consider the two triangular regions defined on either side of e , such that e is the base of each triangle and the base angles are α (see Figure 1). Then at least one of the regions contains no other sites of S .
- ii) *The good polygon property.* For every face of $G(S)$, let A and B be any two sites along its boundary such that the line segment AB lies wholly within the face (such segments are called *chords*). Then the shortest distance from A to B around the boundary of the face is smaller than $d|AB|$.

Intuitively, this property is useful because (i) ensures that an edge cannot act as a long obstacle between a pair of sites and (ii) ensures that the distance “around the boundary” of a graph face can not be too much longer than the distance “across” the face. In fact a formalization of this proves our first result. We will spend the remainder of this section showing that any graph algorithm satisfying this property always produces good approximations of the complete graph. Then in the next section we will show that the Delaunay triangulation, the greedy triangulation and the minimum

weight triangulation satisfy the property. Finally, we will use this property to design an algorithm for constructing sparse nontriangular planar graphs.

For simplicity we assume throughout that π is a multiple of α .

Theorem 1. *For any angle $\alpha < \pi/2$ and any constant d , there exists a constant $c_{\alpha,d}$ such that, if $P_{\alpha,d}(G) = \text{true}$ for some graph algorithm G , then for all S , $\text{measure}(G(S)) \leq c_{\alpha,d}$.*

Proof. Our proof has some flavor of the techniques employed in [DFS]. The latter exploits specific properties of Delaunay triangulations, while ours is a more general approach and is thus more involved.

Let $G'(S)$ be the graph $G(S)$ augmented with all possible chords of faces. Clearly $G'(S)$ may not be planar. We introduce the notion of *pseudo paths*, which are paths in $G'(S)$. We also define *real paths*, which are paths in $G(S)$. The good polygon property ensures that if there exists a pseudo path of length l between sites A and B , then there exists a real path between them of length at most dl . Thus our efforts will be directed towards constructing a short pseudo path between any pair of sites. Henceforth, all paths will be pseudo, unless otherwise mentioned.

The diamond property ensures that every edge e of $G(S)$ is the base of at least one empty triangular region with base angles α . (If both regions are empty, select any one). Let $t(e)$ and $v(e)$ denote this triangle and its third vertex respectively. Note that this vertex need not be a site in S .

Let A and B be two sites such that AB is not in $G(S)$. The two half planes defined by extending AB will be referred to as *Top* and *Bottom* respectively. (Figure 2). We first construct a subgraph of $G'(S)$ within which a short pseudo path between A and B resides. Let e_1, e_2, \dots, e_m be the sequence of edges of $G(S)$ that intersect with AB . Let u_i and l_i refer to the upper and lower sites of e_i . (For simplicity we shall assume these edges do not share vertices). Clearly every adjacent pair of edges e_i and e_{i+1} in this sequence belong to the same face. Let UC_i be the upper convex chain from u_i to u_{i+1} , such that the boundary of the face from u_i to u_{i+1} is above it. The lower convex chain, LC_i , is similarly defined. Let $u_{i,1}l_{i,2}$ and $l_{i,1}u_{i,2}$ be the two common tangents between the two chains. Now consider the graph consisting of the edges of $G(S)$ that intersect AB , along with all upper and lower chains, and all common tangents. Clearly this graph (denoted as $G'(A, B)$) is a subgraph of $G'(S)$. We shall construct a path from A to B in this subgraph that is not too long with respect to AB . Sites in this subgraph may be classified into two groups, those that are end points of edges of $G(S)$ that intersect with AB (called *major sites*), and those that belong to upper and lower

chains (called *minor sites*). Each group may be further classified into upper and lower sites in a natural manner.

Any simple path in this subgraph from A to B consists of alternating portions of the upper and lower boundary, connected by tangents. We now construct a *direct path* of this sort as follows. Begin with A on the path. Suppose u_i (a major site) is the last site of the partially constructed path. If $v(u_{i+1}l_{i+1})$ is inside *Bottom* (lies below AB), then the path is extended to u_{i+1} via the upper chain, UC_i . If $v(u_{i+1}l_{i+1})$ is inside *Top* (lies above AB), then the path is extended to l_{i+1} as follows. The path is first extended via the upper chain to $u_{i,1}$, then along a common tangent to $l_{i,2}$, and eventually along the lower chain to l_{i+1} . If l_i is the last major site, the extension is similar. We denote such direct paths as $DP(A, B)$. The following claim illustrates an important property of direct paths.

Claim 1: Suppose $DP(A, B)$ is *one sided*, that is, the path consists entirely of the upper boundary or entirely of the lower boundary of $G'(A, B)$. Then $\frac{|DP(A, B)|}{|AB|} \leq b_\alpha$, where $b_\alpha = \frac{2\pi}{\alpha \sin \alpha/4}$.

Proof. Without loss of generality, assume the upper boundary is the direct path. Let $R = \text{Bottom} \cup t(e_1) \cup \dots \cup t(e_m)$. Let $R(A, B)$ denote the boundary of R from A to B . Since $DP(A, B)$ is composed of convex chains, $|DP(A, B)| \leq |R(A, B)|$.

