Chapter 1

A New Way to Weigh Malnourished Euclidean Graphs

Gautam Das * Giri Narasimhan † Jeffrey Salowe ‡

Abstract

In this paper, we show that any Euclidean graph over a set V of n points in k-dimensional space that satisfies either the *leapfrog property* or the *isolation property* has small weight, i.e., has weight $O(1) \cdot wt(SMT)$, where SMT is a Steiner minimal tree of V. Both the leapfrog property as well as the isolation property constrain the way the edges of the graph are configured in space. Our main applications are to prove that certain Euclidean graphs known as *t-spanners* can be constructed with optimal weight of $O(1) \cdot wt(SMT)$. This intriguing open problem has attracted much attention recently. The main tool in obtaining the above weight bounds is a theorem that proves the existence of long edges in a Steiner minimal tree on a restricted set of points in k-dimensional space. We also generalize this property for Steiner minimal trees on arbitrary point sets. Since very little is known about high-dimensional Steiner minimal trees, these results are of independent interest.

1 Introduction.

It is often difficult to analyze the total edge weight of an Euclidean graph for the purpose of providing performance bounds. For instance, bounding the length of the Lin-Kernighan heuristic for the traveling salesperson problem appears to be very difficult [16, 5], even for points in Euclidean space. More successful examples include algorithms for the on-line traveling salesperson problem for Euclidean points [13, 5] and the t-spanner problem for Euclidean points [1, 4, 7, 14, 15]. Generally, the analysis centers on relating the lengths of a set of edges to $wt(SMT)$ (or $wt(MST)$ or $wt(TSP)$). Here $wt(SMT)$, $wt(MST)$, and $wt(TSP)$ is the total edge weight of a Steiner minimal tree, minimum spanning tree, and an optimal traveling salesperson tour, resp. of a given set of points. (Note that $wt(TSP) = O(wt(MST)) = O(wt(SMT))$ for a set of points in Euclidean space).

Various tools have been developed in order to perform such analyses. The central idea is to bound the weight of a set of edges that satisfy certain spatial constraints. In [4], it was shown that if a set of edges satisfies the property that the distance between the endpoints of any two edges is at least as large as a constant times the length of the shorter edge (the *gap property*), then the weight of these edges can be bounded by $O(\log n) \cdot wt(SMT)$, where $wt(SMT)$ is the total edge weight of a Steiner minimal tree of the set of endpoints. The gap property found application in two different problems. It is satisfied by the set of edges produced by the greedy algorithm for spanners [4], and also by the set of edges produced by the *2-opt* heuristic for the traveling salesperson problem [5], and consequently helped in their analyses.

The gap property is unfortunately limited in power. It has been shown in [14] that there is a set of edges that satisfies the gap property whose length is $\Omega(\frac{\log n}{\log\log n}) \cdot wt(SMT)$ (very recently this lower bound has been increased to $\Omega(\log n) \cdot wt(SMT)$, [3]). In order to prove stronger results, other properties were considered. In [7], it was shown that if a set of edges in 3-dimensional space satisfies a property called the *leapfrog property*, then the weight of the edges can be bounded by $O(1) \cdot wt(SMT)$. The leapfrog property is a complicated restriction on how a set of edges may be positioned in space, and a precise definition is given later in the paper. This proof was based on a complex charging scheme and employed properties of 2-dimensional Steiner minimal trees. It was left as an open problem to prove the same statement for edges satisfying the leapfrog property in k-dimensional space. The following theorem of this paper solves this open problem. The proof requires a generalization of the properties of the Steiner minimal tree in higher dimensions.

THEOREM 1.1. *If a set of line segments E in k-dimensional space satisfies the leapfrog property, then $wt(E) = O(1) \cdot wt(SMT)$.*

The leapfrog property has found applications in the improved analysis of spanners [7]. Since it is rather diffi-

*The University of Memphis, Memphis, TN 38152. *E-mail*: dasg@next1.msci.memst.edu. Supported in part by NSF RIA CCR-930-6822.

†The University of Memphis, Memphis, TN 38152. *E-mail*: giri@next1.msci.memst.edu. Supported in part by NSF RIA CCR-940-9752 and NSF Grant INT-911-5870.

