Datalogisk Institut
KØBENHAVNS UNIVERSITET

Flexibility of Steiner Trees in Uniform Orientation Metrics

Marcus Brazil, Pawel Winter and Martin Zachariasen

Technical Report no. 2004/09
ISSN: 0107-8283
CR Subject Classification: G.2.2

DIKU
University of Copenhagen • Universitetsparken 1
DK-2100 Copenhagen • Denmark

Flexibility of Steiner Trees in Uniform Orientation Metrics*

M. Brazil † P. Winter‡ M. Zachariasen‡

June 21, 2004

Abstract

We present some fundamental flexibility properties for minimum length networks (known as Steiner minimum trees) interconnecting a given set of points in an environment in which edge segments are restricted to λ uniformly oriented directions. These networks are referred to as λ-SMTs. They promise to play an increasingly important role in the future of optimal wire routing in VLSI physical design, particularly for the next generation of VLSI circuits.

In this paper we develop the concept of a flexibility polygon for a λ-SMT, which is a region representing the union of all λ-SMTs with the same topology on a given set of points. We show that this polygon can be constructed, for a given point set and given topology, in linear time. We discuss some of the future applications of this polygon, which can be thought of as a geometric representation of the amount of flexibility inherent in a given λ-SMT.

Keywords: Interconnection networks, uniform orientation Steiner trees, flexibility, VLSI design.

---

*Partially supported by a grant from the Australia Research Council and by a grant from the Danish Natural Science Research Council (51-00-0336).

†ARC Special Research Centre for Ultra-Broadband Information Networks, Department of Electrical and Electronic Engineering, The University of Melbourne, Victoria 3010, Australia

‡Department of Computer Science, University of Copenhagen, DK-2100 Copenhagen Ø, Denmark

1 Introduction

Interconnects in VLSI design have traditionally used rectilinear (or Manhattan) routing, in which only two perpendicular wiring directions are allowed. Recent technological advances in microchip fabrication have seen an increasing interest in the use of non-rectilinear interconnect architectures in VLSI design. The two alternative architectures that have generated the most interest in recent years are the Y-architecture [4, 5], in which there are three directions for interconnects differing by angles of $2\pi/3$, and the X-architecture [11, 13], in which the rectilinear architecture is supplemented by the pervasive use of diagonal interconnects (at an angle of $\pi/4$ to the rectilinear interconnects). Both traditional rectilinear routing and these new proposed architectures are examples of so-called λ-geometry, in which a fixed set of $\lambda \geq 2$ uniformly oriented directions are allowed.

In VLSI routing, one of the principal objectives is to minimize the total length of interconnections in a net, that is, to compute a λ-geometry *Steiner minimum tree* (or λ-SMT). This is in general an NP-hard problem. However, a surprisingly powerful exact algorithm, GeoSteiner, has been developed for this problem which can find an optimal solution for hundreds of randomly distributed points for arbitrary λ [7]. The key to this algorithm is to exploit the strong geometric structural properties of λ-SMTs. A recent paper [3] has established canonical forms for λ-SMTs which has lead to a further speed-up of GeoSteiner.

In this paper we use these properties to explore the notion of flexibility in a λ-SMT. Informally, this is a measure of the extent to which edges in the minimum length network can be perturbed without increasing the length of the network. This has important applications in solving multi-objective optimisation problems in VLSI physical design, involving minimising the negative effects of properties such as congestion or signal delay as a secondary objective. The concept was introduced for rectilinear Steiner trees in [1] and [9]. Here we provide an effective measure of flexibility for a much wider class of minimum length networks, by defining and constructing the *flexibility polygon* of a λ-SMT.

The paper is organised as follows. In the next section we review some of the important basic properties of any λ-SMT T. We also develop some fundamental properties of 0-shifts, which are special perturbations of the variable points of T which do not increase the length of T. In Section 3 we define the flexibility polygon and show how its boundary can be described in terms of so-called rightmost concave paths, which are paths obtained from paths in T via 0-shifts. A linear time algorithm for constructing the flexibility polygon is given in Section 4. In Section 5 we discuss the flexibility of single Steiner points in a λ-SMT. Concluding remarks, including a discussion of applications, are given in Section 6.

2 Preliminaries

2.1 Properties of λ-SMTs

We begin by establishing some basic definitions and notation, and reviewing some important properties of λ-SMTs. For a more detailed discussion of these properties, see [2, 3].

Let $\lambda \geq 2$ be a given integer. Given λ orientations $i\omega$ $(i = 1, 2, ..., \lambda)$ in the Euclidean plane, where $\omega = \pi/\lambda$ is a unit angle, we represent these orientations by the angles with the x-axis of corresponding straight lines. A line or line segment with one of these orientations is said to be in a *legal direction*. Objects composed of line segments in legal directions are said to belong to a λ-*geometry*.

Since a minimum length network is necessarily a tree, we will only discuss networks in λ-geometry that are trees. We define a λ-tree to be a tree network in λ-geometry interconnecting a given set of points N, also denoted *terminals*. A λ-tree can contain nodes of degree 3 or more that are not terminals. These nodes are called *Steiner points*. Together the terminals and Steiner points are referred to as the *vertices* of the λ-tree.

The graph structure (i.e., pattern of adjacencies of the vertices) of a λ-tree is referred to as its *topology*. In this paper we are concerned with λ-trees T whose total edge length is minimum for a given set of terminals N; these are the λ-*geometry Steiner minimum trees* (λ-SMTs). If the total edge length of T is locally minimum, in that it is minimum for a given topology $\mathcal{T}$, then we say that T is a λ-SMT for $\mathcal{T}$.

Any λ-SMT T can be decomposed into a union of full subtrees meeting only at terminals. These subtrees are referred to as the *full Steiner trees (FSTs)* of T. A λ-SMT T for N is *fulsome* if the number of FSTs is maximized over all λ-SMTs for N. In particular, for a full topology $\mathcal{T}$, a λ-SMT T for N and $\mathcal{T}$ that is fulsome cannot be replaced by two or more FSTs with the same total length (and, hence, fulsomeness is in a sense a property of N and $\mathcal{T}$). For any set of terminals there always exists a λ-SMT in which every full Steiner tree is fulsome.

As in [3], we will focus our attention on full and fulsome λ-SMTs, which form the building blocks for all λ-SMTs. For a set of terminals N and a full topology $\mathcal{T}$ for N, we denote by $S(N, \mathcal{T})$ the *set* of all full and fulsome λ-SMTs (or FSTs) interconnecting N and having topology $\mathcal{T}$. In view of the above, in the remainder

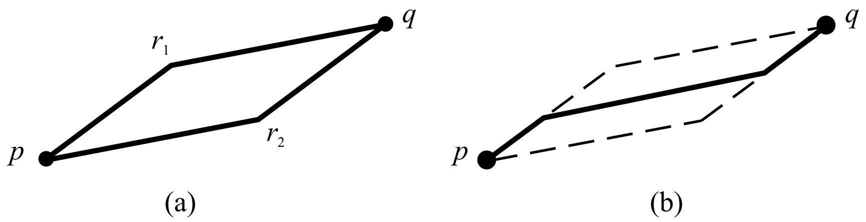

Figure 1: Embeddings of a bent edge in the Euclidean plane. The simplest embeddings (consisting of two half-edges) are shown on the left.

of the paper we assume that $S(N, \mathcal{T}) \neq \emptyset$.

In [2, 10] it was shown that all Steiner points in a λ-SMT have degree 3 or 4. Furthermore, if a Steiner point with degree 4 exists in the λ-SMT, then it spans exactly four terminals (it is a cross); such a cross has no flexibility whatsoever, since any movement of the Steiner point increases the length of the tree [3]. Thus we will assume throughout this paper that every Steiner point has degree 3. The topology of any λ-SMT with degree 3 Steiner points is referred to as a *Steiner topology*.

Edges in a λ-SMT are geodesics (in λ-geometry) between their endpoints. We refer to such an edge as a *straight edge* if it is a single straight line segment (in a legal direction), or else as a *bent edge* if it consists of two or more straight line components. It has been shown in [12] that bent edges are composed of line segments in exactly two legal directions differing by an angle of ω. Furthermore, although there are infinitely many ways of embedding a bent edge pq in the Euclidean plane, there are only two embeddings composed of exactly two straight line segments, as shown in Figure 1(a). Each of these contains two edges of the parallelogram pr_1qr_2 where the interior angles at r_1 and r_2 are $\pi - \omega$. The straight line components in such an embedding are referred to as *half-edges* and the points r_1 and r_2 as *corner points*.

