

Geometric Spanners for Weighted Point Sets

Mohammad Ali Abam^{1*}, Mark de Berg^{2**}, Mohammad Farshi^{3***}, Joachim Gudmundsson⁴,
and Michiel Smid^{5***}

¹ Department of Computer Science, Technische Universität Dortmund, Germany. abam@tu-dortmund.de

² Department of Computer Science, TU Eindhoven, The Netherlands. m.t.d.berg@tue.nl

³ Department of Computer Science, Yazd University, Yazd, Iran. mfarshi@yazduni.ac.ir

⁴ NICTA[†], Sydney, Australia. joachim.gudmundsson@nicta.com.au

⁵ School of Computer Science, Carleton University, Ottawa, ON, K1S 5B6, Canada.

michiel@scs.carleton.ca

Abstract. Let (S, \mathbf{d}) be a finite metric space, where each element $p \in S$ has a non-negative weight $w(p)$. We study spanners for the set S with respect to the following weighted distance function:

$$\mathbf{d}_w(p, q) = \begin{cases} 0 & \text{if } p = q, \\ w(p) + \mathbf{d}(p, q) + w(q) & \text{if } p \neq q. \end{cases}$$

We present a general method for turning spanners with respect to the \mathbf{d} -metric into spanners with respect to the \mathbf{d}_w -metric. For any given $\varepsilon > 0$, we can apply our method to obtain $(5 + \varepsilon)$ -spanners with a linear number of edges for three cases: points in Euclidean space \mathbb{R}^d , points in spaces of bounded doubling dimension, and points on the boundary of a convex body in \mathbb{R}^d where \mathbf{d} is the geodesic distance function.

We also describe an alternative method that leads to $(2 + \varepsilon)$ -spanners for weighted point points in \mathbb{R}^d and for points on the boundary of a convex body in \mathbb{R}^d . The number of edges in these spanners is $O(n \log n)$. This bound on the stretch factor is nearly optimal: in any finite metric space and for any $\varepsilon > 0$, it is possible to assign weights to the elements such that any non-complete graph has stretch factor larger than $2 - \varepsilon$.

1 Introduction

Motivation. Networks play a central role in numerous applications, and the design of good networks is therefore an important topic of study. In general, a good network has certain desirable properties while not being too expensive. In many applications this means one wants a network providing short paths between its nodes, while not containing too many edges. This leads to the concept of spanners, as defined next in the geometric setting.

Let $\mathcal{G} = (S, E)$ be a geometric graph on a set S of n points in \mathbb{R}^d . That is, \mathcal{G} is an edge-weighted graph where the weight of an edge $(p, q) \in E$ is equal to $|pq|$, the Euclidean distance between p and q . The distance in \mathcal{G} between two points p and q , denoted by $\mathbf{d}_{\mathcal{G}}(p, q)$, is defined as the length of a shortest (that is, minimum-weight) path from p to q in \mathcal{G} . The graph \mathcal{G} is called a (geometric) t -spanner, for some $t \geq 1$, if for any two points $p, q \in S$ we have $\mathbf{d}_{\mathcal{G}}(p, q) \leq t \cdot |pq|$. The smallest t for which \mathcal{G} is a t -spanner is called the *stretch factor* (or *dilation*, or *spanning ratio*) of \mathcal{G} and the number of edges of \mathcal{G} is called its *size*.

Geometric spanners have been studied extensively over the past decade. It has been shown that for any set of n points in \mathbb{R}^d and any $\varepsilon > 0$, there is a $(1 + \varepsilon)$ -spanner with only $O(n/\varepsilon^{d-1})$ edges—see the recent book by Narasimhan and Smid [17] for this and many other results on

* MAA was supported by Technische Universität and the MADALGO Center for Massive Data Algorithms, a Center of the Danish National Research Foundation.

** MdB was supported by the Netherlands' Organisation for Scientific Research (NWO) under project no. 639.023.301.

*** MF and MS were supported by Natural Sciences and Engineering Research Council (NSERC) of Canada.

† NICTA is funded by the Australian Government as represented by the Department of Broadband, Communications and the Digital Economy and the Australian Research Council through the ICT Centre of Excellence program

spanners. Instead of considering points in Euclidean space, one can also consider points in some other metric space. Recently the notion of doubling dimension [5, 13, 16] has received considerable attention [2, 15]. The doubling dimension of a metric space (S, \mathbf{d}) is defined as follows. If p is a point of S and $R > 0$ is a real number, then the \mathbf{d} -ball with center p and radius R is the set $\{q \in S : \mathbf{d}(p, q) \leq R\}$. The *doubling dimension* of (S, \mathbf{d}) is the smallest real number d such that the following is true: For every real number $R > 0$, every \mathbf{d} -ball of radius R can be covered by at most 2^d \mathbf{d} -balls of radius $R/2$.

