On Exclusion Regions for Optimal Triangulations

R. L. (Scot) Drysdale
Department of Computer Science
Dartmouth College

Scott McElfresh
Mathematical Sciences Department
Muhlenberg College

Jack Scott Snoeyink*
Department of Computer Science
University of British Columbia
and
Department of Computer Science
University of North Carolina

Abstract

An exclusion region for a triangulation is a region that can be placed around each edge of the triangulation such that the region can not contain points from the set on both sides of the edge. We survey known exclusion regions for several classes of triangulations, including Delaunay, Greedy, and Minimum Weight triangulations. We then show an exclusion region of larger area than was previously known for the Minimum Weight triangulation, which significantly speeds up an algorithm of Beirouti and Snoeyink. We also show that no exclusion region exists for the general class of Locally Optimal Triangulations, in which every triangulation edge optimally triangulates the region determined by its two incident triangles.

1 Introduction

Given a finite set S of points in the plane, a *triangulation* is a maximal set of non-intersecting line segments that join points in S. Triangulations are of interest in areas such as finite element methods and interpolation of numerical bivariate data, where each triangle defines a facet along which a projection of each point can be used in piecewise linear interpolation [20]. Of the many possible triangulations, in this paper we look at four particular classes of triangulations that come from optimizing some criterion: Delaunay, Minimum Weight, Greedy, and Locally Optimal triangulations.

In the rest of this section, we define these classes more precisely. In Section 2, we define exclusion regions and survey those that are known. Section 3 establishes larger exclusion regions for the Minimum Weight Triangulation than were previously known. Section 4 shows that the class of Locally Optimal triangulations does not have an exclusion region.

1.1 Some Optimal Triangulations

The Delaunay Triangulation (DT) of a finite set S of points in the plane is the triangulation of S in which the interior of the circumscribed circle for any triangle is free of points from S. The DT is dual to the Voronoi

*Supported in part by an NSERC Research Grant and a fellowship from the B.C. Advanced Systems Institute.

diagram, and is used in many proximity problems [20]. Computation of the Delaunay Triangulation can be done in $O(n \log n)$ time by a number of algorithms, including divide and conquer methods [20]. When the point set is drawn at random from certain probability distributions, it is possible to compute the DT in $O(n)$ expected time and space.

The DT also has the following interesting property: If the *angle sequence* for a triangulation is the list of all angles of the triangles in increasing order, then the Delaunay triangulation is the triangulation with the maximum angle sequence, where angle sequences are compared lexicographically. Thus, the DT maximizes the minimum angle.

For some applications, it is useful to consider the triangulation that minimizes the total length of all edges in the triangulation (where the length of an edge is the distance between its endpoints measured by the standard Euclidean metric). Such a triangulation is referred to as the Minimum Weight Triangulation (MWT).

The problem of finding the MWT for n-vertex simple polygons can be solved in $O(n^3)$ time [15, 10], but it is not known how to efficiently compute the MWT of a general set of points S. In fact, the complexity of finding an MWT—whether it is polynomial time or NP-hard—is one of the few questions that remains open from Garey and Johnson's classic book on NP-completeness [11]. Some approaches to computing the MWT for general point sets have used the method of finding sets of edges that are subsets of the MWT. These subsets include the Beta-skeleton (a subset of the DT) [14] and the LMT-skeleton [8], which uses a local definition of minimality. The current fastest implementations use exclusion regions as filters before computing a variant of the LMT-skeleton [2, 8, 12].

The Greedy Triangulation (GT) is the triangulation defined procedurally by a greedy algorithm that starts by sorting all $\binom{n}{2}$ pairs of points by lengths of the segments connecting each pair. The algorithm then builds the GT one edge at a time by examining each edge e in increasing order of length: adding edge e if e crosses no previously added edges, or discarding e otherwise.

The greedy algorithm as described takes $\Omega(n^2 \log n)$ time to sort all $O(n^2)$ possible edges. Research on computing the GT by other methods have yielded an algorithm with $O(n \log n)$ worst-case time [21] and an algorithm with expected linear time and space usage [5, 18]. These approaches utilize the constrained Delaunay Triangulation or efficient computation of a superset of the GT of size $o(n^2)$.

The Greedy Triangulation has been proposed as an approximation to the MWT. For point sets on a convex polygon and point sets taken from a uniform distribution, the GT is expected to approximate the MWT within a constant factor [18], although in the worst case it may be off by a factor of $\Theta(\sqrt{n})$ [16, 17].

A triangulation is *a locally optimal triangulation* (LOT) if every quadrilateral formed by two adjacent triangles is optimally triangulated with respect to the Euclidean metric. That is, either the quadrilateral is non-convex, or the other diagonal is at least as long as the diagonal that is in the triangulation.

It is not hard to see that the class of LOTs includes both the Greedy Triangulation and the Minimum Weight Triangulation. (Many examples exist where the Delaunay Triangulation is not locally optimal [19].) Local optimality is one of the primary properties of the Greedy Triangulation that is used in analyzing its usefulness, and in developing algorithms to compute it. Relaxing the greedy edge-ordering requirement, but still requiring local optimality may lead to triangulations that are good MWT approximations, but are much simpler to compute. Empirical evidence indicates that other locally optimal triangulations may be useful in applications [7].

2 Exclusion regions known for these triangulations

Das and Joseph [4] first observed that the edges of these globally-defined triangulations may satisfy locally-defined criteria, which we call "exclusion regions." An exclusion region provides a way to show locally that an edge cannot be part of a global triangulation. Specifically, the edge cannot be in the triangulation if there exist two points in the set to be triangulated that lie within the exclusion region on opposite sides of the segment.

Definition 1 *Let p and q be two points in the plane. The region R is an exclusion region for edge $\overline{pq}$ for a given class of triangulations (e.g. MWT, GT, or LOT) if no triangulation in that class contains $\overline{pq}$ when the set of points to be triangulated contains p, q, and at least one point from region R on either side of $\overline{pq}$.*

Before considering the consequences of having an exclusion region, let us consider some examples.