We now consider some important edge-direction properties in full λ-SMTs. In [3] it was shown that the straight edges and half-edges in a full λ-SMT can be oriented and then partitioned into three equivalence classes, such that each equivalence class contains oriented line segments in at most two directions differing by an angle of ω. In each equivalence class, the right-most edges (or half-edges) are labeled as *primary*, and the left-most ones as *secondary*. Two important properties of this labeling are as follows:

- If λ is a multiple of 3 then there are exactly two feasible directions in each equivalence class. If λ is a not multiple of 3 then one of the equivalence classes contains two feasible directions and the other two classes each contain only one feasible direction, which is said to be both primary and secondary. A primary edge (or half-edge) that is not secondary is said to be *exclusively primary*. Similarly, a secondary edge (or half-edge) that is not primary is said to be *exclusively secondary*. This is illustrated for $\lambda = 6, 7$ and 8 in Figure 2.

- If λ is a multiple of 3 then two exclusively primary or exclusively secondary edges meet at a Steiner point at an angle of $2\pi/3$. If λ is a not multiple of 3 then any pair of edges meet at a Steiner point at an angle that differs from $2\pi/3$ by no more than ω.

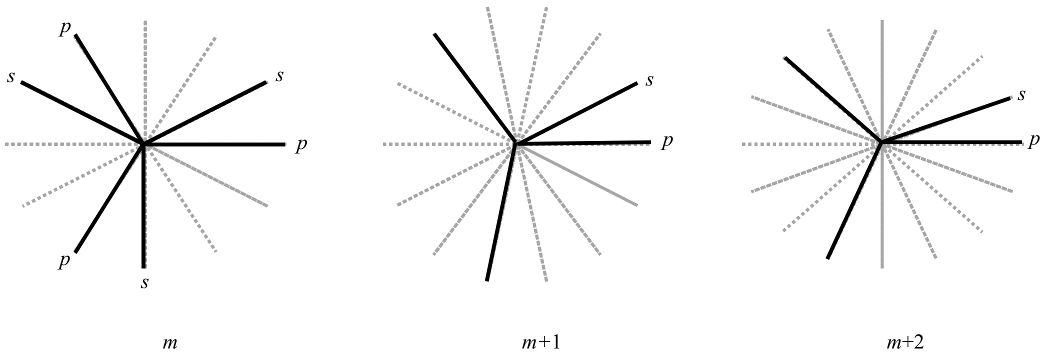

Figure 2: The feasible directions (up to a rotation by a multiple of ω) for edges in a full λ-SMT for $m = 2$. Exclusively primary and secondary directions are indicated by p and s respectively.

Primary and secondary edges of a λ-SMT play a crucial role in determining flexibility. This is due to their connection with the 0-shifts of the paths in the tree, which we define below.

We define a *shift* of a straight edge pq to be a move of p to a new point $p' \neq p$ and a simultaneous move of q to a new point $q' \neq q$ such that $p'q' \parallel pq$. Similarly, a *shift* of a bent edge pq is defined to be a move of p to p' and a simultaneous move of q to q' such that $p'q'$ is either a bent edge with components in the same directions as those in pq or a straight edge whose direction is the same as that of one of the components of pq. The concept of a shift can be generalised to a path P in a full λ-SMT T as follows. A *shift* of $P = ps_1s_2 \dots q$ is a perturbation of T

that moves the internal Steiner points s_i of P to s'_i (and fixes all other nodes of T) such that the following conditions are satisfied:

(1) each internal Steiner point s_i of P moves along the line through s_i containing the straight edge or half-edge of T incident to s_i not lying on P; and

(2) the shift of P induces a shift on each internal edge of P.

Note that the effect of a shift is that it does not change the direction of any straight edge of T except possibly the first and last edges of P.

Given a subpath of a full λ-SMT T, we define a shift on that subpath to be a 0-*shift* if the shift does not increase the length of T. One of the key results on 0-shifts proved in [3] is the following.

Proposition 2.1 *Given an exclusively primary edge or half-edge e_1 and an exclusively secondary edge or half-edge e_2 in a full λ-SMT, there exists a 0-shift on the path between the e_1 and e_2. Furthermore, 0-shifts occur only if e_1 is an exclusively primary edge and e_2 is an exclusively secondary edge, or vice versa.*

2.2 Further properties of 0-shifts

We now derive some new properties of 0-shifts which will form the basis for our later results. A *fundamental* 0-shift is a (non-trivial) 0-shift that moves as few Steiner points as possible. If $\lambda = 3m$ (for some positive integer m) then a fundamental 0-shift moves one Steiner point; if $\lambda = 3m + 1$ or $3m + 2$ then a fundamental 0-shift moves two adjacent Steiner points. The fundamental 0-shifts are illustrated in Figures 3 and 4. Now we have the following elementary decomposition theorem.

Theorem 2.2 *Let T be a given λ-SMT. Any 0-shift in T can be decomposed into a sequence of fundamental 0-shifts in T.*

Proof. Consider a 0-shift between two edges of T, e_1 and e_2, where, without loss of generality, e_1 is exclusively primary and e_2 is exclusively secondary. We prove the theorem by induction on the number of edges of the path P in T from e_1 to e_2. The base case occurs when the 0-shift is fundamental, in which case the theorem is trivially true.

So, suppose the 0-shift is not fundamental. If $\lambda = 3m$, let e_0 be a straight edge or half-edge of P lying strictly between e_1 and e_2. Such an e_0 exists since the 0-shift is not fundamental. If $\lambda \neq 3m$, let s be a Steiner point of P which is not an end point of e_1 or e_2, and let e_0 be an exclusively primary or exclusively

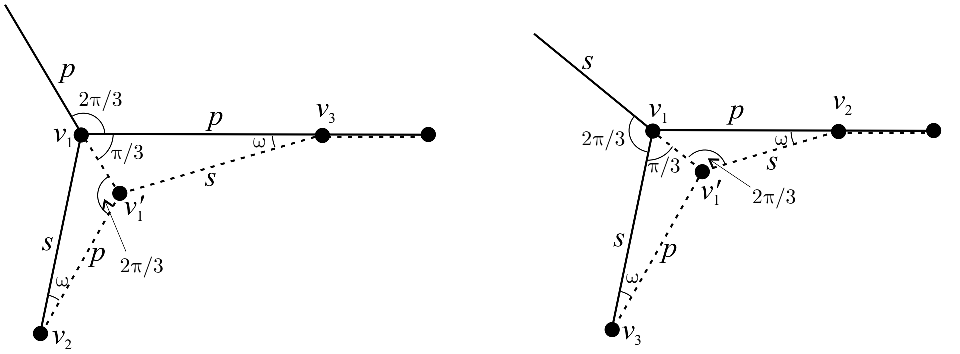

(a) (b)

Figure 3: The fundamental 0-shifts for the case where $\lambda = 3m$. Exclusively primary and exclusively secondary edges are labeled p and s respectively.

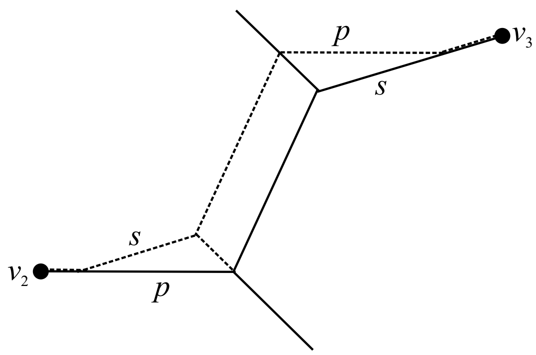

Figure 4: The fundamental 0-shift for the cases where $\lambda \neq 3m$. Exclusively primary and exclusively secondary edges are labeled p and s respectively.

secondary edge or half-edge incident with s. In both cases, note that the paths from e_1 to e_0 and from e_0 to e_2 contain fewer edges than P.

Suppose e_0 is exclusively secondary. Then there exists a 0-shift between e_1 and e_0 which transfers an arbitrarily small amount of exclusively primary material to e_0. One can then perform a 0-shift between e_0 and e_2 transferring all this exclusively primary material in e_0 to e_2. Hence these two 0-shifts give a 0-shift between e_1 and e_2. Thus the 0-shift can be decomposed into two shifts on paths containing fewer edges, completing the induction.

