Steiner Reducing Sets of Minimum Weight Triangulations

Cynthia M. Traub*

Abstract

This paper develops techniques for computing the minimum weight Steiner triangulation of a planar point set. We call a Steiner point P a *Steiner reducing point* of a planar point set X if the weight (sum of edge lengths) of a minimum weight triangulation of $X \cup \{P\}$ is less than that of X. We define the *Steiner reducing set* $St(X)$ to be the collection of all Steiner reducing points of X. We provide here necessary conditions for membership in the Steiner reducing set. We prove that $St(X)$ can be topologically complex, containing multiple connected components or even holes. We construct families of sets X for which the number of connected components of $St(X)$ grows linearly in the cardinality of X. We further prove that $St(X)$ need not be simply connected, and the rank of $H_1(St(X))$ (i.e. the number of holes) can also grow linearly in the cardinality of X.

1 Introduction

We consider minimum weight triangulations of point sets that properly contain an initial input set X of n points. We examine the topology of the collection of Steiner points distinguished by the property that adding one such point to the input set results in a minimum weight triangulation of a set of $n+1$ points with weight less than that of X. We call these distinguished Steiner points *Steiner reducing points*; the collection of all Steiner reducing points is called the *Steiner reducing set* $St(X)$. We present necessary conditions for a point to be a Steiner reducing point. Our two main results prove that the number of connected components of $St(X)$ can grow linearly in the size of the input set, and $St(X)$ may fail to be simply connected, as illustrated in Figure 1. In Section 2 we will present two necessary conditions for constructing Steiner reducing sets. The topology of Steiner reducing sets is studied in Section 3.

1.1 Definitions and Techniques

A *triangulation* of a finite set X of points in $\mathbb{R}^2$ is an inclusion-maximal set of non-intersecting straight line segments between pairs of points in X. The *weight* of a triangulation is defined as the sum of the Euclidean lengths of its line segments, hence, a *minimum weight triangulation* of X is a triangulation (not necessarily unique) with weight less than or equal to that of any other triangulation of X. We denote the weight of a minimum weight triangulation of a point set X by $\mathrm{mwt}(X)$. An edge AB is *unavoidable* if every triangulation of X contains the edge AB. The edges of the boundary of the convex hull are all unavoidable, as are any edges interior to the convex hull which are not properly intersected by any other possible edge. We direct the reader to a standard topology book [9] for formal definitions of the topological terms used in this paper. For our purposes here, the rank of $H_0(Z)$ counts the number of connected components in the point set Z, and the rank of $H_1(Z)$ counts the number of holes or handles in a set Z. A path-connected set without holes is called *simply connected*.

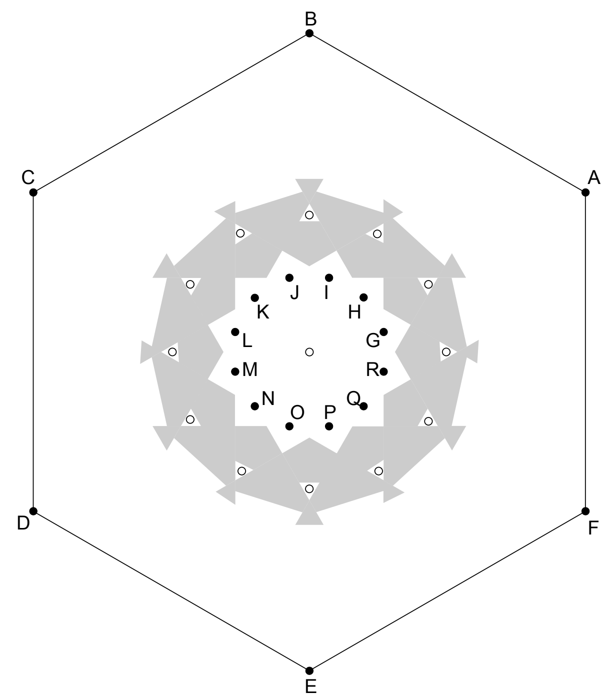

Figure 1: The Steiner reducing set of the point set whose minimum weight triangulation is illustrated in Figure 7; empty circles represent points that are not Steiner reducing points

Calculations of minimum weight triangulations given in this paper were made using integer programming over the universal polytope, or for small examples, verified

