TOPOLOGICAL EFFECTS RELATED TO MINIMUM WEIGHT STEINER TRIANGULATIONS

CYNTHIA M. TRAUB
WASHINGTON UNIVERSITY IN SAINT LOUIS

Abstract. Let $mwt(\mathcal{X})$ denote the sum of the Euclidean edge lengths of a minimum weight triangulation of a point set $\mathcal{X} \in \mathbb{R}^2$. We investigate the conditions under which an n-point set $\mathcal{X}$ will allow an $(n+1)^{\text{st}}$ point P (called a Steiner point) to give $mwt(\mathcal{X} \cup \{P\}) < mwt(\mathcal{X})$. We call the regions of the plane where such a P reduces the length of the minimum weight triangulation *Steiner reducing regions.* We demonstrate by example that these Steiner reducing regions may have many disconnected components or fail to be simply connected. By examining randomly generated point sets, we show that the surprising topology of these Steiner reducing regions is more common than one might expect.

1. Introduction

We are concerned with a long-standing classical problem in computational geometry: that of finding a minimum weight triangulation of a point set. Formally, a *triangulation* of a point set $\mathcal{X} \subset \mathbb{R}^2$ is an inclusion-maximal set of non-intersecting straight line segments connecting pairs of points in $\mathcal{X}$. A triangulation is specified by its *combinatorial type*: a listing of either its edges or its point-empty triangles. The *length* or *weight* of a triangulation of $\mathcal{X}$ is the sum of the Euclidean lengths of the edges used, so a *minimum weight triangulation* of X is a triangulation which has length less than or equal to the length of every other triangulation of $\mathcal{X}$. We note that such a triangulation is not necessarily unique. We denote the length of a minimum weight triangulation of $\mathcal{X}$ by $mwt(\mathcal{X})$, and we denote its set of edges by $\text{MWT}(\mathcal{X})$. For a thorough mathematical treatment of triangulations, consult [8].

Different measures of optimality for triangulations have given rise to useful applications and algorithms. For a survey of optimization with regard to triangulations, see [4] or [8]. Of all the problems of unknown computational complexity collected in [12], the minimum weight triangulation problem is one of the few that remains yet unclassified. There are polynomial-time algorithms for determining the minimum weight triangulation of special classes of point sets, such as polygonal domains [13, 15]. Certain edges and progressively larger subsets of edges have been proven to belong to the minimum weight triangulation. These include the shortest edge [13], all mutual nearest-neighbor edges [22], and two different sets of edges known as the β-skeleton [6, 14] and the LMT-skeleton [1, 9]. Additional work has been done to create and evaluate different methods of finding the exact minimum weight triangulation and also approximating the minimum weight triangulation of point sets in $\mathbb{R}^2$ [2, 10, 16, 17, 18, 19, 20] and higher dimensions [3, 5, 7].

Many thanks to Jesús De Loera for his comments in the creation of this document.

In this paper we consider triangulations whose vertex set contains properly the original input set. In other words, we allow for the addition of a small number of new points, called *Steiner points*, to the input set before triangulating. Reasons for doing this include improving the quality of meshes [4] or approximating the minimum weight triangulation of the input point set [11]. Minimum length triangulations that allow Steiner points are called minimum weight Steiner triangulations.

A perhaps surprising effect of adding Steiner points is that the new point set can have a minimum weight triangulation with length less than that of the original point set. In Figure 1 we show an example.

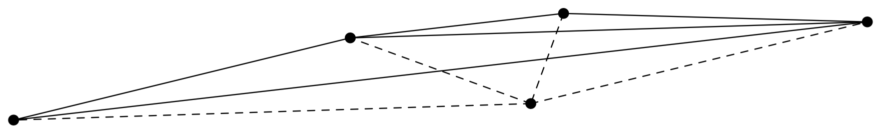

FIGURE 1. A fifth point added to this quadrilateral reduces the length of the minimum weight triangulation of the new larger point set.

We call a point set $\mathcal{X}$ *Steiner reducible* if there exists a point $P = (x, y) \in \mathbb{R}^2 - \mathcal{X}$ such that

$$mwt\big(\mathcal{X} \cup \{P\}\big) < mwt(\mathcal{X}).$$

Such a point P is said to *reduce* the length of the triangulation, and we refer to P as a *Steiner reducing point.* For a point set $\mathcal{X}$, we are concerned with the region of the plane consisting of all reducing points, which we refer to as the *Steiner reducing region.*

In this paper we discuss the topological properties of Steiner reducing regions. Given an input point set, in which regions of the plane can one find a Steiner reducing point? What do these Steiner reducing regions look like? Are there disconnected components or holes? Are the components convex? How does the complexity of the topology increase with the number of points? Our first contribution is to show that, even when connected, the Steiner reducing region can have holes:

Theorem 1. *There exists an 18-point set that admits a connected Steiner reducing region whose first homology group has rank at least 13.*

In Figure 3 we see that a set with as few as five points can have a Steiner reducing region with two disconnected components.

