Geometric Steiner Trees

From the book: "Optimal Interconnection Trees in the Plane"

By Marcus Brazil and Martin Zachariasen

Part 4: Steiner Trees in Minkowski Planes

Marcus Brazil

2015

Part 4: Steiner Trees in Minkowski Planes

1. Properties of Minkowski planes
2. Steiner trees in Minkowski planes
3. Steiner points of degree 3
4. Steiner points of degree 4 or more

Minkowski spaces and norms

A *Minkowski space* is a real vector space V furnished with a *norm*, by which we mean a function $\|\cdot\| : V \to \mathbb{R}$ such that:

- $\|\mathbf{x}\| \leq 0$;
- $\|\mathbf{x}\| = 0$ if and only if $\mathbf{x} = \mathbf{0}$ (positive definiteness);
- $\|\lambda \mathbf{x}\| = |\lambda| \|\mathbf{x}\|$ for all $\lambda \in \mathbb{R}$ (symmetry);
- $\|\mathbf{x} + \mathbf{y}\| \geq \|\mathbf{x}\| + \|\mathbf{y}\|$ (triangle inequality).

A two-dimensional Minkowski space is usually called a *Minkowski plane*. If X is a Minkowski plane with norm $\|\cdot\|$ then the subsets $S_X := \{\mathbf{x} : \|\mathbf{x}\| = 1\}$ and $B_X := \{\mathbf{x} : \|\mathbf{x}\| \leq 1\}$ are called the *unit circle* and *unit disc* of X, respectively. If $\mathbf{x} \in S_X$ then $\mathbf{x}$ is said to be a *unit vector*.

Minkowski planes and unit balls

A *Minkowski plane* X can also be determined by its unit disk, rather than its norm. A unit disk B_X can be any subset of V such that:

- B_X is bounded;
- B_X has a non-empty interior;
- B_X is centrally symmetric;
- B_X is convex.

For a given unit disk B_X the norm can be defined as follows:

$$\|\mathbf{x}\| := \inf\{\mu : \mu > 0,\ \frac{\mathbf{x}}{\mu} \in B_X\}.$$

Examples of norms: ℓ_p norms

An ℓ_p norm ($p \geq 1$), $\|\cdot\|_p$, for any vector $\mathbf{x} = (x, y)$ is defined as follows:

$$\|\mathbf{x}\|_p := (|x|^p + |y|^p)^{1/p}.$$

When $p = 2$ this is the familiar Euclidean norm.

When $p = 1$ the ℓ_p norm is the *rectilinear norm* (or *Manhattan norm*) and has a polygonal unit disk:

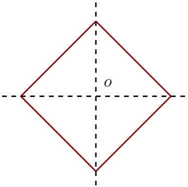

The figure below shows the unit disk when $p = 3/2$:

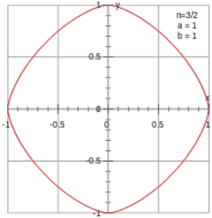

Examples of norms: gradient constrained norms

For any vector $\mathbf{x} = (x, y) \in \mathbb{R}^2$ let $g(\mathbf{x}) = y/x$ denote the *gradient* of $\mathbf{x}$. For a given maximum gradient m, the *gradient constrained norm* $\|\cdot\|_g$ corresponds to the Euclidean length of the shortest path from $(0,0)$ to (x, y) with gradient no greater than m.

$$\|\mathbf{x}\|_g = \begin{cases} \sqrt{x^2 + y^2}, & \text{if } |g(\mathbf{x})| \leq m; \\ \sqrt{1 + m^{-2}}|y|, & \text{if } |g(\mathbf{x})| \geq m. \end{cases}$$

It is easily checked that this defines a norm, as it is the maximum of two norms which are equal exactly when $|g(\mathbf{x})| = m$. Unit disks for $m = 1$ and $m = 1/3$ (respectively) are shown below.

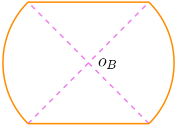

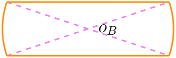

The Steiner tree problem in a Minkowski plane

In formal terms, the Steiner tree problem for a general norm can be stated as follows:

Minkowski Steiner tree problem in the plane

Given: A set of points $N = \{t_1, \ldots, t_n\}$ lying on a Minkowski plane with unit circle $\mathcal{C}$.
Find: A geometric network $T = (V(T), E(T))$, such that $N \subseteq V(T)$, and such that $\sum_{e \in E(T)} \|e\|$ is minimised.