*Department of Mathematics and Statistics, Southern Illinois University Edwardsville, cytraub@siue.edu

by hand. The program universalbuilder written by De Loera and Peterson was used to generate a system of inequalities defining the universal polytope of each input set X, which was then input to the optimization engine CPLEX. This process found the weight of the minimum weight triangulation as well as the triangles used. A Bash shell script was used to test the effects of adding an additional point at each location in an $n \times n$ grid. In order to show that a minimal weight triangulation has a specific structure for each element of a 1- or 2-dimensional set of Steiner points, we appeal to the $\beta-$skeleton, unavoidable edges, and the triangle inequality, as well as the interaction between certain curves called $k-$ellipses and the chambers of the hyperplane arrangement induced by specific pairs of points.

1.2 Background and Previous Work

The minimum weight triangulation decision problem "Given a finite planar point set X and a positive integer b, is there a triangulation of X with weight b or less?" has been a problem of theoretical and computational interest for over thirty years [6]. The problem was shown to be NP-hard in 2006 [12], yet determination of minimum weight triangulations of special classes of point sets, such as polygonal domains, can be done in polynomial time [7, 11]. The shortest edge [7] and specific subsets of edges such as the β-skeleton [2, 10] have been proven to belong to all minimum weight triangulations. The simultaneous addition of many Steiner points to certain point sets can reduce triangulation weight by a significant amount, as shown in [4]. Practical motivations for the study of Steiner points include improving the ability of meshes to approximate fine details [1], and allowing approximation of the minimum weight triangulation of an input point set [4]. Our work presented here is the first to define and study the shape of the Steiner reducing set.

In order to find all edges of a minimum weight triangulation (after perhaps identifying a subset of the edges), one algorithmic approach applies integer programming to the universal polytope. This polytope has vertices corresponding to each triangulation of X. Thus, even though the minimum weight triangulation problem is known to be NP-hard, there are algorithmic means for finding the minimum weight triangulation of planar point sets of up to several hundred points that do not rely on ad-hoc methods [3]. Note that no polynomial time algorithm is known to verify that a proposed triangulation is indeed minimal.

In this paper we reveal topological complexity behind the minimum weight triangulation problem by giving conditions that allow the addition of a Steiner point to a fixed input set to reduce the weight of a minimum weight triangulation. This answers a question posed by Jesús De Loera in 2003 during the MSRI Summer Graduate Workshop on Triangulations of Point Sets.

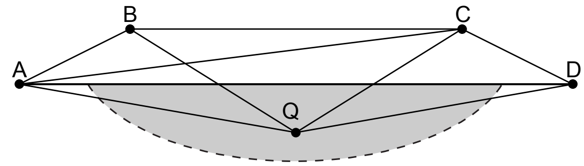

Figure 2: The shaded area is the Steiner reducing set of $Y = \{A, B, C, D\}$, shown with Steiner reducing point Q

2 Steiner Reducing Sets: Necessary Conditions

In this section, we give a means for identifying subsets of the Steiner reducing set via the combinatorial structure of a minimal triangulation and the geometric properties of an associated multi-focal ellipse (known as a k-ellipse). To our knowledge, all known examples of Steiner reducing points are located either outside the interior of the convex hull of X (as in Figure 2), or in the interior of a non-convex polygon formed by edges that belong to a minimum weight triangulation of X (as in Figure 5). The non-convexity of such a polygon will lead to Steiner reducing sets with interesting topology. It remains an open question whether there exist convex polygons which admit Steiner reducing points in the interior of their convex hulls.

It is simple to show that no set of three points admits a Steiner reducing point. Case analysis of four-point sets reveals that no such set will admit a Steiner reducing point in the interior of its convex hull. There are four-point sets, however, which allow Steiner reducing points exterior to and on the edges of their convex hulls. One illustration of such a set is shown in Figure 2.

The minimum weight triangulation of $Y \cup \{Q\}$ for $Q = (x, y)$, $y < 0$, will use edges AQ, BQ, CQ, DQ, when Q is "close" to Y, in addition to edges AB, BC, CD, as seen in Figure 2. Since the new edges replace edges AC and AD from the original triangulation, the Steiner reducing set $St(Y)$ consists of all points Q which satisfy the inequality
$|AQ| + |BQ| + |CQ| + |DQ| < |AC| + |AD| = 10 + \sqrt{65}$.

