On the Configuration Space of Steiner Minimal Trees

Herbert Edelsbrunner · Nataliya P. Strelkova

Abstract Among other results, we prove the following theorem about Steiner minimal trees in d-dimensional Euclidean space: if two finite sets in $\mathbb{R}^d$ have unique and combinatorially equivalent Steiner minimal trees, then there is a homotopy between the two sets that maintains the uniqueness and the combinatorial structure of the Steiner minimal tree throughout the homotopy.

1 Introduction

Shortest networks connecting finite sets of geographic locations were intensely studied when the telephone providers were legally required to charge for phone calls an amount that is proportional to the length of the connection [2]. This motivated the discovery of many mathematical and computational properties of such networks. Nevertheless, many of the fundamental questions are still open. For example, in 1968 Gilbert and Pollak [6] conjectured that the ratio of the length of a minimum spanning tree over the length of a shortest network of the same point set cannot exceed $2/\sqrt{3}$. In 1992, Du and Hwang claimed a proof of the conjecture [3], but 10 years later, Ivanov and Tuzhilin noted that there are gaps [10], and the problem remains open until today [12]. In contrast to the ratio problem, surprisingly little attention was directed toward the configuration space of the shortest networks, a topic we approach in this paper. Specifically, we consider a finite set of points, S, in d-dimensional Euclidean space, which we denote as $\mathbb{R}^d$. A *spanning network* of S is a finite set of rectifyable curves, each connecting two points in $\mathbb{R}^d$ but not necessarily in S, whose union is connected and contains S. We measure the length of the network using the Euclidean metric. A *shortest network* of S is a spanning network of minimum length.

H. Edelsbrunner
IST Austria (Institute of Science and Technology Austria), Klosterneuburg, Austria
E-mail: edels@ist.ac.at

N. P. Strelkova
Moscow State University, Moscow, Russia
E-mail: n.strelk@gmail.com

It is known to exist, and it satisfies the following properties; see e.g. [6]:

- it is a tree whose edges are straight line segments ending at vertices of degree 1, 2, or 3;
- all degree-1 and degree-2 vertices are points in S;
- the angles between the edges meeting at a degree-3 vertex are $120°$, while the angle between the edges meeting at a degree-2 vertex is at least $120°$.
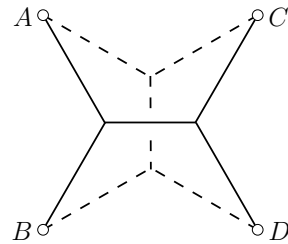

Fig. 1: Two Steiner minimal trees of the four corners of a square.

A shortest network of a finite set of points is usually referred to as a *Steiner minimal tree.* This tree may or may not be unique, and in the latter case, we call the Steiner minimal tree *ambiguous.* For example, the four corners of a square have two Steiner minimal trees; see Figure 1. Note that these two trees have different combinatorial structure if the points are labeled. Just to mention one difference: the points A and B are connected by a path of two edges in the solid tree, and by a path of three edges in the dashed tree. Suppose now that S consists of four (arbitrary) points in the plane, and we need to find a Steiner minimal tree. If we know how many additional vertices there are in the tree, and how they are connected to the four given points and to each other, then we can locate them using the $120°$-condition. Generalizing this idea to n points in the plane gives Melzak's algorithm [13]. But a priori, we do not know which combinatorial structure gives the minimum length, and there is no known way that avoids checking all of them in the worst case. Indeed, the problem of constructing a Steiner minimal tree in $\mathbb{R}^d$ is NP-hard even for $d = 2$; see [5, page 209].

For a given finite set, S, the collection of Steiner minimal trees of S is well defined, and every tree in this collection has a different *combinatorial type* [6], a notion that will be made precise later. We are interested in the following question: **for a given combinatorial type, what does the space of points whose Steiner minimal trees are of this type look like topologically?** To make this more concrete, we order the n points in $S \subseteq \mathbb{R}^d$ and list the coordinates in sequence, mapping S to a point in $\mathbb{R}^{nd}$. For a given combinatorial type, we consider the space of points in $\mathbb{R}^{nd}$ that correspond to sets in $\mathbb{R}^d$ with unambiguous Steiner minimal tree of the given type. Our main result is that this space is path-connected. We have a weaker result if we relax the requirement and allow ambiguous Steiner minimal trees. Here we can prove path-connectivity only for full networks in the plane. Finally, we exhibit an example to show that the space of points with necessarily ambiguous Steiner minimal trees of a given type is generally not connected. While the restricted setting for the result on cells is disappointing, it is consistent with a general lack of understanding of ambiguous Steiner minimal trees. For example, it is not known that the space of ambiguous Steiner minimal trees has measure zero. The only partial result is in $d = 2$ dimensions, where the unambiguous Steiner minimal trees contain an everywhere dense open subset of the configuration space, but the proof in [11] is long and complicated. The approach in [14] leads to a shorter proof.

Outline. Section 2 introduces the main concepts needed to formally state our results. Section 3 presents the key tools used in our proofs. Section 4 presents the proofs. Section 5 concludes the paper.

2 Concepts and Results

The two main concepts in this paper are the spanning networks of a set of n points in $\mathbb{R}^d$, and the decomposition of the nd-dimensional configuration space into cells whose points have combinatorially equivalent shortest networks. We discuss both and introduce the terminology needed to formally state our results.

Graphs, networks, types. Similar to the notion of a curve in differential geometry, we define a network as a map from an abstract topological to a concrete geometric space. In our context, the abstract space is a finite connected graph, $G = (V \sqcup W, E)$. It is *partially ordered*, with V ordered and W unordered. The difference arises because the vertices in V are mapped to the points in S, which are ordered, while the vertices in W can be mapped anywhere. To emphasize the difference, we sometimes call an element of V a *terminal vertex* and an element of W an *interior vertex*; an element of E is an *edge*. Two partially ordered graphs, $G = (V \sqcup W, E)$ and $G' = (V' \sqcup W', E')$, are *combinatorially equivalent* if there exist compatible bijections $V \to V'$, $W \to W'$, and $E \to E'$ of which the first preserves order. Observe that the preservation of order prevents the existence of such bijections for the two trees in Figure 1. We have $V = V' = \{A, B, C, D\}$, and since we are obliged to map A to A and so on, it is impossible to find a compatible bijection between the two edge sets.