‡QuesTech, Inc., 7600A Leesburg Pike, Falls Church, VA 22043. *E-mail*: jsalowe@nvl.army.mil. Supported in part by NSF Grants MIP-91-07717 and CCR-92-24789. Most of this work was performed at the University of Virginia, Charlottesville VA.

cult to appreciate intuitively, in this paper we introduce a more intuitive property called the *isolation property*. Interestingly enough recent spanner algorithms that enforce the isolation property have turned out to be more efficient than those that enforce the leapfrog property [2]. We believe that it will have applicability in other geometric problems. A precise formulation of the property is presented later; for now it suffices to say that a set of edges in k-dimensional space satisfies the isolation property if each edge can be associated with a large enough hypercylindrical region which does not intersect with the other edges. As will be evident later, the isolation property implies the leapfrog property (but not vice versa). However, in this paper we provide a simpler and more direct proof for the following theorem.

THEOREM 1.2. *If a set of line segments E in k-dimensional space satisfy the isolation property, then $wt(E) = O(1) \cdot wt(SMT)$.*

We now define the notion of a *t-spanner*. Let V be a set of n points in k-dimensional Euclidean space, and let $t > 1$. A subgraph of the complete Euclidean graph of V is a t-spanner if for every $u, v \in V$, the shortest path length between any two points is at most t times the Euclidean distance between the two points. Computing a t-spanner with a small weight and other desirable properties has been the focus of much recent research. Although many algorithm exist for producing sparse t-spanners, analyzing the performance of these algorithms turns out to be a difficult task. The earliest results on small weight t-spanners appeared in [1, 15], where it was shown that in 2-dimensional space, such spanners exist with weight at most $O(1) \cdot wt(SMT)$. However, the proofs employed planarity properties and could not be directly extended to higher dimensions. In [14] a spanner based on the *well-separated pairs* construction was described, and its weight was analyzed to be $O(\log^2 n) \cdot wt(SMT)$. In [4] it was shown that in k-dimensional space, t-spanners constructed by a certain greedy algorithm have weight at most $O(\log n) \cdot wt(SMT)$. These results seemed to indicate that eliminating the polylog factors in the weight for higher dimensional spanners will require more sophisticated techniques than had been previously used. The first optimal result for dimensions higher than 2 appeared in [7]. In that paper it was shown that the spanner constructed by the greedy algorithm satisfies the leapfrog property in k-dimensional space. However, the paper only provided a proof of the 3-dimensional version of Theorem 1.1, which implied the existence of spanners with weight $O(1) \cdot wt(SMT)$ in 3-dimensional space.

Our proof of Theorem 1.1 leads us to the generalization stated below. This represents a satisfactory conclusion to the series of previous efforts on constructing small weight t-spanners in higher dimensional space.

THEOREM 1.3. *For a set of points in k-dimensional space, a t-spanner can be constructed with an optimal weight of $O(1) \cdot wt(SMT)$.*

While the leapfrog property has been used here to prove the existence of optimal weight spanners, the less general isolation property has been recently found crucial in devising an *efficient* $O(n \log n)$-time spanner construction with optimal weight [2]. In contrast, the best known algorithm that constructs spanners satisfying the leapfrog property has a running time of $O(n \log^2 n)$ [8].

Our main tool in proving the above theorems is to show properties of higher dimensional Steiner minimal trees. Steiner minimal trees are useful in weight-bounding tasks because they are monotonic in the following sense. If $S_1 \subset S_2$ are sets of points, a SMT on S_2 cannot be shorter than a SMT on S_1. For example, in [7] a SMT is maintained for the endpoints of edges which satisfy the leapfrog property. At every step in the analysis, a suitable edge is removed, and it is shown that the SMT of the endpoints of the remaining edges is shorter than the previous SMT by an amount proportional to the length of the edge removed. Steiner minimal trees are interesting in their own right, and have been studied by many researchers [11] (also see [12]). However, very little is known about Steiner minimal trees in higher dimensions [18, 19].

In this paper we consider Steiner minimal trees that connect $n + 1$ points, n of which lie on the surface of a unit hypersphere in k-dimensional Euclidean space (these points are called *terminals*), while one point lies at the center s of the hypersphere. We call such a tree as a *restricted Steiner minimal tree*. We prove the following theorem.

THEOREM 1.4. *There is a constant $c < 1$ such that for any restricted Steiner minimal tree, the path between s and any terminal u contains an edge of length at least c.*

The existence of these *long edges* in restricted Steiner minimal trees is crucial to the applications that we discuss, as will be described later. The above theorem also provides insight into the local structure of the SMT of an unrestricted set of points. Consider a hypersphere centered at any vertex of an unrestricted SMT with radius equal to the distance to the nearest neighbor of v. The above theorem says that every path from v to the outside of this hypersphere has to contain suitably long edges. Later we also prove several additional properties of unrestricted SMTs.