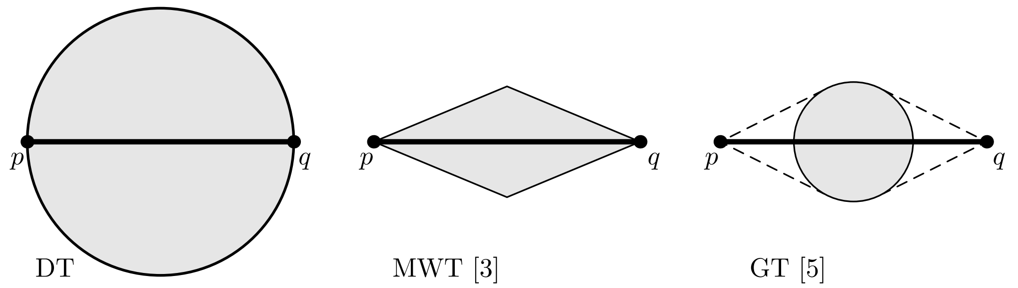

Figure 1: Exclusion regions for Delaunay, Minimum Weight, and Greedy triangulations

The empty-circle definition of the DT naturally leads to the observation that any circle whose diameter is an edge e in the Delaunay Triangulation has at least one of its two semi-circles (separated by e) free of points from the set. In fact, we observe:

Lemma 1 *For the Delaunay Triangulation, the open disk with diameter $\overline{pq}$ is an exclusion region, and no open region including the disk is an exclusion region.*

Proof: For an edge e of the DT, take the empty circumcircle of one of its adjacent triangles. If the center lies on e, then this circumcircle is the exclusion region and both semi-circles are empty. Otherwise, the semi-circle of the transformed exclusion region on the same side of e as the center is contained inside the circumcircle and is therefore empty.

For any larger region R', choose a point $c \in R'$ that is outside the open disk with diameter pq and another point $d \in R'$ just inside the disk so that $\overline{rs}$ goes through the midpoint of $\overline{pq}$. The triangles $\triangle pqr$ and $\triangle pqs$ form a DT, but R' contains points from S on both sides of $\overline{pq}$. Therefore R' is not an exclusion region for edge $\overline{pq}$. ■

Das and Joseph define a diamond-shaped exclusion region for the Minimum Weight Triangulation [4]. We restate their result as the following lemma:

Lemma 2 *The union of two triangles with a shared base $\overline{pq}$ and with base angles $\pi/8$ is an exclusion region for the Minimum Weight Triangulation.*

In [6], Dickerson et al. give an exclusion region for the Greedy Triangulation. We restate it here.

Lemma 3 *A circle of diameter $pq/\sqrt{5}$ centered at the midpoint of $\overline{pq}$ is an exclusion region for the Greedy Triangulation.*

Drysdale, Rote and Aichholzer [5] proved that this region could be enlarged by adding on the regions bounded by tangents from p and q. Das and Joseph showed that their diamond-shaped exclusion region for the MWT is also an exclusion region for the Greedy Triangulation [4].

The fact that a class of triangulations has an exclusion region has two potential benefits: First, for any exclusion region R, we can create a test for edges, where an edge passes the test if at least one of R's two subregions is free of points. Dickerson et al. [6] show that for point sets taken from a uniform distribution over a convex body, only $O(n)$ pairs of points are expected to pass the exclusion region test. This allowed faster algorithms for computing the GT by only considering edges that pass the exclusion region test. Beirouti and Snoeyink [2] utilize the diamond-shaped exclusion region and known subsets of the MWT to compute the Minimum Weight Triangulation on sets of up to 40,000 uniformly-distributed points in five minutes; our improvement of this exclusion region in Section 3 cuts the time in half or lets them double the number of points. The exclusion region helps by allowing them to consider only edges that produce at least one empty semi-region.

Second, if the exclusion region has non-zero area, then for a uniformly distributed set of points, the total length of all edges in the triangulation will be a constant factor of the length of the minimum weight triangulation. This can be proved by extending the approach of [6].

This immediately implies that the Delaunay and Greedy Triangulations are constant factor approximations to the MWT on uniformly distributed points, which was already known [3, 18]. The authors had hoped to use this fact to show that any LOT is expected to be within a constant factor of the MWT for uniformly distributed points. However, we instead showed that no exclusion region can be defined for edges in the class of LOTs.

3 Improving the exclusion region for the MWT

We improve the diamond-shaped exclusion region for the MWT to have base angles $\pi/4.6$, and show a disk-shaped exclusion region with diameter $|e|/\sqrt{2}$. The current and former diamond-shaped exclusion regions for the MWT and the disk-shaped region for the MWT and the disk-shaped region for the GT from [6] are shown in Figure 2. We note that Aichholzer [1] has shown an upper bound on the largest possible disk-shaped exclusion region for the GT that is about 3% smaller than the disk-shaped region proved for the MWT.

As the dashed lines in this figure illustrate, neither the diamond nor the disk contains the other region. We focus on the diamond because it is more amenable to the bucketing used by Beirouti and Snoeyink [2], even though the disk region, when extended to include the regions between the disk and the tangents to the endpoints of e, has slightly greater area.

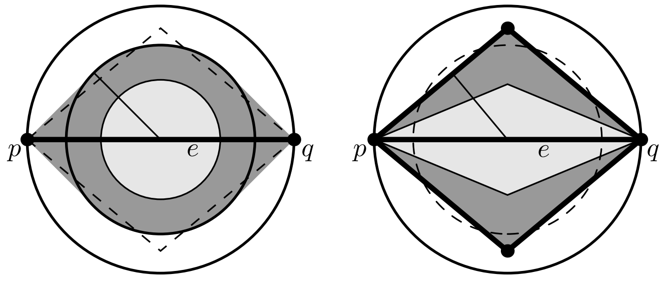

Figure 2: New disk and diamond tests with dark shading; previous tests with light shading

Theorem 4 *The* disk *of diameter $|e|/\sqrt{2}$ centered at the midpoint of e and the* diamond *formed by the two isosceles triangles with base edge e and base angles $\pi/4.6$ are exclusion regions for the Minimum Weight Triangulation.*

Our proof follows the same lines as earlier proofs of tests for the greedy and minimum weight triangulations [4, 6], but tightens many cases. After some definitions and lemmas establishing the structure of the triangulation around e, we show how the assumption that a point from S lies within the exclusion region on each side of e allows us to modify the triangulation to decrease its weight.