We find it convenient to consider a graph as a 1-dimensional topological space, so it makes sense to talk about a point $x \in G$, which may be a vertex or a point on an edge of G. If G and G' are combinatorially equivalent, then there exists a *consistent* homeomorphism $g \colon G \to G'$, that is: g restricts to the bijections establishing the combinatorial equivalence. With this preamble, we introduce the first main concept.

Definition 1 A *network* is a continuous map $\mathcal{G} \colon G \to \mathbb{R}^d$ such that the image of every edge is a rectifyable curve or a point, and the preimage of the image of every vertex is a subgraph of G, possibly together with subsets of points on edges whose images happen to pass through this image. We call G a *parametrization* of $\mathcal{G}$, and a *contracting parametrization* if every subgraph with constant image is connected.

The *(combinatorial) type* of $\mathcal{G}$ is the graph G_0 obtained from G by shrinking each subgraph with constant image to a vertex. Shrinking is similar to the contraction operation defined in Tutte [16, page 32] but more general because we do not require that the subgraphs that shrink to vertices be connected. The type is again a graph although it may have loops and multi-edges even if G does not. The network $\mathcal{G} \colon G \to \mathbb{R}^d$ *spans* an ordered set $S \subseteq \mathbb{R}^d$ if $\mathcal{G}$ restricts to an order-preserving bijection from V to S. To compare networks, we write $\ell(\mathcal{G})$ for the sum of the lengths of the curves. A *straight-line network* maps every edge to a line segment or point. Of particular importance are shortest spanning networks, which are necessarily straight-line but also satisfy the three properties stated in Section 1. A *Steiner minimal tree* is the image of a shortest network. The image of a shortest network is necessarily a tree, so it is natural to parametrize it using a tree. In this case, the type is obtained by shrinking connected subtrees, so the parametrization is necessarily contracting.

The above definitions allow for different degrees of sameness between two networks, $\mathcal{G} \colon G \to \mathbb{R}^d$ and $\mathcal{G}' \colon G' \to \mathbb{R}^d$. We call $\mathcal{G}, \mathcal{G}'$ *identical* if G, G' are combinatorially equivalent and there is a consistent homeomorphism $g \colon G \to G'$ such that $\mathcal{G}(x) = \mathcal{G}'(g(x))$ for every $x \in G$. Note that this is strictly stronger than requiring combinatorially equivalent graphs and equal images. A weaker and perhaps more natural notion calls $\mathcal{G}, \mathcal{G}'$ *the same* if their types, G_0, G_0', are combinatorially equivalent, and there is a consistent homeomorphism $g_0 \colon G_0 \to G_0'$ such that $\mathcal{G}(x) = \mathcal{G}'(g_0(x))$ for every $x \in G_0$. For example, two shortest networks are the same iff they have the same Steiner minimal tree. Yet weaker is the notion that says $\mathcal{G}, \mathcal{G}'$ *have the same type* if G_0, G_0' are combinatorially equivalent. In dimension $d = 2$, we require in addition that there is a homeomorphism $h \colon \mathbb{R}^2 \to \mathbb{R}^2$ such that $\mathcal{G}(x) = h(\mathcal{G}'(g_0(x)))$. For plane embeddings, this is equivalent to requiring that the ordering of the edges around the vertices is the same in both images. While the three notions of sameness are different and relate to each other by a chain of implications, they are all transitive. We get the reverse chain of implications for the negations: having different type implies not being the same implies not being identical.

Configurations, cells, decompositions. Let S be an ordered set of points $p_i = (p_{i1}, p_{i2}, \ldots, p_{id}) \in \mathbb{R}^d$ for $1 \leq i \leq n$. Listing the coordinates in sequence, we obtain a point

$$s = (p_{11}, p_{12}, \ldots, p_{1d}, p_{21}, p_{22}, \ldots, p_{nd}) \in \mathbb{R}^{nd}. \tag{1}$$

We require that the points in S are distinct, which implies that s cannot lie on any of the $(nd - d)$-dimensional planes of the form $\Delta_i^j = \{(\ldots, x, \ldots, x, \ldots) \in \mathbb{R}^{nd} \mid x \in \mathbb{R}^d\}$, where we list the i-th and

the j-th d-tuple of coordinates, both specifying the same point in $\mathbb{R}^d$. Observe that all Δ_i^j contain the diagonal of points $(x, x, \ldots, x)$. We use $\mathbb{R}_\Delta^{nd} = \mathbb{R}^{nd} \setminus \bigcup_{i<j} \Delta_i^j$ as the *configuration space* for sets of n points in $\mathbb{R}^d$.

Definition 2 The *cell* of a type G, denoted as cell(G), is the space of points $s \in \mathbb{R}_\Delta^{nd}$ that have shortest networks of type G. The *unambiguous cell* of G, denoted as ucell(G), is the subspace of points $s \in$ cell(G) such that s has a unique Steiner minimal tree.

The remaining points, in cell(G) \ ucell(G), have at least two shortest networks, one of type G and another of necessarily different type; see [6]. What can we say about the cells? Clearly, they cover the configuration space since every point $s \in \mathbb{R}_\Delta^{nd}$ has at least one shortest network. The cells have possibly non-empty intersections, and removing these intersections leaves us with the unambiguous cells, which are disjoint but do not necessarily cover the configuration space. As mentioned earlier, there is not much known about the space of points with ambiguous shortest networks. To complete the picture, we note that the points in the Δ_i^j correspond to multisets, or after identification to sets of size $m < n$ in $\mathbb{R}^d$. We can therefore complete the decomposition of $\mathbb{R}^{nd}$ by partitioning the Δ_i^j into strata within which m is a constant, and to decompose each stratum according to shortest networks. By definition, the cells in these strata belong to the boundaries of the cells in $\mathbb{R}_\Delta^{nd}$, but not to the cells themselves.
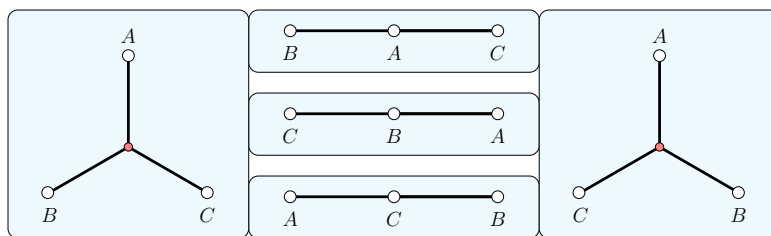

Fig. 2: The five types of shortest networks for three points in the plane. Two boxes touch iff there is a transition between the types that avoids configurations with fewer than three points.