The curved boundary of $St(Y)$ is part of a curve known as a 4-ellipse. More generally, the locus of all points $P \in \mathbb{R}^2$ such that the sum of distances from P to each of k distinguished foci is constant is called a *k-ellipse*. Under this definition, the circle is a 1-ellipse, and the standard ellipse is a 2-ellipse. Let

$$E(\mathcal{Q}; d) = \left\{ P \in \mathbb{R}^2 : \sum_{i=1}^{k} |PQ_i| = d \right\}$$

denote the *k-ellipse* with *foci* in $\mathcal{Q} = \{Q_1, \ldots, Q_k\}$ and corresponding *distance sum* d. Note that $E(\mathcal{Q}; d)$ is a closed curve. For $k > 2$, the interior of this level set does not necessarily contain its foci Q_i, though the curve represented by the level set will be convex for every value of $k \geq 1$, as proven in [15]. Denote by $E_<(\mathcal{Q}; d)$ the points in its interior of the k-ellipse.

The history of the k-ellipse reaches back to Fermat, who posed the following challenge: "Given three points in a plane, find a fourth point such that the sum of its distances to the three given points is as small as possible." (The smallest possible 3-ellipse consists of this single point.) A solution to the problem of Fermat was provided by Evangelista Torricelli around 1640, and the distance minimizing point, termed the Fermat-Torricelli point [8], remains a topic of active research. In the late 1600's, Tschirnhaus generalized the standard string and pins construction of an ellipse, illustrating how to draw a 3-ellipse by hand in [18]. Due to a lack of standardized nomenclature for this object, it has been rediscovered many times throughout the literature, receiving such names as Tschirnhaussche Eiflächen [16], W_n curves [5], polyellipses [19], and egglipses [14]. A survey article of k-ellipses and their basic properties appears in [15]. Noting that the 1- and 2-ellipse both appear as the level-zero set of degree two polynomials, Nie, Parillo, and Sturmfels used semidefinite programming to show that the k-ellipse appears as part of the level zero set of a polynomial of degree 2^k if k is odd and degree $2^k - \binom{k}{k/2}$ if k is even [13]. The author's Ph.D. thesis [17] was the first work detailing the connection between k-ellipses and minimum weight triangulations.

The k-ellipse provides a criterion for measuring proximity to multiple distinct points. For the Steiner reducing set to be non-empty, the sum of lengths of edges incident to new Steiner point(s) must be less than the sum of the lengths of the replaced edges minus the lengths of new edges not incident to the Steiner point. Although it is common for k-ellipses to appear as part or all of the boundary of different Steiner reducing sets, these Steiner reducing sets are not simply unions of sets $E_<\big(X_i; d_i\big)$ for various pairs X_i, d_i of foci and distance sums. This is an issue of *feasibility*; locations of Q interior to quadrilateral $ABCD$ lie within the necessary 4-ellipse, but do not give a complete triangulation. For Q to be a Steiner reducing point, it must lie in the intersection of the 4-ellipse with the appropriate feasibility set, as shown in Figure 3. The intersection of these two sets gives the Steiner reducing set of Y.

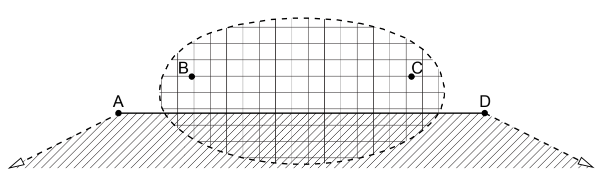

Figure 3: The intersection of $E_<(Y; 10 + \sqrt{65})$ with $F\big(Y, \Pi\big)$

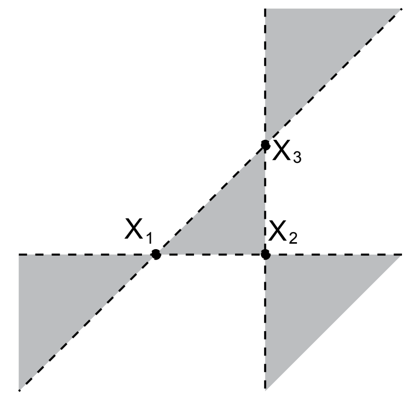

Figure 4: Any triangulation of $X \cup \{V\}$ for V within the shaded feasibility set uses the same set of edges.

We now formally define what we mean by feasibility.