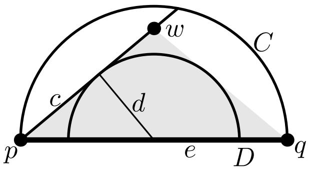

Figure 3: Notation

Assume, for the rest of this paper, that the length of edge e is two, and that e goes horizontally from p on the left to q on the right, as in Figure 3. Let C be the disk with diameter e, and let D be the disk of radius d centered at the midpoint of e. (We will specify values for d later.) Let w be the intersection of tangents to D through p and q; point w is the *apex* of the diamond test.

Let $c = \sqrt{1-d^2}$, which is the length of the segment from p to the point of tangency for the line tangent to D through p. We will use parameters c and d extensively, because they help unify the arguments for the disk and diamond tests. Specifically, for the diamond test we choose $d = 0.631$, which makes the base angles $\pi/4.6$ and apex w be 0.815 above edge e. For the disk test we choose the radius $d = 0.707$.

We use the following lemma frequently.

Lemma 5 *Let $\overline{xv}$ be a chord of circle C above $\overline{pq}$ that intersects the test region—either a disk of radius $d < 1/\sqrt{2}$ or the diamond with $d < \sqrt{\frac{18-3\sqrt{3}}{22}} \approx 0.646$. If x is the endpoint closer to p, then $px < xv$.*

Proof: For the disk test this is easy to establish: The minimum length for any chord of C that crosses disk D is $2c$. The maximum of px occurs when $v = q$ and $\overline{xv}$ is tangent to circle D, so $px \le 2d < 2c \le xv$.

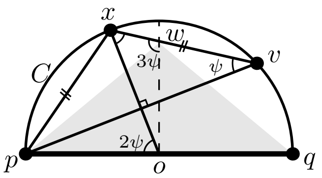

Figure 4: Height of apex w

For the diamond test, we show that x is always left of the perpendicular bisector to $\overline{pv}$ by considering the positions of the apex w that allow x to be on this bisector, as illustrated in Figure 4. Let point o be the midpoint of e and assume that segments $\overline{pv}$ and $\overline{ox}$ are perpendicular, making $\triangle pvx$ isosceles. Apex w must be at or above the intersection of $\overline{xv}$ with the dashed vertical line through point o.

Denote the measure of $\angle pvx$ by ψ. Then $\angle pqx = \psi$ and $\angle pox = 2\psi$. We observe that $\angle wxo = \frac{\pi}{2} - \psi$, $\angle xow = \frac{\pi}{2} - 2\psi$, and $\angle owx = 3\psi$. By the law of sines,

$$ow = \frac{ow}{ox} = \frac{\sin(\frac{\pi}{2} - \psi)}{\sin 3\psi} = \frac{\cos\psi}{\sin 3\psi}$$

The derivative of ow with respect to ψ,

$$\frac{d}{d\psi}\frac{\cos\psi}{\sin 3\psi} \quad = \quad -\frac{\sin\psi\sin 3\psi + 3\cos\psi\cos 3\psi}{\sin^2 3\psi}$$

$$= -\frac{\cos 4\psi + 2\cos 2\psi}{\sin^2 3\psi}$$
$$= -\frac{2\cos^2 2\psi - 1 + 2\cos 2\psi}{\sin^2 3\psi},$$

is zero when $\cos 2\psi = \frac{1}{2}(\sqrt{3} - 1)$, at which point ow attains its minimum value. Using trigonometric relations, we determine that when $ow < \frac{1}{3}\sqrt{3 + 2\sqrt{3}}$, triangle $\triangle pvx$ cannot be isosceles and $px < xv$. ■

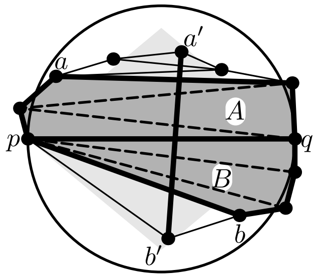

Figure 5: Regions A and B

Now, suppose that the test region—diamond or disk—contains points a' above e and b' below e, as shown in Figure 5. Then segment $\overline{a'b'}$ induces an ordering on the triangles that it intersects; Let A be the set of triangles encountered when tracing $\overline{a'b'}$ toward a', starting from edge e and stopping with the first triangle that has a vertex inside disk C. Let a be the vertex found inside C—if all else fails, then $a = a'$. Similarly, let B be the triangles encountered when tracing from e toward b' until a vertex b is inside C.

The boundary edges of A are grouped naturally into two chains, one from p to a and one from q to a. Although nonadjacent edges can conceivably share endpoints (A need not be a simple polygon), we can still define a sequence of interior angles and vertices for these chains by repeating shared endpoints. We treat a specially, and omit it from both chains of boundary vertices.

We prove the next two lemmas for the chain of A containing p; by symmetry they apply to that containing q and to chains of B. The first lemma is technical, but necessary to prove that triangulations that we define in Lemmas 7 through 9 are valid.