We now show an upper bound for $|R(A, B)|$. Partition the triangles $t(e_1), \dots, t(e_m)$ into $2\pi/\alpha$ groups, $G_0, G_{\alpha/2}, \dots, G_{\pi-\alpha/2}$, where G_θ contains those triangles whose left sides make angles in the range $[\theta, \theta + \alpha/2]$ with AB . Let $R_\theta = \text{Bottom} \cup$ all triangles in G_θ . But $R = R_0 \cup \dots \cup R_{\pi-\alpha/2}$, and hence

$$|R(A, B)| \leq |R_0(A, B)| + \dots + |R_{\pi-\alpha/2}(A, B)|.$$

Consider Figure 3, which illustrates the group G_θ . Clearly $R_\theta(A, B)$ is inside ACB , and is composed of three types of line segments. Segments that go up/down/horizontal make angles in the range $[\theta, \theta + \alpha/2] / [\theta + \alpha, \theta + 3\alpha/2] / [0, 0]$ with AB respectively. It is then easy to see that $|R_\theta(A, B)| \leq |AC| + |BC| \leq \frac{|AB|}{\sin \alpha/4}$, for any θ . Thus $|R(A, B)| \leq \frac{2\pi|AB|}{\alpha \sin \alpha/4}$, and the claim is proved. ■

Since a direct path may not be one sided, we iterate the following modification exhaustively until we obtain a final path, $FP(A, B)$, which is almost one sided. Let u_i and u_j be two upper major sites on the direct path such that the major sites in between, l_{i+1}, \dots, l_{j-1} , are lower sites. Determine if a shortcut is necessary (this will be defined later). If not, retain the path, otherwise replace the portion of the path between $u_{i,1}$ and $u_{j-1,2}$ with the portion of the upper boundary of $G'(A, B)$ between them.

We now define when a shortcut is necessary. Recall that every edge e_k intersecting with AB is associated with an empty triangular region $t(e_k)$. We associate another

empty region with each such edge, denoted as $t'(e_k)$, and defined as follows. Suppose $v(e_k)$ is below AB (Figure 4). Then $t'(e_k)$ is the closed empty region between AB , the common tangents $l_{k-1,1}u_{k-1,2}$ and $u_{k,1}l_{k,2}$, and the path along upper chains from $u_{k-1,2}$ via u_k to $u_{k,1}$. If $v(e_k)$ is above AB , then $t'(e_k)$ is symmetrically defined.

Partition the set of intersecting edges into two sets L and U . The edge e belongs to U (respectively L) if $v(e)$ belongs to *Bottom* (respectively *Top*). Define $R_U = \text{Bottom} \cup_{e \in U} [t(e) \cup t'(e)]$. R_L is symmetrically defined. It is easy to see that $|R_U(A, B)|$ and $|R_L(A, B)|$ are each smaller than $b_\alpha |AB|$.

Let θ be the angle between the tangents $u_{i,1}l_{i,2}$ and $l_{j-1,1}u_{j-1,2}$, when extended. It does not matter whether the extended diagonals meet in *Top* or *Bottom*. Two cases arise.

Case 1: $\theta > \pi/3$. The shortcut is not necessary.

Case 2: $\theta \leq \pi/3$. Intuitively, this means the diagonals are almost parallel. Consider Figure 5. The shaded region is bounded by the tangents $u_{i,1}l_{i,2}$ and $l_{j-1,1}u_{j-1,2}$, $R_U(A, B)$, and $R_L(A, B)$, and is empty of other sites. Let the upper (lower) boundary of this region be T_1 (T_2). We set up co-ordinate axes as shown in the figure. E (F) is the point on T_1 (T_2) with the minimum (maximum) y co-ordinate, and $h = y(E) - y(F)$. PEQ and VFW are parallel to the x axis. Let $w = \max(|PQ|, |VW|)$. If $\frac{h}{w} \leq \frac{1}{2 \tan \alpha/2}$, the shortcut is not necessary, otherwise the shortcut is taken.

The following claims lead to the final stages of the proof of the theorem. We again remind the reader that in this version of the paper, the calculations are rough, and the resulting value of $c_{\alpha,d}$ is not a tight upper bound.

Claim 2: If the shortcut is taken, then the length of the new portion of the path, which is the upper boundary of $G'(A, B)$ from $u_{i,1}$ to $u_{j-1,2}$, is smaller than $b_\alpha(|T_1| + |T_2|)$.