As it turns out, results similar to the Euclidean setting are possible when the *doubling dimension* of the metric space is bounded by a constant d : in this case there is a $(1 + \varepsilon)$ -spanner with $n/\varepsilon^{O(d)}$ edges [15, 19].

Sometimes the cost of traversing a path in a network is not only determined by the lengths of the edges on the path, but also by delays occurring at the nodes on the path: in a (large-scale) road network a node may represent a town and entering and leaving the town will take time, in a computer network a node may need some time to forward a packet to the next node on the path, and so on. The goal of our paper is to study the concept of spanners in this setting.

Problem statement. Let S be a set of n elements—we will refer to the elements as *points* from now on—and let \mathbf{d} be a metric on S . Assume each point $p \in S$ has a non-negative weight, denoted by $w(p)$. We now define a new distance function on S , denoted by \mathbf{d}_w , as follows.

$$\mathbf{d}_w(p, q) = \begin{cases} 0 & \text{if } p = q, \\ w(p) + \mathbf{d}(p, q) + w(q) & \text{if } p \neq q. \end{cases}$$

For a graph $\mathcal{G} = (S, E)$ and two points p and q in S , we denote by $\mathbf{d}_{\mathcal{G}, w}(p, q)$ the length of a shortest path in \mathcal{G} between p and q , where edge lengths are measured using the distance function \mathbf{d}_w ; if $p = q$, then we define $\mathbf{d}_{\mathcal{G}, w}(p, q) = 0$. For a real number $t > 1$, we say that \mathcal{G} is a t -spanner for the weighted point set S , if for any two points p and q in S we have $\mathbf{d}_{\mathcal{G}, w}(p, q) \leq t \cdot \mathbf{d}_w(p, q)$. We want to compute a t -spanner for the weighted point set S having few edges and with a small stretch factor. Unfortunately our metric space (S, \mathbf{d}_w) does not necessarily have bounded doubling dimension, even if the underlying metric space (S, \mathbf{d}) has bounded doubling dimension. A simple example is a set S of n points inside a unit disk in the plane, each having unit weight, and when \mathbf{d} is the Euclidean distance function. Then the doubling dimension of the metric space (S, \mathbf{d}_w) will be $\Theta(\log n)$. This leads us to the main question we want to answer: Is it possible to obtain spanners for weighted point sets with constant stretch factor—that is, stretch factor independent of n , but also independent of the weights of the points—and a near-linear number of edges?

Recently Bose *et al.* [7] also studied spanners for weighted points. More precisely, they consider points in the plane with positive weights and then define the distance between two points p, q as $|pq| - w(p) - w(q)$. The difference between their setting and our setting is thus that they subtract the weights from the Euclidean distance, whereas we add the weights (which in the applications mentioned above is more natural). This is, in fact, a fundamental difference: Bose *et al.* show (under the assumption that the distance between any pair of points is non-negative) that in their setting there exists a $(1 + \varepsilon)$ -spanner with $O(n/\varepsilon)$ edges, while our lower bounds (see below) imply that such a result is impossible in our setting.

Our results. We present two methods for computing spanners for weighted point set. The first method is described in Section 2. It essentially shows that whenever there is a good spanner for the metric space (S, \mathbf{d}) , there is also a good spanner for the metric space (S, \mathbf{d}_w) . This is done by clustering the points in a suitable way, computing a spanner in the \mathbf{d} -metric on the cluster centers, and then connecting each point to its cluster center. We apply our method to obtain, for any $0 < \varepsilon < 1$, $(5 + \varepsilon)$ -spanners for weighted point sets in \mathbb{R}^d and in spaces of doubling dimension d , with $O(n/\varepsilon^d)$ and $n/\varepsilon^{O(d)}$ edges, respectively.

We also apply our method to points on the boundary of a convex body in \mathbb{R}^d , where distances are geodesic distances along the body's boundary. We give a simple and efficient algorithm for computing a well-separated pair decomposition for this metric—we believe this result is interesting in its own right—which proves the existence of a $(1 + \varepsilon)$ -spanner with $O(n/\varepsilon^d)$ edges. When the

points are weighted, we can then use our general method to get a $(5 + \varepsilon)$ -spanner for weighted point sets with $O(n/\varepsilon^d)$ edges.