We illustrate the decomposition of the configuration space with a small example: $n = 3$ points in $d = 2$ dimensions. There are only five different types, each generating a cell in $\mathbb{R}_\Delta^6$; see Figure 2. Factoring out the three dimensions of rigid motions, we may visualize the decomposition in the remaining 3 dimensions. Drawing the diagonal as a point at the origin with three emanating half-lines for the Δ_i^j, $1 \leq i < j \leq 3$, we get a ring of three cells separating the two remaining cells. The origin lies in the boundaries of all five cells, while each half-line belongs to the boundaries of four. Only six of the ten transitions between the five cells avoid the origin and the half-lines, and we visualize them as 2-dimensional membranes, which lie in the common boundaries of the corresponding cells. In this particular example, these membranes belong to the three cells in the ring but not to the other two cells. For example, a point on the membrane separating the cell on the left from the cell at the center in Figure 2 has a Steiner minimal tree consisting of two edges that form an angle of 120° at B. The combinatorial type of this network is different from that on the left but the same as that at the center. Indeed, the two separated cells do not contain any points of their boundary, while the three cells in the ring contain every boundary point that does not belong to the origin and the half-lines.

Main results. We are now ready to give formal statements of our main results. They state that unambiguous cells are connected for all dimensions d, and that cells are connected for full networks in $\mathbb{R}^2$. In contrast, the difference between a cell and the corresponding unambiguous cell is not necessarily connected, which we will show by exhibiting an example of $n = 4$ points in $d = 2$ dimensions.

Theorem 1 (Connectedness of unambiguous cells) *Let G be the type of a shortest network of a finite set of points in $\mathbb{R}^d$, for $d \geq 2$. Then* ucell(G) *is path-connected.*

Theorem 2 (CONNECTEDNESS OF CELLS) *Let G be the type of a shortest network of a finite set of points in $\mathbb{R}^2$ such that all terminal vertices have degree-1. Then* cell(G) *is path-connected.*

3 Tools

The proof of our main result is inductive, applying a procedure similar to Melzak's algorithm; see [6, Section 5] or [13]. We begin by introducing the ingredients of the procedure, proving lemmas in preparation of the proofs of the theorems.

Locally minimal and pairs of codirected networks. A network that satisfies the three properties stated in Section 1 is called *locally minimal*. Every shortest network is locally minimal but there are others; see Figure 3 for an example. As its name suggests, there are no shorter networks nearby, in which the notion of nearness is defined by parametrizing the networks by their interior vertices. More generally, it is not possible to have two locally minimal networks of the same type. We state this result for later reference; see [9, pages 927 and 939].
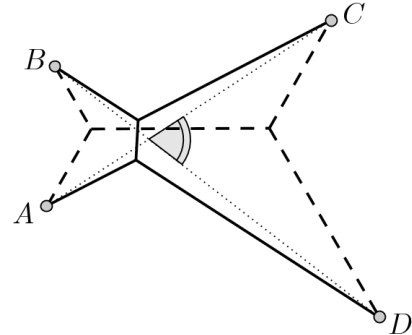

Fig. 3: Two locally minimal networks spanning the same four points in the plane. The dashed tree is shorter because its center edge aligns with the smaller angle pair formed by the diagonals; see [15, Theorem 6].

Proposition 1 (UNIQUENESS OF MINIMAL NETWORK) *Let G be a partially ordered tree with n terminal vertices and S an ordered set of n points in $\mathbb{R}^d$. If there exists a locally minimal network that spans S and is parametrized by G, then it is unique and the shortest such network.*

As a special case, we see that two different shortest networks of the same set S that are parametrized by the same tree are impossible. A related concept is the following. Two straight-line networks, $\mathcal{G}\colon G \to \mathbb{R}^d$ and $\mathcal{H}\colon H \to \mathbb{R}^d$, that span the same set S are *codirected* if they look geometrically the same in the neighborhood of every vertex. More formally, we require that there exists $r > 0$ such that the intersection of the ball with radius r centered at $\mathcal{G}(u) = \mathcal{H}(u)$ with the images of $\mathcal{G}$ and $\mathcal{H}$ are the same for every terminal vertex u. The existence of codirected locally minimal trees that are not the same is open, but in the plane it is known that such pairs do not exist [14].

Proposition 2 (NON-EXISTENCE OF CODIRECTED TREES) *Any two locally minimal trees spanning the same ordered set of finitely many points in $\mathbb{R}^2$ that are codirected are the same.*

Sequences of networks. Another useful tool in the study of cells and their boundaries are converging sequences. We write $\mathcal{G}_i$ for the networks in an infinite sequence, indexing them with the positive integers $i \in \mathbb{N}$. In this paper, we are primarily interested in sequences such that the $\mathcal{G}_i$ have all the same type, G. We say such a sequence $\mathcal{G}_i$ *converges* if the sequence of points $\mathcal{G}_i(u)$ converges for every vertex u of G. For straight edges, the limit of a converging sequence is well defined, namely the straight-line network $\mathcal{G}\colon G \to \mathbb{R}^d$ such that $\mathcal{G}_i(u) \to \mathcal{G}(u)$ for each vertex u. We write $\mathcal{G}_i \to \mathcal{G}$ as well as $\mathcal{G} = \lim_{i \to \infty} \mathcal{G}_i$. Note that $\mathcal{G}$ is not necessarily of type G, but G parametrizes $\mathcal{G}$ by construction. An important method for obtaining converging sequences follows from the Bolzano-Weierstrass Theorem in analysis: starting with an infinite sequence of networks with finitely many types and points within a compact domain, we obtain a converging sequence by repeatedly restricting the choice and continuing the process with an infinite subsequence.