Finally, if e_0 is exclusively primary, a similar argument applies, by swapping the roles of e_1 and e_2 above. ∎

Consider the set $S(N, \mathcal{T})$ of full and fulsome λ-SMTs for a terminal set N and full Steiner topology $\mathcal{T}$. For $T \in S(N, \mathcal{T})$, define $p(T)$ to be the sum of the lengths of all exclusively primary edges in T and $s(T)$ to be the sum of the lengths of all exclusively secondary edges in T. We refer to $p(T)$ and $s(T)$ as, respectively, measures of the amount of exclusively primary material and exclusively secondary material in T. The following theorem shows that $p(T)$ and $s(T)$ depend only on N and $\mathcal{T}$.

Theorem 2.3 *Let $T \in S(N, \mathcal{T})$ be a full and fulsome λ-SMT with terminal set N and topology $\mathcal{T}$. Applying a 0-shift to T does not change the value of $p(T)$ or $s(T)$. Furthermore, $p(T)$ and $s(T)$ are dependent only on N and $\mathcal{T}$, not on the choice of T.*

Proof. For each of the fundamental 0-shifts in Figures 3 and 4 we show directly that $p(T)$ and $s(T)$ are unchanged. The first statement of the theorem then follows immediately from Theorem 2.2. For the second statement of the theorem, we recall from [3, Section 5] that if two λ-SMTs share the same terminal set N and topology $\mathcal{T}$ then there exists a series of 0-shifts to transform one of the trees to the other. Hence, for the second statement it also suffices to show that $p(T)$ and $s(T)$ are unchanged by fundamental 0-shifts.

If $\lambda = 3m$ the two types of fundamental 0-shifts are shown in Figure 3. To see that in each case $s(T)$ is unchanged by this shift, note that

$$\frac{|v_1 v_2|}{\sin 2\pi/3} = \frac{|v_1 v_1'|}{\sin \omega} = \frac{|v_1' v_3|}{\sin \pi/3}.$$

Hence $|v_1 v_2| = |v_1' v_3|$, giving the required result. It also immediately follows that $p(T)$ is unchanged by this shift from the fact that the shift is a 0-shift and $|T| = p(T) + s(T)$.

If $\lambda \neq 3m$ the fundamental 0-shift is shown in Figure 4, and it is clear, from inspection, that $p(T)$ and $s(T)$ are unchanged. ∎

3 The Flexibility Polygon and its Properties

In this section we establish the basic properties of the flexibility polygon, a geometric object which allows us to determine the degree of flexibility available in a given λ-SMT. This polygon indicates the extent to which edges in the tree can be moved using 0-shifts without changing the length of the tree. More precisely, it places tight bounds on the possible positions of the edges of the trees in $S(N, \mathcal{T})$, the set of full and fulsome λ-SMTs for a given terminal set N and full Steiner topology $\mathcal{T}$. Recall that we assume that $S(N, \mathcal{T}) \neq \emptyset$; furthermore, we assume, without loss of generality, that $|N| \geq 3$ and that every Steiner point in $\mathcal{T}$ has degree 3. Note that all λ-SMTs in $S(N, \mathcal{T})$ use the same set of edge directions, that is, every tree can be obtained from another tree using 0-shifts [3].

The *flexibility polygon* $F(N, \mathcal{T})$ for terminal set N and topology $\mathcal{T}$ is defined to be the *union* of the embeddings of all λ-SMTs in $S(N, \mathcal{T})$. We will show that this union forms a simply connected region with a polygonal boundary whose vertices include the terminals N. Some examples of flexibility polygons for the cases where $\lambda = 4$ and $\lambda = 6$ are shown in Figure 5. Notice that in some cases parts of the flexibility polygon may degenerate into single edges (indicating that some edges may exhibit no flexibility at all).

Consider a counter-clockwise outer walk of some $T \in S(N, \mathcal{T})$, beginning and ending at the same terminal. This walk allows us to place a cyclic ordering, $t_1 (= t_{n+1}), t_2, \ldots, t_n$ on the terminals of T. We define the *concave paths* in T to be the paths between t_i and t_{i+1} (for $i = 1, \ldots, n$). In other words, these are paths between terminals where at each intermediate Steiner point we leave using the rightmost outgoing edge. Clearly, the set of all concave paths of T contains every edge of T exactly twice — once in each direction. In fact, up to the starting terminal, the order in which edges of the tree are visited by this outer walk of the tree is unique for a given terminal set N and topology $\mathcal{T}$; this holds since all λ-SMTs in $S(N, \mathcal{T})$ can be obtained from each other using 0-shifts which do not change the ordering of the edges meeting at Steiner points.

In this section we will show that, intuitively, the flexibility polygon $F(N, \mathcal{T})$ can be constructed by pushing every concave path in T as far to the right as possible using 0-shifts. The resulting paths are referred to as *rightmost* concave paths, and in Theorem 3.4 it is shown that the union of these paths forms the boundary

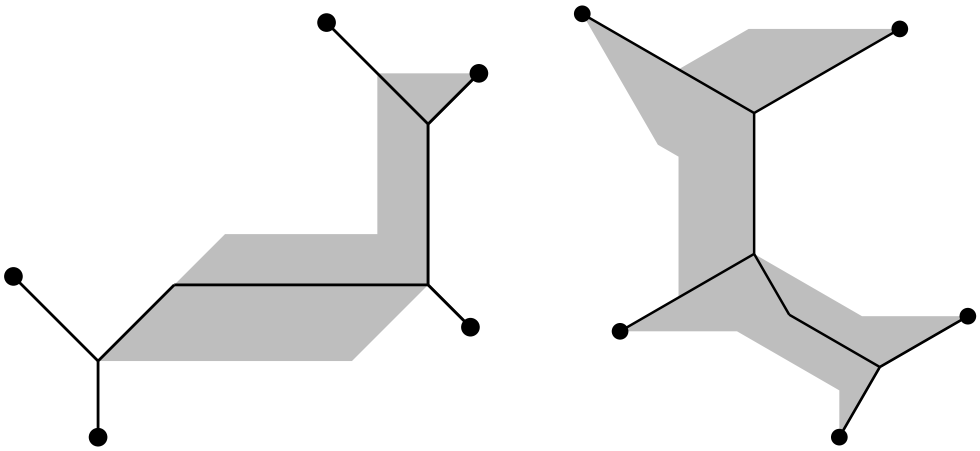

Figure 5: Examples of λ-SMTs and flexibility polygons.

of $F(N, \mathcal{T})$.

3.1 Flexibility of edges

We begin by studying the flexibility of edges of trees in $S(N, \mathcal{T})$. Consider a pair of distinct λ-SMTs $T, T' \in S(N, \mathcal{T})$. Let w be a Steiner point or corner point in T, and let w' be the corresponding node in T' (if w is a corner point and the corresponding edge in T' is a straight edge, then w' is just one of the endpoints of this edge). Let T_w be one of the maximal subtrees of T having w as a leaf, and assume that T_w consists of *straight* edges only; let $T'_{w'}$ be the corresponding subtree in T'. Finally, assume that for every *secondary* edge uv in T_w the corresponding edge $u'v'$ in $T'_{w'}$ is *also* secondary.

Lemma 3.1 *Let T, T_w, T' and $T'_{w'}$ be defined as above. Let the edges in T_w and $T'_{w'}$ be oriented towards their roots w and w', respectively. Let uv be an oriented edge in T_w, and let $u'v'$ be the corresponding oriented edge in $T'_{w'}$. Then no part of $u'v'$ is to the right of the oriented line through u and v.*

Proof. Consider a depth-first traversal of T_w and the associated *post-ordering* of the edges. We prove the statement by induction on this ordering. Consider an oriented edge uv in T_w, where v is closer to the root than u. For the base case of

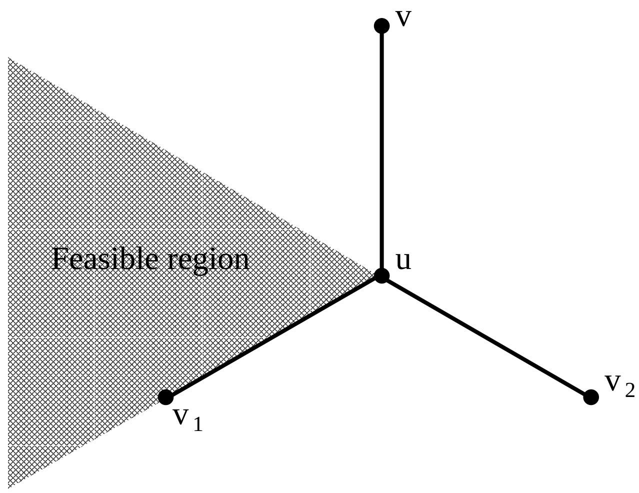

Figure 6: The feasible region for u'.

the induction — the first edge given by the ordering — u must be a terminal. If uv is a primary edge, then the statement trivially holds; if uv is a secondary edge, then $u'v'$ must also be a secondary edge, and the statement holds again.