It is easy to create point sets with many more disconnected components in their Steiner reducing region by scattering small copies of the example in Figure 3 at larger distance from one another. But what is the behavior for random point sets? Or, can one construct non-clustered point sets whose Steiner reducing regions have many disconnected components? Our second contribution is an experimental study of random point sets with respect to Steiner reducing regions. Our experiments indicate that the existence of Steiner reducing regions is relatively common even for random point sets. The random point sets tested also indicated that it is much more likely for a Steiner reduction to occur in the boundary of the convex hull of our input set than in the interior. Finally, multiple disconnected components of the

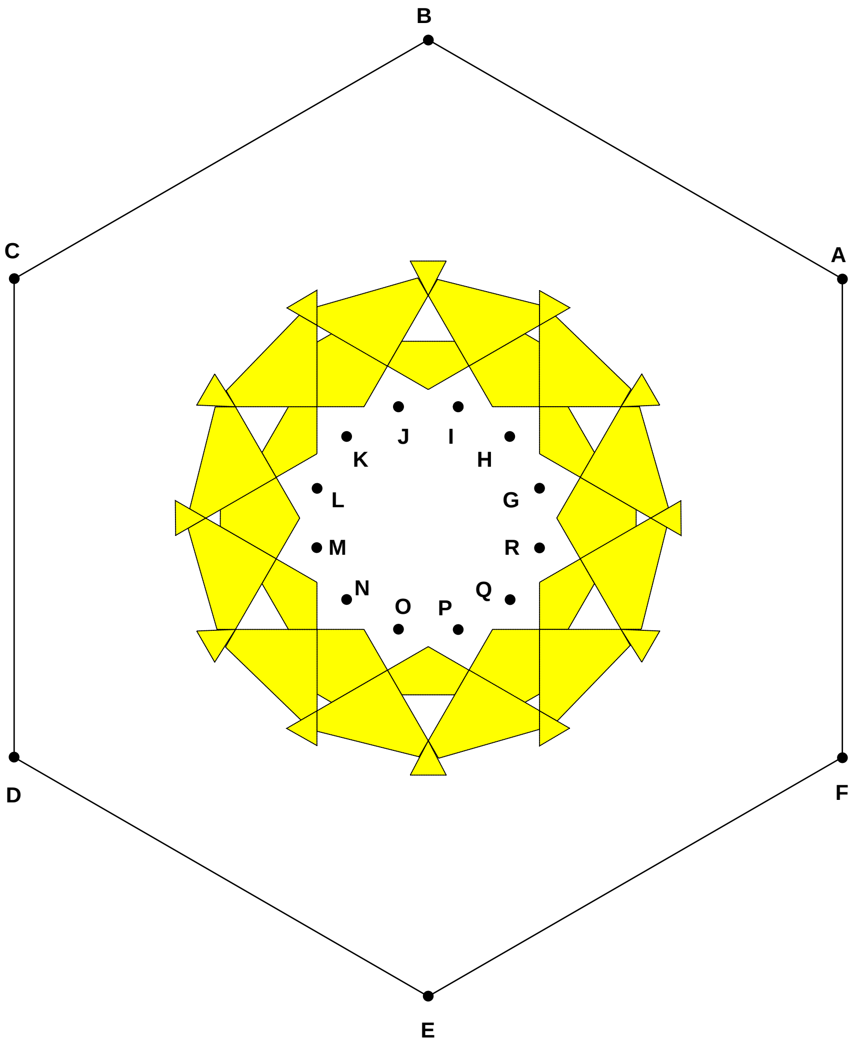

FIGURE 2. This Steiner reducing region has 13 holes.

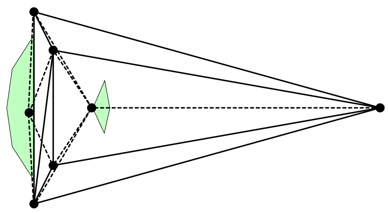

FIGURE 3. This 5-point set has a Steiner reducing region with 2 components.

Steiner reducing region can exist in close proximity to one another without forcing the points of the input set to be clustered far apart from one another: in Figure 4 we show a 15-point set that admits a Steiner reducing region with 20 disconnected components.