For straight-line networks, $\mathcal{G}_i \to \mathcal{G}$ implies that the length of $\mathcal{G}_i$ converges to the length of $\mathcal{G}$. It is therefore easy to see that if all $\mathcal{G}_i$ are shortest networks, then $\mathcal{G}$ is a shortest network with a contracting parametrization. Another important property of converging sequences of shortest networks is that

convergence for the terminal vertices implies convergence for all vertices. In other words, as long as the shortest network does not change its type, it depends continuously on the points it spans.

Lemma 1 (Convergence of shortest networks) *Let G be a tree, and let s_i be a converging sequence of points in* cell(G) *with shortest networks $\mathcal{G}_i\colon G \to \mathbb{R}^d$, all of type G. Then $\mathcal{G}_i$ is a converging sequence of networks, and the limit network is a shortest network of the limit of s_i.*

Proof Since s_i converges, there exists a bounded and convex domain $D \subseteq \mathbb{R}^d$ that contains the images of the terminal vertices of G for all $i \in \mathbb{N}$. The images of the interior vertices lie within the convex hull of the images of the terminal vertices and hence also in D. We can therefore apply the Bolzano-Weierstrass Theorem and obtain a converging subsequence of shortest networks, with limit $\mathcal{G}$. To get a contradiction to the claimed property, we assume that $\mathcal{G}_i$ does not converge to $\mathcal{G}$. Then there exists an interior vertex u of G, a constant $C > 0$, and an infinite subsequence $\mathcal{G}_{i_j}$ such that $\|\mathcal{G}_{i_j}(u) - \mathcal{G}(u)\| > C$ for all j. Applying the Bolzano-Weierstrass Theorem to this subsequence gives another converging sequence of shortest networks, now with limit $\mathcal{G}' \neq \mathcal{G}$. Thus, we get two shortest networks of the same set $S \subseteq \mathbb{R}^d$, both parametrized by the same tree G. Both are locally minimal, which contradicts Lemma 1.   □

Note that Lemma 1 implies that every point s in the boundary of cell(G) has a shortest network parametrized by G, and the parametrization is contracting.

Moustaches. A particularly simple application of infinite sequences is used to shrink and grow edges of Steiner minimal trees that end at degree-1 vertices. Let $\mathcal{G}$ be a shortest network of a set S of $n \geq 2$ points in $\mathbb{R}^d$, let G be its type, and recall that it enjoys the three properties stated in Section 1. Hence, G has at least one of the following two substructures: a *two-sided moustache* consisting of degree-1 vertices p and q both adjacent to a degree-3 vertex v, or a *one-sided moustache* consisting of a degree-1 vertex p adjacent to a degree-2 vertex v. We call v the *anchor* of the moustache in both cases. We *shave* a two-sided moustache by removing p and q together with the edges connecting them to v. This leaves a smaller graph G' and a network $\mathcal{G}'\colon G' \to \mathbb{R}^d$ that spans a smaller set of points, namely $S' = S \setminus \{\mathcal{G}(p), \mathcal{G}(q)\}$ if $\mathcal{G}(v) \in S$, and union $\{\mathcal{G}(v)\}$ if $\mathcal{G}(v) \notin S$. Similarly, we shave a one-sided moustache by removing p together with the edge connecting it to v. In this case, $\mathcal{G}(v)$ is necessarily in S, so we get a spanning network of $S' = S \setminus \{\mathcal{G}(p)\}$. Importantly, the operation preserves minimality in both cases. For later reference, we state and prove a slightly more general claim, namely that trimming a moustache preserves minimality.

Lemma 2 (Trimming a moustache) *Let G be the type of a shortest network $\mathcal{G}$ of a set of $n \geq 2$ points in $\mathbb{R}^d$, let p be a degree-1 vertex of G and v its neighbor. For $0 \leq t \leq 1$, let $\mathcal{G}_t\colon G \to \mathbb{R}^d$ be the same except that it maps p to $\mathcal{G}_t(p) = (1-t)\mathcal{G}(p) + t\mathcal{G}(v)$. Then $\mathcal{G}_t$ is a shortest network of $S_t = S \setminus \{\mathcal{G}(p)\} \cup \{\mathcal{G}_t(p)\}$, and if $\mathcal{G}$ is unambiguous then so is $\mathcal{G}_t$.*

Proof Suppose there exists a network $\mathcal{H}_t\colon G \to \mathbb{R}^d$ that spans S_t with $\ell(\mathcal{H}_t) < \ell(\mathcal{G}_t)$. Append the line segment from $\mathcal{H}_t(p) = \mathcal{G}_t(p)$ to $\mathcal{G}(p)$ to get a network $\mathcal{H}$ that now spans S. Then $\ell(\mathcal{H}) < \ell(\mathcal{G})$, which contradicts the assumption that $\mathcal{G}$ is a shortest network of S. The same argument contradicts the existence of different shortest networks for S_t if there is only one for S.   □

When we shave a moustache, we effectively map a point $s \in \text{cell}(G) \subseteq \mathbb{R}^{nd}_\Delta$ to a point $s' \in \text{cell}(G')$ in $(n-1)d$- or $(n-2)d$-dimensional configuration space, depending on the case. Doing this over all points s and all moustaches of the corresponding shortest networks, we get multiple maps from higher- to lower-dimensional configuration spaces, each skipping either d or $2d$ dimensions. We will exploit the fact that every point in an unambiguous cell has plenty of preimages. In other words, if S' has an unambiguous shortest network, then we can grow a moustache while preserving that the network is shortest and unambiguous, of course now spanning a larger set S. This property does not necessarily hold for sets S'

with ambiguous shortest networks. Indeed, the corners of a regular octagon have eight different shortest networks [4], and none of them permits the growth of a one- or two-sided moustache such that the network remains shortest; see Figure 4.
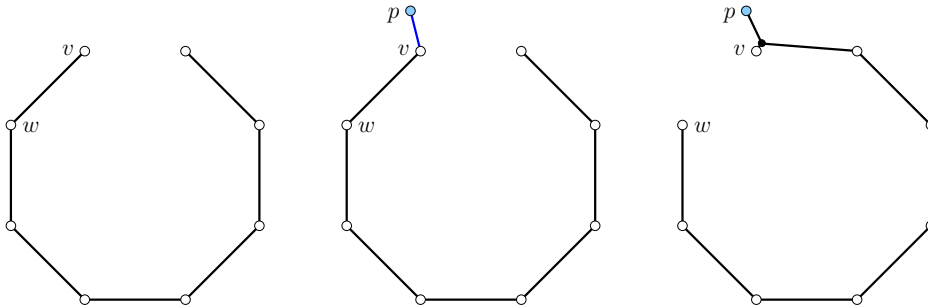

Fig. 4: From *left* to *right*: a Steiner minimal tree of the corners of a regular octagon, the tree with a one-sided moustache attached, a shorter network with a degree-3 vertex connecting the endpoint of the moustache to the rest of the tree.

