

Some Structural and Geometric Properties of Two-Connected Steiner Networks

K. Hvam L. Reinhardt P. Winter M. Zachariasen

Dept. of Computer Science, University of Copenhagen
Universitetsparken 1, DK-2100 Copenhagen Ø, Denmark
e-mail: {pawel,martinz}@diku.dk

Abstract

We consider the problem of constructing a shortest Euclidean 2-connected Steiner network (SMN) for a set of terminals. This problem has natural applications in the design of survivable communication networks. A SMN decomposes into components that are full Steiner trees. Winter and Zachariasen proved that all cycles in SMNs with Steiner points must have two pairs of consecutive terminals of degree 2. We use this result and the notion of reduced block-bridge trees of Luebke to show that no component in a SMN spans more than approximately one-third of the terminals. Furthermore, we show that no component spans more than two terminals on the boundary of the convex hull of the terminals; such two terminals must in addition be consecutive on the boundary of this convex hull. Algorithmic implications of these results are discussed.

Keywords: Survivable networks; Euclidean 2-connected Steiner networks; full Steiner trees.

1 Introduction

The Euclidean Steiner tree problem asks for a shortest possible network spanning a set Z of n terminals in the plane. The solution is a tree, referred to as a *Steiner minimal tree* (SMT). Apart from the terminals, SMTs may contain additional, so-called *Steiner points*, where exactly three edges meet at 120° angles. SMTs are unions of *full Steiner trees* spanning subsets of terminals all having degree 1.

When the objective is to design low cost survivable networks, the problem of constructing *Euclidean 2-connected Steiner minimum networks* in the plane (SMNs) arises. Since a 2-edge-connected minimum-length network necessarily is 2-vertex-connected when the distance function is a metric (Frederickson & Ja'Ja 1982), we use the shorthand 2-connected in the following.

SMNs have been studied by Hsu and Hu (1998), Luebke and Provan (2000), Luebke (2002) and Winter and Zachariasen (2005). Luebke and Provan (2000) proved that the problem is NP-hard and gave a number of structural properties of SMNs. Luebke (2002) introduced the notion of (reduced) block-bridge trees that will play an essential role in this paper. Winter and Zachariasen (2005) proved that all cycles in SMNs with Steiner points must have pairs of consecutive terminals of degree 2.

Copyright copyright 2007, Australian Computer Society, Inc. This paper appeared at Computing: The Australasian Theory Symposium (CATS2007), Ballarat, Australia. Conferences in Research and Practice in Information Technology (CRPIT), Vol. 65. Joachim Gudmundsson and Barry Jay, Eds. Reproduction for academic, not-for profit purposes permitted provided this text is included.

The paper is organized as follows. In Section 2 we formally define the problem and outline some basic properties of SMNs. In Section 3 we discuss reduced block-bridge trees. In Section 4 we show how the properties of the reduced block-bridge trees for SMNs can be used to bound the size of full Steiner trees. Properties of SMNs related to the boundaries of convex hulls of terminals are discussed in Section 5. Algorithmic implications of these results are discussed in Section 6. Concluding remarks are given in Section 7.

2 Basic Properties

Let Z be a set of n terminals in the Euclidean plane. The *2-connected Euclidean Steiner network problem* is to find a minimum length 2-connected network $N(Z)$ spanning Z and possibly additional, so called *Steiner points*. If Z is obvious from the context, we denote $N(Z)$ by N .




Figure 1: SMN with 2 non-trivial FSTs of 3 terminals each.

Hsu and Hu (1998), Luebke and Provan (2000) and Luebke (2002) proved several properties of SMNs. In particular, all Steiner points are incident with three edges meeting at 120° angles (as for the Euclidean Steiner tree problem in the plane). Furthermore, no cycle consists entirely of Steiner points. As a consequence, SMNs are unions of *full Steiner trees* (FSTs), in which all terminals are leaves and all Steiner points are interior vertices (Figure 1). This important result implies that **GeoSteiner**, the 2-phase exact algorithm for the determination of SMTs in the plane (Warme, Winter & Zachariasen 2000) can be adapted to find SMNs. **GeoSteiner** in its first phase generates a superset of FSTs of a SMT. Powerful geometric tests based on non-trivial properties of SMTs permit to keep this superset very small. In the second phase,




Figure 2: Reduced block-bridge tree construction.

