Minimal weight Steiner triangulations need not exist

Christopher J. Bishop

Abstract. A finite planar point set has only finitely many possible triangulations using only the given vertices, so there must be one that attains the minimum weight, i.e., minimum total edge length. When we allow extra vertices, called Steiner points, there are infinitely many possible triangulations, and it is not obvious whether a minimal weight triangulation exists. In this paper, we construct a finite planar point set with no minimal weight Steiner triangulation.

1. INTRODUCTION A finite set in the plane has only a finite number of triangulations, so there is always at least one that is optimal with respect to any given criterion, such as “maximize the minimum angle” or “minimize total edge length”. The first of these is achieved by the well known Delaunay triangulation [1], which can be computed efficiently in time $O(n \log n)$ for n points. However, minimizing total edge length was proven NP-hard by Mulzer and Rote in 2008 [2]. In practice, we are also interested in Steiner triangulations: these allow extra vertices to be added, but now infinitely many triangulations are possible and the existence of an optimal Steiner triangulation for any given criterion is not obvious. The purpose of this paper is to show that the smallest total edge length need not be attained by any Steiner triangulation. Characterizing the point sets for which the infimum is attained remains an open problem. We will state our result more carefully after providing some definitions and illustrating how Steiner triangulations can sometimes give shorter total edge length than non-Steiner triangulations.

A triangulation of a finite set V in the plane is a maximal set of line segments with endpoints in the set and which are pairwise disjoint except for the endpoints. This is equivalent to writing the convex hull of the points as a union of non-degenerate closed triangles with disjoint interiors and satisfying the simplex condition, i.e., distinct triangles are either disjoint, meet at one common vertex or meet at along a common edge. See the left side of Figure 1. An edge of a triangle that lies on the boundary of the convex hull of V is called a boundary edge of the triangulation. Other edges are called internal edges.

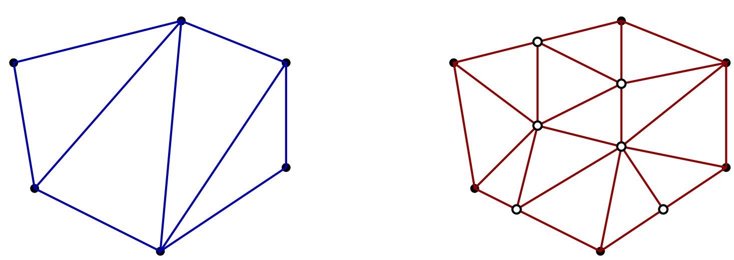

Figure 1. On the left is a triangulation of a given point set V (the black dots). There are six boundary edges and three internal edges. On the right is a Steiner triangulation of the same point set. The white dots are the six Steiner points we have added.

A Steiner triangulation of a set V is a triangulation of some finite set W containing V. In other words, we are allowed to add extra vertices, called Steiner points. See the right-hand side Figure 1. In this picture, we have chosen all the Steiner points to lie inside the convex hull of the original set, but this need not be the case in general. In this paper, we will refer to a triangulation of V without Steiner points as a “triangulation” (or a “non-Steiner triangulation” if we wish to emphasize the point), and one where Steiner points are permitted as a “Steiner triangulation”.

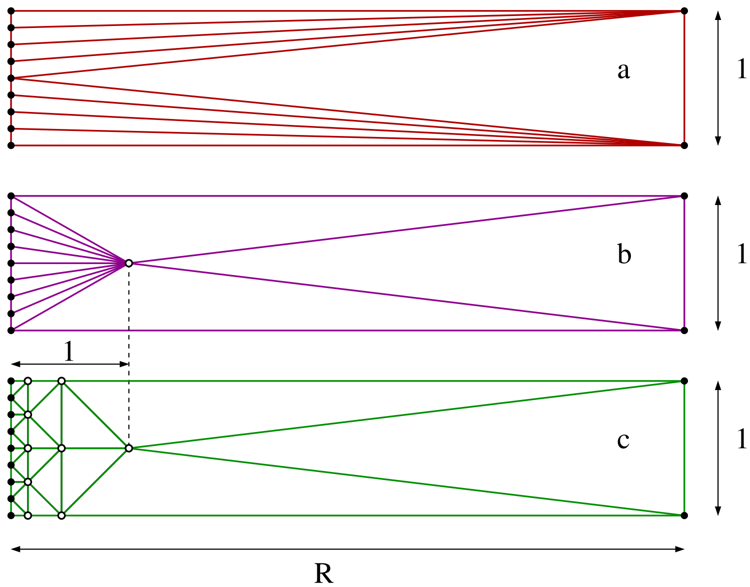

Figure 2. Here the point set V consists of the corners of a $R \times 1$ rectangle plus $n = 2^k - 1$ additional points along one of the shorter edges. The figure shows (a) a triangulation without Steiner points, (b) a triangulation with a single Steiner point, and (c) multiple Steiner points giving an even smaller total edge weight.

Figure 2 shows that allowing Steiner points can lower the total edge length used (we also call this the “weight” of the triangulation). Here the point set V consists of the corners of a $R \times 1$ axis aligned rectangle, plus $n = 2^k - 1$ additional points along the left side of the rectangle. To triangulate V without Steiner points, each of the $n + 1$ points on the left side of the rectangle must be connected by an edge to at least one of the two corners on the right side. Thus any such triangulation must have total length at least $(n + 1)R$; see the top picture in Figure 2.