The rest of the paper is organized as follows. §2 discusses properties of Steiner minimal trees, while §3 discusses applications of these properties in the weight analysis of edge sets satisfying (1) the isolation property, and (2) the leapfrog property. We note that although the results in this paper assume the Euclidean (L_2) metric, they can easily be extended to any arbitrary L_p metric.

2 Properties of Steiner Minimal Trees.

In this section we first derive properties of restricted Steiner minimal trees in 2-dimensions, then generalize them to higher dimensions, and finally consider unrestricted Steiner minimal trees.

2.1 Restricted Steiner Minimal Trees in 2-Dimensions.

Here we prove the 2-dimensional version of Theorem 1.4.

Proof. Let S be the set of terminals, T the Steiner minimal tree under consideration, $\alpha < 1$ a constant, C a circle of radius α and the same center s, N the set of Steiner points of T within C, and $N' \subseteq N$ the set of points that do not have all neighbors in N (these are the *fringe* vertices of the subtree containing vertices from N). We first prove that $|N'|$ is bounded by a constant. For every vertex $v \in N'$ there is at least one path in T of length $(1-\alpha)$ from v to a terminal. Clearly, there exists a constant p such that if $|N'| > p$, then $wt(T)$ would exceed $1 + 2\pi$, which is a simple upper bound for $wt(T)$. All the Steiner points in T have degree 3. Hence $|N| \leq 2|N'|$ and is also bounded by a constant. Thus any path in T from s to a terminal must have a long edge (of length at least $c = \alpha/2p$) which concludes the proof.

The proof also implies that of the $O(n)$ Steiner points in the Steiner minimal tree T, most of them lie far away from the center s and lie close to the unit circle. Corresponding theorems can be trivially proven for the case where terminals lie on a circle with radius r (here α will be replaced by αr), and for the case where terminals lie *outside* the circle. In an unrestricted setting, where the points to be connected do not necessarily lie on a circle, this theorem implies that for any point s, if its nearest neighbor is at distance d away, then the theorem could be applied with $r = d$.

The above proof and the consequent observations do not easily extend to the case when the points are in k-dimensional space. In the 2-dimensional case, the n points on the unit circle could be connected together by a network of length at most 2π (independent of n), which does not hold for n points on the unit hypersphere in k-dimensional space. However, using different significantly different techniques we prove that similar results also hold for the k-dimensional case. This is presented next.

2.2 Restricted Steiner Minimal Trees in k-Dimensions.

We shall now prove Theorem 1.4 for k-dimensions. We need the following result in our proof. It was originally proven by Few [10] and subsequently improved by Smith [17] (see Hwang et al. [12]).

THEOREM 2.1. *The longest possible length of a Steiner minimal tree for n points in the unit hypercube in k-dimensions is $\Theta(n^{1-1/k})$.*

Few's proof of this result is also valid for points on the surface of a unit hypersphere. (The proof is omitted, but the essential idea is to use projections, and the fact that the surface area of a unit k-dimensional hypersphere is $k\frac{\pi^{k/2}}{\Gamma(k/2+1)}$).

COROLLARY 2.1. *The longest possible length of a Steiner minimal tree for n terminals on the surface of a k-dimensional unit hypersphere is $\Theta(n^{1-1/(k-1)})$.*

We first show that restricted Steiner minimal tree T must contain some long edges. We then show that there must be a long edge in each path from s to a vertex on the surface of the hypersphere.

LEMMA 2.1. *T contains an edge of length at least c', for some constant c'.*

Proof. Suppose to the contrary that there are no edges of length c' in T. We show that T must contain greater than n leaves, which contradicts the claim that T is a Steiner minimal tree on n terminals.

To be more specific, consider a sequence of trees $T_1, T_2, \ldots$, where each T_i is a subtree of T. Tree T_1 consists of s and all the portions of T that are within path distance $1/2$ of s. Tree T_2 consists of s and the portions of T within path distance $3/4$ of s. In general, tree T_i consists of s and all the portions of T that are within path distance $1 - 1/2^i$ of s. Note that tree T_j is a subtree of tree T_i, $j < i$. Our goal is to obtain lower bounds on the number of leaves $\ell(T_i)$ in T_i and show that for any n, there is some i such that $\ell(T_i) > n$. This implies that T contains more than n leaves.

We now bound the number of leaves in T_1. Let $c_0 = c'$. By assumption, there are no edges of length c_0 in T, so there are no edges of length c_0 in T_1. Since each Steiner point has degree 3 and each edge has length at most c', the height of T_1 is at least $\frac{1}{2c_0}$, and $\ell(T_1) \geq 2^{\frac{1}{2c_0}}$.