GeoSteiner uses an elaborate branch-and-cut approach to identify FSTs of the superset whose concatenation yields a SMT. Problem instances with up to 10000 terminals can be solved within reasonable amount of time.

The same approach can be used to find SMNs if the connectivity constraints in the integer programming model for the concatenation phase are appropriately modified. While this modification is straightforward, the real problem occurs in the first phase where FSTs are generated. Some of the very powerful geometric test that work for the SMTs, do not apply in the 2-connected case.

The notion of a *chord-path* was introduced in (Winter & Zachariasen 2005). Let G denote an undirected graph and let C be a cycle in G . A chord-path between two distinct vertices u and v on C is a path between u and v in G that shares no interior vertices with C . We proved the following structural results.

Theorem 1 *Any chord-path in a SMN must have a pair of consecutive terminals of degree 2 in its interior.*

Corollary 1 *Suppose that a SMN has vertices of degree 3. Each of its cycles has two pairs of consecutive terminals of degree 2 separated by vertices of degree 3.*

3 Reduced Block-Bridge Trees

Let H be an undirected connected graph. The *block-bridge tree* H_B of H is a tree obtained by contracting each block (i.e., each 2-connected component) to a vertex (Figures 2a and 2b). The *reduced block-bridge tree* H_R of a connected graph H , is obtained from the block-bridge tree H_B of H by replacing each simple path (with interior vertices of degree two) by an edge (Figure 2c). Luebke (2002) introduced reduced block-bridge trees and gave a proof of Theorem 2 in her Ph.D. thesis; here we present a modified and shorter version of the proof.

Given a SMN N of Z and one of its FSTs F , we let $N \setminus F$ denote the graph obtained by deleting from N all interior vertices and all edges of F .

Lemma 1 *Let F be an FST in an SMN N of Z . $H = N \setminus F$ is connected.*

Proof. Suppose that H is disconnected. Let u and v be two vertices of H , each in a different component of H . Since N is 2-connected, it contains two disjoint paths P and P' between u and v . Let C denote the cycle created by P and P' . Both P and P' must go through F . Since F is a tree with no terminals in its interior, C has a terminal-free chord-path. This contradicts Theorem 1. ■

In the remainder of this paper, we let H_B denote the block-bridge tree of $H = N \setminus F$, and H_R the reduced block bridge tree of H .

Consider an edge e of an arbitrary tree T . When e is removed, T breaks down into 2 subtrees, T_L^e and T_R^e . Let S_L^e and S_R^e denote the leaves in T_L^e and T_R^e , respectively. The pair of sets $S^e = (S_L^e, S_R^e)$ is referred to as a *split* of T by the edge e . It is clear that splits generated by different edges of the same tree are different. It is also well-known that two trees with the same set of leaves (and with no vertices of degree 2) are isomorphic if and only if their edges generate the same splits (Semple & Steel 2003).

Lemma 2 *Let F be an FST in a SMN N . There is a one-to-one correspondence between the terminals of F and the leaves of H_R , such that each terminal in F is contracted to a distinct leaf in H_R .*

Proof. First we prove the result for the block-bridge tree H_B and then we extend the result to the reduced block-bridge tree H_R .

Suppose that two distinct terminals t_1 and t_2 of F are contracted to the same vertex v of H_B . Terminals t_1 and t_2 are therefore on a cycle C in H (Figure 3a). The path connecting t_1 and t_2 in F is a terminal-free chord-path of C , contradicting Theorem 1. Thus at most one terminal of F is contracted to a vertex in H_B .

Suppose that v is a leaf in H_B , and no terminal of F is contracted to it (Figure 3b). This is only possible if there is a bridge in N , contradicting the assumption that N is 2-connected.

Suppose that v is a non-leaf of H_B and that a terminal t of F is contracted to v . H_B is a tree. There are therefore two paths in H_B , both starting at v but otherwise disjoint, one through edge f_1 and the other through edge f_2 , ending in leaves v_1 and v_2 (Figure 3c). Consequently, there must be terminals t_1 and t_2 in H contracted to v_1 and v_2 respectively. Since the part of H contracted to v in H_B is 2-connected (with special case where the contracted part consists of t alone), there is a path in H from t_1 to t_2 going through t . Since F is also a tree, there must be a path in F between t_1 and t_2 . The cycle composed by these two paths has a terminal-free chord-path, a contradiction.

