Optimally Sparse Spanners in 3-dimensional Euclidean Space *

(Extended Abstract)

Gautam Das †   Paul Heffernan*‡   Giri Narasimhan*

Abstract

Let V be a set of n points in 3-dimensional Euclidean space. A subgraph of the complete Euclidean graph is a t-spanner if for any u and v in V, the length of the shortest path from u to v in the spanner is at most t times $d(u, v)$. We show that for any $t > 1$, a greedy algorithm produces a t-spanner with $O(n)$ edges, and total edge weight $O(1) \cdot wt(MST)$, where MST is a minimum spanning tree of V.

1 Introduction

Let V be a set of n points in Euclidean space. The *complete Euclidean graph*, $G = (V, E)$ is defined as the complete weighted graph on these points, with edge weights equal to the Euclidean distance. A subgraph $G' = (V, E')$ is a *t-spanner* of G if for every $u, v \in V$, the distance between any two points is at most t times the Euclidean distance between the two points. The minimum value of t such that g' is a t-spanner is called the *stretch factor* of G'. In this paper we consider the greedy algorithm to compute a t-spanner for any complete Euclidean graph in 3-dimensional Euclidean space, and show that it produces a spanner that is *sparse*.

Sparseness of spanners is measured by two criteria, the size and the weight. The *size* of a graph G, $size(G)$, is defined as the number of edges in G and the *weight* of G, $wt(G)$, is defined as the sum of the edge weights of G. The minimum spanning tree (MST) of G is obviously the sparsest possible spanner both in terms of size and weight, but its stretch factor can be as bad as $n - 1$ [1]. The sparseness of a spanner G' is judged by comparing it to the size and weight of the MST.

Spanners for Euclidean graphs as well as for arbitrary graphs have numerous applications in robotics, graph theory, network topology design, distributed system, and communication protocols design. For a recent survey of spanners refer to [2].

We consider the problem of finding sparse spanners for complete Euclidean graphs. Spanners for complete Euclidean graphs in two-dimensional Euclidean space are considered in [1, 3, 4], and it is proved that there exist $O(1)$-spanners with size $O(n)$ and weight $O(1) \cdot wt(MST)$. However, the techniques used in these papers exploit planarity properties, and the results cannot be extended to higher dimensions, for which the only weight bounds known are for arbitrary edge-weighted graphs. The best results for higher dimensional Euclidean space are shown in [2], where it is shown that sparse spanners with size $O(n)$ and weight $O(\log n) \cdot wt(MST)$ can be constructed for a complete Euclidean graph in any dimension.

In this paper, we use substantially different techniques to obtain optimal weight sparseness results for the case of a complete Euclidean graph on a set of points in 3-dimensional Euclidean space. We show how to construct $O(1)$-spanners of $O(n)$ size and weight $O(1) \cdot wt(MST)$. This is clearly an optimal result for the 3-dimensional case.

In [2] the edges are partitioned with respect to their sizes. Since the number of resulting partitions is $O(\log n)$, this introduces a factor of $O(\log n)$ in the

*Appeared in the *Proceedings of the 9th Annual ACM Symposium on Computational Geometry*

†Mathematical Sciences Department, Memphis State University, Memphis, TN 38152.

‡Supported in part by the U.S. Army Research Office, Grant No. DAAL03-92-G-0378

final result. In contrast to [2], we do not group the edges according to their sizes alone. Instead, edges of different sizes are grouped together (to form a constant number of groups), and a complex charging scheme is employed. In this scheme, edges of the spanner are charged to portions of the minimum spanning tree. Another interesting feature of this paper is the cleanup process that is applied to the spanner edges, which greatly facilitates the analysis. We believe that the cleanup process may be useful in other situations where edge lengths need to be analyzed. Another noteworthy point is that one of the cases of our proof uses a locally minimum Steiner tree as an analysis tool.