Admissible extensions. Before proving the property for unambiguous shortest networks, we introduce notation for adding a moustache. Let $\mathcal{G}' \colon G' \to \mathbb{R}^d$ be a shortest network, let v be a degree-1 vertex with neighbor w in G', and let $\varphi \in \mathbb{S}^{d-1}$ be a direction. To add a two-sided moustache with anchor v, we require that the angle between φ and the direction $\omega = (\mathcal{G}'(w) - \mathcal{G}'(v))/\|\mathcal{G}'(w) - \mathcal{G}'(v)\|$ be $120°$. Equivalently, we require $\langle \varphi, \omega \rangle = -\frac{1}{2}$, and if this equation is satisfied, then we call φ an *allowed direction* for v. For such a φ, there exists a unique second direction $\psi \in \mathbb{S}^{d-1}$ such that $\langle \psi, \varphi \rangle = \langle \psi, \omega \rangle = -\frac{1}{2}$. Choosing $r > 0$, we construct $S = S' \setminus \{\mathcal{G}'(v)\} \cup \{\mathcal{G}'(v) + r\varphi, \mathcal{G}'(v) + r\psi\}$. Similarly, we construct G from G' by adding vertices p and q, connecting both with new edges to v, and we construct $\mathcal{G} \colon G \to \mathbb{R}^d$ such that the restriction to G' agrees with $\mathcal{G}'$. We call r *admissible* for v and φ if $\mathcal{G}$ is the unique shortest network of S.

Adding a one-sided moustache is similar, except that φ is an *allowed direction* if $\langle \varphi, \omega \rangle \leq -\frac{1}{2}$, and there is no second direction ψ to worry about.

Lemma 3 (EXISTENCE OF ADMISSIBLE EXTENSIONS) *For each unambiguous shortest network $\mathcal{G}' \colon G' \to \mathbb{R}^d$, each degree-1 vertex v in G', and each allowed direction $\varphi \in \mathbb{S}^{d-1}$, there exists $R > 0$ such that every $0 < r \leq R$ is admissible for v and φ.*

Proof We consider the two-sided case and omit the proof of the easier, one-sided case. Assume without loss of generality that G' is the type of $\mathcal{G}'$, write S' for the ordered set of points spanned by $\mathcal{G}'$, and let $x = \mathcal{G}'(v)$ be the image of the anchor. We prove that the set of admissible extensions is not empty. Given a positive admissible extension, Lemma 2 then implies that all smaller extensions are also admissible, which gives the claimed statement. To get a contradiction, we assume that there is a degree-1 vertex v of G' and an allowed direction φ, such that no $r > 0$ is admissible for v and φ. The plan is to construct two different locally minimal networks that are parametrized by the same tree. By Proposition 1, such networks do not exist, which will furnish the desired contradiction.

The first locally minimal network is $\mathcal{G} \colon G \to \mathbb{R}^d$ obtained by growing a two-sided moustache of sufficiently small size $R = r_1 > 0$ to avoid self-intersections. Assuming G is its type, we get G from G' by adding vertices p and q and edges that connect them to v. To construct the second locally minimal network, we use the assumption that there is no positive admissible extension. Choosing an infinite sequence r_i that vanishes in the limit, we set $S_i = S' \setminus \{x\} \cup \{x + r_i\varphi, x + r_i\psi\}$ and write $\mathcal{H}_i \colon H_i \to \mathbb{R}^d$ for the shortest network of S_i, letting H_i be its type. Since we can select a converging subsequence, we may

assume that $H = H_i$ for all i, and that $\mathcal{H}_i$ converges to a network $\mathcal{H}: H \to \mathbb{R}^d$. Since both $\mathcal{H}$ and $\mathcal{G}'$ are shortest, they are the same, only with different parametrizations. To analyze this difference, we consider the preimage, $H_x = \mathcal{H}^{-1}(x)$. We note that p and q are the only terminal vertices whose images converge
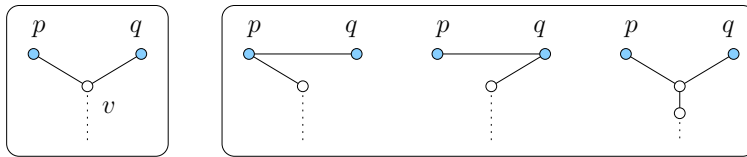

Fig. 5: Possible subtrees H_x. The leftmost is a subtree of G, while the other three exhaust the possibilities for H. The rightmost subtree can be used to reparametrize the others.

to x. All interior vertices of H have degree 3, but one of them can have degree less than 3 in H_x, namely the endpoint of the edge that connects H_x to the rest of the tree. There cannot be a second such vertex for this would contradict that $\mathcal{H}$ has a contracting parametrization as implied by Lemma 1. Since H_x has at most three vertices of degree less than 3, we are left with the choices shown in Figure 5. The tree on the very right can be used to reparametrize the other three. We can therefore reparametrize $\mathcal{G}$ and $\mathcal{H}_1$ with the same tree. Both networks are locally minimal and span the same set, S_1, which contradicts that they are different.   □

Lower semi-continuity. Fixing G', we may consider the shortest networks of all $s' \in \mathrm{ucell}(G')$ and grow a moustache from the fixed anchor, v, in each. Within each such network, we have a direction $\omega(s')$ determined by the image of v and the image of its neighbor, w. By Lemma 1, ω depends continuously on s'. We can therefore choose an allowed direction, φ, that depends continuously on s'. The second direction, ψ, needed to grow a moustache is uniquely determined by ω and φ. Let $R_1: \mathrm{ucell}(G') \to \mathbb{R}$ be defined by mapping s' to the minimum of 1 and the supremum of the thresholds R for which Lemma 3 holds. We do not know whether this function is continuous, but we can prove lower semi-continuity, which is sufficient to prove that the infimum of R_1 over a compact set of points s' is positive.