When going from H_B to H_R , every leaf is preserved. All such vertices contain exactly one terminal from F , finishing the proof of the lemma. ■

Theorem 2 *Let F be an FST in a SMN N . F and H_R are isomorphic.*

Proof. Let N_R denote the union of F and H_R (where the terminals in F have been identified with the corresponding leaves H_R). Consider an arbitrary edge




Figure 3: Impossible contractions.

$e \in F$. We will first show that $N_R \setminus e$ has a unique bridge f located in H_R .

Let u and v denote the end-vertices of e . Assume that there are 2 vertex-disjoint paths between u and v in $N_R \setminus e$ (Figure 4a). Then there exists a cycle in N with e being its terminal-free chord-path, a contradiction.

Assume next that there are 2 edge-disjoint (but not vertex-disjoint) paths R_{uv} and R'_{uv} between u and v in H_R sharing some vertex z (Figure 4b). The vertex z cannot correspond to a single vertex in H ; its degree in H and hence in N would then be at least 4, a contradiction. The vertex z must therefore correspond to a contracted block B of H . Since R_{uv} is a path in H_R , there is a corresponding path P_{uv} in $N \setminus e$. Suppose that P_{uv} enters block B through a vertex u_1 and leaves B through a vertex v_1 . If $u_1 = v_1$, then there exists a path P'_{uv} in $N \setminus e$ corresponding to R'_{uv} which is vertex-disjoint with P_{uv} in B (otherwise $u_1 = v_1$ would have degree 4 in $N \setminus e$). If $u_1 \neq v_1$, let C be a cycle in B through the part of P_{uv} between u_1 and v_1 . If P'_{uv} avoids C , then P_{uv} and P'_{uv} are vertex-disjoint in B . Assume that P'_{uv} enters C through a vertex u'_1 and leaves C through a vertex v'_1 . Since these vertices must have degrees less than 4, u_1, v_1, u'_1 and v'_1 must be mutually different. No matter in what order they appear on C , there always will exist 2 paths between u and v in $N \setminus e$ that are vertex-disjoint in B (follow P_{uv} and P'_{uv} from u until reaching C at u_1 and u'_1 , reach v_1 and v'_1 through disjoint parts of C and continue toward v). The edge-disjoint paths R_{uv} and R'_{uv} can share several vertices. We argued so far that there are paths P_{uv} and P'_{uv} in $N \setminus e$ that are vertex-disjoint in a block corresponding to a shared vertex. Hence, there exist paths P_{uv} and P'_{uv} that are vertex-disjoint in *all* these blocks. Such P_{uv} and P'_{uv} will form a cycle with e as its chord-path, a contradiction.

We have shown so far that $N_R \setminus e$ contains at least one bridge f . Since F and N_R are trees sharing their leaves, f must be located in H_R . We will now show that f is unique. $N_R \setminus \{e, f\}$ is disconnected. As a consequence, the split generated by e in F must be the same as the split generated by f in H_R . Since no pair of splits in a tree (with no vertices of degree 2) is the same (Semple & Steel 2003), it follows that f is unique.

It can be shown in a similar manner that for any $f' \in H_R$, the graph $H_R \setminus f'$ contains a unique bridge e' located in F .

We proved that there is a unique perfect matching between edges of F and edges of H_R , such that for every pair of matched edges $e \in F$ and $f \in H_R$, e is a bridge in $N_R \setminus f$ if and only if f is a bridge in $N_R \setminus e$. Since F and H_R share their leaves, the splits generated by the matched edges are the same. This implies that F and H_R are isomorphic. ■

4 Size of Full Steiner Trees

In this section we use reduced block-bridge trees to obtain an upper bound on the number of terminals in FSTs of a SMN. Such a bound can be used to make the generation phase of the exact algorithm more efficient. It also influences the concatenation phase as the number of FSTs becomes smaller.

Theorem 3 *A SMN N spanning n terminals has no FST spanning more than $\lfloor n/3 \rfloor + 1$ terminals.*

Proof. Let F denote an FST of N . Let m denote the number of terminals in F . Let H , H_R and N_R be as defined in Section 3.