Our second method is described in Section 3. It applies to spaces of bounded doubling dimension for which a semi-separated pair decomposition [1, 20] can be constructed. It leads to spanners with a better stretch factor than our first method, but the size of the spanner is larger. In particular, it leads to $(2 + \varepsilon)$ -spanners with $(n/\varepsilon^{O(d)}) \log n$ edges, for points in \mathbb{R}^d and for points on the boundary of a convex body in \mathbb{R}^d . We also show that the bound on the stretch factor is nearly optimal: in any finite metric space and for any $\varepsilon > 0$, it is possible to assign weights to the points such that any non-complete graph has stretch factor larger than $2 - \varepsilon$.

2 A Spanner Construction Based on Clustering

Let (S, \mathbf{d}) be a finite metric space and let n denote the number of points in S . We assume that each point $p \in S$ has a real weight $w(p) \geq 0$. We will show that if we can find a good spanner for S in the \mathbf{d} -metric, we can also find a good spanner for S in the \mathbf{d}_w -metric.

The main idea is to partition S into clusters, where each cluster has a designated point as its cluster center. The clusters have the following two properties: First, the \mathbf{d} -distances and \mathbf{d}_w -distances between any two centers are approximately equal. Second, for each point p in the cluster with center c , the distance $\mathbf{d}(p, c)$ is $O(w(p))$.

In the second step, a t -spanner, based on the distance function \mathbf{d} , is constructed for the set of cluster centers. Each point p which is not a cluster center is connected by an edge to the center of the cluster that p belongs to. As we will see, the resulting graph is a $O(t)$ -spanner for the weighted point set S .

In this section, we will show that, for some specific metric spaces, we obtain a $(5 + \varepsilon)$ -spanner for weighted point sets having $n/\varepsilon^{O(d)}$ edges.

Clusterings for spanners of weighted point sets. We start by stating more precisely the properties we require from our clustering. Let k_1 and k_2 be two parameters, with $k_1 > 0$ and $k_2 \geq 1$. Define a (k_1, k_2) -clustering of S to be a partitioning of S into a collection $\{C_1, \dots, C_m\}$ of clusters, each with a center denoted by $\text{center}(C_i)$, such that the following three conditions hold:

- (I) for every $1 \leq i \leq m$ and for all $p \in C_i$ we have: $w(\text{center}(C_i)) \leq w(p)$;
- (II) for every $1 \leq i \leq m$ and for all $p \in C_i$ we have: $\mathbf{d}(\text{center}(C_i), p) \leq k_1 \cdot w(p)$;
- (III) for every $1 \leq i, j \leq m$ we have:

$$\mathbf{d}_w(\text{center}(C_i), \text{center}(C_j)) \leq k_2 \cdot \mathbf{d}(\text{center}(C_i), \text{center}(C_j)).$$

Later we will show how to find such a clustering. But first we show how to use such a clustering to obtain a spanner for S in the \mathbf{d}_w -metric.

Let $\{C_1, C_2, \dots, C_m\}$ be a (k_1, k_2) -clustering of S , and let $c_i = \text{center}(C_i)$. Let $\mathcal{C} = \{c_1, c_2, \dots, c_m\}$ denote the set of cluster centers, and let $\mathcal{G}_1 = (\mathcal{C}, E_1)$ be a t -spanner of the set \mathcal{C} in the \mathbf{d} -metric. Finally, let $E_2 = \{(c_i, p) : 1 \leq i \leq m \text{ and } p \in C_i \text{ and } p \neq c_i\}$. In other words, E_2 contains the edges connecting the points in each cluster to the center of that cluster. The next lemma states that augmenting \mathcal{G}_1 with the edges in E_2 gives a spanner in the \mathbf{d}_w -metric.

Lemma 1. *The graph $\mathcal{G} = (S, E_1 \cup E_2)$ is a t' -spanner in the \mathbf{d}_w -metric, where $t' = \max(2 + k_1 + k_1 k_2 t, k_2 t)$.*

Proof. Let p, q be two distinct points in S . We must show that $\mathbf{d}_{\mathcal{G}, w}(p, q) \leq t' \cdot \mathbf{d}_w(p, q)$. Let C_i and C_j be the clusters containing p and q , respectively, and consider $c_i = \text{center}(C_i)$ and $c_j = \text{center}(C_j)$. (It can happen that $i = j$, but this will not invalidate the coming argument.) Note that either $p = c_i$ or (p, c_i) is an edge in \mathcal{G} ; similarly $q = c_j$ or (q, c_j) is an edge in \mathcal{G} . Hence,