Theorem 1.1 *Let $G = (V, E)$ be any n-vertex complete Euclidean graph, where V is a set of points in 3-dimensional space. Let $t > 1$ be any real number. There exists a polynomial time algorithm that constructs a t-spanner of G of size $O(n)$ and weight $O(1) \cdot wt(MST)$. The constant implicit in the O-notation depends only on t.*

This theorem can be easily extended to arbitrary norms. Our current proof does not extend to dimensions higher than three, because at one point we require local properties of 2-dimensional minimum Steiner trees.

2 Preliminaries

We begin by stating some definitions and properties that will be used throughout the paper.

A *complete Euclidean graph* is a complete graph on a vertex set corresponding to a set of points in d-dimensional space with edge weights equal to the Euclidean distances between the points. Let $d(u, v)$ refer to the Euclidean distance (L_2 norm) between points u and v. Let $wt(E)$ refer to the total weight of a set of edges E.

We now summarize some useful results from earlier papers. In [1], a **Greedy Algorithm** is presented that produces a t-spanner for an arbitrary graph. This is a simple generalization of Kruskal's algorithm to build a minimum spanning tree. The algorithm takes as input a weighted graph $G = (V, E)$, and a real number $t > 1$, and produces a t-spanner $G' = (V, E')$ as output. The algorithm builds the spanner incrementally, starting with an empty subgraph. It considers the edges in increasing order of weights. An edge (u, v) is added if and only if the shortest path from u to v in the partially constructed spanner is greater than $t \cdot d(u, v)$. The following two lemmas are proved in [1, 2].

Lemma 2.1 *The output of the greedy algorithm, G', is a t-spanner of G with $O(n)$ edges, for G a complete Euclidean graph.*

Lemma 2.1 states that the spanner of a complete Euclidean graph formed by the Greedy Algorithm satisfies the size bound of Theorem 1.1. The task of this paper is to show that the spanner of the Greedy Algorithm also satisfies the weight bound of the theorem.

Lemma 2.2 *Let (u, v) be any edge on a cycle C in G'. Then $t \cdot d(u, v) < wt(C) - d(u, v)$*

In [2] the following observation is made regarding how any pair of spanner edges are positioned in space.

Lemma 2.3 *Let (u_1, v_1) and (u_2, v_2) be any two spanner edges. Then*

$$t \cdot d(u_1, v_1) < d(u_2, v_2) + t \cdot (d(v_1, u_2) + d(v_2, u_1))$$

Note that (v_1, u_2) and (v_2, u_1) need not be spanner edges. One of the corollaries of this lemma is that, if any two spanner edges are almost parallel, say almost vertical, then the distance between their top (bottom) endpoints is at least proportional to the length of the shorter of the two edges. This single fact is enough to prove the weight result in [2]. However, in order to eliminate the $O(\log n)$ factor found in [2] we will require much more.

We now define what we call the *Leap-Frog Property*. Let E be a set of line segments in space, and let $t > 1$ be a real number. E satisfies the *Leap-Frog* property if, for *every* subset $S = \{(u_1, v_1), \ldots, (u_k, v_k)\}$,

$$\begin{aligned} t \cdot d(u_1, v_1) \; < \; & \textstyle\sum_{i=2}^{k} d(u_i, v_i) + \\ & t \cdot (\textstyle\sum_{i=1}^{k-1} d(v_i, u_{i+1}) + d(v_k, u_1)). \end{aligned}$$

The Leap-Frog property restricts how a set of line segments may be positioned in space, and is a generalization of the property described in Lemma 2.3.

Lemma 2.4 *The spanner edges satisfy the Leap-Frog property.*

The proof of this lemma is fairly straightforward and follows from Lemma 2.2 and Lemma 2.1.

In the rest of the paper, we need not concern ourselves any more with spanners. Instead, for any

$t > 1$, we show that if a set of edges E in 3-dimensional space satisfies the Leap-Frog property, then $wt(E) = O(wt(MST))$, where MST is a minimum spanning tree connecting the endpoints of E. The constant implicit in the O-notation depends only on t.

In fact, what we prove is actually something slightly stronger. Let T be a *minimum Steiner tree* of the endpoints of E. This tree may be shorter than the MST. We prove that $wt(E) = O(wt(T))$. Of course, this fact is not asymptotically significant because in Euclidean space it is well known that $wt(MST) = O(wt(T))$.