In the center of Figure 2, we have used one Steiner point (the white dot). This point is connected to the right-hand corners by two edges each of length between $R - 1$ and R, and the $(n + 1)$ vertices on the left by edges of length between 1 and $\sqrt{2}$. Thus the total edge weight is bounded between $4R + n$ and $4R + 4 + \sqrt{2}n$. The latter quantity is less than $(n + 1)R$ when n and R are both large, so using a Steiner point allows for a triangulation with less edge weight.

The bottom of Figure 2 shows a more intricate construction using more Steiner points. We can divide the construction into $\log_2(n + 1)$ columns so that the total edge length used in each column is $1 + \sqrt{2}$ (a vertical segment plus a union of diagonal segments). Thus the total length used is bounded above by $4R + 4 + (1 + \sqrt{2}) \log_2 n$, which is less than for the construction with one Steiner point, at least when n is large.

As noted above, there are only finitely many possible triangulations of a finite set V, so at least one of them attains the minimum possible edge weight. We let this minimum edge weight be denoted $\mathrm{MWT}(V)$ (for **M**inimum **W**eight **T**riangulation). Similarly, let $\mathrm{MWST}(V)$ be the infimum of edge weights over all Steiner triangulations of V. Eppstein [3] has given an algorithm for computing Steiner triangulations that are within a bounded factor of $\mathrm{MWST}(V)$, but I am not aware of any algorithm for computing $\mathrm{MWST}(V)$ exactly. Some more recent papers related to minimal weight Steiner triangulations include [4], [5] and [6]. The following is the main result of this paper.

Theorem 1.1. *There is a set V with five points so that $\mathrm{MWST}(V)$ is not attained by any Steiner triangulation of V.*

Let $\mathrm{MWT}(V, n)$ denote the infimum of $\mathrm{MWT}(W)$ over all finite sets W containing V and with at most n Steiner points. Clearly this is a non-increasing sequence in n, so $\mathrm{MWST}(V) = \lim_n \mathrm{MWT}(V, n)$ exists and is positive. Theorem 1.1 says that this infimum need not be attained by any Steiner triangulation. More precisely, we will show that $\mathrm{MWST}(V) = \mathrm{MWT}(V, 1) < \mathrm{MWT}(V)$, and $\mathrm{MWT}(W) > \mathrm{MWST}(V)$ for every finite set W containing V.

2. THE EXAMPLE Given two distinct points $x, y \in \mathbb{R}^2$ we let xy denote the closed segment with these points as endpoints, and we let L_{xy} denote the infinite line passing through these two points. Given three distinct, non-co-linear points x, y and z, we let $\triangle_{xyz}$ be the closed triangle with these vertices. Consider the five points shown in Figure 3 where the vertices are labeled $V = \{a, b, c, d, e\}$. In terms of coordinates, we can take

$$a = (0, 0), \quad b = (-1, s + t), \quad c = (-1, s), \quad d = (-1, 0), \quad e = (-r, 0).$$

We shall assume $r > 1$ is large and $s > 0$ is small. We will also assume $t \leq s/8$, $s + t < 1/2$, $s + 2t < 1$ and $r \geq 2/t$. For example, we could take $s = 1/2$, $t = 1/16$ and $r = 32$. Let P be the perimeter of the convex hull of these five points. Clearly

$$P = r + \sqrt{1 + (s + t)^2} + s + t + \sqrt{r^2 + 1},$$

but the exact value will not be important.

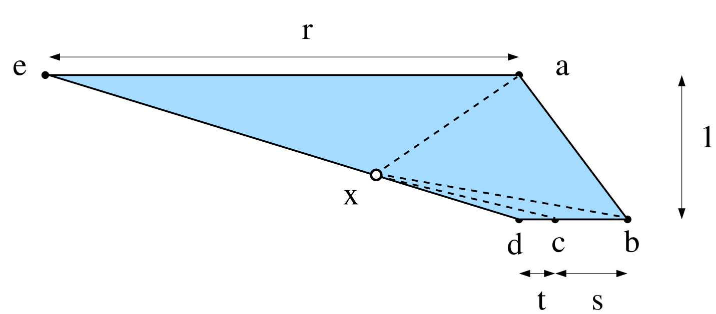

Figure 3. The point set V defined in the text, plus one Steiner point labeled x. As $x \to d$, the total edge length tends to $P + 1 + s + 2t$, where P is the perimeter of the convex hull of V. We will show that every Steiner triangulation has total edge length strictly greater than this, no matter how many extra points are used.

Figure 3 shows V with a Steiner triangulation using one extra point x on the edge de. As x tends to d along this edge, the total length of the triangulation clearly tends to $P+1+s+2t$. Thus $\mathrm{MWST}(V) \le P+1+s+2t$. Figure 4 shows the three possible (non-Steiner) triangulations of these points. Every internal edge of each triangulation in Figure 4 connects one of $\{a,e\}$ to a point in $\{b,c,d\}$, and so has length at least 1. Hence the edge weight of all three triangulations is clearly greater than $P+2 > P+1+s+2t$. Thus for this particular set V, $\mathrm{MWST}(V) < \mathrm{MWT}(V)$. Moreover, we will show that $\mathrm{MWT}(W) > P+1+s+2t$ for any finite set W containing V, and hence that the infimum is never attained; this is exactly what Theorem 1.1 claims.