Let $e \in F$. By Lemma 2, there is a unique edge $f \in H_R$ such that e is a bridge in $N_R \setminus f$ and f is a bridge in $N_R \setminus e$. Furthermore, $N_R \setminus \{e, f\}$ consists of two 2-connected components.

Let H_f denote the subgraph of H contracted to the edge $f \in H_R$. H_f can be viewed as a sequence of 2-connected components. Suppose that H_f has no interior 2-connected components or terminals in this sequence. Hence, f corresponds to an edge in H . The removal of e and f from H_R leaves two 2-connected components. But this implies that the removal of e and f from N also leaves two 2-connected components, contradicting the fact that deleting any pair of edges of an SMN will leave a bridge in one of the resulting components (Monma, Munson & Pulleyblank 1990).

Suppose next H_f has an interior 2-connected component and that it contains no terminal. Then such a component can be replaced in H by a single edge. The new graph remains 2-connected and due to the triangle inequality it is shorter, a contradiction.

It follows that H_f must contain at least one terminal in its interior. F has $2m - 3$ edges. Since F and H_R are isomorphic, H_R also has $2m - 3$ edges. Each of these edges corresponds to a subgraph of H containing at least one terminal. As a consequence, N contains at least $3m - 3$ terminals (m of them belonging to F). Hence, $n \geq 3m - 3$ and it follows that $\lfloor n/3 \rfloor + 1$ is an upper bound on the number of terminals in any FST of N . ■

5 Border Terminals

Let $CH(Z)$ denote the convex hull of Z . A terminal is called a *border terminal* if it is on the boundary of $CH(Z)$. Otherwise it is called an *inner terminal*. Two border terminals are said to be *neighbors* if they share a boundary edge of $CH(Z)$.

Lemma 3 *If N contains edge-disjoint paths P_1 , P_2 and P_3 beginning in a degree 3 vertex v and ending in border terminals z_1 , z_2 and z_3 , then one of these paths contains a chord-path.*




Figure 4: Isomorphism between F and H_R .

Proof. Since N is 2-connected, there exists a path P_{23} connecting z_2 and z_3 , but otherwise disjoint from P_2 and P_3 . P_{23} either avoids P_1 or hits P_1 at a some vertex.

If P_{23} avoids P_1 (Figure 5a), then there must be a path P_{12} from z_1 to z_2 that is disjoint from P_1 and P_2 . If P_{12} avoids P_3 then P_2 is a chord-path for the cycle formed by P_{12} , P_{23} , P_3 and P_1 . If P_{12} hits P_3 , let w be the intersection that is closest to v . The portion of P_3 between v and w is a chord-path for the cycle formed by P_{12} , P_2 and P_1 . (Note that $v \neq w$; otherwise N would have a vertex of degree 4.)

If P_{23} hits P_1 (Figure 5b), let u be the intersection that is closest to v . The portion of P_1 between u and v is a chord-path for the cycle formed by P_{23} , P_3 and P_2 . ■

Lemma 4 *If N contains a path P between two border terminals, and there is at least one terminal on each side of P , then P contains a chord-path.*

Proof. Let z_1 and z_2 be two terminals on opposite sides of P . Since N is 2-connected, there exist two disjoint paths P_{12} and P'_{12} from z_1 to z_2 . Let u and u' denote the intersections of P_{12} and P'_{12} with P ; these intersections can always be selected such that the part of P_1 between u and u' is disjoint from both P_{12} and P'_{12} . This portion of P is a chord-path for the cycle formed by P_{12} and P'_{12} . ■

Theorem 4 *No FST F of N spans more than 2 border terminals. Furthermore, if F spans 2 border terminals, they must be neighbors.*

Proof. The theorem is trivial if $n \leq 3$, so let $n \geq 4$ (where n is the number of terminals spanned by N). Suppose that F spans 3 border terminals z_1, z_2, z_3 . Let v denote the Steiner point of F where the paths P_1, P_2 and P_3 from z_1, z_2 and z_3 all meet. None of these paths contains a terminal and therefore cannot contain a chord-path. This contradicts Lemma 3.

