Euclidean Spanners: Short, Thin, and Lanky

Sunil Arya* Gautam Das† David M. Mount‡ Jeffrey S. Salowe§ Michiel Smid*

Abstract

Euclidean spanners are important data structures in geometric algorithm design, because they provide a means of approximating the complete Euclidean graph with only $O(n)$ edges, so that the shortest path length between each pair of points is not more than a constant factor longer than the Euclidean distance between the points. In many applications of spanners, it is important that the spanner possess a number of additional properties: low total edge weight, bounded degree, and low diameter. Existing research on spanners has considered one property or the other. We show that it is possible to build spanners in optimal $O(n \log n)$ time and $O(n)$ space that achieve optimal or near optimal tradeoffs between all combinations of these properties. We achieve these results in large part because of a new structure, called the *dumbbell tree* which provides a method of decomposing a spanner into a constant number of trees, so that each of the $O(n^2)$ spanner paths is mapped entirely to a path in one of these trees.

1 Introduction

Let $G = (V, E)$ be a weighted graph, and let $d_G(u, v)$ be the length of a shortest path between vertices u and v in G. Let $t > 1$ be any constant. A subgraph G' is a *t-spanner* for G if, for all pairs of vertices u and v, $d_{G'}(u, v)/d_G(u, v) \leq t$. When V is a set of n points in $\mathbb{R}^k$, G is the complete graph, and the length of edge (u, v) is the Euclidean distance between these points, we call G a complete Euclidean graph and G' a *Euclidean t-spanner*. For the purposes of deriving asymptotic bounds, we assume that the dimension k and the spanner factor t are constants independent of n. It is known how to construct a Euclidean t-spanner having $O(n)$ edges in $O(n \log n)$ time [5, 13, 14].

Spanners are important geometrical structures, since they provide a mechanism for approximating the complete Euclidean graph in a much more economical form. Of course, a spanner should have a small number of edges (ideally $O(n)$), but for many applications, it is quite important that the spanner be endowed with other properties. These include the following:

Low weight: The total sum of the edge lengths in the spanner should be as small as possible. The best that can be hoped for is some constant times the weight of the minimum spanning tree, $O(w(MST))$.

*Max-Planck-Institut für Informatik, D-66123 Saarbrücken, Germany. Email: {arya,michiel}@mpi-sb.mpg.de. Supported by the ESPRIT Basic Research Actions Program, under contract No. 7141 (project ALCOM II).

†Math Sciences Dept., The University of Memphis, Memphis, TN 38152. Supported in part by NSF Grant CCR-9306822. E-mail: dasg@next1.msci.memst.edu.

‡Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, Maryland. Partially supported by NSF Grant CCR-9310705. This work was done while visiting the Max-Planck-Institut für Informatik, Saarbrücken. E-mail: mount@cs.umd.edu.

§QuesTech, Inc., 7600A Leesburg Pike, Falls Church, VA 22043. This work was done while visiting the Max-Planck-Institut für Informatik, Saarbrücken. E-mail: jsalowe@nvl.army.mil.

Bounded degree: The number of edges incident to any vertex should be bounded.

Small spanner diameter: The *spanner diameter* (or simply *diameter*) is defined as the smallest integer D such that for any pair of vertices, u and v, there is a t-spanner path between u and v containing at most D edges. For spanners of bounded degree the best that can be hoped for is logarithmic diameter. In some applications even smaller diameters may be desirable, but this comes at the expense of increasing degree.

A natural analogy can be made between spanners and a transportation network of roads connecting a large number of locations. Low weight means that the amount of concrete needed to build the roads is small, bounded degree means that no location in the network has more than a bounded number of roads incident to it, and small diameter means that it is possible to describe any spanner path concisely. Existing work on spanners has focused on achieving one property or the other. However, a transportation network which achieves small diameter by massively increasing total weight is of little practical value. This suggests the important question of whether there exist spanners that *simultaneously* achieve some or all of these properties.

In this paper we present a strong positive answer to this question. We present a number of new constructions for spanners. In almost all cases these constructions are provably optimal from the perspectives of computation time, space, and performance on the properties listed above. The problem is complicated by the fact that there are obvious tradeoffs between these properties. (For example, reducing diameter requires the creation of long edges, which in turn increases total weight, or may increase the number of edges needed in the spanner.) For this reason, we consider all possible combinations of these properties.