$$\begin{aligned} \mathbf{d}_{\mathcal{G}, w}(p, q) &= \mathbf{d}_w(p, c_i) + \mathbf{d}_{\mathcal{G}, w}(c_i, c_j) + \mathbf{d}_w(c_j, q) \\ &= (w(p) + \mathbf{d}(p, c_i) + w(c_i)) + \mathbf{d}_{\mathcal{G}, w}(c_i, c_j) + (w(c_j) + \mathbf{d}(c_j, q) + w(q)) \\ &\leq (2 + k_1) \cdot w(p) + \mathbf{d}_{\mathcal{G}, w}(c_i, c_j) + (2 + k_1) \cdot w(q), \end{aligned}$$

where the last inequality follows from properties (I) and (II) of the clustering. Now consider the shortest path in \mathcal{G}_1 from c_i to c_j in the \mathbf{d} -metric. By property (III) the length of every link on this path—and, hence, its total length—increases by at most a factor k_2 when we measure its length in the \mathbf{d}_w -metric. Since \mathcal{G}_1 is a t -spanner for \mathcal{C} in the \mathbf{d} -metric, we thus have $\mathbf{d}_{\mathcal{G},w}(c_i, c_j) \leq k_2 \cdot \mathbf{d}_{\mathcal{G}_1}(c_i, c_j) \leq k_2 t \cdot \mathbf{d}(c_i, c_j)$. Finally, we observe that

$$\mathbf{d}(c_i, c_j) \leq \mathbf{d}(c_i, p) + \mathbf{d}(p, q) + \mathbf{d}(q, c_j) \leq k_1 \cdot w(p) + \mathbf{d}(p, q) + k_1 \cdot w(q).$$

Combing this with our two earlier derivations, we get

$$\begin{aligned} \mathbf{d}_{\mathcal{G},w}(p, q) &\leq (2 + k_1) \cdot w(p) + \mathbf{d}_{\mathcal{G},w}(c_i, c_j) + (2 + k_1) \cdot w(q) \\ &\leq (2 + k_1) \cdot w(p) + k_2 t \cdot \mathbf{d}(c_i, c_j) + (2 + k_1) \cdot w(q) \\ &\leq (2 + k_1) \cdot w(p) + k_2 t \cdot (k_1 \cdot w(p) + \mathbf{d}(p, q) + k_1 \cdot w(q)) + (2 + k_1) \cdot w(q) \\ &= (2 + k_1 + k_1 k_2 t) \cdot w(p) + k_2 t \cdot \mathbf{d}(p, q) + (2 + k_1 + k_1 k_2 t) \cdot w(q) \\ &\leq \max(2 + k_1 + k_1 k_2 t, k_2 t) \cdot (w(p) + \mathbf{d}(p, q) + w(q)) \\ &= \max(2 + k_1 + k_1 k_2 t, k_2 t) \cdot \mathbf{d}_w(p, q). \end{aligned}$$

□

Computing good clusterings and spanners. The following algorithm takes as input the weighted set S and two real numbers k and $\varepsilon > 0$, and computes a clustering $\{C_1, \dots, C_m\}$ of S .

1. Sort the points of S in non-decreasing order of their weight, and let p_1, p_2, \dots, p_n be the sorted sequence (ties are broken arbitrarily).
2. Initialize the first cluster C_1 : set $C_1 = \{p_1\}$ and $c_1 = \text{center}(C_1) = p_1$. Initialize the set of cluster centers: $\mathcal{C} = \{p_1\}$. Set $m = 1$.
3. For $i = 2$ to n , do the following:
 - (a) Compute an index j with $1 \leq j \leq m$ such that c_j is a $(1 + \varepsilon)$ -approximate nearest-neighbor of p_i in the set \mathcal{C} , in the \mathbf{d} -metric. (Thus, $\mathbf{d}(p_i, c_j)$ is at most $1 + \varepsilon$ times the \mathbf{d} -distance between p_i and the point in \mathcal{C} that is closest to p_i .) Observe that we use approximate nearest neighbor instead of the exact nearest neighbor to improve the efficiency of the algorithm, see Theorem 1.
 - (b) If $\mathbf{d}(c_j, p_i) \leq k \cdot w(p_i)$, then set $C_j = C_j \cup \{p_i\}$. Otherwise, start a new cluster: set $m = m + 1$, set $C_m = \{p_i\}$ and $c_m = \text{center}(C_m) = p_i$, and set $\mathcal{C} = \mathcal{C} \cup C_m$.
4. Return the collection $\{C_1, \dots, C_m\}$ of clusters.