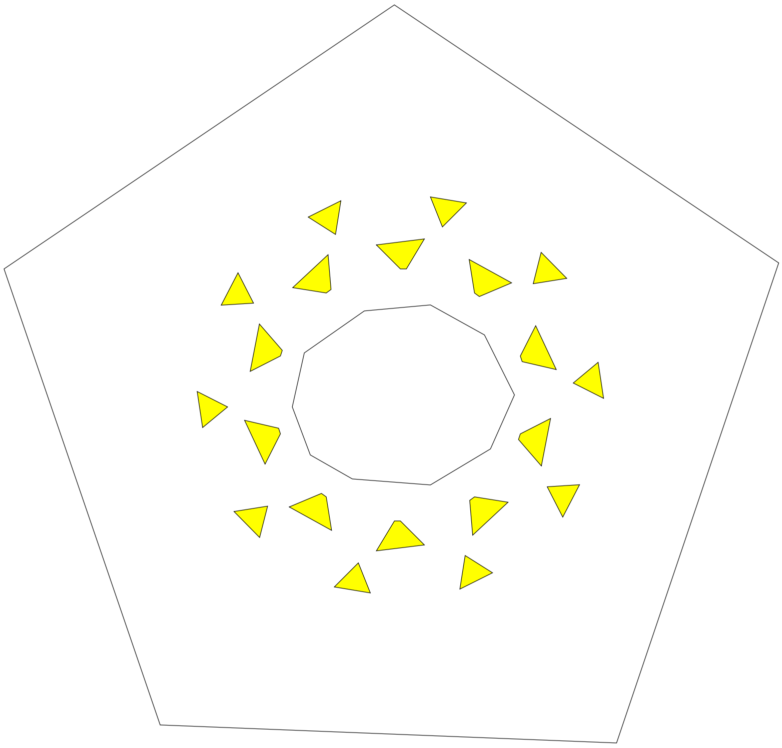

FIGURE 4. This (is a rough sketch of a) 15-point set (which) has a Steiner reducing region with 20 components.

The paper will be organized in the following manner. First, we explore the reasons why Steiner reducing regions exist. We explain why they need not be simply connected in our proof of Theorem 1. Next, we examine the Steiner reducing regions of some random point sets. Lastly, we prove that the Steiner reducing region in Figure 4 exists and is truly disconnected. We close the paper with some questions which remain unresolved.

2. Non-trivial first homology of Steiner reducing regions

The following set of conditions is sufficient to imply that Z is a Steiner reducing point for an input set $\mathcal{X}$:

(1) There exists a non-convex polygon $\mathcal{Q} = \{Q_1, \ldots, Q_k\}$, whose boundary edges are in $\mathrm{MWT}(\mathcal{X})$, and the segments ZQ_i lie in the interior of $\mathcal{Q}$ for all i.
(2) The following inequality holds: $\sum_{i=1}^{k} \mathrm{dist}(Z, Q_i) < L$, where L is the summed length of the interior diagonals which minimally triangulate $\mathcal{Q}$.

The reduction will occur by replacing the interior diagonals which triangulated $\mathcal{Q}$ in $\mathrm{MWT}(\mathcal{X})$ by the edge set $\{ZQ_i\}_{i=1}^{k}$. Condition 1 implies that the new triangulation will be valid. Condition 2 implies that the new triangulation has length less than

$mwt(\mathcal{X})$. Functions of the form $f(Z) = \sum_{i=1}^{k} \mathrm{dist}(Z, Q_i)$ are described as k-ellipses by Sekino in [21]. For a constant C and $k = 1$ or 2, the curve $f(Z) = C$ is a circle or an ellipse, respectively. Moreover, $f(Z) = C$ will be convex for any value of k. Denote by $E(C; Q_1, \ldots, Q_k)$ or $E(C; \mathcal{Q})$ the k-ellipse with foci $\mathcal{Q} = (Q_1, \ldots, Q_k)$ and focal length C. Thus $E(C; \mathcal{Q})$ is the locus of all points whose summed distance to the k fixed points of $\mathcal{Q}$ is C. These k-ellipses will often appear as part or all of the boundary of different Steiner reducing regions.

The Steiner reducing region $\mathcal{S}$ of a point set $\mathcal{X}$ can be written as the union of regions $\mathcal{S}_{\mathcal{Q}}$, where each $\mathcal{S}_{\mathcal{Q}}$ corresponds to a fixed combinatorial type of triangulation of $\mathcal{X} \cup \{Z\}$. It suffices to describe this combinatorial type by a listing $\mathcal{Q} = (Q_1, \ldots, Q_k)$ of the vertices adjacent to Z in the minimum weight Steiner triangulation, since all other edges come from the original MWT$(\mathcal{X})$. Define the *visibility region* $\mathsf{vis}_{\mathcal{Q}}$ corresponding to the possibly non-convex polygon $\mathcal{Q}$ as follows:

$$\mathsf{vis}_{\mathcal{Q}} := \{Z \in \mathbb{R}^2 | \mathrm{int}(ZQ_i) \subset \mathrm{int}(\mathcal{Q}) \text{ for } 1 \leq i \leq k\}.$$

Each $\mathcal{S}_{\mathcal{Q}}$ is the intersection of a visibility region $\mathsf{vis}_{\mathcal{Q}}$ with the interior of the k-ellipse $E\big(\ell_{\mathcal{Q}}; \mathcal{Q}\big)$, where $\ell_{\mathcal{Q}}$ is the sum of the length of diagonals which minimally triangulate $\mathcal{Q}$. The following lemma will allow us to piece together local Steiner reducing regions into the global Steiner reducing region $\mathcal{S}$.