Lemma 4 (LOWER SEMI-CONTINUITY) *Let $s' = \lim_{i \to \infty} s'_i$, all points in a common unambiguous cell. Then $\liminf_{i \to \infty} R_1(s'_i) \geq R_1(s')$.*

Proof We consider the two-sided case and omit the proof of the easier, one-sided case. To get a contradiction, we suppose that $\liminf_{i \to \infty} R_1(s'_i) \leq R_1(s') - 2\delta$ for some $\delta > 0$. Choose a converging sequence r_i such that $R_1(s'_i) < r_i < R_1(s') - \delta$ for each i. Writing S'_i for the set of points that corresponds to s'_i and x_i for the image of the anchor, we introduce $S_i = S'_i \setminus \{x_i\} \cup \{x_i + r_i\varphi, x_i + r_i\psi\}$ and let $\mathcal{G}_i: G \to \mathbb{R}^d$ be the network obtained by adding the moustache to the shortest network of S'_i. Writing G' for the type of that shortest network, we get the type G of $\mathcal{G}_i$ from G' by adding vertices p and q and connecting them with edges to the anchor. Because $R_1(s'_i) < r_i$, the constructed network $\mathcal{G}_i$ is either not a shortest network or it is not unique. In either case, there exists a different shortest network $\mathcal{H}_i$ spanning the same ordered set $S_i \subseteq \mathbb{R}^d$. Since we can choose subsequences, we may assume that the $\mathcal{H}_i$ all have the same type H. By Lemma 1, $\mathcal{H}_i$ converges to a shortest network $\mathcal{H}: H \to \mathbb{R}^d$ spanning $S = \lim_{i \to \infty} S_i$. By construction, $0 < \lim_{i \to \infty} r_i < R_1(s')$, so S has a unique shortest network $\mathcal{G}: G \to \mathbb{R}^d$. This implies that $\mathcal{G}$ and $\mathcal{H}$ are the same. Networks of type G can therefore be parametrized by H. But then $\mathcal{G}_i$ and $\mathcal{H}_i$ are different locally minimal networks parametrized by the same tree, which contradicts Proposition 1.   □

4 Proofs

With the preparations in Section 3, we are now ready to prove Theorems 1 and 2, which have been formally stated at the end of Section 2. In addition, we illustrate the limitation of connectivity with an example.

Proof of Theorem 1. Recall that this theorem asserts that unambiguous cells in the decomposition of the configuration space of n points in $\mathbb{R}^d$ are path-connected for all $d \geq 2$. We use induction over the number of points. For $n = 1$, there is only the trivial Steiner minimal tree, whose unambiguous cell is the entire $\mathbb{R}^d$, which is path-connected. For $n = 2$, there is again only one Steiner minimal tree consisting of two points connected by a line segment. The unambiguous cell is the entire $\mathbb{R}^{2d}$ minus the diagonal, which is a d-dimensional plane. This space has the homotopy type of the $(d-1)$-dimensional sphere. It is path-connected for all $d \geq 2$. This establishes the claim for $n = 1, 2$ and $d \geq 2$.

For the inductive step, we consider two points, s_0 and s_1, in an unambiguous cell $\mathrm{ucell}(G) \subseteq \mathbb{R}^{(n+1)d}_\Delta$, and prove that there exists a path $s\colon [0, 1] \to \mathrm{ucell}(G)$ with $s(0) = s_0$ and $s(1) = s_1$. Choose a moustache in G, shave it, and write G' for the resulting smaller tree. This construction establishes a map from $\mathrm{ucell}(G)$ to $\mathrm{ucell}(G')$, as described in Section 3. Note that $\mathrm{ucell}(G')$ is contained in $\mathbb{R}^{nd}_\Delta$ or in $\mathbb{R}^{(n-1)d}_\Delta$, depending on the choice of moustache. By induction, $\mathrm{ucell}(G')$ is path-connected, which implies the existence of a path $s'\colon [0, 1] \to \mathrm{ucell}(G')$ such that $s'(0)$ and $s'(1)$ are images of s_0 and s_1 under the map. We are going to lift the path to $\mathrm{ucell}(G)$ taking inverses. Lemma 3 states that we can find a preimage of $s'(t)$ for every $t \in [0, 1]$. It remains to show that the preimages can be chosen in such a way that they give a path $s\colon [0, 1] \to \mathrm{ucell}(G)$. But all the work needed to achieve this has already been done. As explained before Lemma 4, we can choose directions φ and ψ that depend continuously on s'. It is easy to modify this direction field so that when we move along s', we connect the needed directions at the two ends. Next, we set $r = \inf_{0 \leq t \leq 1} R_1(s'(t))$ and note that $r > 0$ because $[0, 1]$ is compact, and R_1 is everywhere positive as well as lower semi-continuous; see Lemma 4. Moving along s', we thus get a path $s_r\colon [0, 1] \to \mathrm{ucell}(G)$ consisting of networks with moustaches of uniform size. To get $s\colon [0, 1] \to \mathrm{ucell}(G)$, we first trim the moustache to go from s_0 to $s_r(0)$, we second move from $s_r(0)$ to $s_r(1)$, and we third enlarge the moustache to go from $s_r(1)$ to s_1.