For the inductive step, assume that u is a Steiner point (otherwise, u is a terminal, and the statement holds by the above). Consider the children v_1 and v_2 of u; by induction, we know that the statement holds for the edges v_1u and v_2u. Therefore, the Steiner point u' can neither be to the right of the oriented line through v_1 and u, nor to the right of the oriented line through v_2 and u; the feasible region for u' is shown in Figure 6. Now, if uv is a primary edge, the statement clearly holds for uv; if uv is a secondary edge, then $u'v'$ must also be secondary edge, and the statement holds again. ∎

We say that T_w is a *primary* subtree if all edges in T_w are primary edges, that is, use primary edge directions only. We define a *secondary* subtree analogously.

Corollary 3.2 *Let T_w be any primary (resp. secondary) subtree of T with some Steiner point or corner point w as root and terminals as leaves. Let the edges in T_w be oriented towards the root. Let uv be an oriented edge in T_w, and let $u'v'$ be the embedding of the same edge in any tree $T' \in S(N, \mathcal{T})$. Then no part of $u'v'$ is to the right (resp. left) of the oriented line through u and v.*

3.2 Flexibility of paths

Consider a path $P = v_1 v_2 \dots v_{k-1} v_k$ connecting two *terminals* v_1 and v_k in T. For $i = 2, \dots, k-1$, let T_i denote the maximal subtree of T rooted at Steiner point v_i and not containing any edges of P (see Figure 7).

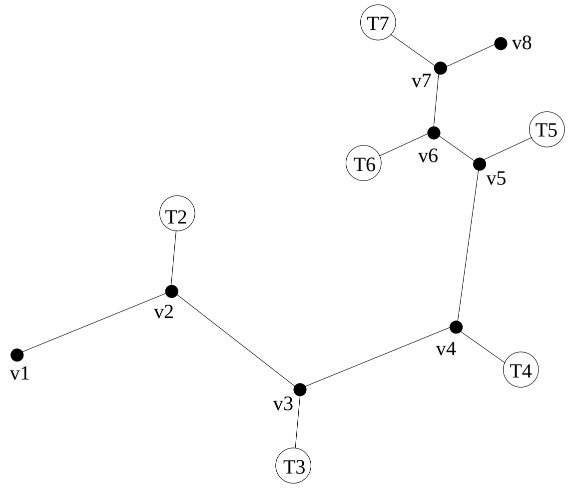

Figure 7: The path P and its associated subtrees T_i.

We number the m edges of T by making a depth-first traversal from v_1. At every Steiner point v_i the subtree T_i is traversed before the edge $v_i v_{i+1}$ is traversed. We call this a depth-first traversal from v_1 along P. Note that the numbering of the edges in P depends only on the topology $\mathcal{T}$. By [3] there exists an integer k, with $1 \leq k \leq m$, and a λ-SMT $T^r \in S(N, \mathcal{T})$ such that all edges in T^r numbered less than k are primary, and all edges numbered greater than k are secondary. The edge numbered k is the (possibly) bent edge.

We note that the tree T^r is in fact *unique* — something that holds for any given numbering of the edges. Assume to the contrary that there exists another integer k', with $1 \leq k' < k \leq m$, and a corresponding λ-SMT $T^{r'} \in S(N, \mathcal{T})$. Let w be the corner point of the bent (k'th) edge in T^r. By applying Lemma 3.1 to the two maximal subtrees of T^r rooted at w, we conclude that $T^{r'}$ cannot have any of its edges to the right of the corresponding edges in T^r. But this makes it impossible to embed the k'th edge of $T^{r'}$, which is a (straight) secondary edge — a contradiction.

Let $P^r = v_1 v_2^r \dots v_{k-1}^r v_k$ be the path in T^r from v_1 to v_k. As shown above, P^r is uniquely defined for any depth-first traversal of T from v_1 along P. In other words, the coordinates of the vertices of P^r do not depend on the choice of depth first traversal in each of the subtrees T_i. We say that P^r is the *rightmost* path from v_1 to v_k.

Theorem 3.3 *Consider an oriented straight edge (or half-edge) uv on a rightmost path P^r from v_1 to v_k; let $u'v'$ be the embedding of the same edge in some λ-SMT in $S(N, \mathcal{T})$. Then no part of $u'v'$ is to the right of the oriented line through u and v.*

Proof. The shape of a rightmost path depends on the location of the bent edge in T^r. If the bent edge is located on P^r, say on edge $v_i^r v_{i+1}^r$, as in Figure 8, all subtrees $T_2^r, \ldots T_i^r$ will clearly be primary subtrees while all subtrees $T_{i+1}^r, \ldots, T_{k-1}^r$ will be secondary subtrees. Alternatively, if the bent edge is located in some sub-

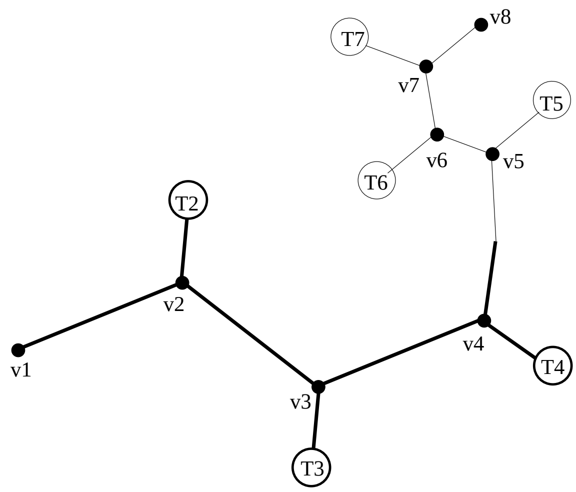

Figure 8: The case where the bent edge lies on P^r. The primary edges and subtrees are indicated by the heavier lines.

tree T_i, all edges in P^r are straight edges, all subtrees $T_2^r, \ldots T_{i-1}^r$ primary subtrees, and all subtrees $T_{i+1}^r, \ldots, T_{k-1}^r$ secondary subtrees (see Figure 9). In both cases there exists on P^r a node w, which may be either a corner point or Steiner point, such that the subtree of T having w as root and containing v_1 is *primary*, while the subtree containing v_k is *secondary*. The theorem follows by applying Corollary 3.2. ∎

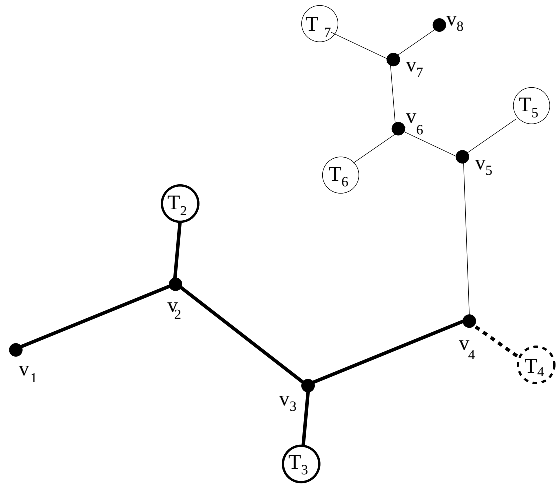

Figure 9: The case where the bent edge does not lie on P^r. Here subtree T_4 may contain a combination of primary and secondary edges.

3.3 Flexibility polygon

The rightmost path of a *concave* path is called a *rightmost concave path*. We now show that the flexibility polygon $F(N, \mathcal{T})$ can be described in terms of the rightmost concave paths for N and $\mathcal{T}$ (Figure 10).