Lemma 1. *If a convex polygon $\mathcal{M} = (M_1, \ldots, M_m)$ lies within the visibility region $\mathsf{vis}_{\mathcal{Q}}$ and M_i is a Steiner reducing point with combinatorial type $\mathcal{Q}$ for all i, then the convex hull of $\mathcal{M}$ is a subset of the Steiner reducing region. Any point $Z \in \mathcal{M}$ will be a Steiner reducing point with combinatorial type $\mathcal{Q}$.*

Proof: Since $\mathcal{M}$ lies within $\mathsf{vis}_{\mathcal{Q}}$, we know that connecting the point $Z \in \mathcal{M}$ to the vertices of $\mathcal{Q}$ will give a valid triangulation. Since each M_i is a Steiner reducing point, we know that $M_i \in E(\ell_{\mathcal{Q}}; \mathcal{Q})$ for all values of i. The convexity of all k-ellipses guarantees that the convex hull of the M_i's will also be contained in $E(\ell_{\mathcal{Q}}; \mathcal{Q})$. Thus every point in $\mathcal{M}$ will be a Steiner reducing point. ■

We can use Lemma 1 to identify polygonal subsets of the Steiner reducing region. By taking the union of these polygonal subsets, we can build approximations to the true Steiner reducing region. We can explicitly describe the Steiner reducing region by considering all polygons formed from the edges in MWT$(\mathcal{X})$.

2.1. **Proof of Theorem 1.** The point set $\mathcal{P}$ considered in Theorem 1 consists of a regular hexagon G_6 of radius 83 containing a smaller regular 12-gon G_{12} of radius 20, with both centered at the origin. We rotate each of these regular n-gons of radius r by an angle of $\frac{\pi}{n}$ from the standard n-gon construction which uses the point $(r, 0)$. Thus coordinates of $\mathcal{P}$ are:

$$G_6 = \left\{ \left(83 \cos \frac{2\pi(2k-1)}{12}, 83 \sin \frac{2\pi(2k-1)}{12} \right) \middle| j = 1..6 \right\}, \text{ and}$$
$$G_{12} = \left\{ \left(20 \cos \frac{2\pi(2k-1)}{24}, 20 \sin \frac{2\pi(2k-1)}{24} \right) \middle| k = 1..12 \right\}.$$

We label the points of G_6 by $A, \ldots, F$, for values of $j = 1..6$. We similarly label the points of G_{12} by $G, \ldots, R$, for values of $k = 1..12$. Notice that our point set is preserved under the standard group action of D_6, the dihedral group of order 12. The symmetries of this point set reduce the number of cases we much consider. The following sets of segments are orbits under the action of D_6, and therefore define equivalence classes based on length.

$$\Gamma := \{AG,\ AH,\ BI,\ BJ,\ CK,\ CL,\ DM,\ DN,\ EO,\ EP,\ FQ,\ FR\}$$

$$\Phi := \{AR,\ AI,\ BH,\ BK,\ CJ,\ CM,\ DL,\ DO,\ EN,\ EQ,\ FP,\ FG\}$$

All segments in Γ have length ≈ 63.8915, and segments in Φ have length ≈ 70.2951.

The minimum weight triangulation of $\mathcal{P}$ includes a minimum weight triangulation of the 12-gon formed by the points of G_{12}. (Notice that the edges of the 12-gon connect mutual nearest neighbors, so they must be in MWT($\mathcal{P}$) by [22].) Eighteen edges are needed to triangulate the region between the convex hull G_6 and the 12-gon G_{12}. A minimal triangulation of $\mathcal{P}$ includes all edges in the set Γ and one edge each from the following six pairs of edges: (AI, BH), (BK, CJ), (CM, DL), (DO, EN), (EQ, FP), (FG, AR). Even without regard to whether these edges form a triangulation (they do), no subset of 18 edges which lie in that region has smaller length. Thus, one example of a minimum weight triangulation of $\mathcal{P}$ uses the following edge set between the convex hulls of G_6 and G_{12}:

$$\Omega := \Gamma \cup \{AI, BK, CM, DO, EQ, FG\}.$$

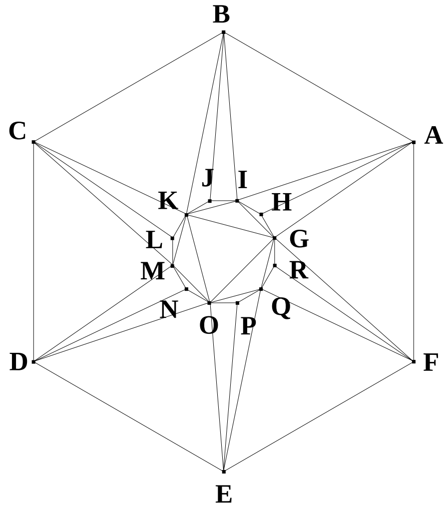

FIGURE 5. $\mathcal{P}$ triangulated minimally with the edges in Ω.

We now establish several polygonal subsets of the Steiner reducing region of $\mathcal{P}$, the union of which will be a connected planar region which we will demonstrate is not simply connected. There are five polygons, up to symmetry, which we will consider. These polygons lie in regions which are bounded by lines extended from the edges of the interior 12-gon.