Proof. Let the lower convex hull of the upper boundary of $G'(A, B)$ from $u_{i,1}$ to $u_{j-1,2}$ be $u_{i,1} = z_1, z_2, \dots, z_k = u_{j-1,2}$. Clearly $|z_1z_2| + |z_2z_3| + \dots + |z_{k-1}z_k|$ is smaller than $|T_1|$. We also observe the following. If H and I are any two points above T_1 , and J is any point below T_2 , such that HJ and IJ intersect the shaded region, then the angle HJI is smaller than α , because of the shortcut condition. Due to this, for $1 \leq q \leq k$, $DP(z_q, z_{q+1})$ is one sided (where $DP(z_q, z_{q+1})$ is defined within $G'(z_q, z_{q+1})$), and the concatenation of all these direct paths is the upper boundary of $G'(A, B)$ from $u_{i,1}$ to $u_{j-1,2}$. Thus the length of the new portion is smaller than $b_\alpha(|T_1|)$, which is smaller than $b_\alpha(|T_1| + |T_2|)$. ■

Claim 3: If the shortcut is not taken due to case 1, then the length of the old portion of the path between $u_{i,1}$ and $u_{j-1,2}$ is smaller than $3(|T_1| + |T_2|)$.

Proof. Clearly, the lower boundary of $G'(A, B)$ is smaller than $|T_2|$. We now compute the total length of the two tangents, $u_{i,1}l_{i,2}$, and $u_{j-1,2}l_{j-1,1}$. Let C be the meeting point of their extensions. Since $\theta > \pi/3$, if C is in *Top* (*Bottom*), their total length is smaller than $2|T_2|$ ($2|T_1|$). Thus their total length is always smaller than $2(|T_1| + |T_2|)$, and the claim is proved. ■

Claim 4: If the shortcut is not taken due to case 2, then the length of the old portion of the path between u_i and u_j is smaller than $a_\alpha(|T_1| + |T_2|)$, where $a_\alpha = \frac{5}{\sin \alpha/2}$.

Proof. As before, the lower boundary of $G'(A, B)$ is smaller than $|T_2|$. We now compute the total length of the two diagonals. Since $\theta \leq \pi/3$, we see that

$$|u_{i,1}P| + |l_{i,2}V| + |u_{j-1,2}Q| + |l_{j-1,1}W| \leq 2(|T_1| + |T_2|)$$

irrespective of where C lies. Also,

$$|PV| + |QW| \leq 4h \leq \frac{2}{\tan \alpha/2}(|T_1| + |T_2|).$$

The claim follows after some manipulations. ■

Thus, whichever case occurs, the length of $FP(A, B)$ from $u_{i,1}$ to $u_{j-1,2}$ is smaller than $[\max(b_\alpha, 3, a_\alpha)](|T_1| + |T_2|) = b_\alpha(|T_1| + |T_2|)$. Summing over all such modifications, we conclude that $|FP(A, B)| \leq 2b_\alpha^2|AB|$. Thus $c_{\alpha,d} = \frac{8\pi^2 d}{\alpha^2 \sin^2 \alpha/4}$, and the theorem is proved. ■

3. LAYOUT PROPERTIES OF CERTAIN TRIANGULATIONS

In this section we investigate the algorithms, Delaunay triangulations, greedy triangulations, and minimum weight triangulations, and show that all three satisfy the property of the previous section. Then in the next section, we consider more general graph algorithms.

Theorem 2. $P_{\pi/4,1}(DT) = \text{true}$.

Proof. Notice that condition (ii) of the property is trivially holds for any triangulation. As for condition (i), it is known that the circumscribing circle of any triangle in a Delaunay triangulation is empty of other sites. Thus every edge is the chord of two empty circles. If we select any one circle of any edge, it will wholly contain at least one of the triangular regions. Thus condition (i) holds. ■

The next theorem describes a similar result for greedy triangulations. We first present a lemma.

Lemma 1. *Let A and B be any two sites such that AB is not an edge in $GT(S)$. Let e be the shortest edge of $GT(S)$ that intersects AB . Then $|e| < |AB|$.*

Proof. Suppose all edges of $GT(S)$ intersecting AB are longer. Consider the state of the algorithm just before the first of these edges is added to $GT(S)$. Clearly AB should be selected instead, because it is shorter, leading to a contradiction. ■

Despite its simplicity this lemma is useful in the proof of the following theorem.

Theorem 3. $P_{\pi/8,1}(GT) = true$.

Proof. Consider Figure 6. AB is an edge of some $GT(S)$. CAB , DAB are triangles with base angles $\pi/8$. O is the center of AB . EAF and KBL are portions of circles centered at B and A respectively. We have to show that either CAB or DAB is empty of other sites.

Consider the left half of the figure. Let a site M be inside CAB . Draw a perpendicular MN onto AB . Let the triangles of $GT(S)$ that intersect with NM be t_1, t_2, \dots, t_k , ordered from right to left along NM . Thus t_1 is the triangle immediately to the left of AB . Let t_i be the *first* triangle in the sequence such that one of its three sites, V , lies inside the region $GABIC$. Without loss of generality, assume V lies in the bottom half, $OBIC$. Consider the polygonal region $P = t_1 \cup t_2 \cup \dots \cup t_i$. The portion of P 's boundary from V to A (V to B) is the *upper chain* (*lower chain*).

We observe that the lower chain degenerates into a single edge, VB . If this was not so, then all edges of $GT(S)$ intersecting VB would be longer, contradicting Lemma 1. Two cases arise for the upper chain.

Case 1: The upper chain is a single edge, VA .