Lemma 2. *The algorithm above computes a $(k, 1 + \frac{2(1+\varepsilon)}{k})$ -clustering of S .*

Proof. Since we treat the points in order of increasing weight and the first point put into a cluster is its center, we have $w(c_j) \leq w(p)$ for every cluster C_j and point $p \in C_j$. Moreover, by step 3 we only put a point p in a cluster C_j if $\mathbf{d}(\text{center}(C_j), p) \leq k \cdot w(p)$. Hence, conditions (I) and (II) are satisfied.

To prove condition (III), consider two distinct cluster centers c and c' . Assume without loss of generality that c was added to \mathcal{C} before c' . Then it follows from the algorithm that $w(c) \leq w(c')$. Consider the iteration of the for-loop in which $p_i = c'$, and consider the set \mathcal{C} at the beginning of this iteration. Observe that $c \in \mathcal{C}$. Let c_j be the $(1 + \varepsilon)$ -approximate nearest-neighbor of c' in \mathcal{C} that is computed by the algorithm. Since c' is added to \mathcal{C} , we have $\mathbf{d}(c_j, c') > k \cdot w(c')$. Let c^* be the exact nearest-neighbor of c' in \mathcal{C} . Then, since $c \in \mathcal{C}$, $\mathbf{d}(c_j, c') \leq (1 + \varepsilon) \cdot \mathbf{d}(c^*, c') \leq (1 + \varepsilon) \cdot \mathbf{d}(c, c')$. It follows that

$$\begin{aligned} \mathbf{d}_w(c, c') &= w(c) + \mathbf{d}(c, c') + w(c') \leq \mathbf{d}(c, c') + 2 \cdot w(c') \\ &< \mathbf{d}(c, c') + \frac{2}{k} \cdot \mathbf{d}(c_j, c') \leq \left(1 + \frac{2(1 + \varepsilon)}{k}\right) \cdot \mathbf{d}(c, c'). \end{aligned}$$

□

By combining Lemmas 1 and 2, we obtain the following result.

Theorem 1. *Let $t > 1$ be a parameter, and let (S, \mathbf{d}) be a metric space with n weighted points such that the following holds:*

- *For any subset $S' \subseteq S$ with m points, we can compute in $T_{\text{sp}}(m)$ time a t -spanner for S' in the \mathbf{d} -metric with $E_{\text{sp}}(m)$ edges, where T_{sp} and E_{sp} are non-decreasing functions.*
- *For any $\varepsilon > 0$ there is a semi-dynamic (insertions-only) data structure for $(1 + \varepsilon)$ -approximate nearest-neighbor queries in the \mathbf{d} -metric for S , such that both insertions and queries can be done in $T_{\text{nn}}(\varepsilon, n)$ time, where the function T_{nn} is non-decreasing in n .*

Then we can construct for any $\varepsilon > 0$ a t' -spanner for S in the \mathbf{d}_w -metric with $O(E_{\text{sp}}(n))$ edges and $t' = 3t + 2 + 2\varepsilon(t + 1)$. The construction can be done in $O(n \log n + T_{\text{sp}}(n) + n \cdot T_{\text{nn}}(\varepsilon, n))$ time.

Proof. We run the clustering algorithm described above with $k = t/(t + 1)$. By Lemma 2 it produces a (k_1, k_2) -clustering with $k_1 = t/(t + 1)$ and $k_2 = 1 + 2(1 + \varepsilon)(t + 1)/t$. Using the data structure for approximate nearest-neighbor queries, the clustering algorithm can be implemented to run in $O(n \log n + n \cdot T_{\text{nn}}(\varepsilon, n))$ time.

Next we compute a t -spanner in the \mathbf{d} -metric on the set \mathcal{C} of cluster centers of the computed clustering. If $|\mathcal{C}| = m$, this takes $T_{\text{sp}}(m) \leq T_{\text{sp}}(n)$ time and produces a spanner \mathcal{G}_1 with $E_{\text{sp}}(m) \leq E_{\text{sp}}(n)$ edges. To this spanner we add the set E_2 of edges connecting each point to its cluster center. This takes $O(n)$ time, and adds another $n - m$ edges. Hence, the total number of edges is at most $E_{\text{sp}}(n) + n - m$, which is $O(E_{\text{sp}}(n))$, since $E_{\text{sp}}(n) \geq n - 1$. This proves the time bound and the bound on the number of edges. The bound on the stretch factor follows by plugging $k_1 = t/(t + 1)$ and $k_2 = 1 + 2(1 + \varepsilon)(t + 1)/t$ in Lemma 1. \square

Applications: Euclidean spaces and spaces of bounded doubling dimension. Theorem 1 can immediately be used to obtain spanners for weighted point sets in Euclidean spaces and metric spaces of bounded doubling dimension.