Region 1 is the chamber bounded by lines HI, GH, KL, and JK. All points in the interior of region 1 are in $\mathsf{vis}_{\mathcal{A}}$ for the set $\mathcal{A} := \{A, B, C, H, I, J, K\}$. Define the pentagon $\mathcal{M}_1 := \operatorname{conv}\{a,b,c,d,e\}$, where

$$\begin{aligned} a &= HI \cap JK \approx (0, 22.30710), \\ b &= GH \cap JK \approx (7.07107, 26.38958), \\ c &= GH \cap (y = 30.675) \approx (4.59688, 30.675), \\ d &= (y = 30.675) \cap KL \approx (-4.59688, 30.675), \text{ and} \\ e &= HI \cap KL \approx (-7.07107, 26.38958). \end{aligned}$$

Then $\operatorname{int}(\mathcal{M}_1) \subset \mathsf{vis}_{\mathcal{A}}$. The edges ZA, ZB, ZC, ZH, ZI, ZJ, ZK will replace edges AI, BI, BJ, BK from the original triangulation, which have a summed length of $\ell_{\mathcal{A}} = 268.374$. Let $d_{\mathcal{A}}(Z)$ be the sum over points $P \in \mathcal{A}$ of the distance from P to Z. Then we have $d_{\mathcal{A}}(a) = 254.103$, $d_{\mathcal{A}}(b) = 264.081$, $d_{\mathcal{A}}(c) = 268.349$, $d_{\mathcal{A}}(d) = 268.349$, and $d_{\mathcal{A}}(e) = 264.081$. Since all five of the above values are less than $\ell_{\mathcal{A}}$, we know by Lemma 1 that any point added within $\mathcal{M}_1 \cap \mathsf{vis}_{\mathcal{A}}$ will reduce the length of the minimum weight triangulation. Note that this result holds for points in interior of the convex hull; the triangulation described is not valid if $Z \in \{a,b,c\}$. In such a case, at least one of the triangles $\triangle ZJK, \triangle ZHI$ will be degenerate. The open segments (d,e) and (b,c) are in $\mathsf{vis}_{\mathcal{A}}$, so Z will be a Steiner reducing point if placed in those intervals.

Region 2 is the chamber bounded by lines GH, IJ, GR, and JK. All points in the interior of region 2 are in $\mathsf{vis}_{\mathcal{B}}$ for $\mathcal{B} := \{A,B,G,H,I,J\}$. Define the pentagon $\mathcal{M}_2 := \operatorname{conv}\{b,f,g,h,i\}$, where

$$\begin{aligned} f &= GH \cap IJ \approx (11.15355, 19.31852), \\ g &= GR \cap IJ \approx (19.31852, 19.31852), \\ h &= GR \cap (y = -0.58307x + 41.77457) \approx (19.31852, 30.51051), \text{ and} \\ i &= JK \cap (y = -0.58307x + 41.77457) \approx (16.77621, 31.99285). \end{aligned}$$

Then $\operatorname{int}(\mathcal{M}_2) \subset \mathsf{vis}_{\mathcal{B}}$. The edges ZA, ZB, ZG, ZH, ZI, ZJ will replace edges AH, AI, BI from the original triangulation, which have a summed length of $\ell_{\mathcal{B}} = 198.079$. Let $d_{\mathcal{B}}(Z)$ be the sum over points $P \in \mathcal{B}$ of the distance from P to Z. Then we have $d_{\mathcal{B}}(b) = 183.697$, $d_{\mathcal{B}}(f) = 173.916$, $d_{\mathcal{B}}(g) = 183.697$, $d_{\mathcal{B}}(h) = 197.097$, and $d_{\mathcal{B}}(i) = 197.124$. Since all five of the above values are less than $\ell_{\mathcal{B}}$, Lemma 1 implies that $\mathcal{M}_2 \cap \mathsf{vis}_{\mathcal{B}} \subset \mathcal{S}_{\mathcal{B}}$, the subset of the Steiner reducing region which corresponds to combinatorial type $\mathcal{B}$. Points on segments $[b,f]$ and $[g,f]$ do not create valid triangulations under this combinatorial type, since those points are not in $\mathsf{vis}_{\mathcal{B}}$. However, the open segments (b,i) and (g,h) are in $\mathsf{vis}_{\mathcal{B}}$ and therefore $\mathcal{S}_{\mathcal{B}}$, as are the points of $\operatorname{int}(\mathcal{M}_2)$.

Region 3 is the chamber bounded by lines GH, KL, GR, and JK. All points in the interior of Region 3 are in $\mathsf{vis}_{\mathcal{C}}$ for $\mathcal{C} := \{A,B,G,H,I,J,K\}$. Let

$$\begin{aligned} j &= GR \cap JK \approx (19.31852, 33.46065) \\ k &= GR \cap (y = -0.24958x + 41.81550) \approx (19.31852, 36.99399), \\ l &= KL \cap (y = -0.24958x + 41.81550) \approx (1.60397, 41.41519), \text{ and} \\ m &= GH \cap KL \approx (0.00000, 38.63703). \end{aligned}$$

Combinatorial type $\mathcal{C}$ will cause the following edges from the original triangulation to be replaced: AH, AI, BI, BJ. The summed length of those four edges is $\ell_{\mathcal{C}} = 261.971$. Let $d_{\mathcal{C}}(Z)$ be the sum over points $P \in \mathcal{C}$ of the distance from P to Z. Then we have $d_{\mathcal{C}}(b) = 208.192$, $d_{\mathcal{C}}(j) = 241.551$, $d_{\mathcal{C}}(k) = 251.262$, $d_{\mathcal{C}}(l) = 259.236$, and