Case 2: The upper chain is not a single edge.

Let the sites along the upper chain be $V = u_0, u_1, \dots, u_m, u_{m+1} = A$.

Since $|VA| < |AB|$, by Lemma 1 at least one of these sites is inside EAG .

We next observe that in fact u_m is inside EAG . If this was not so, let u_p be the last site along the upper chain to lie inside EAG , $1 \leq p < m$. Clearly all edges intersecting $u_p A$ are longer, contradicting Lemma 1.

Now consider the right half of the figure, and assume that DAB also contains a site. A symmetric argument leads to the existence of a site, Q , which corresponds to V of the left half. Clearly Q is inside $HABJD$. Unlike the previous argument however, we shall surrender generality if we assume that Q lies in the bottom half of $HABJD$. Consequently four cases arise.

Case 3: Q lies in the bottom half, and the upper and lower chains are single edges, QA and QB .

Case 4: Q lies in the top half, and the upper and lower chains are single edges, QA and QB .

Case 5: Q lies in the bottom half, and the upper chain is not a single edge.

Case 6: Q lies in the top half, and the lower chain is not a single edge.

We are now ready to prove the theorem, by showing that any combination of the left half cases and right half cases leads to a contradiction. In this version of the paper, only two of the combinations are discussed, the others being similar.

Case 1 and Case 4: AB is the only edge intersecting VQ , and $|AB| > |VQ|$, which contradicts Lemma 1.

Case 2 and Case 6: All edges intersecting u_mQ are longer, which contradicts Lemma 1.

A full analysis of all the combinations proves the theorem. ■

The next theorem describes a similar result for minimum weight triangulations. Here the proof is more involved because we do not have a tool as powerful as Lemma 1. First we present some definitions. A *triangulation* of a n -site simple polygon P is the set of $n - 1$ boundary edges along with $n - 3$ nonintersecting diagonals that partition the interior into $n - 2$ triangles. A *minimum weight triangulation* of P , $MWT(P)$, is a triangulation with the minimum total edge length. Let \mathbf{v} be some direction vector. A *plane sweep triangulation* of P is a triangulation constructed by the following plane sweep algorithm. The sweep line is oriented perpendicular to \mathbf{v} , and moves in the direction \mathbf{v} . On visiting a site, the algorithm constructs all possible nonintersecting diagonals between it and the previously encountered sites.

Theorem 4. $P_{\pi/8,1}(MWT) = true$.

Proof. The proof is very similar in structure to that of Theorem 3. Consider Figure 6 again, and let AB be an edge of some $MWT(S)$. We have to show that either CAB or DAB is empty of other sites.

Consider the left half of the figure, and let a site M be inside CAB . The same argument of Theorem 3 leads to the existence of a polygon P , which has a site V inside $OBIC$. Let the sites along the upper (lower) chain be $V = u_0, u_1, \dots, u_m, u_{m+1} = A$ ($V = l_0, l_1, \dots, l_p, l_{p+1} = B$). The following two claims lead us to conclude that the lower chain degenerates into a single edge, VB .

Claim 5: Either the upper or lower chain is a single edge.

Proof. Suppose this was false. Clearly $MWT(P)$ is a subset of $MWT(S)$. An example is in Figure 7. We now construct a different triangulation of P , called $T(P)$ (see Figure 8). First insert the diagonals VA and VB . Then construct the plane sweep triangulation of the polygon enclosed by the upper chain and VA (lower chain and

VB), with the sweep line moving in the direction **CG** (**CI**). We will now show that $|T(P)| < |MWT(P)|$, thus proving the claim.

Consider Figure 7. $MWT(P)$ has $m + p$ diagonals, and each connects an upper chain site with a lower chain site. Consider any site in u_1, \dots, u_m and l_1, \dots, l_p . The rightmost diagonal terminating at the site is said to *belong* to the site. Thus in our example u_2l_2 belongs to u_2 , while Al_3 belongs to l_3 .

Consider Figure 8. $T(P)$ too has $m + p$ diagonals. We classify the sites u_1, \dots, u_m and l_1, \dots, l_p into three types, *peaks*, *troughs*, and *ordinary*. If both neighbors along the upper (lower) chain of some u_i (l_j) are encountered earlier by the sweep line, u_i (l_j) is a peak. If both neighbors are encountered later, it is a trough, otherwise it is an ordinary site. Thus in our example, u_1 is ordinary, l_3 is a peak, while l_2 is a trough. We observe that in each chain, there is one more of peaks than troughs, and if we ignore the other sites, the peaks and troughs alternate. Now consider all diagonals other than VA and VB . The polygon P is such that, during the plane sweeps, exactly $2/1/0$ diagonals are added on encountering a trough/ordinary site/peak respectively. The diagonal inserted on encountering an ordinary site belongs to the site. The left of the two diagonals added on encountering a trough belongs to the trough, while the right diagonal is donated to the next peak along the chain. Thus in our example, u_1A belongs to u_1 , Vl_2 belongs to l_2 , while Bl_2 belongs to l_3 . Note that the first peaks of each chain (in our case u_2 and l_1) do not possess any diagonals of $T(P)$.