Using Corollary 2.1, we can obtain an even stronger bound on the length of edges in $T_2 - T_1$, the portion of T_2 induced by the removal of tree T_1. There are $\ell = \ell(T_1)$ components in $T_2 - T_1$; in each component, choose a longest edge, and let H be the set of edges chosen. If H is removed from T, the Steiner minimal tree T is broken into $\ell + 1$ components. Choose an arbitrary terminal in each of these components, and let

$H(S)$ be the terminals chosen. Corollary 2.1 implies that the terminals in $H(S) - \{s\}$ can be interconnected using a Steiner minimal tree of length $O(\ell^{1-1/(k-1)})$; an additional edge of length 1 to s interconnects the tree. This implies that the total edge length in H must be $O(\ell^{1-1/(k-1)})$.

Since the total length of edges in H must be $O(\ell^{1-1/(k-1)})$, say $\ell^{1-1/(k-1)}$, at least $\ell/2$ of these edges must have length at most $2/\ell^{1/(k-1)}$. This implies that at least $2^{\frac{1}{2c_0}-1}$ components in $T_2 - T_1$ have longest edge with length at most $c_1 = 2^{1-1/(2c_0(k-1))} \ll c_0$. The number of leaves in T_2 is at least $2^{1/2c_0 - 1} \times 2^{1/4c_1}$.

In general, let c_{i-1} be an upper bound on the length of a longest edge on a particular subtree of T in which the longest root-to-leaf path length is $1/2^i$. This tree T' is intended to be a component of $T_i - T_{i-1}$. ($T_0 = \emptyset$ and $c_0 = c'$). Then the height of T' is at least $\frac{1}{2^i c_{i-1}}$, and the number of leaves is $\ell \geq 2^{\frac{1}{2^i c_{i-1}}}$. From the argument above, at least $\ell/2$ of the subtrees rooted at the leaves of T' have longest edge length at most $2/\ell^{1/(k-1)}$. Consequently, let c_i satisfy the following recurrence:

$$\begin{aligned} c_i &= 2^{1-\frac{1}{2^i(k-1)c_{i-1}}} \\ c_0 &= c'. \end{aligned}$$

Then if a component in $T_i - T_{i-1}$ has longest edge at most c_{i-1}, it must have $2^{\frac{1}{2^i c_{i-1}}}$ leaves, and half of these leaves are roots of subtrees with maximum length bound c_i.

By successively applying the arguments above, we see that the number of leaves in T_i satisfies

$$\ell(T_i) \geq \Pi_{j=1}^{i}(2^{\frac{1}{2^j c_{j-1}}-1}).$$

There is a j such that $\ell(T_j) > n$, completing the proof.

Although Lemma 2.1 guarantees the existence of a long edge in each tree in the forest, it does not guarantee that a long edge lies on the path from s to u. Let $p(s,u)$ be the path from s to u.

LEMMA 2.2. *The total number of Steiner points in $p(s,u) \cap T_1$ is bounded by a constant.*

Proof. Assume the contrary. Suppose there are j such Steiner points. Since the degree of each Steiner point is 3, each Steiner point on $p(s,u) \cap T_1$ is associated with at least one tree connecting that Steiner point to some boundary terminal. By a slight generalization of Lemma 2.1, there is an edge of length at least $c'/2$ in each of j trees. Remove these edges. The length of the removed edges is $c'j/2$. Reconnect these components by choosing one vertex from each component and forming a Steiner minimal tree on the surface of the hypersphere, connected by an edge to s. Corollary 2.1 states that the length of the latter tree is $O(j^{1-1/(k-1)})$, which is less than $c'j/2$ for large enough j.

Lemma 2.2 implies there is an edge of length at most c''' in $p(s,u)$, where c''' is a constant. As a consequence, there is an edge in $p(s,u) \cap T_1$ of length at least c''', for some constant c''', thus proving Theorem 1.4.

2.3 Unrestricted Steiner Minimal Trees.

Here we extend some of our earlier results, by considering Steiner minimal trees of more general point sets. First we consider the case when instead of a single vertex at the center s of the unit hypersphere, we have a set A of m points contained within a hypersphere of radius β centered at s.