The *terminals*, *minimum Steiner tree*, *Steiner points* and *Steiner topology* are all defined as in the Euclidean case.
All non-zero edges in a minimum Steiner tree are geodesics between their endpoints and hence can be embedded as line segments in the given plane. However, there may be many other possible embeddings of a minimum edge, such as a zigzig path, since the Minkowski norm is not necessarily strictly convex.

Full and fulsome Steiner trees

As in the Euclidean case, a given Steiner tree can be uniquely decomposed into *full components*, or *full Steiner trees*. This decomposition is not necessarily unique for a given terminal set. Indeed, the number of full components may not be unique for a given terminal set.

Definition: Fulsome trees

A Steiner tree T is said to be *fulsome* if it has the maximum possible number of full components amongst all Steiner trees with the same length as T for the given terminal set. Hence, a Steiner tree is full and fulsome if there is no Steiner tree with the same length on the same set of terminals with more than one full component.

Restricting our attention to fulsome Steiner trees means that the structure of each full component is as simple as possible.

Steiner configurations

Definition: Configurations

A *Steiner configuration* in a Minkowski plane is defined as a star with centre s and leaves $x_1, \ldots, x_m$ (with $s, x_1, \ldots, x_m$ all distinct) that is part of some minimum Steiner tree with Steiner point s. If a star is not necessarily part of some minimum Steiner tree, then it is simply referred to as a *configuration*.

Note that the word "minimum" is redundant in the above definition (see Problem Sheet 3).

The results in this Part are mainly concerned with local properties of a Steiner tree in a general Minkowski plane; more specifically, with Steiner configurations.

Pointed Steiner configurations

Theorem 22: Pointed Configuration Theorem

Let $\{sx_i : i = 1, \ldots, m\}$ be a configuration about s in a Minkowski plane. If there is a line L through s such that the interior of each segment sx_i is in the same open half-plane bounded by L, then $\{sx_i\}$ is not a Steiner configuration.

Proof (sketch): For the case $m = 3$, there exist collinear points x_i' in the interior of each sx_i, respectively. By the triangle inequality, $\|sx_1'\| + \|sx_3'\| \geq \|x_1'x_3'\|$, and by positive definiteness, $\|sx_2'\| > 0$. Hence, replacing the line segments $\{sx_i' : i = 1, \ldots, 3\}$ by $x_1'x_3'$ strictly reduces the length of the configuration (while maintaining connectivity), showing that $\{sx_i\}$ is not a Steiner configuration.
A similar argument applies if $m > 3$ (see Problem sheet).

Steiner points of degree 3

The following theorem shows that for many norms all Steiner points in a minimum Steiner tree have degree 3.:

Theorem 23

If T is a minimum Steiner tree in a smooth Minkowski plane, then every Steiner point of T has degree 3.

(For the proof, see "G. Lawlor and F. Morgan. Paired calibrations applied to soap films, immiscible fluids, and surfaces or networks minimizing other norms. Pacific Journal of Mathematics, 166: 55–83, 1994".)

Theorem 23 does not hold in general Minkowski planes, where the boundary of the unit ball may have non-differentiable points. However, as we will see later, in these cases higher degree Steiner points only occur in very special circumstances.

The centroid theorem

One of our aims now is to show that for a Steiner configuration in a smooth Minkowski plane, if we know the direction of one of the edges then the norm uniquely determines the directions of the other two edges.

Recall that the *centroid* of a triangle is the point where the three medians of the triangle intersect. It corresponds to the 'center of gravity' of the triangle. Note that the centroid divides each of the medians in a $2:1$ ratio.

Theorem 24: Centroid theorem

Let x_1, x_2 and x_3 be a set of leaves of a Steiner configuration in the Minkowski plane (with unit circle $\mathcal{C}$) with Steiner point s. Let x_1', x_2' and x_3' be the points where $s + \mathcal{C}$ intersects $\overrightarrow{sx_1}$, $\overrightarrow{sx_2}$ and $\overrightarrow{sx_3}$, respectively. Then for each $i \in \{1, 2, 3\}$ there exists a line L_i which is a supporting line of $z + \mathcal{C}$ at x_i', such that L_1, L_2 and L_3 form a triangle whose centroid coincides with s.