Theorem 3.4 *Let N be a set of terminals and $\mathcal{T}$ a full Steiner topology for N. Then the flexibility polygon $F(N, \mathcal{T})$ is a simply connected region whose boundary is the union of the rightmost concave paths for N and $\mathcal{T}$.*

Proof. It immediately follows from Theorem 3.3 that the outer boundary of $F(N, \mathcal{T})$ is the union of the rightmost concave paths for N and $\mathcal{T}$.

Let T be a λ-SMT such that $T \in S(N, \mathcal{T})$. To see that $F(N, \mathcal{T})$ is simply connected, consider, for each pair of terminals t_i, t_{i+1} which are adjacent with respect to boundary order, the region $F^r(P)$ between the concave path $P = t_i \ldots t_{i+1}$ of T and the corresponding rightmost concave path P^r. The region $F^r(P)$ is enclosed by the closed curve formed by P and P^r. We will now show that there exists a λ-SMT in $S(N, \mathcal{T})$ that intersects every point in $F^r(P)$. Clearly, this holds for all points on the boundary, that is, which are on either on P or P^r.

Therefore, consider a point p in the interior of $F^r(P)$. A sequence of 0-shifts that transforms P to P^r can be considered to be a continuous “contraction” of the closed curve given by P and P^r. Since p is not contained in the region obtained

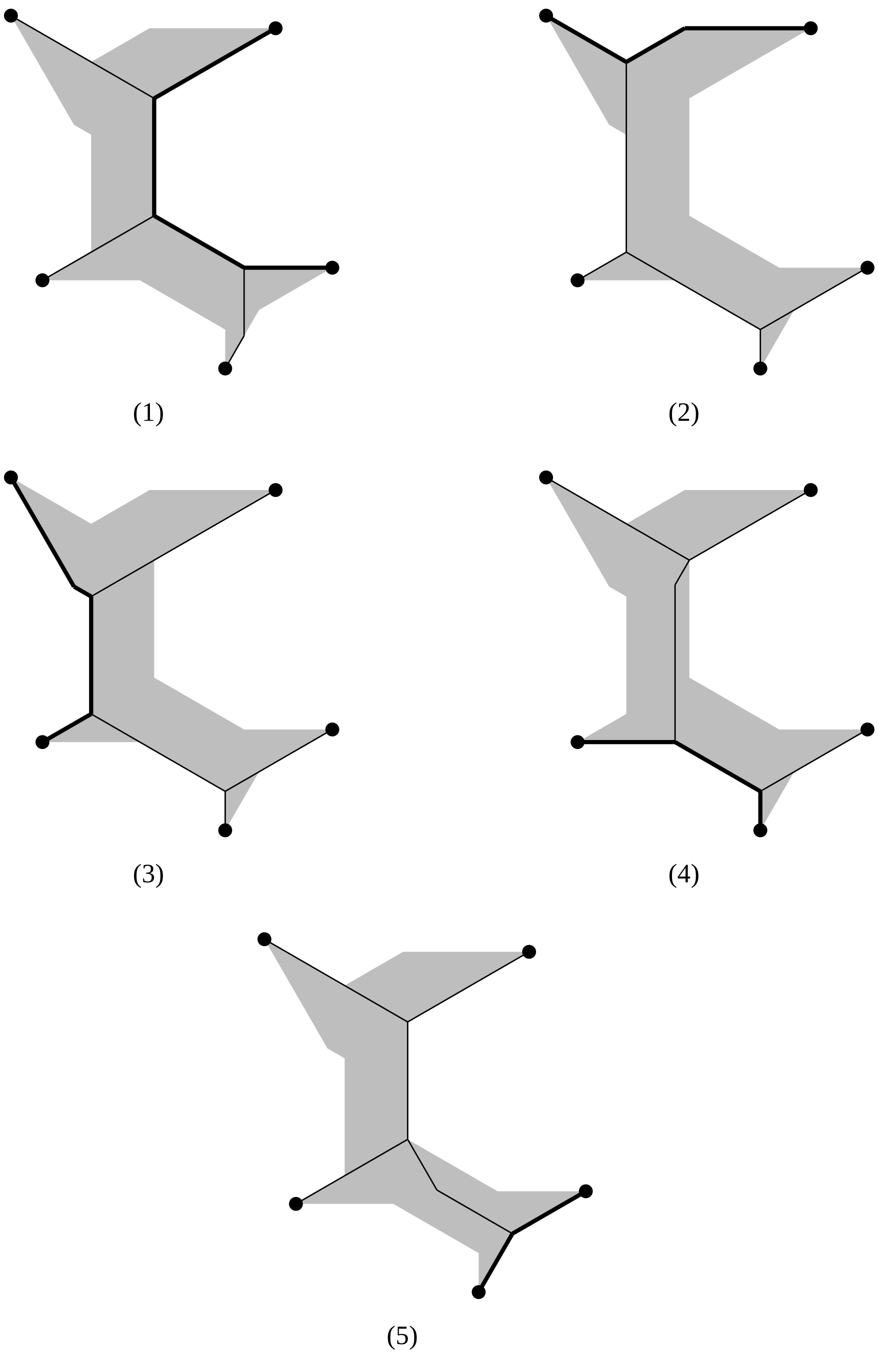

Figure 10: Rightmost concave paths defining a flexibility polygon for $\lambda = 6$. For each rightmost concave path, a λ-SMT realizing this path is shown.

(which is P^r), there must be a point in time where the curve intersects p. Thus there exists a λ-SMT in $S(N, \mathcal{T})$ that intersects every point in $F^r(P)$.

The union of all regions $F^r(P)$ taken over all concave paths P of T is the entire region bounded the rightmost concave paths for N and $\mathcal{T}$. ∎

This characterisation of the boundary of a flexibility polygon in terms of rightmost concave paths immediately implies some nice properties of the flexibility polygon which are summarised in the following corollary.

Corollary 3.5 *Let $F(N, \mathcal{T})$ be a flexibility polygon. Consider a section of the boundary of $F(N, \mathcal{T})$ between two terminals which are adjacent with respect to boundary order. Then:*

1. *Each such section is almost concave – i.e., it would be concave if we flipped the corner point (if any).*

2. *There is a point w on the boundary section such that all edges in the boundary section on one side of w are primary and on the other side of w are secondary.*

3. *Each section contains at most one corner point, and the boundary of $F(N, \mathcal{T})$ contains at most n corner points.*

4. *$F(N, \mathcal{T})$ has a corner point w on the boundary if and only if there exists a λ-SMT $T \in S(N, \mathcal{T})$ passing through w such that T can be split into two subtrees at w, one of which is primary, the other of which is secondary.*

4 Construction of Flexibility Polygon

In this section we give a linear time algorithm for constructing the flexibility polygon $F(N, \mathcal{T})$ for a set of terminals N and a full Steiner topology $\mathcal{T}$. The algorithm to construct the flexibility polygon consists of three steps, each of which can be performed in $O(\lambda n)$ time (where $n = |N|$). In the following we describe each of these steps in detail.

Step 1: Construction of a λ-SMT for N and $\mathcal{T}$

The first step of the algorithm is to construct an arbitrary λ-SMT $T \in S(N, \mathcal{T})$. This can be accomplished in $O(\lambda n)$ time using the algorithm given in [3]. Let

$p(T)$ and $s(T)$ denote the total amounts of respectively exclusively primary and exclusively secondary material in T; by Theorem 2.3 $p(T)$ and $s(T)$ depend solely on N and $\mathcal{T}$ — and not on the particular λ-SMT T.

Let $E(\mathcal{T})$ denote the set of *oriented* edges (or arcs) in $\mathcal{T}$, and consider some edge $[u, v] \in E(\mathcal{T})$. All λ-SMTs in $S(N, \mathcal{T})$ use the same set of (at most) two edge directions for edge $[u, v]$ [3]. By analysing the tree T, we obtain these (at most) two directions, corresponding to the the primary and secondary edge directions. Let $\Theta_p[u, v]$ and $\Theta_s[u, v]$ denote the primary and secondary edge directions for $[u, v]$, respectively. Note that we either have $\Theta_s[u, v] = \Theta_p[u, v] + \omega$ or $\Theta_s[u, v] = \Theta_p[u, v]$ (the latter is only possible when λ is not a multiple of 3).