Definition 1 *Given a set $X \subset \mathbb{R}^2$ of n distinct points, an abstract (yet to be located) vertex $V \notin X$, and a graph $G = (X \cup \{V\}, \Pi)$, the* feasibility set $F(X, \Pi)$ *is the set of points $P \in \mathbb{R}^2 \backslash X$ such that if $V = P$, then the straight line drawing of the graph G is a triangulation of $X \cup \{P\}$.*

In general, a feasibility set $F(X, \Pi)$ is comprised of unions of chambers (together with some boundary segments or rays) of the line arrangement formed by extending the segments of Π between pairs of points of X into lines. Thus, each feasibility set can be described via linear inequalities. Feasibility sets need not be connected, as can be shown by taking X to be a set of three non-collinear points as in Figure 4.

The intersections of feasibility sets with k-ellipses are the basic building blocks of Steiner reducing sets. In the next section we will explore the geometry and topology which occurs as we take unions of such intersections.

3 Topological Properties of Steiner Reducing Sets

3.1 Connectivity

Our first contribution to the study of the topology of Steiner reducing sets is to show that the sets need not be connected. This can be demonstrated with a point set with as few as five points, as shown in Figure 5.

Theorem 2 *There exist sets X of $5n$ points such that the rank of $H_0\big(St(X)\big)$ is at least $2n$.*

Proof. Let $Z = \{(0,0), (2,1), (8,1), (10,0), (5,18)\}$, denoted by A, B, C, D, E, respectively. We will prove that a quadrilateral of non-Steiner reducing points encircles the component of the Steiner reducing set of Z

that lies in the interior of the convex hull. Note first that edges AB, BC, CD, BE and CE are all unavoidable in a triangulation of Z since no other possible triangulation edges intersect these. It follows that Z has exactly two triangulations, and both are minimal. Choose $\mathrm{MWT}(Z)$ to be the minimum weight triangulation of Z that uses edge AC.

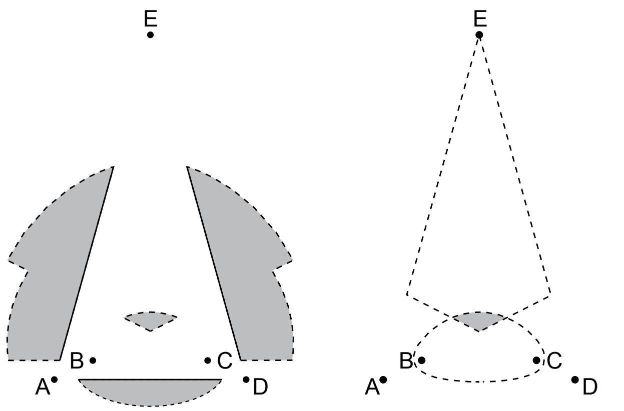

Figure 5: At left, a Steiner reducing set with four connected components; at right, the boundary of the interior component is formed by intersecting a feasibility set and a 5-ellipse

Since the minimum weight triangulation of Z contains the same configuration of edges as the minimum weight triangulation of Y from Figure 2, we know that certain points on or below segment AD will be Steiner reducing points. However, a quick calculation shows $P = (5, 3)$ is also a Steiner reducing point.

Let S_P be the connected component of the Steiner reducing set that contains P. Then $\mathsf{S}_P = E_<(Z; 2\sqrt{298}) \cap F(Z, \Phi)$, where $2\sqrt{298}$ is the length of the edges replaced from $\mathrm{MWT}(Z)$, and the feasibility set is the collection of points within the convex hull that can be connected to all five points of Z without intersecting any edges of a minimal triangulation of $Y = \{A, B, C, D\}$.

In order to establish that the subset S_P of the Steiner reducing set is not path connected to $St(Y)$, consider quadrilateral $EFGH$, where $F = \overleftrightarrow{BE} \cap \overleftrightarrow{CD}, G = \overleftrightarrow{AB} \cap \overleftrightarrow{CD}$, and $H = \overleftrightarrow{CE} \cap \overleftrightarrow{AB}$. This shape surrounds S_P, and its edges contain no Steiner reducing points, a fact which we now prove. To verify that no points of FG or GH are Steiner reducing points, we consider without loss of generality points $R = (x, \frac{1}{2}x)$ within the convex hull of Z and for $x \geq 5$. Such points R lie on the ray $\overrightarrow{GH}$, which makes edges BE, AB, BC, and CD all unavoidable in any triangulation of $Z \cup \{R\}$. Since the triangle inequality implies $|CR| + |ER| \geq |CE|$, the point R is not a Steiner reducing point. A similar argument shows that no point Q on segments FE or EH will be a Steiner reducing point. Therefore there are at least two connected components in $St(Z)$.