THEOREM 2.2. *Let $\beta \leq \alpha_i < \alpha_o < 1$. Let A be a set of m points lying within a hypersphere of radius β centered at a point s. Let B be a set of n points lying outside the unit hypersphere centered at s. Then the number of Steiner points on any Steiner minimal tree T lying within a hyperspherical shell of inner radius α_i and outer radius α_o is bounded by a constant. Also, every path in T connecting a point in A and a point in B must have an edge of length greater than c, for some constant c.*

Proof. Let T be a Steiner minimal tree connecting A and B. Let C_1 be a hypersphere centered at s of radius α_i and let C_2 be a hypersphere centered at s of radius α_o. Let $\mathcal{S}$ be the hyperspherical shell between C_1 and C_2. Let S_A (S_B) be the set of Steiner points beyond (within) radius α_i (α_o) with at least 1 subtree such that all its leaves are in A (B). Let S'_A (S'_B) be the set of fringe vertices of S_A (S_B). In other words, these are the vertices of S_A (S_B) not all of whose neighbors are from S_A (S_B). Also, let R be the set of (remaining) points from T that are not in S_A or S_B.

Consider the subtree of $u \in S'_B$ that has all its leaves from B. By Lemma 2.2 this subtree has at most a constant number of Steiner points inside C_2. We claim that the number of points in S'_B within C_2 is bounded by a constant. To prove this, assume the contrary. Let this number be more than some constant j. Each of these subtrees, by Theorem 1.4 has a long edge of length say c'. By arguments similar to that of the previous section, for large enough j, one can force $c'j > 1 + j^{1-1/(k-1)}$, thus contradicting the assumption that T is a Steiner minimal tree. Similarly, one can also prove that the number of points in S'_A outside C_1 is bounded by a constant. Hence the number of points from $S'_A \cup S'_B$ that lie in the hyperspherical shell $\mathcal{S}$ is bounded by a constant.

Using similar arguments as before, it can be shown that since S'_A and S'_B are the fringe vertices of the subtree of vertices from S_A and S_B, and since their number within the hyperspherical shell $\mathcal{S}$ is bounded by a constant, the number of points from $S_A \cup S_B$ that lie in the hyperspherical shell $\mathcal{S}$ is also bounded by a constant.

In order to show that the number of Steiner points of T that lie within the hyperspherical shell $\mathcal{S}$ is bounded by a constant, we are left to show that the number of points in R that lie within the hyperspherical shell $\mathcal{S}$ is also bounded by a constant. Similar to the approach employed in the previous proofs, we define a set R' of vertices in R, not all of whose neighbors are in R. We will first show that the number of vertices in R' lying in $\mathcal{S}$ are bounded by a constant. Each vertex in R' has degree 3 and has at least one neighbor from $S_A \cup S_B$. By Theorem 1.4, for every vertex in R', the subtree rooted at the neighbor from $S_A \cup S_B$ must contain a long edge since the root is at distance at least α_i or α_o (depending on whether the vertex is in S_A or S_B) away from the leaves of the subtree. Now, if there are more than a constant number of vertices in R', a contradiction to the minimality of the Steiner tree can be arrived at. Having bounded the size of R' lying within $\mathcal{S}$, similar arguments as used before can be used again to prove that the size of R lying within $\mathcal{S}$ is also bounded.

Once the total number of Steiner points in $\mathcal{S}$ are bounded, it is easy to see that every path in T that connects a point in A and a point in B must pass through $\mathcal{S}$, and contains no vertex, one vertex, or a constant number of vertices that lie in $\mathcal{S}$. In each of the cases it is easy to see that there must be an edge on the path of length more than some constant c.

In a more general setting, the results we have shown imply that if for some point s, there exists an *empty* hyperspherical shell between radius βr and r (for some $\beta < 1$), then the Steiner minimal tree will have long edges (of length cr for some constant c), and that the number of Steiner points in a portion of this hyperspherical shell is bounded by a constant. In fact, Theorem 2.2 can be extended to prove the following theorem, which is our most general version.

THEOREM 2.3. *Let $A \cup B$ be a set of points inside the unit hypersphere connected by a SMT. Let the minimum distance between a point in A and a point in B be d. Then there exists an edge of length at least cd, for some constant $c < 1$ on any path in the SMT between a vertex in A and a vertex in B.*

Proof. The main ideas of the proof are as follows. Let $\alpha_i < \alpha_o < 1$. Consider the union of hyperspheres of radius $\alpha_i d$ around every point in A. Consider the union of hyperspheres of radius $(1-\alpha_o)d$ around every point in B. Using arguments similar to the proof of Theorem 2.2 it can be shown that the number of Steiner points in the rest of the region is bounded by a constant. Since any path from a point in A to a point in B must pass through this region, there must be a long edge.

3 Applications to Weight Analysis.

In this section we apply the properties of Steiner minimal trees to bound the weight of edge sets satisfying various spatial constraints. We first prove Theorem 1.2 which concerns the isolation property. Then we outline the proof of Theorem 1.1 which concerns the leapfrog property.