The centroid theorem: sketch proof

By the properties of the gradient of the gauge function in the Minkowski plane, it follows that there exist supporting lines L_i of $s + \mathcal{C}$ at each x_i, such that

$$\frac{\mathbf{u}_1}{h_1} + \frac{\mathbf{u}_2}{h_2} + \frac{\mathbf{u}_3}{h_3} = 0 \tag{1}$$

where each $\mathbf{u}_i$ is the outward normal vector to the supporting line L_i, and each h_i is the (Euclidean) distance from s to L_i.

Since a strictly positive linear combination of the vectors $\mathbf{u}_i$ equals $\mathbf{0}$, it follows that L_1, L_2 and L_3 form a triangle around $s + \mathcal{C}$, which we denote by Δ.

Let L_0 be the line through s parallel to L_1. For $i \in \{2, 3\}$, let p_i be the intersection of L_i and the line perpendicular to L_i through s, and let $y_i = L_i \cap L_0$.

The centroid theorem: sketch proof

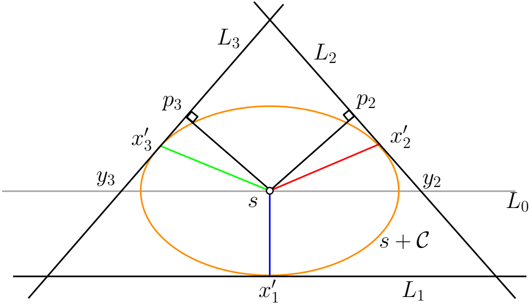

Note that $h_i = |sp_i|$. Equation (1) implies that $\mathbf{u}_2/h_2 + \mathbf{u}_3/h_3$ is perpendicular to L_0, and hence that

$$\frac{cos(\angle p_2 s y_2)}{|sp_2|} = \frac{cos(\angle p_3 s y_3)}{|sp_3|} \tag{2}$$

which implies $|y_2 s| = |y_3 s|$.

The centroid theorem: sketch proof

Hence s lies on the median of Δ through $L_2 \cap L_3$. By symmetric arguments, s also lies on the other two medians of Δ, and hence coincides with the centroid of Δ. **QED**

Definition: Centroid property

Given a Minkowski unit circle $\mathcal{C}$, we say that any set of supporting lines of $\mathcal{C}$ forming a triangle whose centroid is the centre of $\mathcal{C}$ satisfies the *centroid property*.

Theorem 24 shows that for any degree 3 Steiner configuration at a point s there exists a set of lines supporting $s + \mathcal{C}$ at points determined by the Steiner configuration that satisfies the centroid property.

A consequence of the centroid property is that in a smooth Minkowski plane the edges in a full Steiner tree use at most three directions. Before proving this we first establish 3 lemmas.

Centroid property lemmas: 1

The first lemma follows immediately from the observation that the centroid of a triangle divides each of its medians in the ratio $2:1$.

Lemma 25

Let L_1, L_2 and L_3 be three supporting lines of the unit circle $s + \mathcal{C}$, and let L_0 be the line that is parallel to L_1 and contains s. Let $d = \mathbf{d}(L_1, L_0)$, and define L to be the line that is parallel to L_1 at distance $3d$ from L_1, and at distance $2d$ from L_0. Let $y_2 = L_2 \cap L_0$ and let $y_3 = L_3 \cap L_0$. Then supporting lines L_1, L_2 and L_3 satisfy the centroid property if and only if (i) $L_2 \cap L_3$ lies on L, and (ii) $|sy_2| = |sy_3|$.

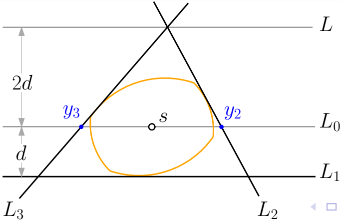

Centroid property lemmas: 2 and 3

In the following lemmas we continue to refer to L_0 and L as defined by Lemma 24.

Lemma 26

Let L_1, L_2 and L_3 be a set of supporting lines of $s + \mathcal{C}$ that fulfils the centroid property, and let L_0 be the line that is parallel to L_1 and contains s. Then neither L_2 nor L_3 supports $\mathcal{C}$ at a point that is strictly between L_1 and L_0.

Lemma 27

Let L_1 be a line that supports a unit circle $\mathcal{C}$ with centre s, and let L_0 be the line that is parallel to L_1 and contains s. If L_0 intersects $\mathcal{C}$ at a differentiable point, then there exists exactly one pair of supporting lines L_2 and L_3, such that L_1, L_2 and L_3 have the centroid property.