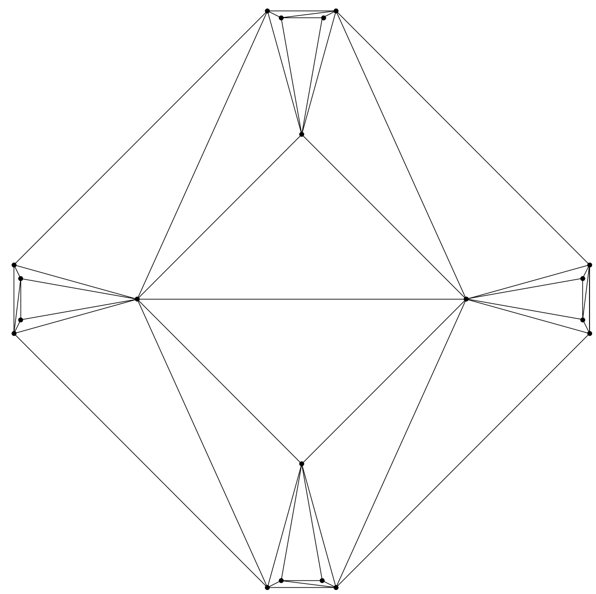

Figure 6: $St\left(Z^{[4]}\right)$ has at least 12 connected components

To prove that $St(Z)$ has four connected components, we verify that the three components outside the convex hull in Figure 5 are not path connected to one another. This can be done, as above, by showing no points of lines $y = 1/2$, $y = 16$ are Steiner reducing points.

By arranging four copies of the point set Z as shown in Figure 6, we can construct a point set with a Steiner reducing set that has at least 8 connected components. The left and right exterior components of the Steiner reducing set from each individual copy of Z are lost, but the the other two components from each copy of Z remain. For this particular example, the midpoint of each boundary edge is a Steiner reducing point, each belonging to its own connected component. By arranging copies of Z so that the images of E from the original set lie on the vertices of an n-gon with sufficiently long edges (relative to the size of Z), we can get a point set denoted $Z^{[n]}$ whose minimum weight triangulation contains all edges of the minimum weight triangulation of each copy of Z. Thus, $X = Z^{[n]}$ is a point set with $5n$ points that has at least $2n$ connected components in its Steiner reducing set. $\square$

In the previous example, we had $Y \subseteq Z$ and the corresponding Steiner reducing sets $St(Y) \subseteq St(Z)$. We note that in general, $X \subseteq W$ does not imply either $\mathrm{MWT}(X) \subseteq \mathrm{MWT}(W)$ or $St(X) \subseteq St(W)$. Even if we have containment of both the point sets $X \subseteq W$ and the edge sets $\mathrm{MWT}(X) \subseteq \mathrm{MWT}(W)$, we may not have $St(X) \subseteq St(W)$. We further note that it is possible for the number of connected components of the Steiner reducing set to exceed the number of points in the set.

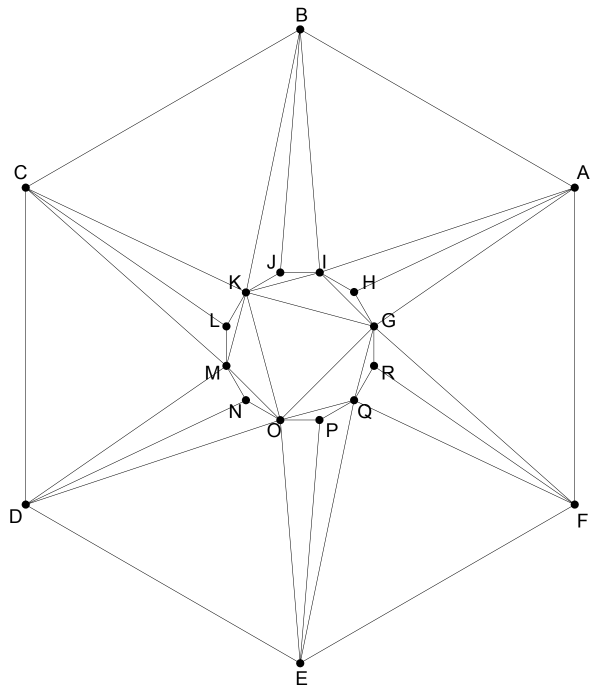

Figure 7: A minimum weight triangulation of Y using the edge set $\mathrm{MWT}(Y)$