The results of this paper arise from a number of improved techniques in spanner constructions, but one deserves particular mention. An important data structure used in the construction of spanners is the well-separated pair decomposition, introduced by Callahan and Kosaraju [4]. This structure represents the $O(n^2)$ pairs of points using only $O(n)$ pairs of geometrically "well-separated" pairs of subsets of points (definitions will be given later). In this paper, we present a novel method of further decomposing a well-separated pair decomposition into a constant number of hierarchically organized sets of well-separated pairs. (The constant depends on the dimension and the separation factor.) Using this decomposition, we show that a class of spanners can be viewed as being the union of a constant number of trees, which we call *dumbbell trees*. Moreover, each of the $O(n^2)$ spanner paths arises as the unique path between two leaves in one of these trees. The fact that the $O(n^2)$ spanner paths can be partitioned among a constant number of trees is a rather remarkable fact in itself, and suggests a great deal about special structure of these graphs. Because of the importance of well-separated pair decompositions to a variety of geometric problems, we suspect that this decomposition may be of use to other geometric problems. The idea of dumbbells has appeared before [7], but their use in decomposing spanner paths is new to this paper.

Here is a summary of the results in this paper. All of the spanner constructions described below run in optimal $O(n \log n)$ time and $O(n)$ space for any fixed dimension k.

Degree: We present an optimal $O(n \log n)$ time construction for spanners of bounded degree. This improves the best known algorithm, due to Arya and Smid [3], which runs in $O(n \log^k n)$ time.

Weight: We present an optimal $O(n \log n)$ time spanner construction that has optimal weight $O(w(MST))$. This improves the best known construction for spanners of low weight, which was due to Das and Narasimhan [8], and which runs in $O(n \log^2 n)$ time.

Diameter: Arya, Mount and Smid [2] give randomized and deterministic constructions of spanners with $O(n)$ edges and $O(\log n)$ spanner diameter. We show that it is possible to achieve diameter $\alpha(n) + 2$ with the same number of edges, where $\alpha(n)$ is the inverse of Ackermann's function. Furthermore, we present a spectrum of tradeoffs between size and diameter. For example, we construct spanners of diameter 2 with $O(n \log n)$ edges, diameter

3 with $O(n \log\log n)$ edges, diameter 4 with $O(n \log^* n)$ edges, and so on. All these spanners have an optimal number of edges for the given diameter.

Degree and weight: The low-weight construction mentioned already has bounded degree, and hence provides an optimal solution to this problem as well. There are no previous results on this problem.

Weight and diameter: By using a new analysis tool, we show that the deterministic low-diameter construction of Arya, Mount and Smid [2] has weight $O(w(MST) \log n)$ as well as diameter $O(\log n)$. This combination is optimal. No simultaneous bounds were previously known.

Degree and diameter: We show how to construct a spanner with bounded degree and $O(\log n)$ diameter. This is optimal with respect to both diameter bound and construction time. No simultaneous bounds were known for this problem.

Degree, weight and diameter: We show how to construct a spanner with bounded degree, weight $O(w(MST) \log^2 n)$, and diameter $O(\log n)$. No simultaneous bounds were previously known.

In summary, all of our results are optimal in terms of providing the best tradeoffs between these properties, except for the spanner having simultaneously bounded degree, low weight, and low diameter, which is possibly suboptimal by at most an $O(\log n)$ factor in weight.

The rest of this paper is organized as follows. In Section 2, we briefly recall the well-separated pair decomposition. In Section 3, we define the dumbbell tree, and show that there exists a spanner that can be decomposed into a constant number of such trees. In Section 4, we give a simple optimal algorithm for constructing a t-spanner of bounded degree. In Section 5, we show that the spanner that results from the well-separated pair decomposition can be pruned in such a way that we get a spanner of weight $O(w(MST))$. Section 6 considers spanners of low diameter. Our results of Section 3 imply that it suffices to add edges to a constant number of bounded degree trees in order to get a spanner of low diameter. This is done by using a technique due to Alon and Schieber[1]. In Section 7, we show how to combine the dumbbell tree with topology trees [10] in order to get a spanner of bounded degree and $O(\log n)$ diameter. In Section 8, we show that spanners that result from the well-separated pair decomposition have weight $O(w(MST) \log n)$. Combining this fact with a result of [2] gives a t-spanner of weight $O(w(MST) \log n)$ and diameter $O(\log n)$. Finally, in Section 9, we consider all properties degree, weight and diameter simultaneously.

2 Split trees and well-separated pairs

Virtually all of our spanner constructions will rely on the notion of a *split tree* and a *well-separated pair decomposition* of a set of points [4, 13, 14]. In this section, we review these data structures.

A *split tree* is a tree that stems from a hierarchical decomposition of a point set into regions that are k-dimensional rectangles of bounded aspect ratio. There are a number of variants on a split tree. We outline the fair split tree, due to Callahan and Kosaraju [4]. Place a smallest-possible k-rectangle R_0 about the point set V. The root of the split tree is R_0. Choose the longest side of R_0 and divide it into two at its bisector. Rectangle R_0 is therefore split into two smaller rectangles, R_1 and R_2. Then the left subtree of R_0 is the split tree for $R_1 \cap V$, and the right subtree is the split tree for $R_2 \cap V$. The process is repeated until a single point remains.