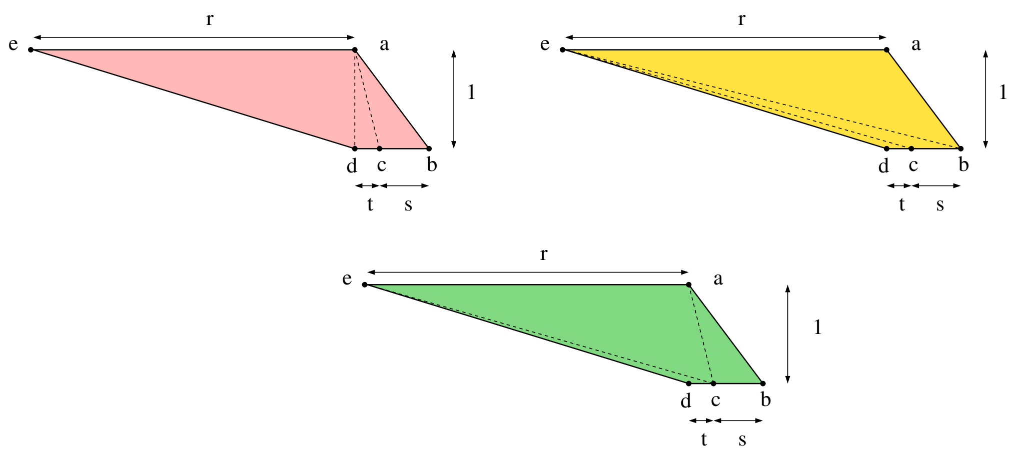

Figure 4. The three possible (non-Steiner) triangulations of V. All the internal edges (the dashed segments) have length at least 1, so all these triangulations have total edge length greater than $P+2 > P+1+s+2t$. Hence they have greater edge length than the Steiner triangulation in Figure 3, at least if x is close enough to d.

We will prove Theorem 1.1 by contradiction. Given a finite set V, we will assume there is point set $W \supset V$ with $\mathrm{MWT}(W) \le P+1+s+2t$, and then prove that the corresponding triangulation must have three internal edges whose lengths sum to strictly more than $1+s+2t$, giving the contradiction. The argument is made precise by the following two lemmas. We say a point $x=(x_1,x_2) \in \mathbb{R}^2$ is to the right of $y=(y_1,y_2) \in \mathbb{R}^2$ if $x_1 \ge y_1$ and strictly to the right if $x_1 > y_1$. We similarly define “to the left of” , “above” and “below”. If $x \in \mathbb{R}^2$ is a point and L is a non-vertical line, we say x is above L, if it is above some point of L, and that x is below L if it is below some point of L. For a non-horizontal line, we can define being left or right of the line is an analogous way.

Lemma 2.1. *Suppose that W is finite point set containing $V=\{a,b,c,d,e\}$, the set defined above, and suppose that $\mathcal{T}$ is a triangulation of W whose total edge length is $\le P+1+s+2t$. Then W contain points v and z so that (see Figure 5)*

1. *v is on or below the line L_{de},*
2. *v is strictly to the left of the line L_{ad},*
3. *v is strictly above the line L_{bd},*
4. *z is to the right of c,*
5. *z is on or below L_{bd},*
6. *the segments va, vb and vz are all internal edges of the triangulation $\mathcal{T}$.*

Lemma 2.2. *For points v and z satisfying the constraints in Lemma 2.1, we have*

$$|v-a|+|v-b|+|v-z| > |d-a|+|d-b|+|d-c| = 1+s+2t.$$

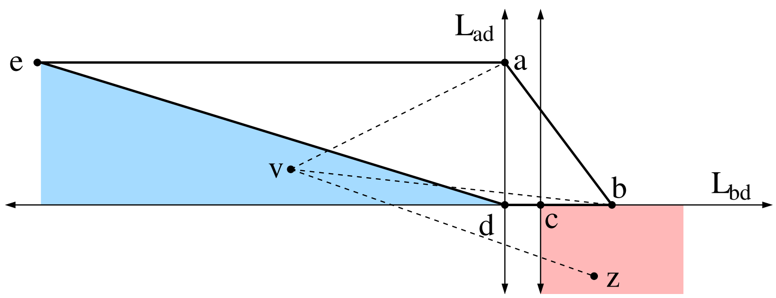

Figure 5. The shaded regions show the restrictions on the points v and z in Lemma 2.1. The dashed segments are the ones whose length sum is estimated in Lemma 2.2.

3. SOME PRELIMINARIES In this section, we present some definitions and elementary results that we will need during the proofs of Lemma 2.1 and 2.2.

Given a finite set W, we let $\mathrm{CH}(W)$ denote its convex hull (the smallest convex set containing W), and we let $\partial\mathrm{CH}(W)$ be the boundary of $\mathrm{CH}(W)$. The convex hull of a finite set is easily seen to be a polygon.