We now consider $MWT(P)$ and $T(P)$ together and observe the following. First, $|VA| + |VB|$ is smaller than the sum of the lengths of the two diagonals of $MWT(P)$ that belong to the first peaks of each chain. In our example,

$$|VA| + |VB| < |u_2l_2| + |u_2l_1|.$$

Second, every site in u_1, \dots, u_m and l_1, \dots, l_p other than the first peaks has a diagonal of $MWT(P)$ and a diagonal of $T(P)$ of smaller length. In our example, for the site l_3 , $|Bl_2| < |Al_3|$. Clearly $|T(P)| < |MWT(P)|$, and the claim is proved. ■

Claim 6: In fact, the lower chain is a single edge, VB .

Proof. Recall that V lies inside $OBIC$ of Figure 6. Suppose the claim was false. By Claim 5, the upper chain is a single edge, VA , and $MWT(P)$ has p diagonals, connecting each of l_1, \dots, l_p with A . As in Claim 5, we construct a different triangulation, $T(P)$. First insert VB , then perform a plane sweep triangulation of the region enclosed by VB and the lower chain, with the sweep line moving in the direction **CI**. Now VB is shorter than the diagonal in $MWT(P)$ which connects the first peak with A . The proof is now similar to that in Claim 5. ■

As in Theorem 3, two cases arise for the upper chain.

Case 1: The upper chain is a single edge, VA .

Case 2: The upper chain is not a single edge. In this case $MWT(P)$ has m diagonals, connecting each of u_1, \dots, u_m with B . We show that all of u_1, \dots, u_m lie within EAG . Suppose some were outside EAG . Then u_1B is the longest diagonal, because if u_iB was the longest diagonal for some u_i in u_2, \dots, u_m , then within the quadrilateral $u_{i-1}u_iu_{i+1}B$, $|u_{i-1}u_{i+1}| < |u_iB|$, which is a contradiction. Now consider a different triangulation, $T(P)$, by first inserting VA , then triangulating the region enclosed by VA and the upper chain by a sweep line moving in the direction CG . Clearly u_1 is the first peak. Since $|VA| < |u_1B|$, the proof follows as in claim . We next observe that at any site u_i in u_2, \dots, u_m , the angle $u_{i-1}u_iu_{i+1}$ internal to P is greater than π , otherwise within the quadrilateral $u_{i-1}u_iu_{i+1}B$, $|u_{i-1}u_{i+1}| < |u_iB|$, which is a contradiction.

Now consider the right half of the figure, and assume that DAB also contains a site. The same argument of Theorem 3 leads to the existence of a site, Q , which corresponds to V of the left half. Similarly four cases for the right half arise. We are now ready to prove the theorem, by showing that any combination of the left half cases and right half cases leads to a contradiction. In this version of the paper, only the most complex combination is discussed.

Case 2 and Case 6: Consider Figure 9. $V, u_1, \dots, u_m, A(Q, l_1, \dots, l_p, B)$ are the sites along the left (right) upper (lower) chain. Consider the polygonal region

$$R = Bu_mAQl_1 \dots l_p.$$

We construct a different triangulation of R by inserting BQ , u_mQ , and l_iQ for each l_i in l_2, \dots, l_p . Since $|BQ| < |AB|$, $|u_mQ| < |Al_1|$, and for each l_i in l_2, \dots, l_p , $|l_iQ| < |l_iA|$, we have a contradiction.

A full analysis of all the combinations proves the theorem. ■

4. A GENERAL PLANAR GRAPH ALGORITHM

Part of the motivation for designing Euclidean planar graphs is because their total edge length is small. Consider the set of all Euclidean planar graphs over some S . Although triangulations are planar, they are usually the longest graphs in this set. Thus it is motivating to look for extremely short (possibly nontriangular) planar approximations of the complete graph. However, a limiting factor is that the length of a *minimum spanning tree* ($MST(S)$) is a lower bound on the length of such graphs, for at the least these graphs should be connected.

We now introduce the following heuristic for constructing such graphs. We start with a triangulation of our choice, for example $GT(S)$. We then select some constant d which specifies the good polygon property of Theorem 1. Now the algorithm removes edges from the triangulation, while retaining both properties as required by Theorem 1. At any iteration, the triangulation has been reduced to some planar graph. All faces of the graph satisfy the good polygon property, and trivially all edges satisfy the diamond property. An edge may be removed if the new face formed by merging its two adjacent faces satisfies the good polygon property. The algorithm terminates if no such edge can be found.

Unfortunately, though this may perform well in practice, in theory we do not know exactly how short the graph is bound to become. To facilitate a more quantitative analysis, we have modified the heuristic into a more complicated algorithm, called A , which is as follows. The first step is to construct $MST(S)$, as well as the *convex hull* of S (denoted as $CH(S)$). This partitions the plane into empty polygons, and the sum of all edges is at most $3|MST(S)|$. The next step is to perform a greedy triangulation of each polygonal face. After that, a constant d is selected (whose value will be specified later) for the good polygon property. Finally, only the greedy triangulation edges are examined for possible removal, in *descending order* of edge length.