An example of a 15-point set X that admits a Steiner reducing set $St(X)$ such that $H_0\big(St(X)\big)$ has rank at least 20 is given in the author's Ph.D. thesis [17].

3.2 Simple Connectivity

Our second topological result is to show that Steiner reducing sets need not be simply connected.

Theorem 3 *There exists a set Y of 18 points such that the rank of $H_1\big(St(Y)\big)$ is at least 13.*

We first describe the structure of a minimum weight triangulation of the point set Y under consideration, then find a connected subset of the Steiner reducing set of Y. We prove this subset is not simply connected by demonstrating 13 curves that lie in the Steiner reducing set of Y and are representatives of linearly independent homology classes within $H_1(St(Y))$.

Proof. Let $Y = G_6 \cup G_{12}$, where

$$G_6 = \left\{ \left(83\cos\frac{\sigma_j}{12},\, 83\sin\frac{\sigma_j}{12} \right) \middle| j = 1, \ldots, 6 \right\}, \text{ and}$$

$$G_{12} = \left\{ \left(20\cos\frac{\sigma_j}{24},\, 20\sin\frac{\sigma_j}{24} \right) \middle| j = 1, \ldots, 12 \right\},$$

where $\sigma_j = 2\pi(2j-1)$. An illustration of Y together with a minimal weight triangulation is given in Figure 7.

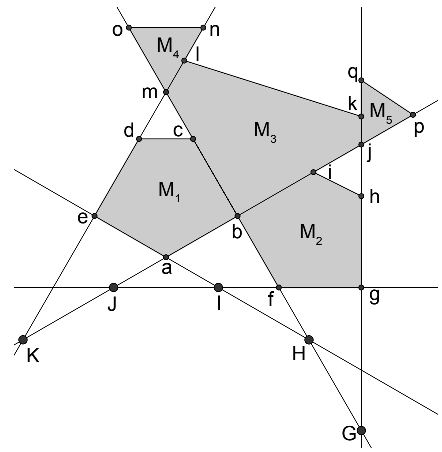

Figure 8: Five subsets M_i of the Steiner reducing set of Y

The set Y is preserved under the standard group action of D_6, the dihedral group of order 12. Let Γ be the orbit of AG under the induced action of D_6 on the edges. The edges of the interior 12-gon formed from the points in G_{12} belong to the β-skeleton of Y, and thus will belong to every minimum weight triangulation of Y. The edges in $\Gamma \cup \{AI, BK, CM, DO, EQ, FG\}$ triangulate the region between the convex hulls of G_6 and G_{12}. No other subset of 18 edges which lie in the annular region between G_6 and G_{12} has smaller weight, so these edges belong to a minimum weight triangulation of Y. Denote by $\mathrm{MWT}(Y)$ a fixed minimum weight triangulation of Y which uses these 18 edges.

Let $\mathcal{H}$ be the line arrangement formed by extending into lines the segments that form the boundaries of $\mathrm{conv}(G_6)$ and $\mathrm{conv}(G_{12})$. In Figure 8 we illustrate five polygons $M_i, 1 \leq i \leq 5$, whose interiors lie in specific chambers of this line arrangement and belong to the Steiner reducing set $St(Y)$. The chamber will determine the available triangulation edges. For example, if a Steiner point Z is added in chamber associated with M_2, Z can be adjacent to any of G, H, I, or J, but cannot be adjacent to K since the edges of the interior 12-gon still belong to a minimum weight triangulation of the augmented point set. We further simplify our search for the Steiner reducing set by utilizing the convexity of k-ellipses. Namely, if $W \subseteq Y$ is a set of d points whose convex hull lies inside a specific feasibility region, and if the points of W all lie in a region bounded by an appropriate k-ellipse, then $\mathrm{conv}(W) \subseteq St(Y)$. By checking the weight of the proposed triangulation at the vertices of the polygons M_i, we use convexity to determine that the interior of each M_i is indeed a subset of a corresponding k-ellipse, and build the Steiner reducing set by verifying minimal triangulations for points that fall on chamber boundary lines.