Lemma 3.1. *If $W \subset W'$ are finite sets, then the length of $\partial\mathrm{CH}(W)$ is less than or equal to the length of $\partial\mathrm{CH}(W')$.*

Proof. Extend an edge of $\partial\mathrm{CH}(W)$ to a line L. This line cuts $\mathrm{CH}(W')$ into two convex bodies, one of which contains W and has boundary length no more than the length of $\partial\mathrm{CH}(W')$ (we are replacing a boundary arc of $\partial\mathrm{CH}(W')$ by a line segment with the same endpoints). See Figure 6. Repeating this for each edge of $\mathrm{CH}(W)$ gives a decreasing, nested sequence of convex bodies that ends with $\mathrm{CH}(W)$, and so that each region has boundary length less than or equal to the previous one. This proves the lemma. ■

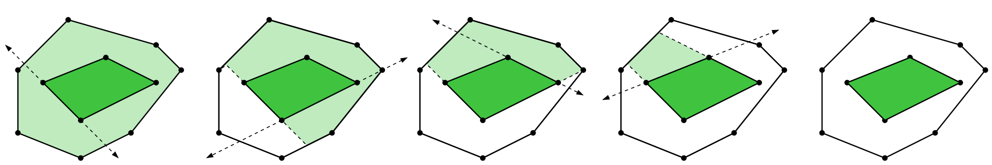

Figure 6. Illustrating the proof of Lemma 3.1.

The previous lemma is true for any pair of nested convex sets (and follows from the above by a limiting argument), but we do not need the general case in this paper.

Lemma 3.2. *In a planar triangulation of a finite point set W, the number of triangles is at least the number of vertices, minus 2.*

Proof. Let $|W|$ be the number of points in W, $|F|$ the number of triangles (= faces), and $|E|$ the number of edges. Let $|B|$ be the number of vertices on $\partial\mathrm{CH}(W)$. Clearly $|B| \leq |W|$. Add a ray connecting each boundary vertex to infinity. This gives a triangulation of the sphere with $|W|+1$ vertices (∞ has been added as a vertex), $|E|+|B|$ edges and $|F|+|B|$ faces. By Euler's formula for the sphere, e.g., [7],

$$2 = (|F|+|B|) - (|E|+|B|) + (|W|+1).$$

Each face has three edges and each edge is counted twice, so $3(|F|+|B|) = 2(|E|+|B|)$, which gives $2 = -(|F|+|B|)/2 + (|W|+1)$ or $|F|+|B| = 2|W|-2$. Since $|B| \leq |W|$, we deduce $|F| \geq |W|-2$, as claimed. ■

An open strip $S \subset \mathbb{R}^2$ of width d consists of all points that are strictly within distance $d/2$ of some line, called the axis of the strip. By a line inside S, we mean any line parallel to the axis, and that is contained in S.

Lemma 3.3. *Suppose $W \subset \mathbb{R}^2$ and that $\mathcal{T}$ is a triangulation of W. Suppose S is a strip of width d so that every line L inside S hits the interior of the convex hull of W. Then either*

1. *the internal edges of $\mathcal{T}$ have total length at least d or,*
2. *there is a line L in S that separates a point $w \in W$ from all other points of W.*

In the second case, the point w is on the boundary of the convex hull of W and the line L intersects the two boundary edges e_1, e_2 of $\mathcal{T}$ that are adjacent to w. L may be chosen to hit interior points of each edge. If x and y are the endpoints of e_1 and e_2 that are not w, then $\triangle_{wxy}$ is an element of $\mathcal{T}$. If $|W| \geq 4$, then removing this triangle from $\mathcal{T}$ leaves a triangulation of the set $V \setminus \{v\}$. In particular, the line $L_{x,y}$ separates w from all the other points of W.

Proof. If every line in S hits an internal edge of T, then the internal edges have length at least d. This is the first alternative of the lemma.

Otherwise, there is a line L in the strip that hits the interior of convex hull of W, but that does not hit any internal edge of the triangulation. Edges are closed segments, so the set of lines missing every internal edge is open, and hence by perturbing the line, if necessary, we may assume the line also misses all points of W. Thus we can choose L so that it hits two distinct boundary edges of the triangulation, e_1 and e_2 at interior points of those edges. Moreover, e_1 and e_2 must be edges of the same triangle in the triangulation, since they are connected by a sub-segment of L that does not cross any other triangulation edges. Let $w \in W$ be the common endpoint of e_1 and e_2. Then w is clearly on the convex hull boundary of W (since it is the endpoint of two boundary edges), and the line L separates w from all other points of W. If x, y are the other endpoints of e_1 and e_2, then the segment xy is an edge of $\mathcal{T}$ and becomes a boundary edge of the new triangulation when we remove the triangle $\triangle_{wxy}$ from $\mathcal{T}$. ■

4. PROOF OF LEMMA 2.1 Suppose W is a finite set containing V and that $\mathrm{MWT}(W) \leq P+1+s+2t$. Let $\mathcal{T}$ be a minimal weight triangulation of W attaining this value. We will find the points v and z described in Lemma 2.1 using three applications of Lemma 3.3.