Step 2: Construction of primary and secondary subtrees

The second step of the algorithm is to construct all primary and secondary subtrees of $\mathcal{T}$. For an edge $[u, v] \in E(\mathcal{T})$ the maximal subtree of topology $\mathcal{T}$ rooted at u and not containing v is denoted by $\mathcal{T}[u, v]$. The *primary* subtree for $\mathcal{T}[u, v]$ (if it exists) is the geometric embedding of $\mathcal{T}[u, v]$ such that every edge of $\mathcal{T}[u, v]$ uses its *primary* edge direction. In addition, we require that the amount of exclusively primary edge material $p[u, v]$ used by the primary subtree is less than $p(T)$, the total amount of exclusively primary material available. If the primary subtree for $\mathcal{T}[u, v]$ exists, we let $\Phi_p[u, v]$ denote the coordinates of the node u in this subtree; otherwise $\Phi_p[u, v] = \text{NIL}$.

The algorithm ConstructSubtrees given in Figure 11 computes $\Phi_p[u, v]$ for every subtree $\mathcal{T}[u, v]$ in $O(n)$ time; note that there are $O(n)$ subtrees, since there are $O(n)$ oriented edges in $E(\mathcal{T})$. In the algorithm N and S denote the set of terminals and Steiner points in $\mathcal{T}$, respectively. Furthermore, $(\Phi_p[u, v], \Theta_p[u, v])$ denotes the *ray* with source $\Phi_p[u, v]$ and direction $\Theta_p[u, v]$. The function $d^*(q_1, q_2)$ returns the Euclidean distance between points q_1 and q_1 provided that the direction from q_1 to q_2 is an *exclusively* primary direction — otherwise $d^*(q_1, q_2)$ returns zero.

In the first phase of the algorithm (lines 1–12), we initialize $\Phi_p[u, v]$ for every subtree $\mathcal{T}[u, v]$. Also, every subtree $\mathcal{T}[u, v]$ such that u is a Steiner point with two terminals (other than v) as neighbours is inserted into a queue Q. The queue Q holds all subtrees $\mathcal{T}[u, v]$ that *can* be constructed at given point in time. A subtree $\mathcal{T}[u, v]$ can be constructed if the children of u in $\mathcal{T}[u, v]$ already have been constructed. Since any Steiner topology has at least one Steiner points with two neighbouring terminals, the queue is non-empty when the initialization phase finishes.

```
CONSTRUCTSUBTREES(N, 𝒯, T)
// Initialization phase
Q = ∅    // empty queue of oriented edges (=subtrees)
forall [u, v] ∈ E(𝒯) do
    if u ∈ N then
        Φ_p[u, v] = u    // coordinates of terminal u
        p[u, v] = 0
    else
        Φ_p[u, v] = NIL
        p[u, v] = ∞
        Let v_1 and v_2 be the two neighbours of u other than v
        if v_1 ∈ N and v_2 ∈ N then
            ENQUEUE(Q, [u, v])
// Construction phase
while Q ≠ ∅
    [u, v] = DEQUEUE(Q)
    Let v_1 and v_2 be the two neighbours of u other than v
    Let r be the intersection (if any) between
    the rays (Φ_p[v_1, u], Θ_p[v_1, u]) and (Φ_p[v_2, u], Θ_p[v_2, u])
    if r exists then
        p[u, v] = p[v_1, u] + p[v_2, u] + d*(Φ_p[v_1, u], r) + d*(Φ_p[v_2, u], r)
        if p[u, v] < p(T) then
            Φ_p[u, v] = r    // subtree 𝒯[u, v] has now been constructed
            if v ∈ S then
                Let u_1 and u_2 be the two neighbours of v other than u
                if Φ_p[u_1, v] ≠ NIL then ENQUEUE(Q, [v, u_2])
                if Φ_p[u_2, v] ≠ NIL then ENQUEUE(Q, [v, u_1])
```

Figure 11: Construction of primary subtrees of topology $\mathcal{T}$.

In the second phase of the algorithm (lines 13–25) we construct the subtrees that have been inserted into Q. The subtree is only constructed if the total amount of exclusively primary material is less than the total amount of exclusively primary material in T (lines 19–20). If the construction succeeds and $v \in S$, we investigate if the newly constructed subtree can be used to construct some larger subtree (lines 22–25). This is done by checking if either of the neighbours of v (other than u) already have had their subtree constructed; if this is the case then the subtree rooted at v is inserted into Q.

Lemma 4.1 *Algorithm* CONSTRUCTSUBTREES *correctly constructs all primary subtrees in $O(n)$ time.*

Proof. The running time analysis is straightforward, since each oriented edge in $E(\mathcal{T})$ is inserted into Q at most once. Processing an edge (=subtree) takes $O(1)$ time.

The correctness follows by induction on the depth of the constructed subtrees. The base cases are the subtrees consisting of terminals only (that have depth 0), and Steiner points with terminals as children (that have depth 1); the latter are inserted into the queue Q in the initialization phase and therefore obviously constructed. For the induction step, assume that all subtrees with depth up to $k \geq 1$ have been constructed. It is then clear that subtrees of depth $k + 1$ will also be constructed, since these are inserted into Q when the smaller depth trees are constructed. ∎

By using an analogous algorithm, we can also construct all *secondary* subtrees in $O(n)$ time. Here we let $\Phi_s[u, v]$ denote the coordinates of the node u in the secondary subtree $\mathcal{T}[u, v]$. This tree only exists if the amount of exclusively secondary material is less than $s(T)$, the total amount of exclusively secondary material available.

Step 3: Construction of the boundary of the flexibility polygon

Consider the section of the flexibility polygon between a pair of consecutive terminals t_i and t_{i+1} (with respect to the terminal ordering from a counter-clockwise outer walk of the tree). This is a rightmost concave path. It is constructed iteratively by building from t_i a path of edges of primary subtrees — constructed in step 2 — until we have a sequence of consecutive nodes u, v and w such that $\Phi_p[u, v] \neq$ NIL and $\Phi_p[v, w] =$ NIL (or $\Phi_p[u, v] \neq$ NIL and $v = t_{i+1}$). We distinguish between two cases:

- The secondary subtree $\mathcal{T}[v,u]$ exists (i.e., $\Phi_s[v,u] \neq$ NIL). By Theorem 3.3 the intersection between the rays ($\Phi_p[u,v]$, $\Theta_p[u,v]$) and ($\Phi_s[v,u]$, $\Theta_s[v,u]$) must exist and defines a corner point of a bent edge. The boundary of the flexibility polygon consists of primary edges from t_i up to the corner point, and secondary edges from the corner point to t_{i+1} (Figure 8).

- The secondary subtree $\mathcal{T}[v,u]$ does *not* exist. By Theorem 3.3 the secondary subtree $\mathcal{T}[w,v]$ must exist. The intersection between the rays ($\Phi_p[u,v]$, $\Theta_p[u,v]$) and ($\Phi_s[w,v]$, $\Theta_s[w,v]$) defines the position of Steiner point v on the boundary of the flexibility polygon. Thus the boundary of the flexibility polygon consists of primary edges from t_i up to the Steiner point v, and secondary edges from v to t_{i+1} (Figure 9).

Therefore, by using the information computed in steps 1 and 2, we can construct the complete boundary of the flexibility polygon in one counter-clockwise outer walk of T.

Theorem 4.2 *The flexibility polygon for a set of terminals N and topology $\mathcal{T}$ (where $n = |N|$) can be constructed in time $O(\lambda n)$.*

The presented algorithm is easy to implement and involves small constants in the asymptotic running time. The algorithm has already been used as a corner stone in the design of a new VLSI router for nets in λ-geometry [8]. The time to construct the flexibility polygon was negligible compared to other parts of the router algorithm.

5 The Flexibility Polygon of a Steiner Point

Throughout this section, let T be a full and fulsome λ-SMT for a given terminal set N and topology $\mathcal{T}$. Let $F(N,\mathcal{T})$ be the flexibility polygon for N and $\mathcal{T}$. Here we address the following problem: given a Steiner point v in T what are all the possible positions for the Steiner points corresponding to v in the λ-SMTs for N with topology $\mathcal{T}$? We refer to the union of all such points as the *flexibility polygon of the Steiner point* v, denoted $F_v(N,\mathcal{T})$. In this section we will show that $F_v(N,\mathcal{T})$ is a simply connected convex region with polygonal boundary, and we will show that, similarly to $F(N,\mathcal{T})$, its boundary can be constructed from rightmost paths. Furthermore, if $F(N,\mathcal{T})$ has already been constructed then $F_v(N,\mathcal{T})$ can be constructed for any Steiner point v in constant time.