Lemma 6 *If a chain of boundary vertices from p has no three consecutive vertices r, s, and t with $sq > tq$ that form an internal angle of less than π, then the clockwise (cw) limit on the directions of the boundary edges is perpendicular to $\overleftrightarrow{qw}$, the tangent from q to D.*

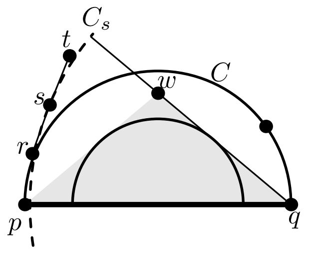

Figure 6: Limiting direction $\overrightarrow{st}$

Proof: If the internal angle at s is less than π—in other words, points r, s, t form a right turn—then, by the hypothesis of the lemma, point t is outside of the circle C_s centered at q through s, as illustrated in

Figure 6. Since all edges are below the line $\overleftrightarrow{qw}$ tangent to D, the limiting direction cw is perpendicular to $\overleftrightarrow{qw}$.

The first edge is similarly limited, since its endpoints lie below $\overleftrightarrow{qw}$ and outside of circle C. Finally, when the internal angle at s is $\geq \pi$, then the direction is limited by that of the previous edge. This establishes the lemma. ∎

Lemma 7 *If the chain of boundary vertices from p has three consecutive vertices r, s, and t with $sq \geq tq$ that form an internal angle of less than π, then the weight of triangulation $T(S)$ can be decreased by retriangulating A.*

Proof: Let us assume that r, s, and t are the first triple that satisfy the conditions of the lemma when adding triangles to form region A starting from edge e. This may involve swapping the rôles of p and q.

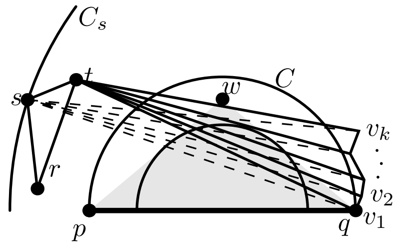

Figure 7: Retriangulating the fan

Let $v_1, v_2, \ldots, v_k$ be the opposite endpoints of chords from s, listed in order along the boundary chain from q to a; the chords are shown dashed in Figure 7. (Recall that point a is not in either chain; in particular, $t \neq a$.) We retriangulate the *fan* of triangles incident on s by replacing these chords with $\overline{rt}$, $\overline{tv_1}$, $\overline{tv_2}$, $\ldots$, $\overline{tv_{k-1}}$, shown solid in Figure 7. We must show that this gives a valid triangulation of the fan having smaller weight.

Lemma 6 applies to the chain from q and limits the direction of the fan's boundary edges so that none have direction between $\overline{sv_i}$ and $\overline{tv_i}$. Thus, the fan boundary does not obstruct chord $\overline{tv_i}$. Moreover, chord $\overline{rt}$ remains in a convex region bounded by the lowest chord out of t, the highest chord out of r, the edges $\overline{rs}$ and $\overline{st}$, and a segment of C on the right side. Thus, we have a valid retriangulation.

To see that weight decreases, notice that the direction of st is also constrained so that each vertex v_i is closer to t than to s—the perpendicular bisector of $\overline{st}$ intersects $\overline{pq}$. Thus, $sv_i \geq tv_i$ for all $i < k$, and the lemma will be established once we show that $sv_k > rt$.

In the rest of the proof, we investigate the worst placements of r, s, t, and v_k, and show that $sv_k > rt$ holds. Since all chords cross the test region, points r, s and t lie left of C between the rays $\overrightarrow{v_k p}$ and $\overrightarrow{v_k w}$, where the later is the tangent to the test region. Let v be the intersection of $\overline{wv_k}$ with C. If v_k is not v, and we replace v_k with v, then the region in which r, s and t may lie becomes strictly larger and $sv_k \geq sv$. (See Figures 8 and 9.)

Next, fix the circle C_s through s and centered at q; note that point t is inside C_s. Let x and y be the points where $\overrightarrow{vw}$ intersects the circles C and C_s, as shown in Figure 8. When $sq > 2$, let z denote the intersection of $\overrightarrow{vp}$ and C_s, as shown in Figure 9.

If point r is outside of C_s, as in Figure 8, the choice of rst as the first triple that satisfies the hypothesis implies that the entire chain from p to s has all internal angles $\geq \pi$ and lies inside C_p, the circle of radius 2 centered at q.

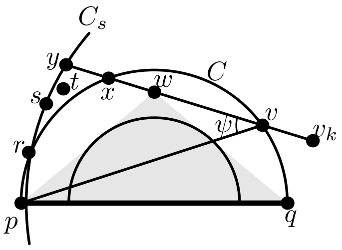

Figure 8: rst and v when $sq \leq 2$.

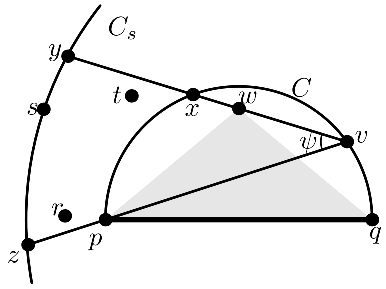

Figure 9: Maximum zx/yx when $z = p$.

Therefore, we can upper bound rt by the maximum of lengths yx, px, py, and, when $sq > 2$, pz, zy, zx. We therefore show that sv is greater than each of these. In many cases we use the observation that $sv \geq yv$.

Case $\overline{yx}$: This is easy, $sv \geq yv > xy$.

Case $\overline{px}$: From Lemma 5, $sv \geq yv > xv > px$.

Case $\overline{py}$: $sv \geq yv = yx + xv > yx + px > py$ by Lemma 5 and the triangle inequality.

Case $\overline{pz}$: $zx > pz$ since $\angle zpx$ is obtuse. Hence, this case reduces to the $\overline{zx}$ case.

Case $\overline{zy}$: If angle $\angle vyz$ is obtuse, then $zx > zy$, so we need not consider this case. But when $\angle vyz$ is not obtuse, then the fact that $\angle zvy < \pi/4$ implies $sv > yv > zy$.

Case $\overline{zx}$: We first show that the zx/yz is smallest when $z = p$; then we show that it is less than one.