Corollary 1. (i) *Given a set S of n points in \mathbb{R}^d , each having a non-negative weight, and given a real number $0 < \varepsilon < 1$, we can construct a $(5 + \varepsilon)$ -spanner of the weighted point set S having $O(n/\varepsilon^d)$ edges in $O((n/\varepsilon^d) \log n)$ time.*

(ii) *Given a metric space (S, \mathbf{d}) of constant doubling dimension d , where S is a set of size n , and in which each point of S has a non-negative real weight, and given a real number $0 < \varepsilon < 1$, we can construct a $(5 + \varepsilon)$ -spanner of the weighted point set S having $n/\varepsilon^{O(d)}$ edges in $O(n \log n) + n/\varepsilon^{O(d)}$ time.*

Proof. Callahan and Kosaraju [9] have shown that for any set of n points in \mathbb{R}^d and any $0 < \varepsilon < 1$, one can compute a $(1 + \varepsilon)$ -spanner with $E_{\text{sp}}(n) = O(n/\varepsilon^d)$ edges in $T_{\text{sp}}(n) = O(n \log n + n/\varepsilon^d)$ time. Moreover, Arya *et al.* [4] presented a data structure for $(1 + \varepsilon)$ -approximate nearest-neighbor queries in \mathbb{R}^d that has $O((1/\varepsilon^d) \log n)$ query time, and in which insertions can be done in $O(\log n)$ time. Part (i) of the theorem now follows by applying Theorem 1, replacing ε by $\varepsilon/10$ and setting $t = 1 + \varepsilon/10$.

Roditty [18] (see also Gottlieb and Roditty [12]) have shown that for any metric space (S, \mathbf{d}) with n points and doubling dimension d and any $0 < \varepsilon < 1$, one can compute a $(1 + \varepsilon)$ -spanner with $E_{\text{sp}}(n) = n/\varepsilon^{O(d)}$ edges in $T_{\text{sp}}(n) = O(n \log n) + n/\varepsilon^{O(d)}$ time. Moreover, Cole and Gottlieb [10] presented a data structure for $(1 + \varepsilon)$ -approximate nearest-neighbor queries in (S, \mathbf{d}) that has $2^{O(d)} \log n + 1/\varepsilon^{O(d)}$ query time, and in which insertions can be done in $2^{O(d)} \log n$ time. Part (ii) now follows by applying Theorem 1, replacing ε by $\varepsilon/10$ and setting $t = 1 + \varepsilon/10$. \square

More applications: the geodesic metric for a convex body. Let S be a set of n points on the boundary $\partial\mathcal{B}$ of a convex body \mathcal{B} in \mathbb{R}^d . For any two points $p, q \in S$, let $\mathbf{d}_{\mathcal{B}}(p, q)$ be the geodesic distance between p and q along $\partial\mathcal{B}$, and let $\mathbf{d}(p, q)$ denote their Euclidean distance. In order to apply Theorem 1 to the metric space $(S, \mathbf{d}_{\mathcal{B}})$, we need a sparse $(1 + \varepsilon)$ -spanner for a set $S' \subseteq S$ based on the distance function $\mathbf{d}_{\mathcal{B}}$. We will obtain such a spanner using a so-called well-separated pair decomposition (WSPD).

Well-separated pair decompositions were introduced by Callahan and Kosaraju [9] for the Euclidean metric and by Talwar [19] for general metric spaces. They are defined as follows. Let (S, \mathbf{d}) be a finite metric space. The *diameter* $\text{diam}_{\mathbf{d}}(A)$ of any subset A of S is defined as $\text{diam}_{\mathbf{d}}(A) = \max\{\mathbf{d}(a, b) : a, b \in A\}$, and the *distance* $\mathbf{d}(A, B)$ of any two subsets $A, B \subseteq S$ is defined as $\mathbf{d}(A, B) = \min\{\mathbf{d}(a, b) : a \in A, b \in B\}$. For a real number $s > 0$, we say that the subsets A and B of S are *well-separated with respect to s* , if $\mathbf{d}(A, B) \geq s \cdot \max(\text{diam}_{\mathbf{d}}(A), \text{diam}_{\mathbf{d}}(B))$.

Definition 1. *Let (S, \mathbf{d}) be a finite metric space and let $s > 0$ be a real number. A well-separated pair decomposition (WSPD) for (S, \mathbf{d}) , with respect to s , is a set $\{(A_1, B_1), \dots, (A_m, B_m)\}$ of pairs of non-empty subsets of S such that*

1. *for each i , A_i and B_i are well-separated with respect to s , and*
2. *for any two distinct points $p, q \in S$, there is exactly one index i with $1 \leq i \leq m$ such that (i) $p \in A_i$ and $q \in B_i$ or (ii) $p \in B_i$ and $q \in A_i$.*

The following lemma, due to Callahan and Kosaraju [8], shows how a spanner can be obtained from a WSPD.