The following theorems and lemmas prove that $A(S)$ approximates the complete graph, and also lead to a quantitative upper bound on its length. Specifically, we prove that $|A(S)| = O(|MST(S)|)$.

Theorem 5. $P_{\pi/8,d}(A) = true$.

Proof. During edge removal, this property is maintained by the algorithm. Thus we only have to ensure that the edges of $MST(S)$, $CH(S)$, and the greedy triangulations of the polygonal faces satisfy the diamond property. It is known that if e is an edge of $MST(S)$, then the circle with e as diameter is empty of other sites. Thus all such edges satisfy the diamond property. If e is an edge of $CH(S)$, the exterior triangular region is empty, thus satisfying the diamond property.

Let $e = AB$ be a greedy triangulation edge. As in Theorem 3, consider the left half of Figure 6, and let a site M be inside CAB . This leads to the existence of a polygon P , which has a site V inside $OBIC$. It is easy to see that none of the diagonals of P belong to $CH(S)$. Similarly, none of them belong to $MST(S)$ either, otherwise the circle with the diagonal as diameter will contain M , which is a contradiction. Thus the diagonals are greedy triangulation edges, and so the proof of the theorem is identical to Theorem 3. ■

The above theorem ensures that the output of the algorithm indeed approximates the complete graph. We now turn our attention to estimating the length of $A(S)$. Clearly the $MST(S)$ and $CH(S)$ edges add up to $O(|MST(S)|)$. We have to show that the greedy triangulation edges remaining in $A(S)$ also add up to $O(|MST(S)|)$. Recall that the former edges are first added, which results in a planar graph with polygonal faces, hereafter referred to as *regions*. The greedy triangulation edges that are later added can be considered as *diagonals*, which split regions into *subregions*. Note that subregions are essentially faces of $A(S)$. We will consider each region in isolation, and it will be sufficient to prove that the length of all the diagonals is proportional to the length of its boundary. The following lemmas will be useful for the purpose. They describe local properties of $A(S)$, in particular how an edge relates to its neighboring subregions.

Lemma 2. *Let e be a diagonal of some region R . Then there exists a pair of sites A and B along the boundaries of its two adjacent subregions respectively, such that if e is removed, (i) AB will be a chord of the merged subregion, (ii) the shortest distance between A and B along the merged subregion boundary will be greater than $d|AB|$, and (iii) $|AB| \geq |e|$.*

Proof. The merged subregion clearly violates the good polygon property. Consider all pairs of sites where the violation occurs. Let A and B be the pair with the largest separation, $|AB|$ (Figure 10). To contradict the lemma, let us assume that $|AB| < |e|$. Clearly AB intersects e . Thus AB does not belong to the original greedy triangulation. But by Lemma 1, in the original greedy triangulation, there existed an edge f intersecting AB such that $|f| \leq |AB|$. Thus $|f| < |e|$. But since the candidates for removal were in decreasing length, while e was being examined, f was present in the graph. However, at that point f itself separated A and B into different subregions, so they could not have been a violating pair during e 's examination. This contradiction proves the lemma. ■

The next lemma describes how an edge relates to its next-to-adjacent subregions. Let $e = AB$ be a diagonal of some region R . Consider any side of e , and let the polygon P be the union of its adjacent subregion and next-to-adjacent subregions on that side (Figure 11). Let P 's boundary without the edge AB be denoted as $P(A, B)$. It is composed of diagonals, as well as boundary edges of R . Let the set of diagonals along $P(A, B)$ be referred to as $diag(A, B)$.

Lemma 3. *Suppose algorithm A is designed with a (suitably large) constant $d > 1$ specifying the good polygon property. Then there exist constants $b > 1$ and $c > 1$ such that for any S , and any diagonal $e = AB$,*

$$\frac{|P(A,B)|}{|AB|} \leq c \text{ implies } \frac{|AB|}{|\text{diag}(A,B)|} \geq b.$$

Proof (sketch). The lemma in words is, if $P(A,B)$ is very close in length to AB , then the contribution to its length from diagonals is small. We shall give an intuitive sketch of the proof. Details are in [DJ].

Assume the contrary, that is even though $P(A,B)$ is short, most of it is formed by diagonals. Let $Q(A,B)$ refer to the boundary of the subregion adjacent to AB , apart from AB itself. $Q(A,B)$ is shorter than $P(A,B)$, and even more of its length is due to diagonals. Let the length of all its diagonals be l . By Lemma 2, every diagonal e on $Q(A,B)$ is associated with a portion of $P(A,B)$ of length at least $d|e|$. If each such portion is disjoint, we can conclude that $P(A,B)$ is at least as long as dl . By selecting a large enough d , we can arrive at a contradiction, because $P(A,B)$ is initially assumed to be small.

The complication is due to the fact that the above portions are *not* disjoint. However, we show in [DJ] that the overlap is not too much. ■

We assume henceforth that the algorithm A is designed with the constant d as in Lemma 3. We are now ready for the final theorem.