In order to simplify some of our arguments, it is convenient to think of a fair split tree in an ideal form, which we call the *idealized box split tree*. In this tree, rectangles are k-dimensional hypercubes, each split recursively into 2^k identical hypercubes of half the side length. Actual constructions will be carried out using the fair split tree, but the idealized box split tree provides a clean way of conceptualizing the fair split tree for purposes of analysis.

Next we consider well-separated pair decompositions. Let $s > 0$ be a constant. Two point sets A and B are *well separated* if they can be enclosed in k-spheres of radius r, whose distance of closest approach is at least sr. A *well-separated pair* decomposition is a set of pairs of nonempty subsets of S, $\{\{A_1, B_1\}, \{A_2, B_2\}, \ldots, \{A_m, B_m\}\}$, such that (1) the sets A_i and B_i are disjoint, (2) for each pair $a, b \in S$, there is a unique pair $\{A_i, B_i\}$ such that $a \in A_i$ and $b \in B_i$, and (3) A_i and B_i are well-separated. Callahan and Kosaraju use a split tree to compute a set of $O(s^k n)$ well-separated pairs in $O(n \log n + s^k n)$ time.

Given these well-separated pairs, Callahan and Kosaraju show that a spanner can be constructed easily. For each pair $\{A_i, B_i\}$ in the well-separated pair decomposition, choose arbitrary points, called *representatives*, $a_i \in A_i$ and $b_i \in B_i$, and connect a_i and b_i with an edge in the spanner. Similar constructions were previously given by Vaidya[14] and Salowe[13].

3 The dumbbell tree

One of the major difficulties in establishing the results of this paper is the lack of structure in well-separated pair decompositions and the spanners that are derived from them. Unlike the split tree, well-separated pair decompositions do not possess any obvious hierarchical structure. One of the major innovations of this paper is the observation that well-separated pair decompositions, and hence the spanners derived from them, can be decomposed into a constant number of hierarchically organized structures. This greatly simplifies the analysis and construction of spanners, by reducing problems on general graphs to much simpler problems on trees. This decomposition may have applications to a number of other problems where sparse geometric graphs are used.

Space does not permit a complete presentation of the decomposition, but the intuition is relatively straightforward. First observe that each well-separated pair $\{A_i, B_i\}$ can be viewed as a geometrical object, consisting of two rectangles containing A_i and B_i, respectively, joined by a line segment. The resulting shape, is called a *dumbbell* and the rectangles (or in fact, small perturbations of these rectangles) are called the *heads* of the dumbbell. The *length* of a dumbbell is defined as the distance between the centers of its heads. The *size* of a head is defined to be half its diameter.

Das, Heffernan and Narasimhan [7] introduced the concept of the dumbbell. We claim that it is possible to partition the set of dumbbells arising from the well-separated pair decomposition into a constant number of groups, such that within each group, dumbbell heads are either disjoint, or one dumbbell is nested entirely within the head of the other dumbbell. In particular, we can show the following (proofs will appear in the full paper):

Theorem 1 *Consider the dumbbells resulting from a well-separated pair decomposition of a set of n points in dimension k with separation factor s. In $O(n)$ time it is possible to partition these dumbbells into $O(s)^k$ classes, such that within each class:*

(1) *two dumbbells either have lengths that are within a factor of 2 of one another, or else they differ by a factor of at least s,*

(2) *any two dumbbells within the same length interval $[x, 2x]$, are separated by a distance greater than $2x/s$, and*

(3) *we may deform the heads of each dumbbell (forming pseudo-dumbbells) such that a dumbbell of length x has a head of size at most $4x/s$, and such that the heads of any two pseudo-dumbbells are either disjoint or else one is nested within a head of the other.*

The nesting of dumbbells provides us with a tree structure, which we call a *dumbbell tree*. The important fact about the dumbbell tree decomposition is that spanners can be derived from the well-separated pair decomposition which inherit this structure. Thus, they can be viewed as consisting of the union of a constant number of trees. Furthermore, we show that each spanner path is mapped entirely to one tree. Our main result is summarized in the following theorem:

Theorem 2 *Given a set V of n points in dimension k, and given $t > 1$, a forest consisting of $O(1)$ rooted binary trees can be built in $O(n \log n)$ time and $O(n)$ space, having the following properties:*

(1) *For each tree in the forest, there is a 1–1 correspondence between the leaves of this tree and the points of V.*

(2) *Each internal node has a unique* representative point, *which can be selected arbitrarily from the points in any of its descendent leaves.*

(3) *Given any two points $u, v \in V$, there is a tree T of the forest, so that the path formed by walking from representative to representative along the unique path in T between these nodes, is a t-spanner path for u and v.*

The constant factors for the number of trees, preprocessing time and space are $O(s^k \log s)$, where s is $O(k/(t-1))$. With the addition of an augmenting data structure of size $O(n)$, we can compute a t-spanner path between any two points in $O(p + \log n)$ time, where p is the number of edges on the path.