$d_{\mathcal{C}}(m) = 251.505$. Since all five of those values are less than $\ell_{\mathcal{C}}$, any point added within the convex hull of $\mathcal{M}_3$ will indeed reduce the length of the minimum weight triangulation since $\mathcal{M}_3 \cap \mathsf{vis}_{\mathcal{C}} \subset \mathcal{S}_{\mathcal{C}}$. The segments $[b, m]$ and $[b, j]$ are not in $\mathsf{vis}_{\mathcal{C}}$ so they will not make a valid triangulation with this particular combinatorial type. The open segments (l, m) and (j, k) are in $\mathcal{M}_3 \cap \mathsf{vis}_{\mathcal{C}}$, so they will be subsets of $\mathcal{S}_{\mathcal{C}}$.

Define region 4 to be the bounded chamber formed by lines KL, GH and convex hull edges AB, BC. All points in the interior of region 4 are in $\mathsf{vis}_{\mathcal{D}}$ for $\mathcal{D} := \{A, B, C, G, H, I, J, K, L\}$. Define the triangle $\mathcal{M}_4 := \mathrm{conv}\{m, n, o\}$, where

$$\begin{aligned} n &= KL \cap (y = 44.6) \approx (3.12136, 44.6), \text{ and} \\ o &= GH \cap (y = 44.6) \approx (-3.12136, 44.6). \end{aligned}$$

Combinatorial type $\mathcal{D}$ replaces the following edges from the original triangulation: AH, AI, BI, BJ, BK, CK. The summed length of those six edges is $\ell_{\mathcal{D}} = 396.158$. Let $d_{\mathcal{D}}(Z)$ be the sum over points $P \in \mathcal{D}$ of the distance from P to Z. Then we have $d_{\mathcal{D}}(m) = 362.079$, $d_{\mathcal{D}}(n) = 389.779$, and $d_{\mathcal{D}}(o) = 389.779$. Since all three of the above values are less than $\ell_{\mathcal{D}}$, we have $\mathcal{M}_4 \cap \mathsf{vis}_{\mathcal{D}} \subset \mathcal{S}_{\mathcal{D}}$. Any point added to $\mathrm{int}(\mathcal{M}_4)$ will reduce the length of the minimum weight triangulation. Points on segments $[m, o]$ and $[m, n]$ do not create valid triangulations under combinatorial type $\mathcal{D}$, since those intervals are not in $\mathsf{vis}_{\mathcal{D}}$.

Region 5 is the chamber bounded by lines JK, GR and convex hull edge AB. All points in the interior of region 5 are in $\mathsf{vis}_{\mathcal{E}}$ for $\mathcal{E} := \{A, B, G, H, I, J, K, R\}$. Define the triangle $\mathcal{M}_5 := \mathrm{conv}\{j, p, q\}$, where

$$\begin{aligned} p &= JK \cap (y = -0.56463x + 50.38075) \approx (24.58335, 36.50030), \text{ and} \\ q &= GR \cap (y = -0.56463x + 50.38075) \approx (19.31852, 39.47297). \end{aligned}$$

The following edges will be replaced from the original triangulation: AG, AH, AI, BI, BJ. The summed length of those five edges is $\ell_{\mathcal{E}} = 325.863$. Let $d_{\mathcal{E}}(Z)$ be the sum over points $P \in \mathcal{E}$ of the distance from P to Z. Then $d_{\mathcal{E}}(j) = 280.188$, $d_{\mathcal{E}}(p) = 303.605$, and $d_{\mathcal{E}}(q) = 303.332$. Since all three of the above values are less than $\ell_{\mathcal{E}}$, any $Z \in \mathcal{M}_5 \cap \mathsf{vis}_{\mathcal{E}}$ will be a Steiner reducing point. Points on segments $[j, q]$ and $[j, p]$ do not create valid triangulations under combinatorial type $\mathcal{E}$, since they are not in $\mathsf{vis}_{\mathcal{E}}$.

The interior of polygons $\mathcal{M}_i$ for $i = 1..5$ are subsets of the Steiner reducing region. Moreover, the collection of Steiner reducing boundary segments will glue these polygons together into one connected region. Segment (b, c) connects $\mathcal{M}_1$ to $\mathcal{M}_3$, segment (b, i) connects $\mathcal{M}_2$ to $\mathcal{M}_3$, segment (l, m) connects $\mathcal{M}_3$ to $\mathcal{M}_4$, and segment (j, k) connects $\mathcal{M}_3$ to $\mathcal{M}_5$. Segments (d, e) and (g, h) provide the remaining necessary glue for symmetric images of $\mathcal{M}_3$ to connect to the other sides of $\mathcal{M}_1$ and $\mathcal{M}_2$. It follows that the subset of the Steiner reducing region established so far is connected. To finish the proof of Theorem 1, we must prove the existence of 13 holes within this reducing region. We will do so by finding points in the interior of the holes that do not reduce.