Lemma 27: sketch proof

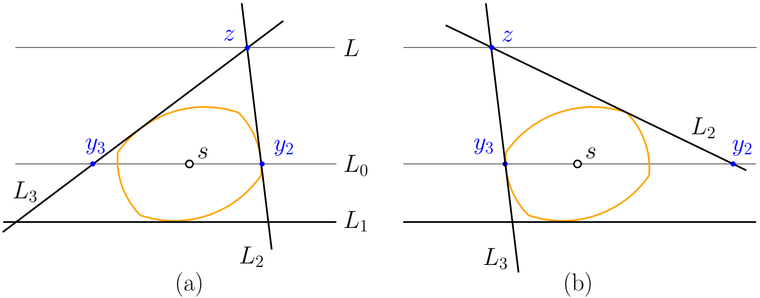

Illustration of proof of Lemma 27. (a) The initial positions of supporting lines L_2 and L_3. (b) The final positions of the supporting lines. The position for which the centroid property holds lies between these two extremes at the point where $|sy_2| = |sy_3|$.

Steiner trees in smooth Minkowski planes

The following corollary is an immediate consequence of Lemma 27.

Corollary 28

In a smooth Minkowski plane the directions of the edges in a degree 3 Steiner configuration are uniquely determined by the direction of any one edge.

Theorem 29

In a smooth Minkowski plane the edges of a full Steiner tree use at most three distinct directions.

The proof of Theorem 29 follows easily from Theorem 23 and Corollary 28, by contradiction.

Steiner points of degree ≥ 4

In this section we consider higher degree Steiner points in Minkowski planes which are not smooth. All of the results in this section are based on a *replacement principle*.

This replacement principle operates in one of two ways: either we replace certain line segments in a minimum tree T by new line segments with the same length and direction (and hence the same cost), or we replace a set of line segments by a new set of line segments which by minimality we can show have the same cost.

Theorem 30

Given a set of terminals N in a Minkowski plane, there exists a minimum Steiner tree T for N in which every vertex has degree at most 4.

Proof of Theorem 30

Let s be a Steiner point of degree 5 or more. There exists a set of three adjacent points to s, say x_1, x_2, x_3, such that all three lie in the interior of a half-plane induced by a line through s.

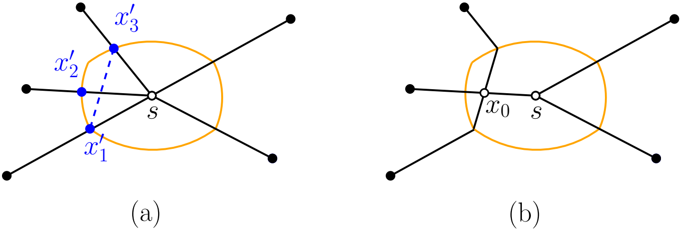

By the triangle inequality we can reduced the degree of s (by 2) by introducing a new degree 4 Steiner point x_0, without increasing the length of T.

Opposite pairs of edges at a Steiner point

In light of Theorem 30, we can now restrict our attention to Steiner points of degree 4.

Definition: Opposite pairs of edges

Let $\{sx_i : i = 1, \ldots, 4\}$ be a degree 4 Steiner configuration around s where the neighbours of s are indexed in counter-clockwise order around s. We say that such a Steiner configuration consists of two *opposite pairs of edges*, $\{sx_1, sx_3\}$ and $\{sx_2, sx_4\}$.

Lemma 31

In a degree 4 Steiner configuration in a Minkowski plane one of the opposite pairs of edges is collinear.

Proof: Let $\{sx_i : i = 1, \ldots, 4\}$ be a degree 4 Steiner configuration around a Steiner point s. We assume that neither of the opposite pairs of edges is collinear, and obtain a contradiction by the replacement principal.