4 Spanners of bounded degree

In this section, we prove the following general result, which will be used to construct in $O(n \log n)$ time a t-spanner of bounded degree.

Theorem 3 *Let V be a set of n points in $\mathbb{R}^k$ and let $t' > t > 1$. Let G be a t-spanner for V and assume that the edges of G can be directed such that each point has outdegree at most α. In $O(n \log n)$ time, we can construct a t'-spanner for V in which each point has degree bounded by $O(\alpha\,(ct/(t'-t))^{k-1})$, for some constant c.*

In order to prove this result, we need the notion of single-sink spanner. Let V be a set of points in $\mathbb{R}^k$, let x be a point of V, and let $t > 1$. A directed graph having the points of V as its vertices is called an *x-single-sink t-spanner* for V if for every point y in V there is a t-spanner path from y to x.

Let θ be a fixed angle such that $0 < \theta < \pi/4$ and $1/(\cos\theta - \sin\theta) \leq t$. Let $\mathcal{C}$ be a collection of k-dimensional cones such that (i) each cone has its apex at the origin, (ii) each cone has angular diameter at most θ, and (iii) the union of these cones covers $\mathbb{R}^k$. For each point $p \in \mathbb{R}^k$ and $C \in \mathcal{C}$, let C_p be the cone $C + p := \{a + p : a \in C\}$.

Now consider the set V and the point x. Let n be the size of V. For each $C \in \mathcal{C}$, let V_C be the set of all points of $V \setminus \{x\}$ that are contained in the cone C_x. If a point is contained in more than one cone, then we put it in only one subset. If a subset V_C contains more than $n/2$ points, then we partition it (arbitrarily) into two subsets $V_{C,1}$ and $V_{C,2}$, each of size at most $n/2$.

The x-single-sink t-spanner for V is obtained as follows. For each subset V_C—or in case this set contains more than $n/2$ points, for each subset $V_{C,i}$, $i = 1, 2$—we take a point y in this subset that is closest to x, and we add an edge from y to x. Then we recursively construct a y-single-sink t-spanner for this subset. The recursion stops if a subset has size one.

Using exactly the same analysis as in [12], it follows that the graph is a single-sink t-spanner.

Lemma 1 *Let V be a set of n points in $\mathbb{R}^k$, let $x \in V$, and let $t > 1$. In $O(n \log n)$ time, we can construct an x-single-sink t-spanner for V, such that each point has outdegree at most 1 and indegree bounded by $O((c/(t-1))^{k-1})$, for some constant c.*

Now we are ready to give the transformation that will prove Theorem 3. Let V be a set of n points in $\mathbb{R}^k$ and let $t' > t > 1$. Let G be a t-spanner for V and assume that the edges of G can be directed such that each point has outdegree at most α. We denote this directed version of G by $\vec{G}$.

For each point x of V, we do the following. Consider all points of V that have an edge in $\vec{G}$ towards x. Let W be the set of these points. We replace all edges from W to x by an x-single-sink (t'/t)-spanner for the set $W \cup \{x\}$.

This gives a directed graph $\vec{G}_0$. We remove the direction from each edge and call the resulting graph G_0. We claim that G_0 is a t'-spanner for V in which each point has a degree bounded by a constant.

To prove this, let p and q be any two points of V. There is a t-spanner path $p = p_0, p_1, p_2, \ldots, p_m = q$ in G between p and q. Consider any edge $\{p_i, p_{i+1}\}$ on this path. Assume w.l.o.g. that in $\vec{G}$ this edge is directed from p_i to p_{i+1}. The directed graph $\vec{G}_0$ contains a p_{i+1}-single-sink (t'/t)-spanner with p_i as one of its vertices. Hence, in the graph G_0 there is a (t'/t)-spanner path between p_i and p_{i+1}. The concatenation of all these paths has length at most $\sum_{i=0}^{m-1} (t'/t)\,|p_i p_{i+1}| \leq (t'/t)\,t\,|pq| = t'\,|pq|$.

Consider the directed graphs $\vec{G}$ and $\vec{G}_0$. It follows from Lemma 1 that the outdegrees of both these graphs are the same. Hence, each point in $\vec{G}_0$ has outdegree at most α. Let x be any point of V. We bound the indegree of x in $\vec{G}_0$. This graph contains an x-single-sink spanner having $O((\frac{c}{t'/t-1})^{d-1}) = O((\frac{ct}{t'-t})^{d-1})$ edges with sink x. Now let y be any point such that $\vec{G}$ contains an edge from x to y. (There are at most α such points y.) Then x occurs in a y-single-sink spanner, and it has indegree bounded by $O((ct/(t'-t))^{d-1})$ in this spanner. Hence, in the directed graph $\vec{G}_0$, point x has indegree bounded by $O((1+\alpha)(ct/(t'-t))^{d-1})$. This proves that in the undirected t'-spanner G_0, each point has a degree bounded by a constant.