Claim 1. *None of the points* $X = (0, 35.08709)$, $Y = (18.47521, 32)$, *or* $W = (0, 0)$ *will reduce.*

Proof: A minimum weight triangulation of $\mathcal{P} \cup \{X\}$ has combinatorial type $\mathcal{F} := \{B, H, I, J, K\}$, and $d_{\mathcal{F}}(X) = 154.2164$. Combinatorial type $\mathcal{F}$ replaces edges BI and BJ, which have a summed length of 127.78. Therefore X is not a Steiner reducing point.

A minimum weight triangulation of $\mathcal{P} \cup \{Y\}$ has combinatorial type $\mathcal{B} = \{A, B, G, H, I, J\}$. We note that $d_{\mathcal{B}}(Y) = 198.912$, which is greater than the summed lengths of edges BI, AH, and AI : 198.079. It follows that Y is not a Steiner reducing point.

A minimum weight triangulation of $\mathcal{P} \cup \{W\}$ has combinatorial type $\mathcal{G} := \{G, I, K, M, O, Q\}$, and replaces edges GK, KO, and GO. The triangle inequality implies that W is not a Steiner reducing point. ■

Since none of the points X, Y, W were on the boundary of $\mathsf{vis}_{\mathcal{Q}}$ for their respective combinatorial types $\mathcal{Q}$, there is a non-reducing neighborhood around each of the points. The symmetry of our input set $\mathcal{P}$ implies that X and Y are each in one of six symmetric non-reducing holes. The thirteenth hole comes from W. Thus Theorem 1 is proven.

We note that the high degree of symmetry in this point set was only utilized to shorten the above proof. By slightly deforming our original input set $\mathcal{P}$, we can find a point set in general position with no symmetries that also will have 13 holes in its Steiner reducing region. It is likely that similarly constructed point sets will give Steiner reducing regions with larger numbers of holes. An upper limit on the number of holes that can be obtained with a construction of this type, if the limit exists, is not known.

3. Connectivity of Steiner reducing regions

An examination of 30 uniformly distributed sets of 10 points revealed k_1 input sets with a non-empty Steiner reducing region. Of these, k_2 had Steiner reducing regions with more than one component, and k_3 admitted Steiner reducing points interior to their convex hulls.

Increasing the number of points in the input set to 20 gave the following results: (list them here)

Tests of 30-point sets indicate that the complexity of the topology continues to grow with the size of the input: (list results here)

The number of disconnected components of a Steiner reducing region can be larger than the number of components of the input set. Consider the point set $\mathcal{N}$ consisting of a regular pentagon G_5 of radius 32 containing a smaller regular 10-gon G_{10} of radius 8. Explicitly, we again rotate each of the regular n-gons by an angle of $\frac{\pi}{n}$ from the standard n-gon construction. The coordinates of $\mathcal{N}$ are:

$$G_5 = \left\{ \left(32 \cos \frac{2\pi(2j-1)}{10}, \, 32 \sin \frac{2\pi(2j-1)}{10} \right) \middle| j = 1..5 \right\}$$

$$G_{10} = \left\{ \left(8 \cos \frac{2\pi(2k-1)}{20}, \, 8 \sin \frac{2\pi(2k-1)}{20} \right) \middle| k = 1..10 \right\}$$

We label the points of G_5 by $A, \ldots, E$, for values of $j = 1..5$. We similarly label the points of G_{10} by $F, \ldots, O$, for values of $k = 1..10$. We note that the dihedral group of order 10, D_5, will act on the point set $\mathcal{N}$ and create many symmetries which we shall exploit in the course of our proof.

As was the case for the similarly constructed $\mathcal{P}$ in the previous section, any minimum weight triangulation of $\mathcal{N}$ will contain a minimum weight triangulation of G_{10}. Moreover, any minimum weight triangulation of $\mathcal{N}$ contains the edges in the set $\{AF, AG, BH, BI, CJ, CK, DL, DM, EN, EO\}$, plus one edge each from the following five pairs of edges: (AH, BG), (BJ, CI), (CL, DK), (DN, EM), and (AO, EF).

Claim 2. *The point set $\mathcal{N} = G_5 \cup G_{10}$ described above has 20 disconnected Steiner reducing regions within its convex hull.*

Proof: There are two triangular subsets of the Steiner reducing region of $\mathcal{N}$, the union of the symmetric images of which will be a disconnected planar region with 20 components. These triangles lie in regions which are bounded by lines extended from the edges of the interior 10-gon.

Region 1 is bounded by lines FG, FO, HI, and IJ. Points in the interior of region 1 are in $\mathsf{vis}_{\mathcal{J}}$ for $\mathcal{J} := \{A, B, F, G, H, I\}$.

$\because$ (continue proof in style of thm 1)

Region 2 is bounded by lines KL, NO, and convex hull edges CD and DE. Points in the interior of region 2 are in $\mathsf{vis}_{\mathcal{K}}$ for $\mathcal{K} := \{D, E, K, L, M, N, O\}$.