Proof of Theorem 2. Our second result asserts path-connectivity for cells, but only in dimension $d = 2$ and for a full network in which all terminal vertices have degree 1. Suppose $\mathcal{G}\colon G \to \mathbb{R}^2$ is a shortest such network for a set S, let $G = (V \sqcup W, E)$ be its type, and let $s_0 \in \mathrm{cell}(G) \setminus \mathrm{ucell}(G)$. Since we know that $\mathrm{ucell}(G)$ is path-connected, it suffices to construct a path $s\colon [0, 1] \to \mathrm{cell}(G)$ that starts at $s(0) = s_0$ and ends at a point $s(1) \in \mathrm{ucell}(G)$. We prove the existence of such a path by considering the derivative of the length function. Write $\mathbf{m}(u) \in \mathbb{R}^2$ for the *motion vector* of $u \in V \sqcup W$, and define the *direction vector* by taking the sum of unit vectors in the directions of the incident edges:

$$\mathbf{d}_\mathcal{G}(u) = \sum_{\{u,v\} \in E} \frac{\mathcal{G}(v) - \mathcal{G}(u)}{\|\mathcal{G}(v) - \mathcal{G}(u)\|}. \tag{2}$$

The first derivative of the length function in the direction of the motion is the sum of scalar products between motion and direction vectors. Since the direction vector of every interior vertex vanishes, this gives

$$\frac{\partial \ell(\mathcal{G})}{\partial t} = \sum_{u \in V} \langle \mathbf{m}(u), \mathbf{d}_\mathcal{G}(u) \rangle; \tag{3}$$

see [8, page 64]. We are interested in the special case in which $\mathbf{m}(u) = \mathbf{d}_\mathcal{G}(u)$ for every $u \in V$. Then $\partial \ell(\mathcal{G})/\partial t = n$, the number of terminal vertices. Returning to the task at hand, we consider a second shortest network, $\mathcal{H}\colon H \to \mathbb{R}^2$ of S, and we compare the two derivatives. The vectors $\mathbf{d}_\mathcal{H}(u)$ have at most

unit length. This is immediate if the degree of u in $\mathcal{H}$ is 1, and it follows from the angle condition if the degree is 2. Hence, $\partial\ell(\mathcal{G})/\partial t - \partial\ell(\mathcal{H})/\partial t \geq 0$, with equality iff $\mathbf{d}_\mathcal{G}(u) = \mathbf{d}_\mathcal{H}(u)$ for all terminal vertices. If the difference is positive, then we move the points $\mathcal{G}(u)$ to $\mathcal{G}(u) + r\mathbf{d}_\mathcal{G}(u)$ for $r > 0$ smaller than the length of the shortest edge of $\mathcal{G}$. After this motion, the adjusted network $\mathcal{G}$ is shorter than the adjusted network $\mathcal{H}$.

It remains to show that vanishing difference is not possible. As mentioned, this is equivalent to $\mathbf{d}_\mathcal{G}(u) = \mathbf{d}_\mathcal{H}(u)$ for all $u \in V$. If all terminal vertices have degree 1, then this implies that $\mathcal{G}$ and $\mathcal{H}$ are codirected, which according to Proposition 2 is impossible for two shortest networks that are not the same. Otherwise, we grow a new edge from every vertex u whose degree in $\mathcal{H}$ is 2 in the direction $-\mathbf{d}_\mathcal{G}(u)$. We grow these edges both in $\mathcal{G}$ and in $\mathcal{H}$. The two trees may no longer be shortest, but for sufficiently short extensions, they are both locally minimal. We thus get a contradiction from the stronger version of Proposition 2 which asserts that no two locally minimal trees can be codirected.

Example. To complement our theorems, we show that ambiguous subsets of cells are not necessarily connected. Specifically, we exhibit an ambiguous shortest network of type G such that $\mathrm{cell}(G) \setminus \mathrm{ucell}(G)$ is not connected. The network spans the ordered set of $n = 4$ points in $\mathbb{R}^2$ drawn in Figure 6. The two trees superimposed on the left are the only two Steiner minimal trees of the four points. Indeed, there cannot be a full Steiner minimal tree with two interior vertices because the four points do not form a strictly convex quadrilateral; see [15, Lemma 5], and the remaining types with one or zero interior vertices give longer networks.
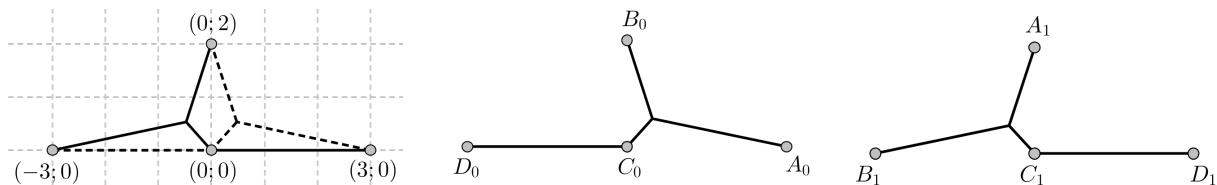

Fig. 6: *Left:* the two Steiner minimal trees of four points in the plane. *Middle* and *right:* the two ambiguous shortest networks that cannot be connected by a path of like networks.

For the next step, we choose two orderings of the four points such that the first Steiner minimal tree for the first ordering has the same type as the second Steiner minimal tree for the second ordering; see Figure 6 in the middle and on the right. We claim that there is no path from the first to the second tree along which the Steiner minimal tree remains of the same type and ambiguous at all times. To derive a contradiction, we suppose there is such a path $s\colon [0, 1] \to \mathrm{cell}(G) \setminus \mathrm{ucell}(G)$. Consider the counterclockwise angle, α_t, from the edge connecting C_t to D_t to the second edge incident on C_t. Since we have a Steiner minimal tree at all times, we have $\frac{2\pi}{3} \leq \alpha_t \leq \frac{4\pi}{3}$ for all t. Taking a look at Figure 6, we see that $\alpha_0 > \pi$ and $\alpha_1 = 2\pi - \alpha_0 < \pi$. Since α_t depends continuously on t, there is a value $0 < \tau < 1$ such that $\alpha_\tau = \pi$. At $t = \tau$, the Steiner minimal tree is determined by the points A_τ, B_τ, D_τ. However, the Steiner minimal tree of three points cannot be ambiguous. This contradicts that S is a path inside $\mathrm{cell}(G) \setminus \mathrm{ucell}(G)$. We note that this example does not extend to $d \geq 3$ dimensions.

5 Discussion

This paper continues the study of the decomposition of the configuration space of n points in $\mathbb{R}^d$ as defined by the combinatorial types of the Steiner minimal trees, which was initiated in [11]. We consider the ordered setting and our main result is a proof that the cells consisting of configurations with unambiguous Steiner minimal trees are path-connected. There are many questions that remain open.