We begin with a lemma which shows that the position of a Steiner point v on a rightmost path of T is a local property, requiring knowledge only of the two neighbours of v on the path and not the entire path.

Lemma 5.1 *Let v be a Steiner point of T and let u and w be two vertices of T adjacent to v. Suppose P is a path between terminals in T containing vertices u, v, w in that order (i.e., $P = \cdots uvw \cdots$). Let v^r be the vertex corresponding to v in P^r, the rightmost path of P, and let T^r be a tree realizing P^r. Then the position of v^r and the direction of each straight edge or half edge of P incident with v^r is independent of the choice of P (i.e., is determined by u, v and w).*

Proof. Let T_1, T_2 and T_3 be the three maximal subtrees of T containing v as a leaf, such that T_1 contains u and T_2 contains w. It follows from the proof of Theorem 3.3 that the effect of transforming T to T^r is to transfer as much exclusively primary material as possible to T_1 and as much exclusively secondary material as possible to T_2. Hence, *independent* of the choice of P, the tree T^r — which consists of the subtrees T_1^r, T_2^r and T_3^r — has one of the following three forms:

1. The bent edge is in T_1^r, and both T_2^r and T_3^r are secondary subtrees.
2. The bent edge is in T_2^r, and both T_1^r and T_3^r are primary subtrees.
3. The bent edge is in T_3^r, and T_1^r is a primary subtree while T_2^r is a secondary subtree.

Since two of the subtrees of v^r are primary/secondary subtrees the position of v^r is uniquely determined (by bottom-up construction). Furthermore, if T_1^r contains any secondary material then the straight edge or half edge of T_1^r incident with v^r will be secondary. Similarly, if T_2^r contains any primary material then the straight edge or half edge of T_2^r incident with v^r will be primary, completing the proof of the lemma. ∎

Let v be a Steiner point of T and let u_1, u_2 and u_3 be the three vertices of T adjacent to v (in counterclockwise order around v). Then for each pair i, j, such that $i, j \in \{1, 2, 3\}$ and $i \neq j$, let P_{ij} be a path in T connecting two terminals of T such that $P_{ij} = \cdots u_i v u_j \cdots$. We define v_{ij} to be the vertex corresponding to v in P_{ij}^r, the rightmost path of P_{ij}. Consider the edge in P_{ij}^r corresponding to $u_i v$. If this edge is a straight edge then we define $u_{i(ij)}$ to be the vertex on P_{ij}^r corresponding to u_i. On the other hand, if the edge is a bent edge then $u_{i(ij)}$ is defined to be its corner point (under the embedding of P_{ij}^r). Similarly, $u_{j(ij)}$ is defined to be the vertex or corner point in P_{ij}^r corresponding to u_j (Figure 12).

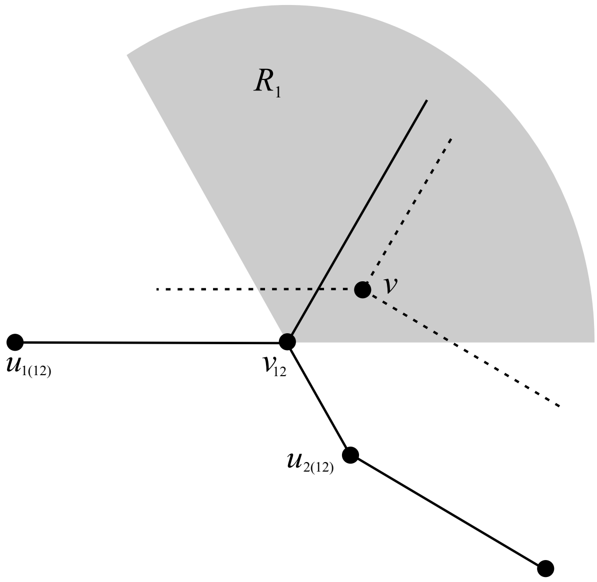

Figure 12: An example of the construction in Section 5 for $\lambda = 6$. In Theorem 5.2 we show that v_{13} lies on the line extending $u_{1(12)}v_{12}$.

By Lemma 5.1, the point v_{ij} and the directions of $u_{i(ij)}v_{ij}$ and $v_{ij}u_{j(ij)}$ are independent of the specific choice of P_{ij}. The set $\{v_{ij}\}$ consists of at most six distinct points.

For each $i \in \{1, 2, 3\}$ let $k = k(i) = (i \bmod 3) + 1$. Let R_i be the intersection of the closed half-plane to the left of the directed line through $u_{i(ik)}v_{ik}$ with the closed half-plane to the left of the directed line through $v_{ik}u_{k(ik)}$. (For the example in Figure 12 R_1 is indicated by the shaded region.) Define the region R_v as follows:

$$R_v = \bigcap_{i \in \{1,2,3\}} R_i.$$

The following theorem shows that R_v is the flexibility polygon for v.

Theorem 5.2 *Let v be a Steiner point of T, and let R_v be defined as above. Then $F_v(N, \mathcal{T}) = R_v$, and hence $F_v(N, \mathcal{T})$ is a simply connected region with polygonal boundary. Furthermore, the set $\{v_{ij} : i, j \in \{1, 2, 3\}$ and $i \neq j\}$ maps surjectively (under the trivial map) onto the vertices of the boundary of $F_v(N, \mathcal{T})$.*

Proof. Note that $F_v(N, \mathcal{T}) \subseteq R_v$, by Theorem 3.3. In order to prove the reverse subset inequality, we first prove the second half of the theorem.

Clearly the points v_{ik} (where $k = k(i) = (i \bmod 3) + 1$) are vertices of the boundary of $F_v(N, \mathcal{T})$. We need to consider the points v_{ki}. We begin by studying the case $i = 3$ and showing that v_{13} lies on the line extending $u_{1(12)}v_{12}$.

For a given path $P_{12} = \cdots u_1 v u_2 \cdots$ in T let T' be a λ-SMT containing P_{12}^r obtained from T by 0-shifts. Let T_1, T_2 and T_3 be the three maximal subtrees of T' containing v_{12} as a leaf, such that T_1 contains $u_{1(12)}$ and T_2 contains $u_{2(12)}$. Since P_{12}^r is a rightmost path, it follows (from the proof of Theorem 3.3) that T_1 contains as much exclusively primary material as possible and T_2 contains as much exclusively secondary material as possible. Hence, T_2 and T_3 are either both secondary subtrees, or T_1 is a primary subtree. In the former case it is clear that $v_{13} = v_{12}$ (by bottom-up construction). In the latter case observe that the subtree corresponding to T_1 in a λ-SMT containing v_{13} must also be a primary subtree. Hence v_{13} lies on the (primary) line extending $u_{1(12)}v_{12}$. By a similar argument (in which the roles of primary and secondary edges are reversed), v_{13} also lies on the line extending $u_{3(23)}v_{23}$. Applying the same argument to each of the vertices v_{ki} shows that $\{v_{ij} : i, j \in \{1, 2, 3\} \text{ and } i \neq j\}$ coincides with the vertices of the boundary of $F_v(N, \mathcal{T})$.

Finally, we show that $F_v(N, \mathcal{T}) \supseteq R_v$ by convexity of the distance function in λ-geometry. Since all vertices of R_v are possible positions of v under 0-shifts, it follows that all points on the boundary of the convex hull correspond to positions of v by convexity of the distance function (see [3, Section 3.3]). Similarly all points inside R_v correspond to positions of v by convexity. ∎

Examples of flexibility polygons for Steiner points are given in Figure 13. Note that one or more of the points v_{ij} will often coincide, resulting in a polygon with less than six vertices. It is possible, however, for all six points v_{ij} to be distinct, as shown in the flexibility polygon on the left.

In the following corollary, we show that if $\lambda = 3$ then each $F_v(N, \mathcal{T})$ has a triangular boundary. Note that this corollary is a generalisation of a result in [6], where the same result was shown for Steiner trees on three terminals.