3 Overview

As mentioned earlier, for any $t > 1$ we analyze the weight of a set of edges E in 3-dimensional space satisfying the Leap-Frog property. Let T always refer to a minimum Steiner tree of the endpoints of E.

We first provide an overview of the proof. The various steps are elaborated in later sections. The proof employs a complex *charging* technique, where edges of E are charged to portions of T in such a way that 1) each edge of E is charged to a portion of T proportional to its length, and 2) the same portion of T is never charged twice.

The approach is as follows. To facilitate the analysis, the edges of E go through a *Cleanup* phase, which partitions the edges into a constant number of groups. For each group we prove that the sum of the lengths of the edges is at most a constant times the weight of T. Since the number of groups is itself a constant, this proves the necessary result. The cleanup process, which is described in some detail in Section 4, ensures that the set of edges in a group satisfy the nice properties listed below:

1. **Near-Parallel Property**: Any pair of edges in a group are nearly parallel.

2. **Length-Grouping Property**: In a group, any two edges have lengths that are either nearly equal or differ by a large amount.

3. **Empty-Region Property**: In a group any two edges of nearly equal length are far apart.

4. **Nested-Dumbbell Property**: As explained later, this is a crucial property used in the proof. Each edge in a group defines a *dumbbell*. A dumbbell consists of an edge attached at each endpoint to a cylinder, or *dumbbell head*, of suitable size. These dumbbells are properly *nested*, that is, either one is completely contained in the head of the other, or the two are completely disjoint (see Figure 3.1).

5. **Isolated-Centers Property**: The grouping is such that, given two edges of different sizes, the endpoint of the longer edge is not inside a head of the smaller dumbbell.

After the Cleanup process, each group is analyzed separately, and we show that the weight of each group is $O(wt(T))$. Since T is a minimum Steiner tree of the endpoints of E, it is also a (not necessarily minimum) Steiner tree of each group. To achieve our analysis we need to show that for any edge in the group, there is enough portion of T within the two dumbbell heads such that all the edges completely contained in the heads *as well as* the edge itself can be charged to this portion.

Without loss of generality we assume that the edges in a group are parallel to the z axis (vertical). First transform T in the following manner. Replace each edge (u, v) of T by a pair of line segments (u, w) and (w, v), where w is a new Steiner point, (u, w) is parallel to the z axis, and (w, v) is parallel to the xy plane. The weight of the new tree T is still within a constant multiple of the weight of the old tree.

T must connect all the endpoints of the edges in the group, hence it must pierce all the dumbbell heads. We partition the set of dumbbells into two sets based on the location of the points at which T intersects the surface of the cylinders of the dumbbell. Let (u, v) be an edge in a group. Consider the fragments of T inside the two heads of (u, v). Let T_u (respectively T_v) be the fragment containing u (respectively v). If neither T_u nor T_v pierces any of the flat faces of the two dumbbell heads, then we call such a dumbbell a *Lateral* dumbbell with respect to T, and the corresponding edge a *Lateral* edge. All other dumbbells (and edges) are termed as *Non-Lateral* with respect to T. Figure 3.2 describes these notions.

Accordingly, each group is further partitioned into a Lateral and a Non-Lateral group. We shall show separately that the weight of the edges in each group is $O(wt(T))$. This is tackled in Sections 5 and 6, respectively.

The global techniques used in both cases are similar, because in both we perform a bottom-up analysis in the implicit forest defined by the nesting of the dumbbells. The local techniques used, however, are quite different. The non-lateral case uses the generalized version of the Leap-Frog property, while the lateral case does not use the Leap-Frog property at all. Instead, it uses local properties of minimum Steiner trees in 2-dimensions. We have not been able to extend this part of the proof to arbitrary dimensions. However, what is interesting and surprising is that minimum Steiner trees are used as a tool for analyzing the lengths of edges.

4 Cleanup Phase

In this section we show that a set of edges satisfying the Leap-Frog property can be partitioned into a constant number of groups, each having the nice properties listed below.