This proves Theorem 3. It turns out that several known spanners have the property that their edges can be directed such that each point has bounded outdegree. For example, for any $0 < \theta < \pi/4$, the θ-graph (see [12, 2]) is a t-spanner for $t \geq 1/(\cos\theta - \sin\theta)$. This spanner is directed already and each point has outdegree bounded by $O((c/\theta)^{d-1})$. It can be constructed in $O(n \log^{k-1} n)$ time.

Spanners based on well-separated pair decompositions also have the property we need. Essentially, the construction is to enumerate $O(n)$ sets of "box pairs." For each well-separated pair of boxes $\{A, B\}$, choose an arbitrary point $a \in A$ and $b \in B$ and add an edge $\{a, b\}$ to the spanner. This edge is directed from a to b if the parent box of A is not larger than the parent box of B, then the resulting graph has bounded outdegree. (For details, see Callahan and Kosaraju [5].) The entire graph can be constructed in $O(n \log n)$ time.

Now we can prove the main result of this section. Let V be any set of n points in $\mathbb{R}^k$ and let $t_0 > 1$. To construct a t_0-spanner for V having bounded degree, we set $t = \sqrt{t_0}$ and $t' = t_0$. In $O(n \log n)$ time, we construct a t-spanner G satisfying the condition of Lemma 3. Then we apply the given transformation and obtain the desired t_0-spanner. This proves:

Theorem 4 *Let V be a set of n points in $\mathbb{R}^k$ and let $t > 1$. In $O(n \log n)$ time, we can construct a t-spanner for V in which each point has a degree that is bounded by a constant only depending on t and k.*

5 Spanners of low weight

In this section, we give an $O(n \log n)$ time construction of a t-spanner that has weight $O(w(MST))$. In order to bound the weight of this graph, we use a theorem from Das, Narasimhan and Salowe[9].

Let $c > 0$ be a constant, let A be a set of edges, and let $e \in A$ be an edge of weight l. If it is possible to place a cylinder B of radius and height $c \cdot l$ each, such that the axis of B is a subedge of e and $B \cap (A \setminus \{e\}) = \emptyset$, then e is said to be *isolated*. The set A has the *isolation property* if all edges are isolated.

Theorem 5 ([9]) *If A has the isolation property, then $w(A) = O(w(MST))$, where MST is a minimum spanning tree with respect to the endpoints of A.*

It is easy to see that in the definition of the isolation property, one can replace the cylinder with a sphere, box, etc., without affecting the above theorem.

The low-weight spanner is constructed in the following way. Let C be a cone, and let $E(C)$ be the set of edges in the box well-separated pair construction, that, when translated such that one of their endpoints coincide with the apex of C, lie inside of C. We change the endpoints of an edge to ensure that the edge does not intersect the interior of the convex hull of the points within the respective boxes. These endpoints are chosen from among the points with maximum or minimum coordinates in a particular dimension. For each point p, mark edge $e \in E(C)$ if it is the shortest edge in $E(C)$ with one endpoint in a box ancestor of p. The spanner G_1 consists of the union of the marked edges.

We claim that G_1 is a spanner and that its edges satisfy the isolation property. The fact that G_1 is a spanner can be proved by a straightforward induction proof. To show that G_1 has the isolation property, we use some of the pruning techniques of Das, Heffernan, and Narasimhan[7]. We may assume that edges have been placed into a constant number of groups so that each edge has either approximately the same length or differs in length by a sufficiently large amount (but bounded by a constant factor).

We now show that edge $e = (a, b)$ has the isolation property. Edge e corresponds to some well-separated pair, say $\{A, B\}$, in the idealized box split tree. Note that the length $w(e)$ of e is related by a constant factor to the diameter d of these boxes. Place a box β of diameter d about the center point of e; we first claim that β does not contain any points.

To show this, edge e is present because there is some point $p \in A$, say, for which (a, b) is a shortest well-separated pair in direction C. If there was a point q in β, this would imply that p and q would be in a well-separated pair. However, the length of this well-separated pair would be smaller than $w(e)$, and it would be in the direction of C. (We note that this new well-separated pair edge may be in a nearby cone as well; this detail can be fixed using some pruning techniques.)

We now claim that at most a constant number of edges intersect a slightly-shrunken version of β. Suppose that an edge that is significantly shorter than $w(e)$ intersects β. Then we can shrink β by a small amount. Note there are no points inside of β, so β will not be shrunken by more than a small percentage.