Since we have seen that all (non-Steiner) triangulations of V have edge weight $> P+1+s+2t$, we know that W must have at least six points, and hence $\mathcal{T}$ has at least four triangles by Lemma 3.2. The convex hull of W contains the convex hull of

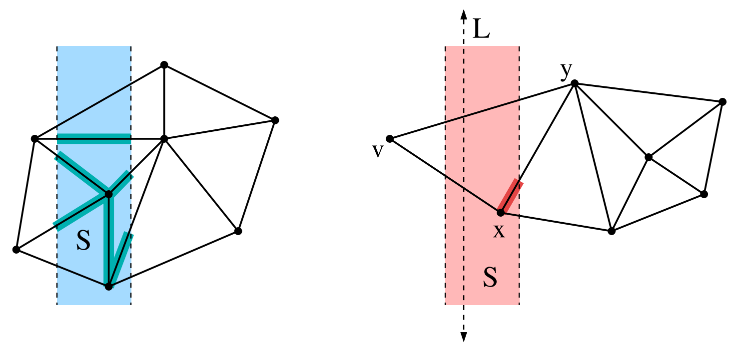

Figure 7. The two cases of Lemma 3.3. On the left every line in the strip hits an interior edge, so the total length of the internal edges is at least the width of the strip. On the right, there is some line L in S that does not hit any interior edge, and it specifies a triangle $\triangle_{vxy}$ with two boundary edges.

V, so its perimeter is at least P by Lemma 3.1. Therefore the internal edges of $\mathcal{T}$ have total length $\leq 1 + s + 2t$.

Now we make the first application of Lemma 3.3 to $\mathcal{T}$ using the vertical strip S_2 that has width 2 and whose right edge is L_{ad}. See Figure 8. Since $1 + s + 2t < 2$, the internal edges of $\mathcal{T}$ have length strictly less than 2, and hence Lemma 2.1 says there is some line L_1 inside S_1 that hits two boundary edges of $\mathcal{T}$ and that separates a single point of W from the rest of W.

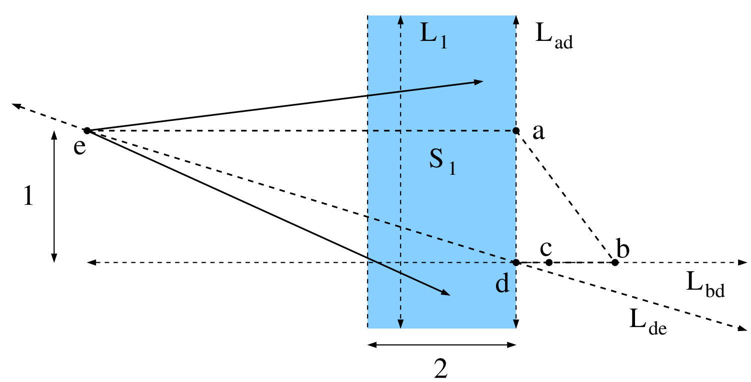

Figure 8. The convex hull of W might surround the convex hull of V. However, every line in the vertical strip S_1 of width 2 hits the interior of the convex hull of V and hence also hits the interior of the hull of W. Thus Lemma 3.3 applies.

Since $e \in V \subset W$ is separated from a, b, c, d by L, the point e must be the same as the point given by the lemma. Let w and v be the vertices of $\partial\mathrm{CH}(W)$ that are adjacent to e, and assume w is the point on or above L_{ae} and v is the point on or below L_{de}. This verifies Conclusion (1) of Lemma 2.1. See Figure 9.

Let L_2 be the line parallel to L_{bd} and distance s above L_{bd}. We claim v is below this line. To see this, note that the line L_{de} has slope between 0 and $-1/r$, and that v lies to the right of the left edge of S_1; this edge is distance 2 from L_{ad}, and this implies that v is at most distance $2/r$ above L_{bd}. Since we have assumed $2/r < t < s/2$ and know that w is above L_{ae}, we see that $|v - w| > 1 - t$. See Figure 9.

We will now apply Lemma 3.3 a second time. Let S_2 be the horizontal strip of width $2s$ and lower edge L_2. Let $\mathcal{T}_2$ be the triangulation $\mathcal{T}$ with the triangle $\triangle_{evw}$ removed.

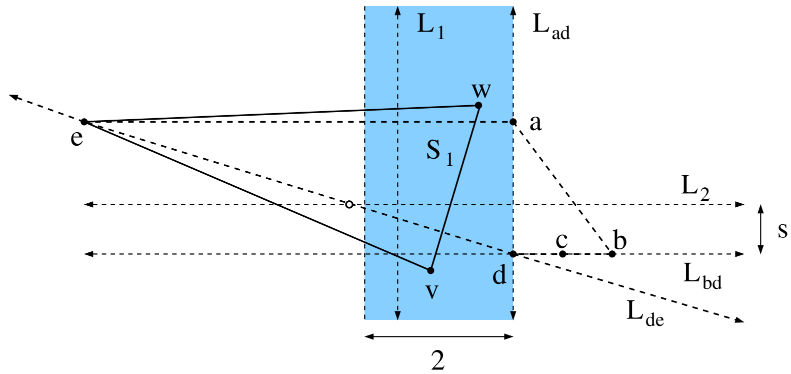

Figure 9. The triangulation $\mathcal{T}$ must contain a triangle $\triangle_{ewv}$ as shown. The point v lies below the line L_{de}, and also below the line L_2, as explained in the text.