A number of constants are introduced in this section. Each constant is actually a function dependent only on t. In this version of the paper we do not derive the functions; instead, we provide intuitive explanations.

1. **Near-Parallel Property**: Any pair of edges in a group is *Near-Parallel*, i.e., the angle between the two edges is less than a small constant, which depends on t.

   The technique to achieve this grouping is borrowed from [2]. Cover the half-space above the xy-plane into a constant number of cones $C_1, \ldots C_{O(1)}$, such that the angle of each cone is a small constant. Each edge of E is assigned to a group E_i in such a manner that, upon translating the endpoint with lesser z-coordinate to the origin, the edge lies completely within cone C_i. The union of the E_i's is E. In the rest of the paper, we assume that the edges in any E_i are actually parallel; a more rigorous proof allowing Near-Parallel edges appears in the full version.

2. **Length-Grouping Property**: Let ϵ be a small constant, and s a large constant, both dependent upon t. For the ease of notation, we define $\delta = 1-\epsilon$. In a group, any two edges with lengths x and y (say $x \geq y$) are either of *Near-Equal* length ($x \geq y \geq x\delta$), or have lengths that differ by a large amount ($y \leq x\delta^{s-1}$).

   To achieve this grouping, consider a group E_i of Near-Parallel edges. Let the longest edge be of length x. Partition the interval $[0, x]$ into an infinite set of intervals $\{[x\delta, x], [x\delta^2, x\delta] \ldots\}$. Define the subgroup $E_{i,j}$ for $j \in \{0, \ldots, s-1\}$ as having those edges whose lengths belong to the intervals $\{[x\delta^{j+1}, x\delta^j], [x\delta^{j+s+1}, x\delta^{j+s}], [x\delta^{j+2s+1}, x\delta^{j+2s}], \ldots\}$.

   Clearly, $E_{i,j}$ satisfies both Properties 1 and 2, and there is only a constant number of such groups. Again, we may imagine that Near-Equal edges are actually equal; a more rigorous proof allowing Near-Equal edges appears in the full version.

3. **Empty-Region Property**: A group $E_{i,j}$ may be further subdivided into groups $E_{i,j,k}$ such that any two points on two Near-Equal edges (of length x each) are separated in space by a distance which is a large (but constant) multiple of x. In other words, two Near-Equal edges cannot be close to each other.

   To achieve this grouping, construct a graph where the nodes are edges of $E_{i,j}$, and two Near-Equal nodes e_1 and e_2 (of length x) have a graph edge between them if e_1 intersects a large cylinder of radius cx and height cx (c is a large constant), centered at the center of e_2. This graph has a constant degree, because by Lemma 2.3, there can be only a constant number of similar Near-Equal edges whose endpoints can be packed into the cylinder. Thus this graph has a constant chromatic number, and consequently a constant number of independent sets, and each set forms a group satisfying Properties 1, 2, and 3.

4. **Nested-Dumbbell Property**: Each group $E_{i,j,k}$ also satisfies the Nested-Dumbbell property, which is the most crucial property in the entire proof. First we describe dumbbells more precisely. For each edge of length x, the dumbbell head around an endpoint is a cylinder of radius ax and height hx ($1 \gg h \gg a > 0$). The axis of the cylinders lie along the edge. Observe that the radius of each cylinder is much smaller than its height, which in turn is much smaller than the edge length. Any two dumbbells in the same group $E_{i,j,k}$ are nested, i.e., either one dumbbell is completely contained within one of the cylinders of the other dumbbell, or the two dumbbells are completely disjoint (as in Figure 3.1).

   As we shall now see, this nesting is achieved by relaxing the requirement that each dumbbell head be a cylinder. Instead, each dumbbell head is actually a closed region in space *resembling* a cylinder, in that the minimum radius is nearly equal to the maximum radius, and the minimum height is nearly equal to the maximum height. The construction of the dumbbells proceeds from the smallest to the largest edge in the group. Recall that edges are either Near-Equal, or vary greatly in length.