Let ψ denote the angle $\angle zvy$. The law of cosines gives us the following equations:

$$\begin{aligned} sq^2 &= zv^2 + qv^2 \\ sq^2 &= yv^2 + qv^2 + 2yv \cdot qv \sin\psi \\ zx^2 &= xv^2 + zv^2 - 2xv \cdot zv \cos\psi \end{aligned}$$

To know whether $zx < yv$, we check whether their squared ratio is less than unity, applying the previous equations to simplify the ratio. In the last line we use the abbreviation $\beta = xv/2\cos\psi$.

$$\begin{aligned} \frac{zx^2}{yv^2} &= \frac{xv^2 + zv^2 - 2xv \cdot zv \cos\psi}{zv^2 - 2yv \cdot qv \sin\psi} \\ &= 1 + \frac{xv^2 - 2xv \cdot zv \cos\psi + 2yv \cdot qv \sin\psi}{zv^2 - 2yv \cdot qv \sin\psi} \\ &= 1 + \frac{2xv(\beta - zv)\cos\psi + 2yv \cdot qv \sin\psi}{zv^2 - 2yv \cdot qv \sin\psi}. \end{aligned}$$

We can show that the numerator of this last fraction is negative by showing that

$$\frac{zv - \beta}{yv} > \frac{qv \sin \psi}{xv \cos \psi}.$$

For a fixed v, the right hand side is constant; it is enough to check where the left hand side takes on its minimum value. The derivative

$$\begin{aligned} \frac{d}{d\,yz} \frac{zv - \beta}{yv} &= \frac{d}{dyz} \frac{\sqrt{yv^2 + 2yv \cdot qv \sin \psi} - \beta}{yv} \\ &= \frac{\frac{-yv \cdot qv \sin \psi}{\sqrt{yv^2 + 2yv \cdot qv \sin \psi}} + \beta}{yv^2} \\ &= \frac{\beta zv - yv \cdot qv \sin \psi}{yv^2 zv} \end{aligned}$$

is positive over its range, since $zv > yv$ and $xv > qv$ from Lemma 5 implies that $xv > qv \sin 2\psi$ or $\beta > qv \sin \psi$. Thus, the minimum occurs when the point $z = p$.

Now, $\frac{yv}{px} > \frac{xv}{px} > 1$ by Lemma 5. Therefore, $sv \geq yv > zx$. This completes the $\overline{zx}$ case.

To conclude, in all configurations $sv_k \geq sv > rt$, so the weight of the triangulation decreases. ■

Now, we can retriangulate regions A and B and, unless they are fairly special fan-like triangulations, decrease the weight.

Lemma 8 *Suppose that region A contains the triangle $\triangle pqv$ and that, without loss of generality, vertex a is not adjacent to q. Then A can be retriangulated so that its weight decreases or, when A is a fan with all chords incident to vertex p, its weight increases by at most $qa - pv$. Furthermore, when A is a fan incident to p, triangles $\triangle pqa$ and $\triangle qva$ are in the new triangulation and $\overline{va}$ intersects the exclusion region.*

Proof: We can apply Lemma 7 to retriangulate and decrease the weight unless no boundary chain has three consecutive vertices r, s, and t with $|sq| > |tq|$ that form an internal angle of less than π.

When Lemma 7 does not apply, we retriangulate by removing all chords from region A and retriangulating with a fan from vertex a to all vertices of the boundary chains. Lemma 6 implies that each new chord from a remains in the interior of A, so this retriangulation is valid.

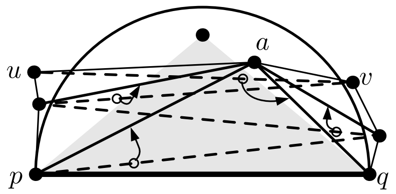

Figure 10: Pairing chords A

To bound the change in weight, we pair removed chords with inserted chords of equal or smaller length. (See Figure 10.) List the original chords of A in the order that they intersect the line through a that is perpendicular to $\overline{pq}$, with those that intersect the line closer to $\overline{pq}$ coming before those that intersect further from $\overline{pq}$. Except for $\overline{qa}$, each new chord $\overline{au'}$ is paired with the last original edge incident on u'. This pairs every original chord except for the last one, $\overline{uv^*}$. We thus pair $\overline{qa}$ with $\overline{uv^*}$.

For each original chord $\overline{u'v'}$, we have the weight $u'v' > u'a$ since $\overline{u'v'}$ crosses C above $\overline{pq}$ and a is inside C above $\overline{u'v'}$. The remaining pairing of $\overline{qa}$ with $\overline{uv^*}$ may increase weight but we claim $qa - uv^* < pv - pa$; that is, the increase is less than the amount saved by pairing $\overline{pa}$ with the first removed chord $\overline{pv}$.

To justify this claim, note that the following movements neither decrease $pa+qa$ nor increase $pv+uv^*$: First, move a up until w lies on $\overline{ua}$ or $\overline{v^*a}$, then move a to an intersection of C with $\overleftrightarrow{wa}$. Next, move u and v^* to the intersections of C with $\overleftrightarrow{wa}$, then move v to v^*. Thus, the worst configuration has $v = v^*$, both u and v on the circle C, and a at u or v, and the inequality is reduced to either $qv < uv$ or $pu < uv$, both of which were established by Lemma 5.

Finally, when A is a fan, the unpaired edges are the last triangle's base, $\overline{pv^*}$, and the first triangle's chord, $\overline{qa}$. Since Lemma 7 does not apply, the chords of the fan cannot decrease in length, thus the increase in weight $qa - pv^* \leq qa - pv$.

To complete the proof of the lemma, we need to verify the last statement - that va intersects the region. In order to have $qa - pv > 0$, we must have va intersect the exclusion region. ■

There remain cases in which both regions A and B are fans and retriangulating either does not decrease the cost. In a final lemma, we reduce these to the cases of pentagons and hexagons and, with the assistance of the algebra package Mathematica, establish that they can be retriangulated at lower cost.