Then the internal edges of $\mathcal{T}_2$ are the internal edges of $\mathcal{T}$, except that vw has been omitted. Thus the internal edges of $\mathcal{T}_2$ have total length at most

$$(1 + 2t + s) - (1 - t) = 3t + s < 2s.$$

Therefore, Lemma 3.3, applied to $\mathcal{T}_2$ and the horizontal strip S_2, implies that there is a horizontal segment in S_2 connecting vw to a boundary edge wy of $\mathcal{T}_2$ without passing through any other edge of $\mathcal{T}_2$. See Figure 10.

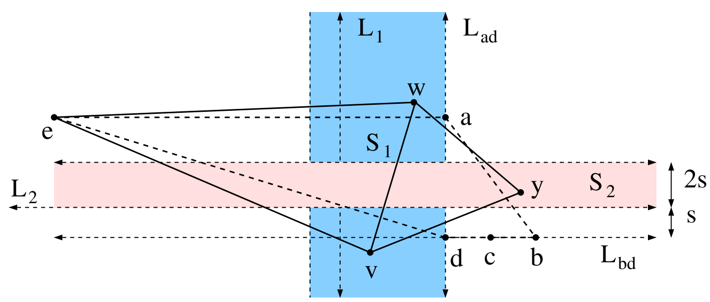

Figure 10. The triangulation of W must contain a triangle $\triangle exw$. We remove this triangle and apply Lemma 3.3 using the horizontal strip S_2 shown.

Since w is separated from the other vertices of $W \setminus \{e\}$ by this line and a also has this property, we also deduce $w = a$. This implies that v is to the left of L_{ad}, for otherwise $\triangle_{eva}$ would contain the point d, which it cannot since elements of $\mathcal{T}$ contain no points of W except their corners. This is Conclusion (2) of Lemma 2.1.

We also claim that $|v - y| > (s + t)/2$. To see this, let q be the point on ab that is closest to d. See Figure 11. Since $\triangle_{abd}$ is similar to $\triangle_{dbq}$, we deduce that

$$|v - y| \geq |d - b| = (s + t)/|a - b| > (s + t)/(|a - d| + |d - b|) = |s + t|/2.$$

We can now make our third application of Lemma 3.3. We use a vertical strip S_3 of width s, with the point c on its left boundary edge and b on its right boundary edge. See Figure 12. Let $\mathcal{T}_3$ be the triangulation $\mathcal{T}_2$ with the triangle $\triangle_{ayv}$ removed. Since $|v - y| > (s + t)/2$, the internal edges of $\mathcal{T}_3$ have total length at most

$$3t + s - (t + s)/2 = (5/2)t + s/2 < s.$$

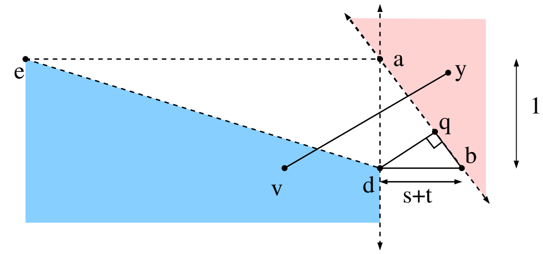

Figure 11. Proof that $|v - y| \geq |d - q| > (s + t)/2$.

Therefore Lemma 3.3 implies that there is a vertical line in S_3 that connects vy to a boundary edge of $\mathcal{T}$. As before, we can deduce that $y = b$ and that the boundary edge is of the form zb for some $z \in W$. See Figure 12.

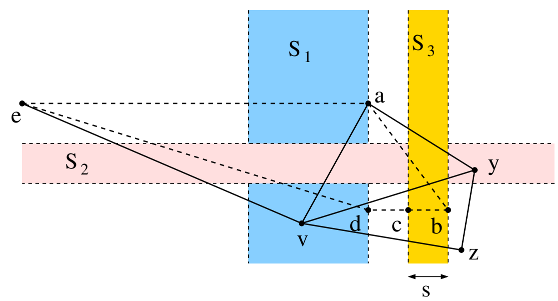

Figure 12. We apply Lemma 3.3 using the vertical strip S_3 of width s, and deduce that $y = b$ and there is a point $z \in W$ so that zb is a boundary edge of the triangulation.

However, there is a slight problem with the way Figure 12 is drawn. Since bz is a boundary edge of the triangulation, z must be below the line L_{bd}; otherwise the points c and d would be outside the convex hull. This is Condition (5) of Lemma 2.1. Furthermore, since $\triangle_{vbz}$ cannot contain the point c, z must be the right of c and v must be above L_{bd}; these are Conditions (3) and (4) of lemma 2.1. See Figure 13.

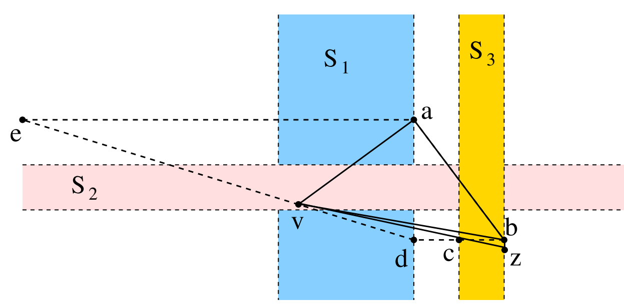

Figure 13. The corrected version of Figure 12. Since the triangle $\triangle_{vbz}$ can't contain the points c or d and since z must be to the right of c, we can deduce that v is above L_{bd}.