3.1 The Isolation Property. Let $c > 0$ be a constant. Let E be a set of edges in k-dimensional space, and let $e \in E$ be an edge of length l. If it is possible to place a hypercylinder B of radius and height $c \cdot l$ each, such that the axis of B is a subedge of e and $B \cap (E - \{e\}) = \emptyset$, then e is said to be *isolated.* If all the edges in E are isolated, then E is said to satisfy the *isolation property.* We now prove Theorem 1.2.

Proof. Let E be a set of edges satisfying the isolation property. Let T be a Steiner minimal tree connecting the endpoints of E. We first partition E into a constant number of groups, $E_1, E_2, \ldots E_{O(1)}$, such that within each group the edges are almost parallel to each other. This can be done by using the standard technique of covering the space with cones, and selecting edges that lie within a particular cone (see [7]). We will now prove that for each i, $wt(E_i) = O(1) \cdot wt(T)$, where T is now a Steiner (though not necessarily minimal) tree connecting the endpoints of E_i.

Let the k co-ordinate axes be named $x_1, x_2, \ldots x_k$, and w.l.o.g. let the x_1 axis be parallel to the edges of E_i. If a line segment is parallel (perpendicular) to the x_1 axis, it is said to be *vertical* (*horizontal*). Furthermore, if a point u has a larger x_1 co-ordinate than a point v, u is said to be *above* v (and v is said to be *below* u). Since E_i satisfies the isolation property, each edge e is associated with a hypercylinder B which does not intersect with the rest of the edges. However the hypercylinders may intersect each other. To eliminate this problem, we shrink each hypercylinder horizontally by halving its radius. Each hypercylinder's surface is composed of three pieces: each end is a $(k-1)$-dimensional hypersphere such that one is above the other, and connecting the two is a $(k-1)$-dimensional surface representing the *vertical wall.*

We first transform T to T' by replacing each edge (u,v) of T by a vertical edge (u,w) and a horizontal edge (w,v). Let $hwt(T')$ denote the weight of the horizontal edges of T'. Clearly $hwt(T') = O(wt(T')) = O(wt(T))$. Partition E_i further into two groups L_i and N_i. If there is a connected portion of T' completely contained within a hypercylinder B and touching both the upper and lower surfaces, then the corresponding edge belongs to N_i, otherwise it belongs to L_i. We shall prove that $wt(N_i) = O(wt(T'))$ and $wt(L_i) = O(hwt(T'))$, which will suffice to prove the theorem.

The first part is easy, because each edge in N_i can be charged to the portion of T' contained within its hypercylinder. To prove the second part, consider the structure consisting of the tree T' and the edges in L_i. We proceed to remove (suitably selected) edges from L_i one by one, and at each iteration restructure T' such that, (1) it is still a Steiner tree for the endpoints of the remaining edges in L_i, (2) it is composed of horizontal and vertical edges, (3) it does not have a connected portion completely contained in any remaining hypercylinder which touches both the upper and lower surfaces, and (4) the weight of its horizontal edges is smaller than the corresponding weight of the horizontal edges of the previous tree by an amount at least a constant times the length of the edge removed.

At any iteration, select the B whose upper hypersphere has the largest x_1 co-ordinate, and remove the corresponding edge (say $e = (u, v)$ where u is above v). Consider the unbounded hypercylindrical region whose surface is formed by extending the vertical wall of B in both directions until infinity. Let T'_u be the maximal connected portion of T' that includes u and that lies completely within this region. Notice that the interior vertices of T'_u other than u have to be Steiner points, and cannot be endpoints of the remaining edges in L_i. Let T'_u pierce the surface of this region at points $U = \{u_1, u_2, \ldots, u_m\}$. Let u_1 be the unique point in U which lies along the path between u and v in T'.

Our restructuring of T' will proceed as follows. For each u_j in U there is a corresponding u'_j whose x_1 co-ordinate is the same as that of u, but the rest of the co-ordinates agree with those of u_j. Let U' be the set of these corresponding points. Remove T'_u. Construct a $(k-1)$-dimensional restricted Steiner minimal tree S_u connecting U' with u. Add vertical edges between each u_j and its corresponding u'_j. Add the vertical edge (u, v). Remove the long edge (Theorem 1.4) along the path in S_u between u and u_1. It is easy to see that the new T' satisfies all four properties.

Thus at each iteration we can charge the removed edge to the difference in the horizontal weights of the successive trees. The process is then repeated for the next edge. We can therefore conclude that $wt(L_i) = O(hwt(T'))$, which proves Theorem 1.2.