Suppose that an edge $e' = (a', b')$ that is significantly longer than $w(e)$ intersects β. Then, if the idealized box A' containing a' is sufficiently far away, any point in A' would be in a well-separated pair with a, and the edge corresponding to this well-separated pair would be shorter than $w(e')$ and in approximately the same direction. (Again, the proof is only sketched. Note that this is where we need to choose the box representatives in a careful way.)

Finally, consider an edge that has approximately the same length as $w(e)$. Then these edges must correspond to idealized boxes within distance $w(e)$ of A or B. Packing arguments that use the fact that $w(e)$ is related to the width of A imply that there are only a constant number of idealized boxes in this area. Therefore, there are only at most a constant number of edges that can intersect β. Again, using the decomposition technique of Das et al.[7], one can partition the group of edges into a constant number of edge sets, each possessing the isolation property. The following theorem is proved.

Theorem 6 *In any dimension, for any $t > 1$, t-spanner G_1 can be constructed in $O(n \log n)$ time, and it has weight $O(w(MST))$.*

Our construction actually has bounded degree as well. This is because any set of edges possessing the isolation property has bounded degree (a straightforward proof). We therefore have the following:

Corollary 6.1 *In any dimension and for any $t > 1$, t-spanner G_1 can be constructed in $O(n \log n)$ time, and it has weight $O(w(MST))$ and bounded degree.*

6 Spanners of small diameter

We first consider the case of the 1-dimensional spanner, and then we show that all the higher-dimensional cases are closely related to the 1-dimensional case.

In the 1-dimensional case, the input is a set of n points on a line, and the output is a graph with small diameter. Surprisingly, a useful construction has already been discovered. It was devised by Alon and Schieber[1]. Among other results, this construction implies that there is a linear-sized 1-spanner with diameter $\alpha(n) + 2$.

Here are the essential aspects of the Alon and Schieber construction (they are tailored somewhat to enhance the analogy to our problem). Suppose we want a spanner of diameter d that contains as few edges as possible. Alon and Schieber divide up the point set into ℓ pieces, each piece of size n/ℓ. For each piece, recursively construct a spanner of diameter d; this accounts for spanner paths within the pieces. In order to account for the spanner paths between the pieces, select the points in each group with smallest and largest values. Each of the points in a particular group are connected directly with the two group representatives; the representatives themselves are connected with a spanner of diameter $d - 2$.

The number of edges $T(n, d)$ used in the Alon and Schieber construction is given by the recurrence:

$$\begin{aligned} T(n,d) &= O(n) + T(2\ell, d-2) + \ell T(n/\ell, d) \\ T(n,1) &= O(n^2). \end{aligned}$$

By choosing the values of ℓ appropriately, it is possible to show that $T(n, 2) = O(n \log n)$, $T(n, 3) = O(n \log \log n)$, $T(n, 4) = O(n \log^* n)$, and so on. It is also possible to show that the diameter is $\alpha(n)+2$ if one allows only $O(n)$ edges.

In order to generalize this idea to the k-dimensional case, we use the fact that there exists a spanner which can be represented as the union of a constant number of bounded degree trees. (See

Theorem 2.) Let T be one of these trees. We construct a modified version of T whose degree is bounded by a constant, and we endow it with some additional geometric properties. Specifically, this modified dumbbell tree T' is a tree whose vertices are original points and whose edges are Euclidean edges. The important geometric property of T' is the following: if a pair of points a and b are in a well-separated pair that actually appears as a dumbbell in T, then the path between a and b in T' is a t-spanner path.

The actual construction of T' is done in the following way. A dumbbell Δ in T contains several children dumbbells $\Delta_1, \Delta_2, \ldots, \Delta_m$ and several isolated points $p_1, p_2, \ldots p_j$. Consider these isolated points to be degenerate dumbbells and therefore children of Δ. This possibly large set of edges will be replaced by a tree T''' whose degree is bounded by a constant, described below.

For each box in Δ_i, choose a representative point. Let tree T'' be the minimal tree (with respect to edge inclusion) in the fair-split tree that connects these representative points. This tree T'' is a Steiner tree: it consists of original points (the representative points), Steiner points (degree-3 vertices), and paths connecting these two types of points. T''' is the tree that results from replacing each of the paths in T'' with a single edge.

The proof that T' has the t-spanner path property stems from the fact that the children dumbbells are much smaller than the parent dumbbell and the fact that the diameter of a box is halved in the fair-split tree within a constant number of levels. A detailed proof is omitted.

We apply a construction akin to the one of Alon and Schieber to shortcut the paths. Let $P = x_1, x_2, \ldots, x_j$ be a path in T'. By the triangle inequality, any path $P' = x_1, x_{p(2)}, x_{p(3)}, \ldots, x_j$, where $1 < p(2) < p(3) < \cdots < j$ has length less than or equal to the length of P. Appropriate shortcuts, therefore, have spanner properties. The details of how these shortcuts are constructed is omitted.