The segment va is an internal edge of $\mathcal{T}$ since it is on both $\triangle_{eva}$ and $\triangle_{vab}$. Similarly, vb is and edge of both $\triangle_{abv}$ and $\triangle_{vbz}$ so it is also an internal edge. Finally, vz

is internal for otherwise $\mathcal{T}$ would consist of only the three triangles $\triangle_{eav}$, $\triangle_{abv}$ and $\triangle_{bzv}$, which contradicts that it contains at least four triangles. This verifies Condition (5) of Lemma 2.1, and completes the proof of that lemma.

5. PROOF OF LEMMA 2.2 It is immediate from Conditions (2) and (4) in Lemma 2.1 that $|v - z| \geq |d - c|$, so in order to prove Lemma 2.2, it suffices to prove that $|v - a| + |v - b| \geq |d - a| + |d - b|$, given the assumed constraints on v, i.e., that it is strictly above L_{bd} and on or below L_{de}.

If v is strictly below the line L_{de}, move it to the right until it hits this line. This decreases both $|v - a|$ and $|v - b|$, so decreases their sum. Therefore we may assume v is on the line L_{de} and strictly to the left of d. We want to show that $|v - a| + |v - b|$ decreases as v approaches d along this line. This is clear if v is to the left of p, the closest point on L_{de} to a. See Figure 14. To the right of this point $|v - a|$ increases as v approaches d, but we claim the sum $|v - a| + |v - b|$ still decreases. One could demonstrate this using explicit formulas and computing derivatives, but we give a more geometric proof.

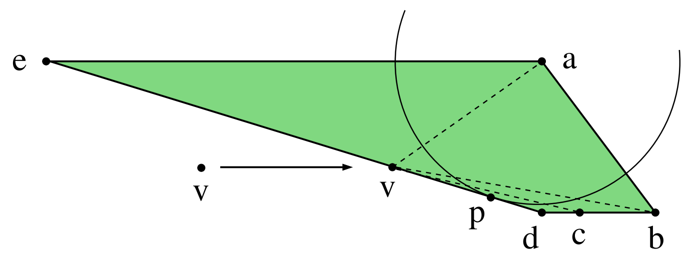

Figure 14. Let p be the point on L_{de} closest to a. To the left of p moving v towards d along L_{de} clearly decreases both $|v - a|$ and $|v - b|$ and hence deceases the sum. We will show the sum continues to decrease to the right of p, even though $|v - a|$ is increasing.

Lemma 5.1. *If $L \subset \mathbb{R}^2$ is a line and if α and β are points on opposite sides of L, then the point γ where the segment $\alpha\beta$ crosses L minimizes $|\gamma - \alpha| + |\gamma - \beta|$ over all $\gamma \in L$. If α and β are on the same side of L, let β^* be the reflection of β across L. Then the point γ where $\alpha\beta^*$ crosses L minimizes $|\gamma - \alpha| + |\gamma - \beta|$. The function $f(\sigma) = |\sigma - \alpha| + |\sigma - \beta|$ is monotone increasing as we move along L away from γ in either direction.*

Proof. If α and β are on opposite sides of L, this is just the triangle identity. See Figure 15. The case when the points are on the same side of L reduces to the first case since $|\gamma - \alpha| + |\gamma - \beta| = |\gamma - \alpha| + |\gamma - \beta^*|$. For the final claim, it suffices to assume α and β are on opposite sides of L. The conclusion then follows from Lemma 3.1 and the fact that for $\mu, \tau \in L$ on the same side of γ, we have $\triangle_{\sigma\alpha\beta} \subset \triangle_{\tau\alpha\beta}$ if $|\gamma - \sigma| < |\gamma - \tau|$. ■

To finish the proof of Lemma 2.2, apply Lemma 5.1 to the line $L = L_{de}$, and the points $\alpha = a$ and $\beta = b$. Let b^* be the reflection of b across L_{de}. Since a and b lie on the same side of L_{de}, the sum $|a - x| + |b - x|$ is minimized (as a function of $x \in L_{de}$) at the point where ab^* crosses L_{de}. We claim this point is to the right of d. This, in turn, holds if b^* is to the right of d, and the proof of this fact is given in Figure 16. Thus $|x - a| + |x - b|$ attains its minimum on L_{de} to the right of d, and this function of x is decreasing on L_{de} as we approach d from the left. This completes the proof of Lemma 2.2, and also the proof of Theorem 1.1.

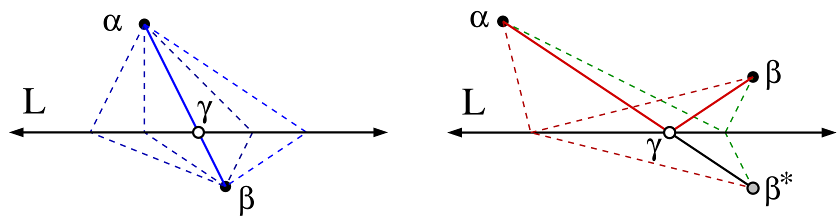

Figure 15. Minimizing the sum of distances of two points $\alpha, \beta \in \mathbb{R}^2$ to a third point on a line. The sum is minimized where the segment $\alpha\beta$ crosses the line (if α and β are on opposite sides) or where $\alpha\beta^*$ crosses the line (otherwise).