$\because$ (continue proof in style of thm 1) ∎

(insert concluding paragraph here - why these results are interesting)

4. Unresolved questions

- Can the number of holes in a connected Steiner reducing region of an n-point set grow without bound as n increases?
- Can the number of disconnected components of the Steiner reducing region of an n-point set grow with n at a rate which is faster than linear?
- Is there a convex polygon which admits an interior Steiner reducing point?
- (The teamwork question.) Are there finite point sets $\mathcal{X}, \mathcal{R}$ such that
$$mwt(\mathcal{X} \cup \mathcal{R}) < mwt(\mathcal{X}),$$
but we have
$$mwt(\mathcal{X} \cup \mathcal{Q}) > mwt(\mathcal{X})$$
for every possible set $\mathcal{Q}$ of smaller cardinality than $\mathcal{R}$? (A related question: can points work together to create reductions none of their subsets could produce?)
- As a function of n, what is the expected number of components in the Steiner reducing region of a uniformly distributed set of n points? What is the expected number of interior components?

References

[1] O. Aichholzer, F. Aurenhammer, and R. Hainz. New results on MWT subgraphs. *Inform. Process. Lett.*, 69(5):215–219, 1999.

[2] O. Aichholzer, F. Aurenhammer, G. Rote, and Y.-F. Xu. Constant-level greedy triangulations approximate the MWT well. *J. Comb. Optim.*, 2(4):361–369, 1999.

[3] B. Aronov and S. Fortune. Approximating minimum-weight triangulations in three dimensions. *Discrete Comput. Geom.*, 21(4):527–549, 1999.

[4] M. Bern and D. Eppstein. Mesh generation and optimal triangulation. In *Computing in Euclidean geometry*, volume 1 of *Lecture Notes Ser. Comput.*, pages 23–90. World Sci. Publishing, River Edge, NJ, 1992.

[5] S.-W. Cheng and T. K. Dey. Approximate minimum weight Steiner triangulation in three dimensions. In *Proceedings of the Tenth Annual ACM-SIAM Symposium on Discrete Algorithms (Baltimore, MD, 1999)*, pages 205–214, New York, 1999. ACM.

[6] S.-W. Cheng and Y.-F. Xu. On β-skeleton as a subgraph of the minimum weight triangulation. *Theoret. Comput. Sci.*, 262(1-2):459–471, 2001.

[7] J. A. De Loera, S. Hoşten, F. Santos, and B. Sturmfels. The polytope of all triangulations of a point configuration. *Doc. Math.*, 1:No. 04, 103–119 (electronic), 1996.

[8] J. A. De Loera, J. Rambau, and F. Santos. *Triangulations: Applications, Structures, Algorithms.* To appear.

[9] M. T. Dickerson, J. M. Keil, and M. H. Montague. A large subgraph of the minimum weight triangulation. *Discrete Comput. Geom.*, 18(3):289–304, 1997. ACM Symposium on Computational Geometry (Philadelphia, PA, 1996).

[10] D. Eppstein. Approximating the minimum weight triangulation. In *Proceedings of the Third Annual ACM-SIAM Symposium on Discrete Algorithms (Orlando, FL, 1992)*, pages 48–57, New York, 1992. ACM.

[11] D. Eppstein. Approximating the minimum weight Steiner triangulation. *Discrete Comput. Geom.*, 11(2):163–191, 1994.

[12] M. R. Garey and D. S. Johnson. *Computers and intractability: A guide to the theory of NP-completeness.* W. H. Freeman and Co., San Francisco, Calif., 1979.

[13] P. Gilbert. New results in planar triangulations. Master's thesis, University of Illinois, 1979.

[14] J. M. Keil. Computing a subgraph of the minimum weight triangulation. *Comput. Geom.*, 4(1):13–26, 1994.

[15] G. T. Klincsek. Minimal triangulations of polygonal domains. *Ann. Discrete Math.*, 9:121–123, 1980.

[16] C. Levcopoulos and D. Krznaric. A linear-time approximation scheme for minimum weight triangulation of convex polygons. *Algorithmica*, 21(3):285–311, 1998.

[17] C. Levcopoulos and D. Krznaric. Quasi-greedy triangulations approximating the minimum weight triangulation. *J. Algorithms*, 27(2):303–338, 1998. 7th Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996).

[18] C. Levcopoulos and A. Lingas. On approximation behavior of the greedy triangulation for convex polygons. *Algorithmica*, 2(2):175–193, 1987.

[19] C. Levcopoulos and A. Lingas. Greedy triangulation approximates the optimum and can be implemented in linear time in the average case. In *Advances in computing and information—ICCI '91 (Ottawa, ON, 1991)*, volume 497 of *Lecture Notes in Comput. Sci.*, pages 139–148. Springer, Berlin, 1991.

[20] C. Levcopoulos, A. Lingas, and J.-R. Sack. Heuristics for optimum binary search trees and minimum weight triangulation problems. *Theoret. Comput. Sci.*, 66(2):181–203, 1989.

[21] J. Sekino. n-ellipses and the minimum distance sum problem. *Amer. Math. Monthly*, 106(3):193–202, 1999.

[22] B. T. Yang, Y. F. Xu, and Z. Y. You. A chain decomposition algorithm for the proof of a property on minimum weight triangulations. In *Algorithms and computation (Beijing, 1994)*, volume 834 of *Lecture Notes in Comput. Sci.*, pages 423–427. Springer, Berlin, 1994.