Theorem 6. *Let R be a region. Then the length of all diagonals is of the order of the length of R 's boundary.*

Proof (sketch). The details are in [DJ]. Figure 12 shows a region with all its subregions. Consider a subregion which has only one diagonal, PQ , along its boundary. We remove this subregion from R . Construct the *accounting tree* of edges as follows. Each node contains an edge, and PQ is at the root. The boundary edges are at the leaves, and the diagonals are at the internal nodes. The tree may be built conceptually in a top down manner by "exploring" subregions (Figure 13). First, the subregion adjacent to PQ is explored, and all its edges (other than PQ itself) are made children of PQ . Nodes containing diagonals are expanded in a natural manner, by exploring their adjacent subregions.

Our objective is to show that the total length of all edges is of the order of the total length of leaf edges. The internal nodes may be classified into two types, (a) and (b). Consider an internal node $e = AB$. If we start from e and explore two levels below into the tree, we shall effectively come across its adjacent and next-to-adjacent subregions. The leaf nodes at the first level along with all nodes at the second level

compose $P(A, B)$, as in Lemma 3. If the condition before the implication in Lemma 3 is false, then e is a type (a) node, otherwise e is a type (b) node.

We omit further details and claim that such a condition throughout the tree is sufficient to prove the theorem. As an intuition, the theorem is easily seen to be true if the tree had only one of either type of nodes. ■

5. CONCLUSIONS

We have identified a sufficiency condition such that, graph algorithms satisfying it produce good approximations of the complete graph. We have shown that this condition is quite general because it is satisfied by Delaunay triangulations, greedy triangulations, and minimum weight triangulations. We also use it to design an algorithm for constructing sparse nontriangular graphs. In network design this has some significance because these algorithms can sometimes be better alternatives to the Delaunay triangulation. Furthermore, we have increased our knowledge of minimum weight triangulations, of which little is currently known.

Since the upper bounds in our general results are not tight, our future research is aimed at improving them, possibly by discovering additional properties of the specific algorithms considered. An intriguing open problem is to design an algorithm which produces graphs with the *least measure*. It may be that Delaunay triangulations is the answer here. Finally, other sparse Euclidean graphs need to be investigated, including graphs that are not even planar.

6. ACKNOWLEDGEMENTS

We thank Meera Sitharam and Jorg Peters for several useful discussions.

7. BIBLIOGRAPHY

- [C] Chew: There is a Planar Graph Almost as Good as the Complete Graph: ACM Symposium on Computational Geometry, 1986, 169-177.
- [DJ] Das, Joseph: Planar Euclidean Graphs and their uses in Network Design: Technical Report, UW-Madison, in preparation.
- [DFS] Dobkin, Friedman, Supowit: Delaunay Graphs are Almost as Good as Complete Graphs: IEEE Symposium on Foundations of Computer Science, 1987, 20-26.
- [K] Klincsek: Minimal Triangulations of Polygonal Domains: Annals of Discrete Mathematics 9, 1980, 121-123.

- [L] Lingas: On Approximation Behavior and Implementation of the Greedy Triangulation for Convex Planar Point Sets: ACM Symposium on Computational Geometry, 1986, 72-79.
- [MZ] Manacher, Zobrist: Neither the Greedy nor the Delaunay Triangulation of a Planar Point Set Approximates the Optimal Triangulation: Information Processing Letters, July 1979, 31-34.




Figure 1





Figure 2

Figure 3
(The group G_θ)