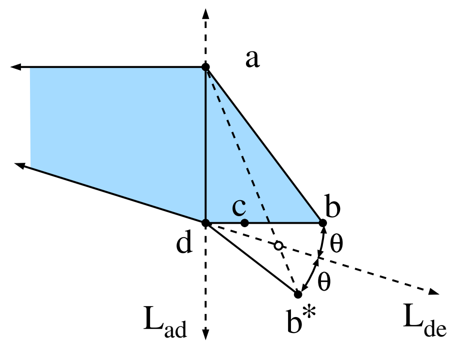

Figure 16. Let θ denote the angle L_{de} makes with the horizontal. Since $r > 1$ we know $0 < \theta < \pi/4$. If b^* is the reflection of b across L_{de}, then the x coordinate of b^* is $|b - d| \cos(2\theta) > 0$. Thus the intersection of the segment ab^* and the line L_{de} (the white dot) occurs to the right of d, as desired.

6. QUESTIONS AND REMARKS In our example, the five points form the vertices of a convex polygon with four extreme points, i.e., three of the points are co-linear, and the minimizing sequence of Steiner triangulations has two triangles that degenerate into line segments. What happens if there are no co-linear points?

Question 6.1. *If $V \subset \mathbb{R}^2$ is a finite set with no three points co-linear, does some Steiner triangulation attain the minimum possible edge weight?*

Many other problems remain open. For example, is there a finite set V so that each of the finite minimums $\mathrm{MWT}(V, n)$ is attained, but $\mathrm{MWST}(V)$ is not, e.g., $\{\mathrm{MWT}(V, n)\}_{n=1}^{\infty}$ is not eventually constant? In other words, is there a finite set V where arbitrarily many Steiner points are needed to approximate $\mathrm{MWST}(V)$? If this can happen, it would be interesting to see if a sequence of nearly optimal Steiner triangulations exhibit some sort of self-similar structure.

The points of our example form the vertices of a convex polygon. Is it easier to decide the problem in the previous paragraph when we restrict to such sets? What about the analogous problem when we consider Steiner triangulations of the interior of a polygon? Is there any characterization of polygons for which the infimal edge weight for Steiner triangulations is attained?

In this paper we considered Steiner triangulations that minimize total edge length, but one can ask analogous questions for other measurements. A well studied example is to minimize the maximal angle used in the triangulation, or maximize the minimum angle. It is fairly easy to see that the Steiner triangulation of a point set can have angles as close to 60° as desired, but that this bound can only be attained for special point sets that are subsets of some equilateral triangular lattice. Minimizing the maximum angle for Steiner triangulations of polygons is harder, but it was recently proven that the optimal bound is always attained except in one special case: polygons where all interior angles are multiples of 60° and with some pair of edges whose lengths have an irrational ratio. In this case, the minimum is 60°, but it is not attained by any Steiner triangulation. See Corollary 1.2 of [8]. Are there other interesting measurements (besides minimizing the maximum angle or minimizing total edge length) for which we can characterize when an optimal Steiner triangulation exists?

REFERENCES

[1] Lawson CL. In: Software for C^1 surface interpolation. New York: Academic Press [Harcourt Brace Jovanovich Publishers]; 1977. p. ix+388. Publication of the Mathematics Research Center, No. 39.
[2] Mulzer W, Rote G. Minimum-weight triangulation is NP-hard. J ACM. 2008;55(2):Art. 11, 29. Available from: `http://dx.doi.org/10.1145/1346330.1346336`.
[3] Eppstein D. Approximating the minimum weight Steiner triangulation. Discrete Comput Geom. 1994;11(2):163–191. Available from: `http://dx.doi.org/10.1007/BF02574002`.
[4] Cheng SW, Lee KH. Quadtree, ray shooting and approximate minimum weight Steiner triangulation. Comput Geom. 2002;23(2):99–116.
[5] Eppstein D, Hadizadeh Z. Minimum-weight Steiner triangulation of convex polygons requires interior Steiner points. 2026;Preprint, arxiv.org/abs/2606.25302, accessed 09/05/2026.
[6] Traub CM. Steiner reducing sets of minimum weight triangulations: structure and topology. Comput Geom. 2015;49:24–36.
[7] Eppstein D. Twenty-one proofs of Euler's formula;. Accessed: 09/05/2026. The geometry junkyard. Available from: `https://ics.uci.edu/~eppstein/junkyard/euler/`.
[8] Bishop CJ. Optimal triangulation of polygons. Discrete Comput Geom. 2026;Available from: `https://doi.org/10.1007/s00454-026-00831-z`.

CHRISTOPHER J. BISHOP received a bachelors degree from Michigan State University, did Part III of the Mathematical Tripos as a Churchill Scholar at Cambridge University, and earned a PhD from the University of Chicago, advised by Peter W. Jones. He is currently Distinguished Professor of Mathematics at Stony Brook University, was an invited speaker at the 2018 ICM, and was awarded the Senior Berwick Prize by the London Mathematical Society in 2024. His research interests include various aspects of analysis: complex function theory, holomorphic dynamics, hyperbolic geometry, geometric measure theory, fractals, and computational geometry.
Department of Mathematics, Stony Brook University, Stony Brook NY, 11794
bishop@math.stonybrook.edu