First construct dumbbells for the smallest Near-Equal class. Since they are pairwise sufficiently far apart, we can construct heads which are actual cylinders. At any stage, suppose we are considering the edge e of length x. As an induction hypothesis, assume that all the dumbbells of the previous edges satisfy the Nested-Dumbbell property. As a first attempt, construct a dumbbell D, by attaching cylindrical heads of radius and height ax and hx respectively. If the surface of D does not intersect any of the previous dumbbells, we can go on to the next edge. Suppose the surface of D intersects some of the previous dumbbells, say $D_1, \ldots D_m$. However, the sizes of the previous dumbbells are much smaller than even the radius of D (this can be achieved by a very large constant s in the Length-Grouping property). Construct the actual dumbbell as $D \setminus (D_1 \cup D_2 \cup \ldots D_m)$. Clearly the nested property is maintained. The shapes of the dumbbells may get quite complicated as the process continues, but the heads still *resemble* cylinders, because the minimum radius and height are always very close to their maximum values. Some of these notions are illustrated in Figure 4.1.

From now on, it will help to imagine that all the heads are actually cylinders; a more rigorous proof allowing the general shapes appears in the full version of the paper.

5. **Isolated-Centers Property**: The grouping is such that even though a smaller dumbbell is completely contained in a larger dumbbell head, the endpoint of the longer edge is not inside a cylinder of the smaller dumbbell.

   Each group $E_{i,j,k}$ is further broken up into groups $E_{i,j,k,l}$ to achieve this property. Techniques similar to Property 3 are used here (application of Lemma 2.3, as well as graph coloring), and we omit details in this version.

5 The Non-Lateral Case

In this section we prove that the sum of the weights of the edges in a non-lateral group with respect to the tree T is $O(wt(T))$. Recall that T is a Steiner tree of the endpoints of the edges in the group, transformed in such a way that each line segment of T is either parallel to the z axis or the xy plane. It is easy to see that from T we can construct a traveling salesman cycle C of the endpoints of the group edges, by doubling the line segments. Henceforth, we shall show that the weight of the non-lateral group is $O(wt(C))$.

The dumbbells in the group define a forest of rooted trees as follows. If dumbbell D_1 is contained within a head of dumbbell D_2, then D_2 is an ancestor of D_1. Thus the dumbbells which do not contain other dumbbells are the leaves of the forest, and likewise the dumbbells which are not contained by other dumbbells are the roots of the forest. We shall employ a bottom-up analysis in this forest. At any stage, we shall remove some leaf dumbbells, and restructure C to connect the remaining dumbbells in such a way that several invariants remain true.

We describe the weight of the cycle C. During the bottom-up analysis, each segment of the current cycle C has weight equal either to its Euclidean length or to zero. The segments with weight zero are called *null segments*, and are crucial to the charging scheme. All null segments are vertical, and initially there are no null segments. The weight of the current C, denoted $wt(C)$, is defined as the sum of its segment weights.

We employ the following notation. Let $e = (u, v)$ be an edge in the group. Let D_e refer to the dumbbell of e, and H_u and H_v its two heads. Let C_u (C_v) be the maximal portion of the current cycle C that contains u (v) and is fully contained within H_u (H_v). Let $\{u_1, u_2\}$ ($\{v_1, v_2\}$) be the two points of C_u (C_v) on the surface of H_u (H_v). The non-lateral property dictates that at least one of u_1, u_2, v_1, v_2 lies on a flat face of the heads.

Basis Step: Perform the following restructuring at all the leaf dumbbells. Let D_e be a leaf dumbbell with edge weight x, and with cylindrical dumbbell heads of height hx and radius ax. W.l.o.g., let u_1 lie on a flat face of D_e. Replace C_u by a horizontal segment (u_1, w) and a vertical segment (w, u_2), where w is a new Steiner point inside H_e, and (w, u_2) is a null segment.

Do not remove the leaf dumbbells as yet. Observe that within each leaf dumbbell the weight of C has reduced by at least $hx/2$ (because of the use of a null segment). Since $h \gg a$, the edge corresponding to the leaf dumbbell can be charged to this difference.