Lemma 3. [8] *Let (S, \mathbf{d}) be a finite metric space and let $t > 1$ be a real number. Furthermore, let $\{(A_1, B_1), \dots, (A_m, B_m)\}$ be a WSPD for (S, \mathbf{d}) , with respect to $s = \frac{2(t+1)}{t-1}$ and, for each i with $1 \leq i \leq m$, let a_i be an arbitrary point of A_i and b_i be an arbitrary point of B_i . Then the graph $G = (S, E)$ where $E = \{(a_i, b_i) : 1 \leq i \leq m\}$ is a t -spanner for S with m edges.*

Lemma 3 tells us that if we have a WSPD for S in the $\mathbf{d}_{\mathcal{B}}$ -metric, we can get a spanner for S in the $\mathbf{d}_{\mathcal{B}}$ -metric. Using Theorem 1 we can then also get a spanner for the weighted point set S . As we show in Lemma 10, the metric space $(S, \mathbf{d}_{\mathcal{B}})$ has bounded doubling dimension. Using the algorithm of Har-Peled and Mendel [15] we can thus construct a WSPD for this metric space. Unfortunately, their algorithm needs an oracle that returns, for any two points p and q , the geodesic distance $\mathbf{d}_{\mathcal{B}}(p, q)$ in $O(1)$ time, and computing geodesic distances on a convex body is not so easy. We therefore describe a more direct method for computing a WSPD for points on a convex body. The basic idea behind our method is to compute a WSPD for the Euclidean space (S, \mathbf{d}) , and then refine this WSPD in a suitable way to obtain a WSPD for $(S, \mathbf{d}_{\mathcal{B}})$. For the refinement, we only need to know the normal vectors of all points $p \in S$; we do not need any distance computations in the $\mathbf{d}_{\mathcal{B}}$ -metric. An additional advantage of our method over Har-Peled and Mendel's method is that the dependency on ε will be better.

For any point p on $\partial\mathcal{B}$, we denote by $\mathcal{N}_{\mathcal{B}}(p)$ the (outer) normal vector of \mathcal{B} at p . If the tangent plane of p at \mathcal{B} is not unique, then we choose for $\mathcal{N}_{\mathcal{B}}(p)$ the normal vector of an arbitrary tangent plane. We fix a real number σ such that $0 < \sigma < \pi/2$. The following lemma states that $\mathbf{d}_{\mathcal{B}}(p, q)$ and $\mathbf{d}(p, q)$ are approximately equal, provided the angle between the normals of p and q is at most σ . Similar observations have been made in papers on approximate shortest paths on polytopes; see e.g. [3].

Lemma 4. *Let p and q be two points on $\partial\mathcal{B}$ such that $\angle(\mathcal{N}_{\mathcal{B}}(p), \mathcal{N}_{\mathcal{B}}(q)) \leq \sigma$. Then $\mathbf{d}(p, q) \leq \mathbf{d}_{\mathcal{B}}(p, q) \leq \frac{\mathbf{d}(p, q)}{\cos \sigma}$.*

The normal vector of each point of $\partial\mathcal{B}$ at \mathcal{B} can be considered to be a point on the sphere of directions, denoted \mathbb{S}^{d-1} , in \mathbb{R}^d . We partition \mathbb{S}^{d-1} into $O(1/\sigma^{d-1})$ parts such that the angle between any two vectors in the same part is at most σ . Based on this, we partition $\partial\mathcal{B}$ into *patches*: A σ -patch is the set of all points of $\partial\mathcal{B}$ whose normals fall in the same part of the partition of \mathbb{S}^{d-1} .

Let $s > 0$ be a real number, and let $\{(A_1, B_1), \dots, (A_m, B_m)\}$ be a WSPD for the Euclidean metric space (S, \mathbf{d}) , with respect to s , where $m = O(s^d n)$. We refine the WSPD by partitioning

each A_i and B_i into subsets A_i^1, \dots, A_i^ℓ and $B_i^1, \dots, B_i^{\ell'}$, respectively, where $\ell = O(1/\sigma^{d-1})$ and $\ell' = O(1/\sigma^{d-1})$. The partitioning is done such that the points in each subset belong to the same σ -patch. Define $\Psi = \{(A_i^j, B_i^k) : 1 \leq i \leq m, 1 \leq j \leq \ell \text{ and } 1 \leq k \leq \ell'\}$.