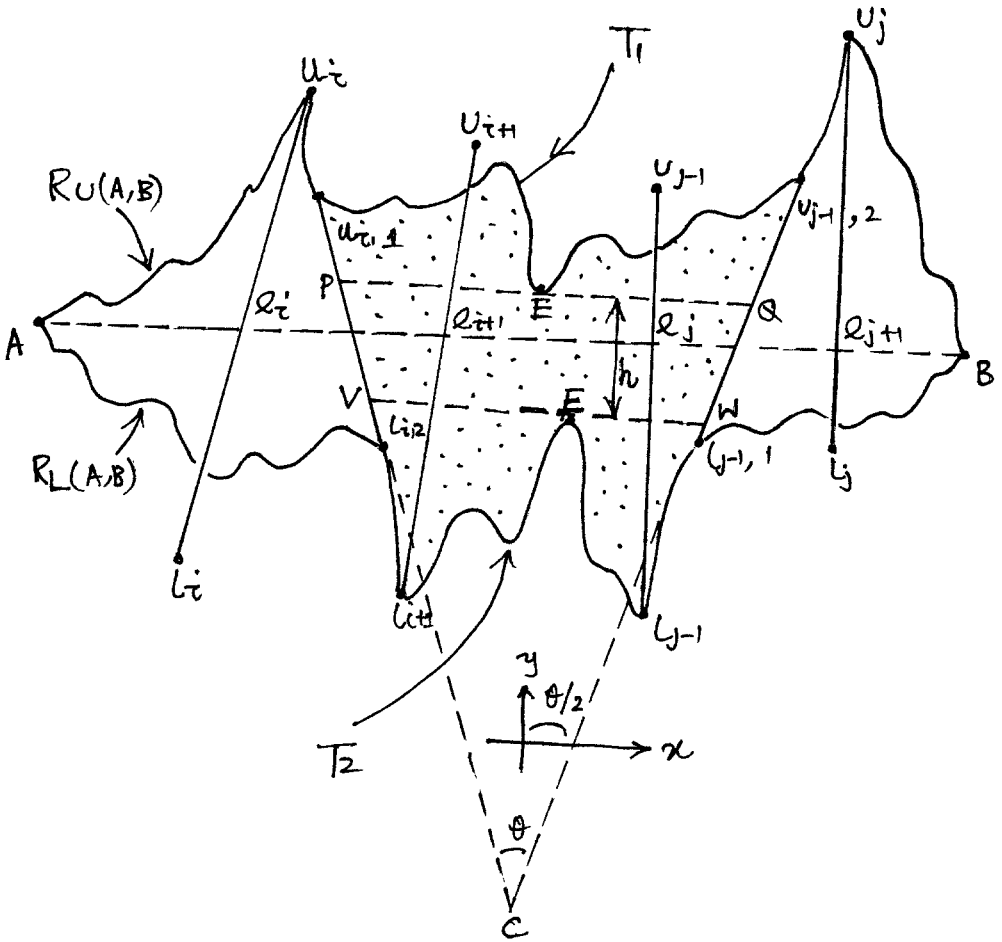


Figure 4

Figure 5






Figure 6




Figure 7




Figure 8






Figure 9

Figure 10



parts of
 $P(A,B)$



- diagonals
- == MST, CH edges
- - - rest of R's boundary

Figure 11

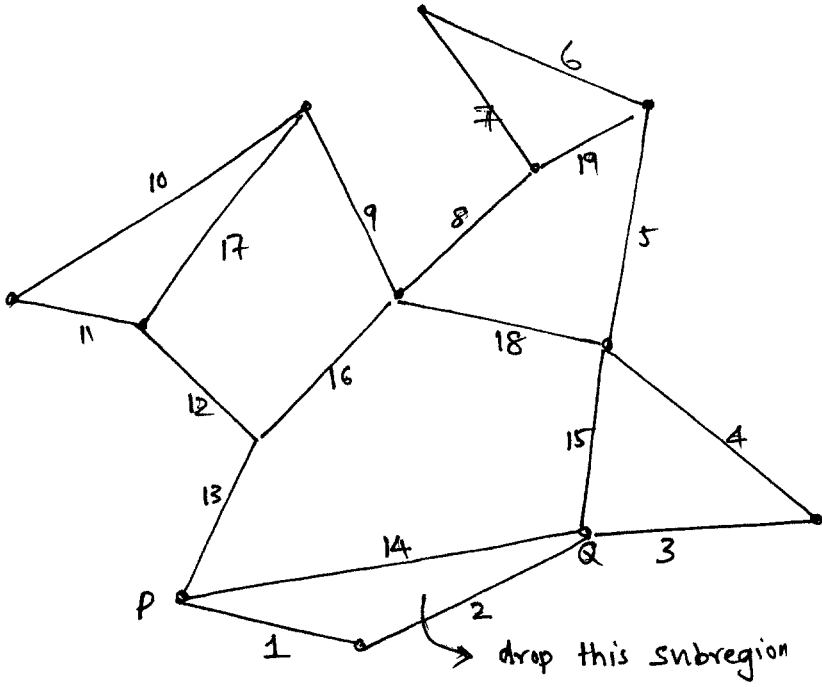


Figure 12

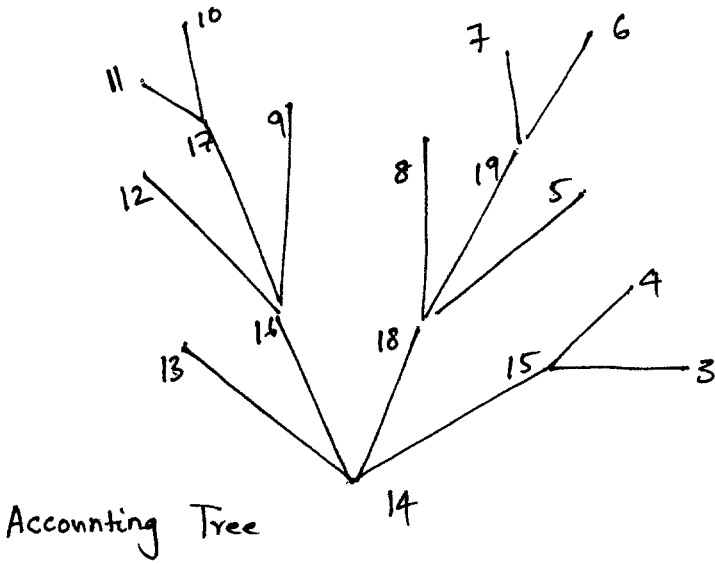


Figure 13