Inductive Hypothesis: The following properties hold at all stages of the bottom up traversal. Let w

be the weight of the original C before the traversal started, w_1 be the weight of the current C (where null segments have weight zero), and w_2 be the weight of all dumbbells charged so far.

1. $w - w_1 > c \cdot w_2$, for some constant c.
2. In the current forest, the leaf dumbbells have been charged already, while the nonleaf ones have not yet been charged.
3. The current cycle C connects all the endpoints of the nonleaf dumbbells.
4. The portions of the current C outside the leaf dumbbells are identical to the corresponding portions of the original C.
5. Some of the vertical line segments of the current C may be null segments. However, all such segments are confined within the heads of the leaves.

Inductive Step: Select an edge $e = (u, v)$ such that all its children are leaves. Let its length be x. W.l.o.g. assume u_1 lies on a flat face of D_e. We remove the children, thus making D_e a leaf. At the same time, we restructure C so that the inductive hypothesis is maintained. Two cases arise.

Case 1 : $wt(C_u) > 3ax$

In this case, restructure C by replacing C_u by a horizontal segment (u_1, w) and a vertical null segment (w, u_2). Remove all children of D_e from the forest. Since the length of (u_1, w) is at most $2ax$, the difference between the weights of the previous C and the new C is at least ax. The edge e can be charged to this difference. It is easy to see that the inductive hypothesis is still true.

Case 2 : $wt(C_u) \leq 3ax$

We will show that this case is impossible. This is where where the Leap-Frog property is used, and the analysis is more difficult. In this version we simply provide an intuitive outline.

Ignore the null labels on vertical segments of C_u for the moment. The portion of C_u from u_1 to u has a vertical displacement of $hx/2$ and a *total* horizontal displacement of at most $3ax$. (By total, we mean the sum of all horizontal segments on this portion). Since $h \gg a$, this intuitively implies that this portion of C_u is vertical most of the way, occasionally straying a little bit in the horizontal direction. The proof will show that, due to the Leap-Frog property, null segments can comprise only a small fraction of the vertical sections of this portion of C_u.

In this version, we describe a special situation which provides insight into why Case 2 cannot occur. Suppose the portion of C_u from u_1 to u is completely vertical (see Figure 5.1). We need to show that a large fraction of this path is not composed of null segments. Recall that null segments are confined within the dumbbell heads of the children of D_e. We will show that, given a vertical segment (u_1, u) of length $hx/2$, it is impossible to arrange children dumbbells alongside it with their heads intersecting the segment, such that the *exposed* portions (portions not contained within the dumbbell heads of the children of D_e) of (u_1, u) is $\leq 3ax$.

Consider the example in Figure 5.1 again. The leaf dumbbells that intersect (u_1, u) are those of $e_1 \ldots e_5$, where $e_i = (p_i, q_i)$, and whose corresponding edge lengths are, say, $x_1 \ldots x_5$. Here p_i is the top endpoint while q_i is the bottom endpoint. In what follows, based on the context, e_i may refer to an edge, or its corresponding dumbbell. By selecting suitable constants in the Length-Grouping and Empty-Region properties, it can be assured that these dumbbells cannot *straddle* one another, that is, the endpoints of e_i and e_j cannot be vertically ordered as either p_i, p_j, q_i, q_j or p_j, p_i, q_j, q_i.

This situation defines a natural hierarchy of the dumbbells. For example, in Figure 5.1, e_1 and e_2 are at the top of the hierarchy. At the next level, e_3 and e_4 are confined between the two heads of e_1, while e_5 is confined between the two heads of e_2. In general there could be further levels in the hierarchy.

Consider the top level. Clearly the portions of (u_1, u) above the top head of e_1, between the bottom head of e_1 and the top head of e_2, and below the bottom head of e_2 cannot contain null segments. The total length of these portions is $hx/2 - (x_1 + x_2) - (hx_1 + hx_2)$.