To finish the proof of Theorem 3, we prove the existence of 13 holes within the Steiner reducing set $St(Y)$. To do so, we find generators of linearly independent homology classes in $H_1(St(Y))$. The technique requires finding 13 points in the interior of $\text{conv}(St(Y))$ that are not Steiner reducing points, together with 13 closed curves γ_i such that $\gamma_i \subset St(Y)$ for $1 \le i \le 13$ and each bounds a compact subset of $\mathbb{R}^2$ that contains a point which is not a Steiner reducing point.

None of the points $(0, 35)$, $(18, 32)$, or $(0, 0)$, as well as rotations of these by multiples of $\frac{\pi}{3}$ are Steiner reducing points. It is not difficult to use sets M_i to construct three generators $\gamma_1, \gamma_2, \gamma_3$ of linearly independent homology classes in $H_1(St(Y))$ such that $(0, 35)$ lies interior to curve γ_1, $(18, 32)$ lies interior to γ_2, and $(0, 0)$ lies interior to γ_3. It follows that no path homotopy exists within the set $St(Y)$ from γ_i to γ_j, $i \neq j$, so the holes detected are distinct. Define γ_i for $4 \le i \le 13$ to be the distinct images of γ_1 and γ_2 under rotations about the origin by integer multiples of $\pi/3$ radians. □

As in the proof of the first theorem, by aligning copies of Y with the vertices of an n-gon with sufficiently long edges (relative to the size of Y), we can get a point set $Y^{[n]}$ whose minimum weight triangulation contains all edges of the minimum weight triangulation of each copy of Y, thus giving a point set with $18n$ points that has at least $13n$ holes.

4 Future Work

These results have implications for the design of algorithms to search for minimum weight Steiner triangulations. Questions remain about whether two points can be better than one: in a case when the Steiner reducing set $St(X)$ is empty, is it possible that two Steiner points P and Q can team up to cause the weight of a minimal weight triangulation of $X \cup \{P, Q\}$ to be less than that of X? It may be possible to use the structures of the k-ellipse and feasibility sets to design faster algorithms to find the minimal weight triangulation of certain classes of point sets in an efficient manner.

References

[1] 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.

[2] 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.

[3] J. A. De Loera, J. Rambau, and F. Santos. *Triangulations: Structures for Algorithms and Applications*, volume 25 of *Algorithms and Computation in Mathematics*. Springer-Verlag, Berlin, 2010.

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

[5] P. Erdős and I. Vincze. On the approximation of convex, closed plane curves by multifocal ellipses. *J. Appl. Probab.*, (Special Vol. 19A):89–96, 1982.

[6] 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, 1979.

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

[8] C. Groß and T.-K. Strempel. On generalizations of conics and on a generalization of the Fermat-Torricelli problem. *Amer. Math. Monthly*, 105(8):732–743, 1998.

[9] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

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

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

[12] W. Mulzer and G. Rote. Minimum weight triangulation is NP-hard. In *Computational geometry (SCG'06)*, pages 1–10. ACM, New York, 2006.

[13] J. Nie, P. A. Parrilo, and B. Sturmfels. Semidefinite representation of the k-ellipse. In *Algorithms in algebraic geometry*, volume 146 of *IMA Vol. Math. Appl.*, pages 117–132. Springer, New York, 2008.

[14] P. V. Sahadevan. The theory of the egglipse—a new curve with three focal points. *Internat. J. Math. Ed. Sci. Tech.*, 18(1):29–39, 1987.

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

[16] G. Sz.-Nagy. Tschirnhaus'sche Eiflächen und Eikurven. *Acta Math. Acad. Sci. Hungar.*, 1:36–45, 1950.

[17] C. Traub. *Topological Effects Related to Minimum Weight Steiner Triangulations*. PhD thesis, Washington University in St. Louis, 2006.

[18] E. W. v. Tschirnhaus. *Medicina mentis, sive artis inveniendi praecepta generalia*. J. Thomann Fritsch, Leipzig, 1695.

[19] A. Varga and C. Vincze. On a lower and upper bound for the curvature of ellipses with more than two foci. *Expo. Math.*, 26(1):55–77, 2008.