Lemma 9 *When both A and B are fans, then we can retriangulate and decrease the cost.*

Proof: If A and B each consist of a single triangle, then we can replace $\overline{pq}$ by $\overline{ab}$, decreasing the cost. Thus, assume without loss of generality that A is a fan with at least two triangles incident on p. From Lemma 8, we can retriangulate A at a cost of $qa - pv$. If this cost is negative, then we are done; if the cost is zero, then retriangulating turns A into a single triangle, and we can exchange the rôles of A and B.

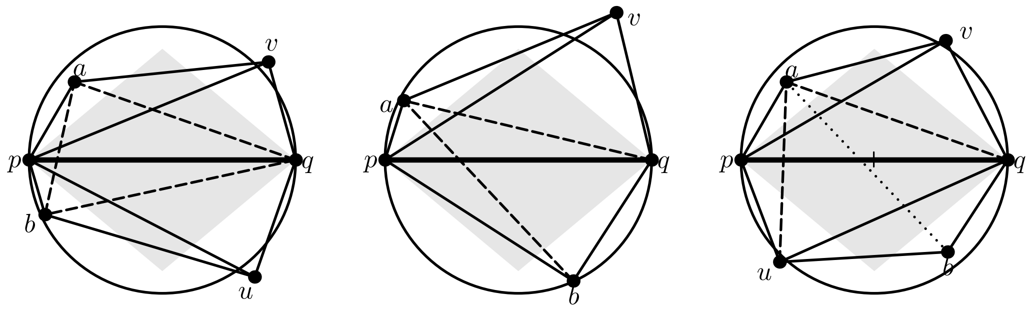

Figure 11: Three cases for retriangulating fans

Therefore, we can assume that A is a fan with $qa > pv$. There are now three cases to consider for B, illustrated in Figure 11: Either (i) B is a fan on vertex p with $qb > pu$, or (ii) B is a single triangle, or (iii) B is a fan on vertex q with $pb > qu$. For each case, we have a pentagon or hexagon for which some retriangulation will decrease the total weight.

(i) B is a fan on vertex p with $qb > pu$: Apply Lemma 8 to retriangulate A and B, then replace $\overline{pq}$ with $\overline{ab}$. The weight increases by at most

$$\begin{aligned} qa - pv + qb - pu + ab - pq &\leq qa - pv + qb - pu + (pa + pb) - 2 && (1) \\ &= (pa + qa - pv - 1) + (pb + qb - pu - 1). && (2) \end{aligned}$$

These terms are symmetric, so it suffices to prove that $pa + qa < pv + 1$.

For simplicity, we would like a to lie outside of the exclusion region, with either $\overline{pa}$ or $\overline{av}$ crossing the region. We thus move a up perpendicular to $\overline{pq}$ until a is outside of the region. Note that this increases $pa + qa$, so will not make our inequality easier to prove. However, this movement may cause a situation where neither $\overline{pa}$ nor $\overline{av}$ crosses the region. This happens precisely when a passes through w, the apex of the region. To avoid this, if a lies directly below w, we instead first move a a small amount to the left or to the right (whichever direction increases $pa + qa$) and then directly up as before.

To minimize pv, we can assume that v is at the intersection of $\overrightarrow{aw}$ with circle C; to maximize $pa+qa$ we can move a along $\overrightarrow{vw}$ until either $qa = pv$ or a is also on C. In the former case, the inequality reduces to $pa < 1$, which holds for the chosen w: the maximum of pa, which occurs when $\overline{va}$ is horizontal through w, is less than 0.92. In the later case, the maximum of $pa + qa - pv$ when a is constrained to lie on the circle actually occurs at the same location.

(ii) B is a single triangle: This is the case that determines our choice of angle $\pi/4.6$. Assuming that $bv > pq = 2$, we show that $ba + qa < pv + 2$ so that segments $\overline{ba}$ and $\overline{qa}$ can replace $\overline{pq}$ and $\overline{pv}$.

As in case (i), we would like a to lie outside of the region, so we move a perpendicularly away from $\overline{bq}$, adjusting analogously to above if a passes through w.

Consider the positions of a, b, and v that make the inequality as tight as possible. We can assume that $bv = 2$ with v on $\overrightarrow{aw}$, because this minimizes pv. We can assume that b is on the circle C; if not, moving b while preserving $bv = 2$ will only increase ba.

The point a on $\overrightarrow{vw}$ that maximizes $ba+qa$ either has $qa = pv$, or is the intersection of C and $\overrightarrow{vw}$. In the former case, the inequality reduces to $pq > ba$, which is true. In the later case, we must verify that $pv+2-ba-qa$ is positive, where a and b are parameterized by angle and v is determined by the locations of a and b. When the diamond is chosen to have angle $\pi/4.6$, calculations in Mathematica show that the minimum value of this function, which occurs with a at 2.85 radians and b at 4.92 radians, is indeed positive.

(iii) B is a fan on vertex q with $pb > qu$: We have saved the most slippery case for last. We will show that $qa + au < 2 + pv$ so that segments $\overline{qa}$ and $\overline{au}$ can replace $\overline{pq}$ and $\overline{pv}$.

As in the previous cases, we would like a to lie outside of the region, so we move a perpendicularly away from $\overline{qu}$, adjusting as above if a passes through w.

We can assume, without loss of generality, that the angle from a to the x-axis is not more than the angle from b to the x-axis—equivalently, that the origin is on or above $\overline{ab}$. To minimize pv, move v to the intersection of circle C with $\overrightarrow{aw}$. (Possibly ignoring the constraint $uv \geq 2$, which does not play a role in this case.)

Now, the placement of b restricts the placement of u since $\overline{bu}$ intersects the test region. Placing b on $\overrightarrow{ao}$ makes the region for u as large as possible; placing u on $\overrightarrow{bw}$ with $qu = pb$ makes au as large as possible. (Note that rays $\overrightarrow{bu}$ and $\overrightarrow{va}$ diverge.)