It is easy to see that in the definition of the isolation property, we can replace the hypercylinder by a hypersphere, hypercube etc, without affecting Theorem 1.2. The theorem can also be strengthened so that all edges of E do not have to be isolated; only certain long edges need be. The details are omitted from this version of the paper. This property was the linchpin in devising an $O(n \log n)$ time spanner construction with weight $O(1) \cdot wt(SMT)$ [2]; this latter result serves to prove Theorem 1.3. Extrapolating from our experience with the gap property, we expect that the isolation property too should prove useful in analyzing other Euclidean graphs.

3.2 The Leapfrog Property. Here we strengthen the previous result by showing that a set of edges satisfying the leapfrog property has a small weight. The leapfrog property is defined as follows. Let $d(u, v)$ denote the Euclidean distance between points u and v. Let E be a set of edges in k-dimensional space, and let $t > 1$ be a real number. E satisfies the leapfrog property if, for *every* subset $S = \{(u_1, v_1), \ldots, (u_m, v_m)\}$,

$$t \cdot d(u_1, v_1) \;<\; \sum_{i=2}^{m} d(u_i, v_i) + t \cdot \Big(\sum_{i=1}^{m-1} d(v_i, u_{i+1}) + d(v_m, u_1)\Big)$$

The isolation property is a special case of the leapfrog property. This is because for every subset of edges $\{(u_1, v_1), \ldots, (u_m, v_m)\}$, the empty hypercylinder of (u_1, v_1) ensures that the leapfrog condition holds.

We sketch the proof of Theorem 1.1. It is a generalization of the 3-dimensional proof that appears in [7], and we shall assume the reader is familiar with that proof.

Proof. Essentially, every edge is associated with a pair of small hypercylindrical regions at either ends. These regions together with the edge define a *dumbbell.* It is shown that the set of edges can be partitioned in such a way that the dumbbells in each group E_i are almost parallel, and in addition satisfy the *nested* property, i.e either two dumbbells are disjoint, or one dumbbell is completely contained in one of the hypercylinders of the other dumbbell.

Let the k co-ordinate axes be named $x_1, x_2, \ldots x_k$, and w.l.o.g. let the x_1 axis be parallel to the edges of E_i. We generalize the proof in [7] as follows. Construct a SMT of the endpoints of E_i and transform this to a tree T' by replacing each tree edge with a vertical edge and a horizontal edge. Partition E_i into two groups, the *lateral group* L_i and the *non-lateral group* N_i, exactly as in [7]. The weight analysis of the non-lateral group does not need generalization. However, we need to generalize the analysis of the lateral group to higher dimensions.

Consider the analysis of the lateral group as described in [7]. The proof is by induction, where at every iteration the current shortest dumbbell $e = (u, v)$ is removed, and T' is restructured to connect the remaining endpoints. It is shown that the removed edge can be charged to the difference in weight between the two trees. Using the terminology in [7], let T'_u be the maximal connected portion of T' that includes u and that lies wholly within the dumbbell head centered at u. A similar piece T_v lies around vertex v. Let the piece T_u pierce the dumbbell head at points $U = \{u_1, u_2, \ldots, u_k\}$. Let u_1 be the unique point in U which lies along the path between u and v in T'.

Instead of doing the restructuring as described in [7] (which only works for 3-dimensions), our restructuring of T' will proceed as follows (which is very similar to our proof of Theorem 1.2). For each u_j in U there is a corresponding u'_j whose x_1 co-ordinate is the same as that of u, but the rest of the co-ordinates agree with those of u_j. Let U' be the set of these corresponding points. Remove T'_u. Construct a $(k-1)$-dimensional restricted Steiner minimal tree S_u connecting U' with u. Add vertical edges between each u_j and its corresponding u'_j. Add the vertical edge (u, v). Remove the long edge (Theorem 1.4) along the path in S_u between u and u_1.

Thus at each iteration we can charge the removed edge to the difference in the horizontal weights of the successive trees. The process is then repeated for the next edge. We can therefore conclude that $wt(L_i) = O(hwt(T'))$, which proves Theorem 1.1.

Since the greedy algorithm ([4, 7]) produces t-spanners with the leapfrog property, this also serves to prove Theorem 1.3.

4 Conclusions and Open Problems.

In this paper, we have proved the existence of long edges in Steiner minimal trees connecting restricted sets of points. We have demonstrated the usefulness of these results in the geometric analysis of the weight of a set of edges in k-dimensional space satisfying some spatial constraints. As an application, we solve the open problem of proving that small weight t-spanners can be constructed for points in k-dimensional space. Some interesting open problems follow.