Theorem 7 *For any $t > 1$, and any dimension k, there is a t-spanner containing $O(n)$ edges and constructible in $O(n \log n)$ time with diameter $\alpha(n)+2$.*

If one allows more space, the diameter can be reduced. Here are some of our results.

Theorem 8 *For any $t > 1$, and any dimension k,*

1. *there is a t-spanner containing $O(n \log n)$ edges and constructible in $O(n \log n)$ time with diameter 2,*

2. *there is a t-spanner containing $O(n \log\log n)$ edges and constructible in $O(n \log n)$ time with diameter 3,*

3. *there is a t-spanner containing $O(n \log^* n)$ edges and constructible in $O(n \log n)$ time with diameter 4.*

7 Spanners of bounded degree and small diameter

Theorem 9 *For any $t > 1$, and any dimension k, in $O(n \log n)$ time, a t-spanner whose degree is bounded by a constant and whose diameter is at most $O(\log n)$ can be constructed.*

Our low-diameter constructions of the previous section have high degree. On the other hand, it is difficult to bound the diameter of our bounded-degree constructions. Note, however, that our diameter results are closely related to the one-dimensional results.

Consider the following strategy to produce a spanner of $O(\log n)$ diameter and bounded degree for a set of n points on a horizontal line. Without loss of generality, assume that n is a power of 2 and that they are numbered 0 through $n-1$ from left to right.

Include an edge $(i, i+1)$, $0 \leq i < n$. The resulting graph is a spanner, but its diameter is $n-1$. Select the set of even-numbered points, $0, 2, 4, \ldots$, and connect them by a set of edges, $(2i, 2i+2)$, $0 \leq i < n/2$. Repeat this process. The resulting set of edges has $\log n$ diameter, but several of the points have degree $\log n$ as well.

In order to reduce the degree, note that $O(\log n)$ diameter would have been preserved if the odd-numbered points were chosen at the second "level," or if either $2i$ or $2i+1$ was "promoted" to the second level. A similar statement can be made at the ℓ-th level ($2^\ell i$ through $2^\ell(i+1)-1$ can be promoted). If one is careful about alternating "promotions," the resulting structure, reminiscent of a bounded degree skip-list, has bounded degree and logarithmic diameter.

In order to generalize this proof to all dimensions, we need to apply the same strategy to trees, specifically the modified dumbbell tree of Section 6. Here, the appropriate analogue to the leveling idea seems to be Frederickson's topology trees[10].

We provide a rough outline of the method and the properties we need to maintain. Suppose we have a rooted tree T whose degree is bounded by a constant. Furthermore, assume that every leaf node has a unique label and that any internal node can be labeled with the label of an arbitrary leaf node. The first step is to choose representatives for the nodes in T. To do this, we propagate leaf labels. A node chooses one of the propagated labels and propagates the other up the tree. Each label is used at most twice, once at a leaf, and once at an internal node.

We then perform a layering approach, grouping sets of nodes into a single node at the next layer. An important issue is the maintenance of a tree whose maximum degree is bounded by a constant at every level.

Given this layered tree, labels are again distributed so that no label is used more than a constant number of times. Roughly, the labeling procedure ensures that points (corresponding to the labels) have degree bounded by a constant, and the leveling process ensures that the path has link-distance $O(\log n)$. Full details will be included in the final paper.

8 Spanners of low weight and small diameter

We use the following spanner construction, due to Arya, Mount, and Smid[2]: Start with a fair-split tree, and designate some nodes as heavy and some as light. A node is heavy if it contains more points in its subtree than its sibling, and it is light otherwise (if both subtrees contain an equal number of points, the left child is heavy and the right child is light). We use this designation to determine box representatives for the well-separated pairs; specifically, a parent box inherits the representative of its heavy child. Arya et al.[2] show that if the representatives are chosen in this way, then the resulting spanner has diameter $O(\log n)$.

We now show that the weight of well-separated pair constructions is $O(w(MST)\log n)$, which is tight[11]. This improves the results of Lenhof et al.[11], who prove that the sum D of the diameters of the boxes in a box split tree is $O(w(MST)\log^2 n)$ and that the length of the well-separated pair edges is $O(D)$. Our techniques can be used to show that $D = O(w(MST)\log n)$.

Rather than focus on the split tree, we focus on the dumbbell tree. Recall the gap property of Chandra et al.[6]: A set of edges E has the *gap property* if for every pair e_1 and e_2, the distance between the closest endpoints of e_1 and e_2 is at least the length of the smaller edge. Chandra et al. prove that if E has the gap property, $w(E) = O(w(MST)\log n)$.

In our case, let E be the set of well-separated edges represented by a dumbbell tree. We show that there is a set of edges $E' \subseteq E$ such that E' has the gap property and $w(E') = \Theta(w(E))$. This proves that $w(E) = O(w(MST)\log n)$.