Let us now go to the next level; in fact, two levels will be enough to bring about the contradiction. Consider e_1. In what follows, we shall show that the exposed portions between the two heads of e_1 can be made as close to x_1 as possible, by suitably selecting the constants in the Cleanup stage, and applying the Leap-Frog property. Let the length of the exposed portions between the top heads of e_1 and e_3, between the bottom head of e_3 and the top head of e_4, and between the bottom heads of e_4 and e_1 be z. Let $d = d(p_1, p_3) + d(q_3, p_4) + d(q_4, q_1)$. Clearly, $d < z + (h + 2a)(x_1 + x_3 + x_4)$. But $z \leq x_1 - (x_3 + x_4)$. Let $h + 2a = \beta$. Substituting we get $d < z(1-\beta) + x_1 2\beta$.

The Leap-Frog property says that $tx_1 < (x_3 + x_4) + td$. Substituting for $(x_3 + x_4)$ and d, and simplifying, we get

$$x_1\left(\frac{1-2\beta-1/t}{1-\beta-1/t}\right) < z$$

In the above example, by suitably selecting a small β (by making h and a suitably small), we can make z as close to x_1 as we like. We offer an intuitive description for this phenomena. Given any edge e_1, we cannot "leap frog" from one endpoint to the other a lot, via other parallel edges close by. Now consider e_2 at the same level of the hierarchy. As above, we can make the exposed portions between the top heads (and bottom heads) of e_2 and e_5 as close to x_2 as we like.

At this stage we see that the exposed portion of (u_1, u) between the dumbbell heads is at least

$$hx/2-(x_1+x_2)-(hx_1+hx_2)+\frac{1-2\beta-1/t}{1-\beta-1/t}(x_1+x_2)$$

But an upper bound for $x_1 + x_2$ is $hx/2$. Substituting above and simplifying, we see that the total exposed portion is at least

$$(hx/2)\left(\frac{1-2\beta-1/t}{1-\beta-1/t} - h\right)$$

By suitably selecting a small β (by making h and a suitably small), we can make this as close to $hx/2$ as we like. Since $h \gg a$, clearly we can violate the condition that this exposed portion be $\leq 3ax$, which disallows Case 2.

The entire argument can be applied to any hierarchy of dumbbells intersecting with (u_1, u). In a more rigorous proof, we have to take into account the fact that the portion of C_u from u_1 to u is "almost vertical". This is omitted from this version of the paper.

In the next section, we consider the case of lateral dumbbells.

6 The Lateral Case

In this section we prove that the sum of the lengths of the edges in a group that are lateral with respect to T is bounded by a constant factor times the weight of T. Recall that T is a minimum Steiner tree of the endpoints of the edges in the group, transformed such that each segment of T is either parallel to the z axis or to the xy plane.

In our analysis a horizontal segment has weight equal to its length, while a vertical segment has zero weight. Intuitively, what we will be proving is that it is sufficient to charge the edges of the group only against the horizontal segments of T.

To prove this claim, we use induction on the number of dumbbells. For the inductive step, we pick the shortest remaining edge $e = (u, v)$. The corresponding dumbbell is then eliminated and the tree is restructured. The length of e is charged to the difference in weight between the original and the restructured tree. The inductive hypothesis states that all eliminated edges have already been charged to horizontal portions of T. The portion of T inside only one of the two dumbbell heads of e is restructured; the portion within the other dumbbell head is left unchanged. The restructured tree is a Steiner tree that connects the remaining points, but has a total length that is less than the previous tree by an amount proportional to the length of the eliminated dumbbell. The eliminated dumbbell can be charged against this reduction in length. The remaining dumbbells continue to be lateral with respect to the new tree, since the restructured portion of T lies within the eliminated dumbbell.

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. Note that once the dumbbell associated with e is eliminated, the restructured tree need not visit the vertices u and v. So if T_u is a simple segment from the surface of the cylinder to vertex u, then this entire piece can be eliminated, to which the edge e can be charged. However, things may not be that simple. In general, the piece T_u may be a complex tree structure. We will show that regardless of its precise structure, we can restructure it and shorten it sufficiently.