Finally, to maximize $qa + ua$, we can move a on $\overrightarrow{vw}$ until $qa = pv$ or a is also on the circle C. In the former case, the inequality reduces to $ua < 2$. Here we can be fairly sloppy. The angle for v determines the position of a. The angle for a determines and angle for b and thus a maximum $pb > qu$. The maximum position for u with qu intersecting the test region has $au < 2$.

For the later case, even dropping the constraint that qu intersect the test region, we get a function $pv+2-qa-ua$ parameterized by the angle of a and length of b that we can demonstrate is positive.

This completes the case analysis for the lemma. ■

3.1 Upper Bounds on Exclusion Region Size

We suspect that the analysis of the previous section is not tight—for one thing, it doesn't necessarily use the points inside the test region when constructing a smaller triangulation.

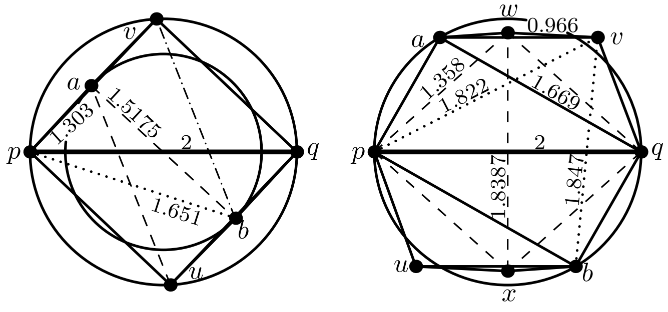

Figure 12: Upper bounds on disk and diamond test regions

Figure 12 shows two examples of minimum weight triangulations using edge $\overline{pq}$ that give upper bounds on the test region size: The diagram at left shows an MWT of six points with one inside a disk of radius $d = 0.759$. Points $puqv$ form a rectangle inscribed in circle C so that sides $\overline{pv}$ and $\overline{qu}$ just enter the circle D at their midpoints. Thus, a can be considered as lying infinitesimally above $\overline{pv}$. The minimum weight triangulation uses $\overline{pq}$, $\overline{pv}$, and $\overline{qu}$; its competitors are symmetric to the triangulation using $\overline{pv}$, $\overline{vb}$, and $\overline{pb}$, which is greater whenever $d > \sqrt{(1+\sqrt{13})/8} \approx 0.759$, and the triangulation using $\overline{vb}$, $\overline{ab}$, and $\overline{au}$, which is greater whenever $d > (15+\sqrt{17})/26 \approx 0.736$.

The diagram at right shows an MWT of eight points with one inside a diamond of angle $\pi/4.23$. Points a and v are chosen on a horizontal line just below w, which is 0.91933 above $\overline{pq}$; a is chosen on the circle C, and v on a circle of radius 1.0833. The minimum weight triangulation uses edges $\overline{av}$, $\overline{aq}$, $\overline{pq}$, and their reflections through the origin. Edges for competing triangulations are labeled with their lengths.

In contrast, the best upper bound known for a diamond-shaped exclusion region for the GT is $\pi/4.51$ [1].

We summarize these results as a theorem.

Theorem 10 *No disk-shaped exclusion region for the MWT centered on the midpoint of an edge e can have radius larger than $0.759|e|$. No diamond-shaped exclusion region for the MWT can have base angles larger than $\pi/4.23$.*

4 Locally Optimal Triangulations and Exclusion Regions

In this section, we show that no exclusion region exists for the general class of Locally Optimal Triangulations (LOTs). Given a segment and two points on opposite sides of the segment, we can create an LOT including the segment and those points.

Lemma 11 *For any segment $\overline{pq}$ and any pair of points a and b lying on opposite sides of $\overline{pq}$ it is possible to create a set of points S including p, q, a, and b and a locally optimal triangulation of S that includes edge $\overline{pq}$. Therefore no exclusion region can be defined for locally optimal triangulations.*

Proof: Let $e = \overline{pq}$, and let R be a region satisfying the conditions of the lemma. Orient the plane so that e is horizontal and p is to the left of q. Let C be the circle with $\overline{pq}$ as diameter. Consider points a

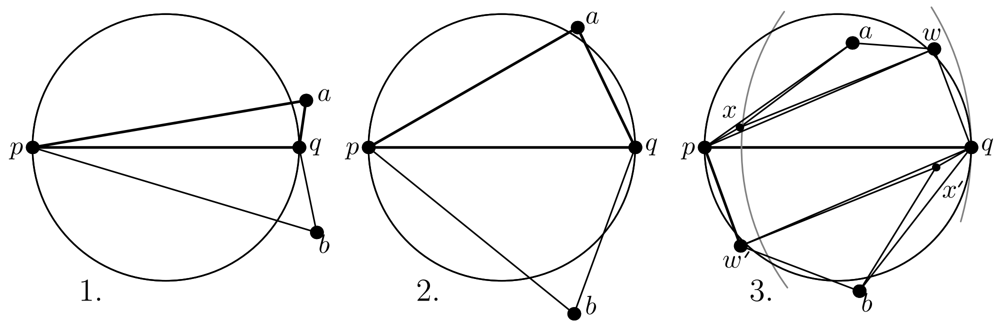

Figure 13: Constructed counterexample

and b in the region R that lie above and below $\overline{pq}$, respectively. If $\angle apq \geq \angle bqp$, then rotate to swap the roles of a and b, and of p and q.

If the quadrilateral $paqb$ is not convex, as in Figure 13.1, or if $pq \leq ab$, as in Figure 13.2, then the triangulation using $\triangle pqa$ and $\triangle pqb$ is locally optimal.

Otherwise, $pq > ab$. We add four points to the set. Choose a point w that lies in $\angle apq$, outside of C, but inside the circle centered at p with radius pq. Choose a point x lying in $\angle apw$, and outside of the pw-radius circle centered at q. This guarantees that $pw < qx$. Define w' and x' to be reflections of w and x through the midpoint of $\overline{pq}$. The placement of these relative to circles and angles involving b and q is the same as the points above the line relative to a and p, since a was initially assumed to make the smaller angle with $\overline{pq}$.