A recent paper ([13]) studies the problem of constructing on-line algorithms for the traveling salesperson problem. It is unknown whether there exists a constant-competitive algorithm for the on-line traveling salesperson problem under the *fixed graph* scenario for points in k-dimensional space. Perhaps some of the techniques in this paper can be applied in resolving this problem.

Some of the off-line heuristics for the TSP problem, most notably the Lin-Kernighan heuristic, have not been satisfactorily analyzed in the Euclidean setting. It is intriguing whether the suboptimal tour is no more than a constant factor longer than the optimal. It would be interesting to study whether our results give more insight into the heuristic.

In the construction of sparse t-spanners, very tight results have been obtained in reducing the number of edges. In [6] it has been shown that, given any $\delta > 1$, spanners can be constructed which have at most $n \cdot \delta$ edges. The corresponding problem in weight is still open, i.e. given $\gamma > 1$, are there spanners with weight $wt(MST) \cdot \gamma$?

We believe that other applications of the Steiner tree properties as well as the isolation property are very likely. In addition, Steiner tree properties are interesting in their own right, and need to be explored further.

References

[1] I. Althöfer, G. Das, D. P. Dobkin, D. A. Joseph, and J. Soares, *On sparse spanners of weighted graphs*, Discrete and Computational Geometry, 9(1993), pp. 81-100.

[2] S. Arya, G. Das, D. M. Mount, J. S. Salowe, and M. Smid, *Euclidean spanners: short, thin, and lanky*, manuscript, (1994).

[3] S. Arya, J. S. Salowe, private communication, (1994).

[4] B. Chandra, G. Das, G. Narasimhan, and J. Soares, *New sparseness results on graph spanners*, Proc. of ACM Symposium on Computational Geometry, (1992), pp. 192-201.

[5] B. Chandra, H. Karloff, and C. Tovey, *New results on the old k-opt algorithm for the TSP*, Proc. of SIAM-ACM Symposium on Discrete Algorithms, (1993), pp. 150-159.

[6] G. Das and P. J. Heffernan, **Constructing degree-3 spanners with other sparseness properties**, Proc. of Intl. Symposium on Algorithms and Computation, (1993), LNCS, Springer-Verlag, pp. 53-62.

[7] G. Das, P. J. Heffernan, and G. Narasimhan, *Optimally sparse spanners in 3-dimensional Euclidean space*, Proc. of ACM Symposium on Computational Geometry, (1993), pp. 53-62.

[8] G. Das and G. Narasimhan, A fast algorithm for constructing sparse Euclidean spanners, Proc. of ACM Symposium on Computational Geometry, (1994), pp. 132-139.

[9] D.-Z. Du and F. K. Hwang, *An approach for proving lower bounds: solution of Gilbert-Pollak's conjecture on Steiner ratio*, Proc. of IEEE Symp. on Foundations of Comp. Science, (1990), pp. 76-85.

[10] L. Few, *The shortest path and the shortest road through n points*, Mathematika, 2(1955), pp. 141-144.

[11] E. N. Gilbert and H. O. Pollak, *Steiner minimal trees*, SIAM Journal of Appl. Math., 16(1) (1968), pp. 1-29.

[12] F. K. Hwang, D. S. Richards, and P. Winter, *The Steiner tree problem*, North-Holland, (1992).

[13] B. Kalyanasundaram, and K. Pruhs, *Constructing competitive tours from local information*, Proc. of

ICALP, (1993).

[14] H.-P. Lenhof, J. S. Salowe, and D. E. Wrege, *Mixed spanning trees*, manuscript, (1993).

[15] C. Levcopoulos and A. Lingas, *There are planar graphs almost as good as the complete graphs and as short as the minimum spanning trees*, Proc. of Intl. Symposium on Optimal Algorithms, (1989), LNCS, Springer-Verlag, pp. 9-13.

[16] S. Lin and B. W. Kernighan, *An effective heuristic algorithm for the Traveling-Salesman problem*, Operations Research, 21(1973), pp. 498-516.

[17] W. D. Smith, *Studies in computational geometry motivated by mesh generation*, Ph.D. Thesis, Princeton University, (1988).

[18] W. D. Smith, *How to find Steiner minimal trees in Euclidean d-space*, Algorithmica, 7(1992), pp. 137-177.

[19] T. L. Snyder, *On the exact location of Steiner points in general dimensions*, SIAM Journal of Computing, 21(1992), pp. 163-180.