To select E', initially let $E' = E$, and consider any pair of edges e_1 and e_2 in E'. If they violate the gap property, remove the shorter one, say e_1, and continue with $E' \setminus \{e_1\}$. Eventually, E' will have the gap property.

In order to show that $w(E') = \Theta(w(E))$, build the following directed forest: when e_1 is eliminated because of e_2, direct an edge from e_2 to e_1. Note that only the root e of a tree $t(e)$ in the forest will remain in E', so we want to show that $w(e) = \Theta(w(t(e)))$.

Consider the children of edge e' in $t(e)$. Recall the length grouping property of Theorem 1. The children of e' are of length approximately $c^i \cdot w(e')$, where $i > 0$ indicates the length group, and $0 < c \ll 1$ is a constant. From dumbbell tree properties, only a constant number ξ of edges of length ℓ may be within a distance ℓ of a fixed point, so the number of children of e' in group i is at most 2ξ, ξ for each endpoint. The parameter c can be chosen independently of ξ, so that $\xi c < 1/2$. It follows that the total weight of the children of e' is at most $\frac{2c}{1-c}\xi w(e')$. This implies that $w(t(e)) \leq \frac{w(e)}{1-\delta}$, where $\delta = \frac{2c}{1-c}\xi$, which in turn implies the main theorem in this section.

Theorem 10 *For any $t > 1$, and any dimension k, there is a t-spanner, constructible in $O(n \log n)$ time, with $O(\log n)$ diameter and weight $O(w(MST) \log n)$.*

9 Spanners of bounded degree, low weight and small diameter

It turns out that our bounded degree, $O(\log n)$ diameter spanner also possesses some interesting weight properties. Our analysis above shows that the sum of the diameters D of the boxes in an appropriate box split tree construction is $O(w(MST) \log n)$, so one layer of the construction has weight $O(w(MST) \log n)$. Since there are $O(\log n)$ layers, the weight is $O(w(MST) \log^2 n)$. We conclude with the following result:

Theorem 11 *For any $t > 1$, and any dimension k, there is a t-spanner, constructible in $O(n \log n)$ time, with bounded degree, $O(\log n)$ diameter, and weight $O(w(MST) \log^2 n)$.*

Conjecture 1 *For any $t > 1$, and any dimension k, there is a t-spanner, constructible in $O(n \log n)$ time, with bounded degree, $O(\log n)$ diameter, and weight $O(w(MST) \log n)$.*

References

[1] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. Tech. report 71/87, Tel-Aviv University, 1987.

[2] S. Arya, D. M. Mount, and M. Smid. Randomized and deterministic algorithms for geometric spanners of small diameter. In *Proc. 35th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 703–712, 1994.

[3] S. Arya and M. Smid. Efficient construction of a bounded degree spanner with low weight. In *Proc. 2nd Annu. European Sympos. Algorithms (ESA)*, volume 855 of *Lecture Notes in Computer Science*, pages 48–59, 1994.

[4] P. B. Callahan and S. R. Kosaraju. A decomposition of multi-dimensional point-sets with applications to k-nearest-neighbors and n-body potential fields. In *Proc. 24th Annu. ACM Sympos. Theory Comput.*, pages 546–556, 1992.

[5] P. B. Callahan and S. R. Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *Proc. 4th ACM-SIAM Sympos. Discrete Algorithms*, pages 291–300, 1993.

[6] B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. In *Proc. 8th Annu. ACM Sympos. Comput. Geom.*, pages 192–201, 1992.

[7] G. Das, P. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proc. 9th Annu. ACM Sympos. Comput. Geom.*, pages 53–62, 1993.

[8] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. In *Proc. 10th Annu. ACM Sympos. Comput. Geom.*, pages 132–139, 1994.

[9] G. Das, G. Narasimhan, and J. S. Salowe. A new way to weigh malnourished Euclidean graphs. In *Proc. 6th ACM-SIAM Sympos. Discrete Algorithms*, pages 215–222, 1995.

[10] G. N. Frederickson. A data structure for dynamically maintaining rooted trees. In *Proc. 4th ACM-SIAM Sympos. Discrete Algorithms*, pages 175–194, 1993.

[11] H.-P. Lenhof, J. S. Salowe, and D. E. Wrege. New methods to mix shortest-path and minimum spanning trees. Manuscript, 1994.

[12] J. Ruppert and R. Seidel. Approximating the d-dimensional complete Euclidean graph. In *Proc. 3rd Canad. Conf. Comput. Geom.*, pages 207–210, 1991.

[13] J. S. Salowe. Constructing multidimensional spanner graphs. *Internat. J. Comput. Geom. Appl.*, 1(2):99–107, 1991.

[14] P. M. Vaidya. A sparse graph almost as good as the complete graph on points in K dimensions. *Discrete Comput. Geom.*, 6:369–381, 1991.