– Can our partial results for cells be extended to all configurations in $\mathbb{R}^2$ and to higher dimensions?
– Is it true that the set of points in $\mathbb{R}^{nd}_\Delta$ with ambiguous shortest networks has measure zero?
– Is the non-empty intersection of two cells in $\mathbb{R}^{nd}_\Delta$ necessarily path-connected?
– Is the union of sets $\mathrm{cell}(G) \setminus \mathrm{ucell}(G)$ in $\mathbb{R}^{nd}_\Delta$ over all types G path-connected?

Besides the ordered setting considered in this paper, it would be interesting to address the same questions for unordered point sets. The additional symmetries complicate the global topology [1], and more sophisticated notions of connectivity as offered by homology and homotopy groups seem appropriate.

References

1. J. S. Birman. *Braids, Links, and Mappings.* Princeton Univ. Press, Princeton, New Jersey, 1974.
2. X. Cheng, Y. Li, D.-Z. Du and H. Q. Ngo. Steiner trees in industry. In *Handbook of Combinatorial Optimization, Vol. 5*, eds.: D.-Z. Du and P. M. Pardalos, 193–216, Kluwer Academic Publ., 2004.
3. D.-Z. Du and F. K. Hwang. A proof of Gilbert-Pollak conjecture on the Steiner ratio. *Algorithmica* **7** (1992), 121–135.
4. D.-Z. Du, F. K. Hwang and J. F. Weng. Steiner minimal trees for regular polygons. *Discrete Comput. Geom.* **2** (1987), 65–84.
5. M. R. Garey and R. L. Johnson. *Computers and Intractability. A Guide to the Theory of NP-Completeness.* Freeman and Company, San Francisco, California, 1979.
6. E. N. Gilbert and H. O. Pollak. Steiner minimal trees. *SIAM J. Appl. Math.* **16** (1968), 1–29.
7. F. K. Hwang. A linear time algorithm for full Steiner trees. *Oper. Res. Lett.* **4** (1986), 235–237.
8. A. O. Ivanov and A. A. Tuzhilin. *Minimal Networks: the Steiner Problem and Its Generalizations.* CRC Press, Boca Raton, Florida, 1994.
9. A. O. Ivanov and A. A. Tuzhilin. The space of parallel linear networks with a fixed boundary. *Izvestiya: Mathematics* **63** (1999), 923–962.
10. A. O. Ivanov and A. A. Tuzhilin. The Steiner ratio: the current state. *Mat. Vopr. Kibern.* **11** (2003), 27–48 [in Russian].
11. A. O. Ivanov and A. A. Tuzhilin. Steiner minimal tree uniqueness for boundaries in general position. *Matem. Sbornik* **197** (2006), 55–90. [English translation in *Sb. Math.* **197**, 1309–1340.]
12. A. O. Ivanov and A. A. Tuzhilin. The Steiner ratio Gilbert-Pollak conjecture is still open: clarification statement. *Algorithmica* **62** (2012), 630–632.
13. Z. A. Melzak. On the problem of Steiner. *Canad. Math. Bull.* **4** (1961), 143–148.
14. K. L. Oblakov. Non-existence of distinct codirected locally minimal trees on a plane. *Moscow Univ. Math. Bull.* **64** (2009), 62.
15. H. O. Pollak. Some remarks on the Steiner problem. *J. Combin. Theory* **Ser. A 24** (1978), 278–295.
16. W.,T. Tutte. *Graph Theory.* Cambridge Univ. Press, Cambridge, England, 2001.

A Notation

$S \subseteq \mathbb{R}^d$	ordered set of n points in Euclidean space
$G = (V \sqcup W, E)$	partially ordered connected graph
$\mathcal{G} : G \to \mathbb{R}^d, G_0$	network, graph; type
$\mathcal{H} : H \to \mathbb{R}^d, H_0$	network, graph; type
g, g_0, h	homeomorphisms
$s \in \mathbb{R}^{nd}_\Delta$	configuration, configuration space
$\mathrm{ucell}(G) \subseteq \mathrm{cell}(G)$	unambiguous cell, cell
$s : [0, 1] \to \mathrm{ucell}(G)$	path
$t, \tau \in [0, 1]$	time parameter, moment
$\mathcal{G}_i, i \in \mathbb{N}$	infinite sequence, index
$\mathcal{G}_i \to \mathcal{G}, \mathcal{G} = \lim \mathcal{G}_i$	convergence, limit
$C; D \subseteq \mathbb{R}^d$	constant; bounded convex region
$u, v, w; p, q$	vertices
$\varphi, \psi, \omega \in \mathbb{S}^{d-1}$	directions
r_i, r, R	positive extensions
$R_1 : \mathrm{ucell}(H) \to \mathbb{R}$	supremum extension function
$\mathbf{m}(u), \mathbf{d}_\mathcal{G}(u)$	motion vector, direction vector
$\ell(\mathcal{G}), \partial\ell(\mathcal{G})/\partial t$	length, derivative
α, α_t	angles

Table 1: Notation for geometric concepts, sets, functions, vectors, variables used in this paper.
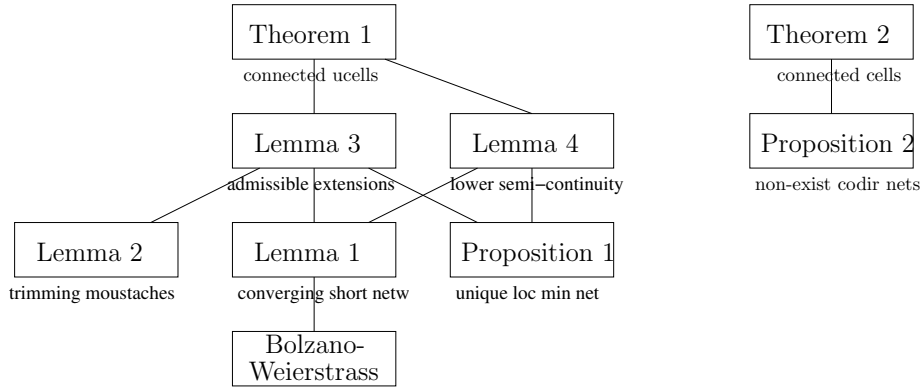
B Dependence

Fig. 7: Dependence of Lemmas and Theorems from top to bottom.