Corollary 5.3 *Suppose $\lambda = 3$, and let v be a Steiner point of T. Then $F_v(N, \mathcal{T})$ has a triangular boundary.*

Proof. Consider again the proof of Theorem 5.2. There we show that either $v_{13} = v_{12}$ or v_{13} lies on a primary line extending $u_{1(12)}v_{12}$. Similarly, we show that either $v_{13} = v_{23}$ or v_{13} lies on a secondary line extending $u_{3(23)}v_{23}$. However, for $\lambda = 3$, the primary line through $u_{1(12)}v_{12}$ and the secondary line through $u_{3(23)}v_{23}$ are parallel, and hence if they meet they do so at an angle π. It follows that v_{12}, v_{13},

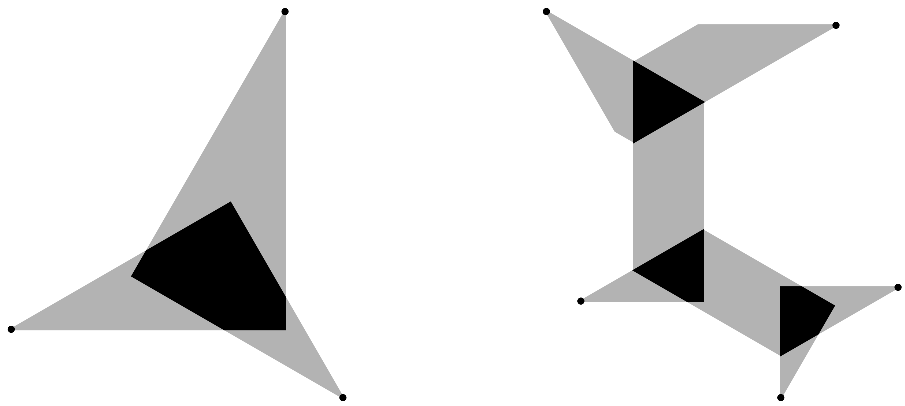

Figure 13: Examples of flexibility polygons (shaded dark) for Steiner points ($\lambda = 6$). The flexibility polygon on the left has 6 vertices, while the flexibility polygons on the right have 3, 4 and 5 vertices, respectively.

v_{23} are collinear. After applying the same argument to v_{21} and v_{32} the result easily follows. ∎

We make a few final observations about the flexibility polygon of a Steiner point.

- The flexibility polygon of a Steiner point v can be constructed in constant time once the flexibility polygon $F(N, \mathcal{T})$ for T has been constructed. This follows immediately from Theorem 5.2, since the directions of the adjacent edges to v in the relevant rightmost paths have already been computed.

- The flexibility polygon for an edge of T can also be defined in a similar way. It can be constructed by taking the union of the two flexibility polygons for the edge's endpoints and the part of $F(N, \mathcal{T})$ lying between those two polygons. It is easy to see that this is again a simply connected region with polygonal boundary, and one that can be constructed in constant time once $F(N, \mathcal{T})$ has been constructed.

6 Conclusion

The flexibility polygon is a compact description of the region in which λ-SMTs for a given set of terminals may be embedded. In VLSI routing a huge number of Steiner trees must be routed simultaneously on the surface of the chip. By placing the flexibility polygons for all nets on top of each other, congested regions are identified where many polygons overlap (Figure 14). Furthermore, this gives the basis for a completely new routing paradigm in VLSI routing (in arbitrary λ-geometry): Congested regions in the overlay graph associated with flexibility polygons give an indication of which nets are in congested regions. These nets should preferably be routed first in such a way that they avoid these regions.

Flexibility polygons have recently been used in a prototype of a new VLSI router for nets in λ-geometry [8]. Although computational results are limited, and the router is only a prototype, the usability of flexibility polygons is remarkable when it comes to avoiding congested areas while keeping nets at minimum length with very few vias (changes of routing layers).

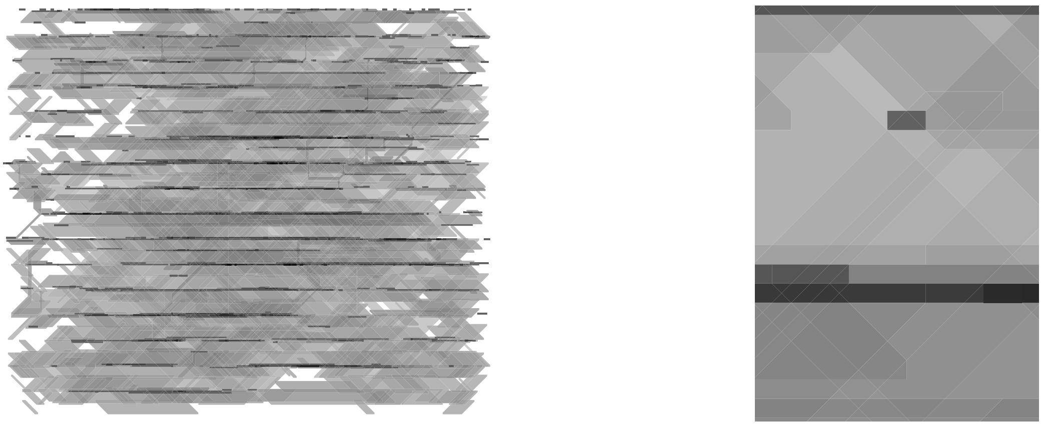

Figure 14: Overlay of flexibility polygons in VLSI routing. One the left is a full picture of a small chip, and on the right a zoom-in of a region of the picture of the left. Darker regions indicate more congested regions.

Our descriptions of the regions in which individual *Steiner points* in a λ-SMT may be placed immediately give flexibility regions for individual edges and bounds on their lengths under 0-shifts. We believe that these results will be important in improving delay-related measures for nets in VLSI routing. Finally, our result on the invariability of exclusively primary/secondary edge material in a full and fulsome λ-SMT has the potential to improve existing exact algorithms for computing λ-SMTs.

Acknowledgments

This work was partially supported by a grant from the Australia Research Council and by a grant from the Danish Natural Science Research Council (51-00-0336).

References

[1] E. Bozorgzadeh, R. Kastner, and M. Sarrafzadeh. Creating and Exploiting Flexibility in Steiner Trees. *IEEE Transactions on Computer-Aided Design of Integrated Circuits and Systems*, 22:605–615, 2003.

[2] M. Brazil, D. A. Thomas, and J. F. Weng. Minimum Networks in Uniform Orientation Metrics. *SIAM Journal on Computing*, 30:1579–1593, 2000.

[3] M. Brazil, D. A. Thomas, J. F. Weng, and M. Zachariasen. Canonical Forms and Algorithms for Steiner Trees in Uniform Orientation Metrics. Technical Report 02/22, DIKU, Department of Computer Science, University of Copenhagen, 2002.

[4] H. Chen, C. K. Cheng, A. B. Kahng, I. I. Mandoiu, Q. Wang, and B. Yao. The Y-Architecture for On-chip Interconnect: Analysis and Methodology. In *Proceedings ACM/IEEE International Conference on Computer-Aided Design (ICCAD)*, pages 13–19, 2003.

[5] H. Chen, B. Yao, F. Zhou, and C. K. Cheng. The Y-Architecture: Yet Another On-chip Interconnect Solution. In *Proceedings Asia and South Pacific Design Automation Conference*, pages 840–846, 2003.

[6] Y. Y. Li, S. K. Cheung, K. S. Leung, and C. K. Wong. Steiner Tree Constructions in λ_3-Metric. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 45(5):563–574, 1998.

[7] B. K. Nielsen, P. Winter, and M. Zachariasen. An Exact Algorithm for the Uniformly-Oriented Steiner Tree Problem. In *Proceedings of the 10th European Symposium on Algorithms, Lecture Notes in Computer Science*, volume 2461, pages 760–772. Springer, 2002.

[8] M. Paluszewski, P. Winter, and M. Zachariasen. A New Paradigm for General Architecture Routing. In *ACM Great Lakes Symposium on VLSI (GLSVLSI)*, pages 202–207, 2004.

[9] S. Peyer, M. Zachariasen, and D. G. Jørgensen. Delay-related Secondary Objectives for Rectilinear Steiner Minimum Trees. *Discrete Applied Mathematics*, 136:271–298, 2004.

[10] M. Sarrafzadeh and C. K. Wong. Hierarchical Steiner Tree Construction in Uniform Orientations. *IEEE Transactions on Computer-Aided Design*, 11:1095–1103, 1992.

[11] S. Teig. The X Architecture: Not Your Father's Diagonal Wiring. *SLIP*, pages 33–37, 2002.

[12] P. Widmayer, Y. F. Wu, and C. K. Wong. On Some Distance Problems in Fixed Orientations. *SIAM Journal on Computing*, 16(4):728–746, 1987.

[13] X Initiative Home Page. `http://www.xinitiative.com`, 2001.