Assume therefore that F spans 2 border terminals z_1 and z_2 and that they are not neighbors. Let P denote the unique path between z_1 and z_2 in F . Since z_1 and z_2 are not neighbors, then there are two other border terminals separated by P . By Lemma 4, P contains a chord-path. In particular, it contains 2 consecutive terminals. But this contradicts the assumption that P is in the FST F . ■

Let F be an FST in $N(Z)$. Let z_F denote the number of terminals in F . Let b_F denote the number of border terminals in F and let i_F denote the number of inner terminals in F . Finally, let i be the number of terminals in the interior of $CH(Z)$.

Theorem 5 $b_F + 2i_F - i \leq 2$.

Proof. If $z_F = 2$, then by a straightforward analysis of three cases ($b_F = 0, 1, 2$), the theorem holds.

We therefore assume that $z_F \geq 3$. By Theorem 4, $b_F \leq 2$ and therefore $i_F \geq 1$. Let z_1, z_2, \dots, z_{i_F} denote the inner terminals of F . Let F_1, F_2, \dots, F_{i_F} denote FSTs, other than F , containing z_1, z_2, \dots, z_{i_F} , respectively. These FSTs must be pairwise disjoint and they each share exactly one terminal with F . Otherwise N would contain a cycle with 2 terminals, contradicting Corollary 1.

Let G denote the union of $F, F_1, F_2, \dots, F_{i_F}$. G is a tree. The number of terminals spanned by G is at least $2i_F$. The number of border terminals spanned by G is therefore at least $b_F + 2i_F - i$. Suppose that G has 3 border terminals z_1, z_2 and z_3 . Since G is a tree, there must be a vertex v in G where otherwise disjoint paths from z_1, z_2 and z_3 meet. By Lemma 3, one of these paths contains a chord-path in N . In particular, such a path must contain 2 consecutive interior terminals of degree 2. But this contradicts the manner in which G was constructed. The claim follows. ■

Theorem 6 $z_F \leq 2 + \lceil i/2 \rceil$

Proof. Since $z_F = i_F + b_F$, Theorem 5 gives $z_F \leq 1 + b_F/2 + i/2$. By Theorem 4, $b_F \leq 2$ and $z_F \leq 2 + \lceil i/2 \rceil$. ■

Corollary 2 *SMNs for problem instances with all but at most one terminal on the boundary of $CH(Z)$ contain no Steiner points.*

6 Pruning Applications

Equilateral points are defined recursively as follows. A terminal t is an equilateral point based on t . Let a be an equilateral point based on a subset of terminals Z_a and let b be an equilateral point based on a subset of terminals Z_b . If $Z_a \cap Z_b = \emptyset$, then the equilateral point e_{ab} based on $Z_a \cup Z_b$ is defined as the third corner of the equilateral triangle with ab as one of its sides and with e_{ab} to the right of the line through a and b when looking from a (Figure 6a). Note that $e_{ab} \neq e_{ba}$.

Equilateral points are building blocks of FSTs during the generation part of the two-phase algorithm for SMTs and SMNs. Every FST F spanning a subset of terminals Z_e and a closing terminal $t, t \notin Z_e$, is completely determined by an appropriate equilateral point e based on Z_e . Furthermore, the length of F is equal to the length of the Simpson line connecting e with t . The locations of the Steiner point adjacent to t is the intersections of the Simpson line with the circle circumscribing the equilateral triangle defining e (Figure 6b). The locations of other Steiner points is determined recursively with the earlier determined Steiner points acting as closing terminals.

It follows that FSTs can be generated by the enumeration of all equilateral points closed by appropriate terminals. Many of these equilateral points can be




Figure 5: Paths to three border terminals.

pruned away as they violate 120° angle requirement between edges meeting at Steiner points or they lead to long FSTs that cannot appear in an optimal solution. The reader is referred to (Warne et al. 2000) for a more detailed discussion of equilateral points and their role in the construction of FSTs.