We now triangulate this set of points with the following edges, as illustrated in Figure 13.3: $\overline{pq}$, $\overline{pw}$, $\overline{px}$, $\overline{pa}$, $\overline{xw}$, $\overline{xa}$, $\overline{aw}$, $\overline{qw}$, and the symmetric edges below $\overline{pq}$: $\overline{qw'}$, $\overline{qx'}$, $\overline{qb}$, $\overline{x'w'}$, $\overline{x'b}$, $\overline{bw'}$, and $\overline{pw'}$. We may need to complete the triangulation with edges needed on the convex hull, if w or w' lie within quadrilateral $paqb$. As necessary, add edges $\overline{aw'}$, $\overline{pb}$, $\overline{bw}$, and $\overline{qa}$. Note that at most two of these edges will be added, since $pq > ab$.

The edges in this triangulation are on the convex hull or are diagonals of a non-convex quadrilateral, with the exceptions of $\overline{pq}$, $\overline{pw}$, and $\overline{qw'}$. By construction, each of these three edges is shorter than the other diagonal of its quadrilateral. Hence, the triangulation that was constructed is locally optimal. ■

Thus, we find that for any region, we can place points on both sides of an edge pq such that neither semi-region is empty, but there still exists a locally optimal triangulation that contains edge $\overline{pq}$.

5 Conclusion

This paper increases the size of the known exclusion regions for the Minimum Weight Triangulation. These larger exclusion regions help algorithms that compute the MWT by allowing them to eliminate more edges from consideration. For example, simply enlarging the angle of the diamond-shaped region from $\pi/8$ to $\pi/4.6$ cut the running time of Beirouti and Snoeyink's algorithm in half and allowed them to handle twice as many points in core memory—the MWT for 80,000 points chosen unformly at random could be computed in under five minutes. We have also shown by counterexample that there exists no exclusion region for the more general class of locally optimal triangulations.

6 Acknowledgements

The authors wish to thank Matt Dickerson for helpful discussions during the process of coming up with these results, and the anonymous referees for their careful reading and helpful comments.

References

[1] O. Aichholzer, "Local Properties of Triangulations", CG'95, 11th European Workshop on Computational Geometry, Linz (1995).

[2] R. Beirouti and J. Snoeyink, "Implementations of the LMT Heuristic for Minimum Weight Triangulations", *Proceedings of the Fourteenth Annual ACM Symposium on Computational Geometry* (1998) 96–105.

[3] R. C. Chang and R. C. T. Lee, "On the average length of Delaunay triangulations." *BIT* **24** (1984) 269–273.

[4] G. Das and D. Joseph, "Which triangulations approximate the complete graph?", *Optimal Algorithms, Lecture Notes in Computer Science* **401**, Springer Verlag, (1989) 168–192.

[5] R.L. Drysdale, G. Rote, and O. Aichholzer, "A Simple Linear Time Greedy Triangulation Algorithm for Uniformly Distributed Points", IIG-Report-Series 408, Technishe Universitaet Graz, 1995.

[6] M. Dickerson, R.L. Drysdale, S. McElfresh, and E. Welzl, "Fast Greedy Triangulation Algorithms." *Computational Geometry: Theory and Applications* (1997) 67–86.

[7] M. Dickerson, S. McElfresh, and M. Montague, "New algorithms and empirical findings on minimum weight triangulation heuristics." *Proceedings of the Eleventh Annual ACM Symposium on Computational Geometry* (1995) 238–247.

[8] M. Dickerson, J.M.Keil, and M. Montague, "A Large Subgraph of the Minimum Weight Triangulation." *Discrete and Computational Geometry* 18 (1997) 289–304.

[9] M. Dickerson and M. Montague, "A (Usually?) Connected Subgraph of the Minimum Weight Triangulation." *Proceedings of the Twelth Annual ACM Symposium on Computational Geometry* (1996) 204–213.

[10] P. Gilbert, "New results in planar triangulations." MS Thesis, University of Illinois, Urbana, IL, 1979.

[11] M. Garey and D. Johnson, *Computers and Intractability: A Guide to the Theory of NP-Completeness* (1979).

[12] R. Hainz, O. Aichholzer, and F. Aurenhammer. "New results on minimum weight triangulations and the LMT-skeleton (extended abstract)." *CG'97, 13th Euro. Workshop on Comp. Geom.*, Würzburg Germany, 1997.

[13] L. Heath and S. Pemmaraju, "New results for the minimum weight triangulation problem." *Algorithmica* (1994) 533–552.

[14] M. Keil, "Computing a subgraph of the minimum weight triangulation," *Computational Geometry: Theory and Applications* (1994) 4:13-26.

[15] G. Klincsek, "Minimal triangulations of polygonal domains." *Ann. Discrete Math.* **9** 121–123.

[16] C. Levcopoulos. "An $\Omega(\sqrt{n})$ lower bound for the nonoptimality of the greedy triangulation." *Inform. Process. Lett.*, 25:247–251, 1987.

[17] C. Levcopoulos and D. Krznaric, "Quasi-Greedy Triangulations Approximating the Minimum Weight Triangulation", *J. Algorithms* **27** (1998) 303–338.

[18] C. Levcopoulos and A. Lingas, "Greedy Triangulation Approximates the Minimum Weight Triangulation and Can be Computed in Linear Time in the Average Case." Tech. Report LU-CS-TR:92-105, Dept. of Computer Science, Lund University, 1992.

[19] G. Manacher and A. Zobrist, "Neither the greedy nor the Delaunay triangulation of the planar set approximates the optimal triangulation." *IPL* **9** (1979) 31–34.

[20] F. Preparata and M. Shamos, "Computational Geometry: an introduction." Springer-Verlag, 1985.

[21] C. A. Wang, "An optimal algorithm for greedy triangulation of a set of points." *Proceedings of the 6th Canadian Conference on Computational Geometry* (1994) 332–338.