Let the piece T_u pierce the dumbbell head at points $U = \{u_1, u_2, \ldots, u_k\}$. Let u_1 be the point in U that is "closest" to bottom portion T_v. More precisely, let u_1 be the point in U such that any path from a point in U to T_v must pass through u_1; such a u_1 exists. Let $U' \subseteq U$ be such that $u_i \in U'$ if the path from u_i to u_1 does not pass through u. Let T'_u be the portion of T_u that connects U' and u.

Our approach will be to restructure T'_u so that it still connects all the points in U', but is disconnected from u. Since this may disconnect T, we can reinstate connectivity by adding a vertical null segment from u to v. Vertical segment can be added "free of charge" since they have zero weight. Consequently, it is sufficient to imagine the points inside a dumbbell head to be on a plane. The first step of the restructuring will be achieved by replacing T'_u by the minimum Steiner tree connecting U' and u. Call this tree S_u. Clearly, since a minimum Steiner tree is a minimum connecting network, $wt(S_u) \leq wt(T'_u)$.

Now consider an imaginary circle C centered at u and of radius $ax\beta$ where x is the length of the edge e, ax is the radius of the cylindrical dumbbell head around u, and $\beta < 1$ is an appropriately small constant. Several cases arise for the restructuring of S_u. These cases depend on the number of Steiner points, s, that are present on S_u and that lie inside circle C. In fact, we will show that the only possible cases are $s = 0$ and $s = 1$, and that $s \geq 2$ is impossible.

Case 1 : $s = 0$

If S_u does not have any Steiner points inside C, then the portion of S_u inside C is a straight line of length at least $ax\beta$. To restructure T, simply remove this piece of S_u, and add a vertical null segment from u to v. Clearly the weight is reduced by at least a constant factor times the length of e, and the inductive hypothesis is preserved. This case is shown in Figure 6.1.

Case 2 : $s = 1$

In this case, S_u touches C at two points, let us call them u' and u''. To restructure T, remove the entire portion of S_u inside C, connect u' and u'' by a straight line, and add a vertical null segment from u to v. Since the two straight lines that pierce C form an angle of 120 degrees, this restructuring involves a weight reduction of at least $(2 - \sqrt{3})ax\beta$, to which e can be charged (see Figure 6.2).

Case 3 : $s \geq 2$

We show that this case is impossible.

Assume that S_u starts off from u with segment e_u. Since $s \geq 2$, the segment e_u must terminate inside C at a Steiner point w. Since every Steiner point must have degree 3, assume that segments e_w and e_y are the other two segments incident at w. Assume that the segment e_y terminates inside C at another Steiner point y. Let e_z be another segment incident on y, as shown in Figure 6.3.

Now consider the segments e_w and e_z. We will first show that both e_w and e_z cannot be too long. In other words, it cannot happen that both $wt(e_w)$ and $wt(e_z)$ are larger than $ax\beta/4$. If that does happen then S_u cannot be a minimum Steiner tree since it can be shortened as shown in Figure 6.4.

If either e_w or e_z is not too long, say $wt(e_z) \leq ax\beta/4$, then consider the Steiner point z into which segment e_z terminates. Since z is a Steiner point, there must be another segment (say e_x) that is parallel to e_w. Now again, if both e_w and e_x are very long (of length more than $ax\beta/4$), then as shown in Figure 6.5, S_u cannot be a minimum Steiner tree.

However, if either e_w or e_x is short, then their subtrees will eventually intersect before hitting the surface of the cylinder, contradicting the fact that S_u is a minimum Steiner tree.

This completes the proof that Case 3 ($s \geq 2$) is impossible.

7 Open Problems

The most interesting and immediate open question we would like to answer is whether our techniques can be used without much modification for the case of higher dimensional point sets. It would also be interesting to see whether the constant that appears in the weight bound of Theorem 1.1 can be reduced.

References

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

[2] Chandra, B; Das, G.; Narasimhan, G.; Soares, J., New Sparseness Results on Graph Spanners, *ACM Symposium on Computational Geometry*, (1992), 192-201.

[3] Das, G.; Joseph, D., Which triangulations approximate the complete graph? *International Symposium on Optimal Algorithms* (1989), (LNCS, Springer-Verlag).

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