The upper bound on the number of terminals in FSTs of an SMN obtained in Section 5 can be used to prune many other equilateral points. Let a and b denote two equilateral points based on disjoint subsets of terminals Z_a and Z_b . Let e be an equilateral point based on $Z_e = Z_a \cup Z_b$. Let b_{Z_e} and i_{Z_e} denote the number of border terminals and interior terminals, respectively, that e is based on. It follows immediately from Theorem 5 that e can be pruned away if $b_{Z_e} > 2$. Furthermore, if $b_{Z_a} = 1$ and $b_{Z_b} = 1$, then e needs to be constructed only if the two border terminals are neighbors on $CH(Z)$.

Lemma 5 *If $b_{Z_e} = 2$ and $i_{Z_e} > i/2 - 1$ then e_{ab} can be pruned away.*

Proof. Any FST F involving e must include a terminal t not in Z_e . The terminal t cannot be a border terminal since Z_e already contains two border terminals. Hence, t must be an inner terminal, and $i_F \geq i_{Z_e} + 1$. Then

$$b_F + 2i_F - i \geq 2 + 2i_{Z_e} + 2 - i > 2 + i - i = 2$$

where we use the assumption $2i_{Z_e} + 2 > i$ in the second inequality — contradicting Theorem 5. ■

Lemma 6 *If $b_{Z_e} + 2i_{Z_e} - i > 1$ then e_{ab} can be pruned away.*

Proof. Closing such equilateral point e by a terminal (border or inner terminal) will cause the inequality of Theorem 5 to be violated. ■

The above pruning tests are applied to equilateral points. Each such equilateral point can be closed by any terminal $t \in Z \setminus Z_e$. Such closures may permit pruning away some of the FSTs (even though the final equilateral point could not be pruned away). This is for example the case if t is a border terminal and $b_e = 2$ or if t is a border terminal, $b_e = 1$ and the two border terminals are not neighbors on $CH(Z)$. This latter test in particular indicates that edges connecting border terminals can appear in SMNs only if the terminals are neighbors on $CH(Z)$. Finally, we remark that every FST F has to satisfy the inequality of Theorem 5.

7 Concluding Remarks

We presented some new structural properties of SMNs. The bound on the size of full Steiner trees can be used in a straightforward manner. FSTs are generated during the first phase of the algorithm in non-decreasing order of the number of terminals they span. So the generation can be cut off when the number of terminals gets beyond $\lfloor n/3 \rfloor + 1$. The properties based on the convex hull of the terminals can be used to obtain new criteria for pruning equilateral points and FSTs. These criteria will in particular be very powerful in connection with problem instances where the number of interior terminals is small. For example, problem instances with 8 inner terminal will only involve FSTs spanning at most 6 terminals. Many FSTs for such problem instances will be pruned away as they can contain at most 2 border terminals. Furthermore, if FSTs contain 2 border terminals, then they have to be neighbors on $CH(Z)$.

References

- Frederickson, G. N. & Ja'Ja, J. (1982), 'On the Relationship between the Biconnectivity Augmentation and Traveling Salesman Problem', *Theoretical Computer Science* **13**, 189–201.
- Hsu, D. F. & Hu, X.-D. (1998), 'On Shortest Two-Edge Connected Steiner Networks with Euclidean Distance', *Networks* **32**, 133–140.
- Luebke, E. L. (2002), *k*-Connected Steiner Network Problems, PhD thesis, University of North Carolina, USA.
- Luebke, E. L. & Provan, J. S. (2000), 'On the Structure and Complexity of the 2-Connected Steiner Network Problem in the Plane', *Operations Research Letters* **26**, 111–116.
- Monma, C. L., Munson, B. S. & Pulleyblank, W. R. (1990), 'Minimum-Weight Two-Connected Spanning Networks', *Mathematical Programming* **46**, 153–171.
- Semple, C. & Steel, M. (2003), *Phylogenetics*, Vol. 24 of *Oxford Series in Mathematics and Its Applications*, Oxford University Press.
- Warne, D. M., Winter, P. & Zachariasen, M. (2000), Exact Algorithms for Plane Steiner Tree Problems: A Computational Study, in D.-Z. Du, J. M. Smith & J. H. Rubinstein, eds, 'Advances




Figure 6: Constructing a FST using equilateral points.

in Steiner Trees', Kluwer Academic Publishers, Boston, pp. 81–116.

Winter, P. & Zachariasen, M. (2005), 'Two-Connected Steiner Networks: Structural Properties', *Operations Research Letters* **33**, 395–402.