Proof of Lemma 31

We show non-minimality in two cases, first where there is a meeting angle of π:

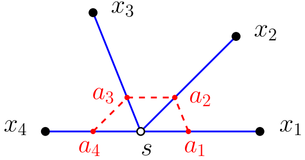

and second where each meeting angle is $< \pi$:

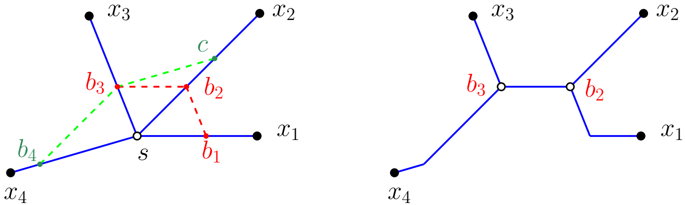

Properties of the second opposite pair of edges

Definitions: First and second opposite pairs

Let $\{sx_i : i = 1, \ldots, 4\}$ be a degree 4 Steiner configuration around s where the neighbours of s are indexed in counter-clockwise order around s. One of the opposite pairs of edges, say $\{(s, x_1), (s, x_3)\}$, must be collinear, by Lemma 31; we refer to this as the *first* opposite pair of edges. The other pair of edges, $\{(s, x_2), (s, x_4)\}$, is called the *second* opposite pair of edges.

Note that above classification into pairs is not necessarily unique.

Lemma 32

Suppose the second opposite pair of edges $\{(s, x_2), (s, x_4)\}$ around a degree 4 Steiner point s in a Steiner configuation $\{sx_i : i = 1, \ldots, 4\}$ is not collinear. Then there exists a point s' in the interior of sx_1 or sx_3 such that for every point $s_0 \in ss'$ we have $\|sx_2\| = \|s_0x_2\|$ and $\|sx_4\| = \|s_0x_4\|$.

The proof of this lemma is straightforward (see Problem sheet).

Splitting degree 4 Steiner points

A consequence of Lemma 32 is that if the second opposite pair of edges is not collinear, then we can split the Steiner point s into a pair of adjacent degree 3 Steiner points:

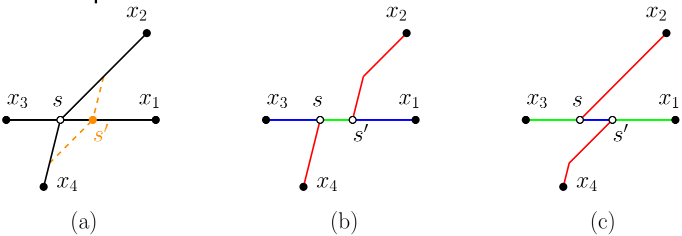

Definition: Cross

A *cross* is a degree 4 Steiner point where both the first and second opposite pairs of edges are collinear.

Thus far we have shown that unless a degree 4 Steiner point is a cross, we can always split it into two adjacent degree 3 Steiner points.

The Sliding Lemma

Lemma 33: Sliding lemma

Let $e = (s_1, s_2)$ be an edge connecting two Steiner points (s_1 and s_2) in a fulsome minimum Steiner tree T. Let $e_1 = (s_1, v_1)$ be the next edge incident with s_1 (in counter-clockwise order around s_1 from e), and let $e_2 = (s_2, v_2)$ be the next edge incident with s_2 (in clockwise order around s_2 from e. Then θ, the angle at s_1 between e and e_1, and ϕ, the angle at s_2 between e and e_1, satisfy $\theta + \phi > \pi$.

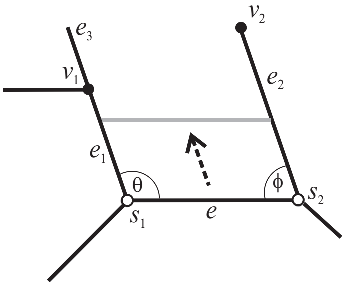

The idea is to show that if $\theta + \phi \leq \pi$, e can slide as shown until it meets a terminal, contradicting fulsomeness.

Degree 4 Steiner points can almost always be split

Theorem 34

In a fulsome minimum Steiner tree, a degree 4 Steiner point s can always be split into two adjacent degree 3 Steiner points *unless* it is a cross and is adjacent to terminals only.

It follows that in a GeoSteiner-type algorithm for constructing a minimum Steiner tree one can, for the most part, limit the construction of candidate full Steiner components to full and fulsome minimum Steiner trees where all Steiner points have degree 3. For many Minkowski norms, including the *rectilinear norm* and the *gradient constrained norm*, there exist efficient methods for constructing such candidates. (The construction of a cross with terminals as neighbours can easily be handled separately.)

Proof of Theorem 34

By the comments on slide 25, we only need to consider the case where s is a cross. Assume that one of the neighbours of s is a Steiner point v with degree 3. Then, the local perturbation shown splits s into two adjacent Steiner points without changing the length of the tree.