Lemma 5. *The set of pairs in Ψ forms a WSPD with respect to $s \cos \sigma$ for the metric space $(S, \mathbf{d}_{\mathcal{B}})$. The number of pairs in this WSPD is $O((s^d/\sigma^{2d-2})n)$.*

Proof. It is clear that Ψ contains $O((s^d/\sigma^{2d-2})n)$ elements. It is also clear that condition 2. in Definition 1 is satisfied. It remains to show that condition 1. is satisfied. Consider a pair $(A_i^j, B_i^k) \in \Psi$. We have to show that

$$\mathbf{d}_{\mathcal{B}}(A_i^j, B_i^k) \geq s \cos \sigma \cdot \max(\text{diam}_{\mathbf{d}_{\mathcal{B}}}(A_i^j), \text{diam}_{\mathbf{d}_{\mathcal{B}}}(B_i^k)). \quad (1)$$

We first show that

$$\text{diam}_{\mathbf{d}}(A_i^j) \geq \text{diam}_{\mathbf{d}_{\mathcal{B}}}(A_i^j) \cos \sigma. \quad (2)$$

To show this, let a and a' be two arbitrary points in A_i^j . Using Lemma 4, we obtain $\mathbf{d}_{\mathcal{B}}(a, a') \leq \frac{\mathbf{d}(a, a')}{\cos \sigma} \leq \frac{\text{diam}_{\mathbf{d}}(A_i^j)}{\cos \sigma}$, from which (2) follows. By a symmetric argument, we obtain

$$\text{diam}_{\mathbf{d}}(B_i^k) \geq \text{diam}_{\mathbf{d}_{\mathcal{B}}}(B_i^k) \cos \sigma. \quad (3)$$

Let a be an arbitrary point of A_i^j and let b be an arbitrary point of B_i^k . Since $A_i^j \subseteq A_i$ and $B_i^k \subseteq B_i$, and since A_i and B_i are well-separated with respect to s (in the Euclidean metric \mathbf{d}), we have $\mathbf{d}_{\mathcal{B}}(a, b) \geq \mathbf{d}(a, b) \geq s \cdot \max(\text{diam}_{\mathbf{d}}(A_i), \text{diam}_{\mathbf{d}}(B_i)) \geq s \cdot \max(\text{diam}_{\mathbf{d}}(A_i^j), \text{diam}_{\mathbf{d}}(B_i^k))$. Combining this with (2) and (3), it follows that $\mathbf{d}_{\mathcal{B}}(a, b) \geq s \cos \sigma \cdot \max(\text{diam}_{\mathbf{d}_{\mathcal{B}}}(A_i^j), \text{diam}_{\mathbf{d}_{\mathcal{B}}}(B_i^k))$. This proves that (1) holds. \square

Lemmas 3 and 5 now imply the following result (take for instance $\sigma = \pi/3$, so that $\cos \sigma = 1/2$).

Theorem 2. *Let S be a set of n points on the boundary of a convex body \mathcal{B} in \mathbb{R}^d , and let $0 < \varepsilon < 1$ be a real number. If we can determine for any $p \in S$ an outward normal of \mathcal{B} at p in $O(1)$ time then we can compute in $O(n \log n + n/\varepsilon^d)$ time a $(1 + \varepsilon)$ -spanner of S in the $\mathbf{d}_{\mathcal{B}}$ -metric, with $O(n/\varepsilon^d)$ edges.*

Corollary 2. *Let S be a set of n points on the boundary of a convex body \mathcal{B} in \mathbb{R}^d , each with a non-negative weight. For any $0 < \varepsilon < 1$, there is a $(5 + \varepsilon)$ -spanner of the weighted point set S having $O(n/\varepsilon^d)$ edges.*

3 A $(2 + \varepsilon)$ -Spanner for a Weighted Point Set

In each of the applications considered in the previous section, our method generated a $(5 + \varepsilon)$ -spanner for weighted point sets. The goal of this section is to see if we can obtain spanners for weighted point sets with a smaller stretch factor. We start with a lower bound.

Lemma 6. *For any finite metric space (S, \mathbf{d}) and any real number $\varepsilon > 0$, there exists a set of weights for the points of S , such that every non-complete graph with vertex set S has stretch factor larger than $2 - \varepsilon$.*

Proof. Let $D = \text{diam}_{\mathbf{d}}(S)$. Assign each point in S a weight D/ε . Consider a non-complete graph G with vertex set S , and let p and q be two points in S that are not connected by an edge in G . We have $\mathbf{d}_w(p, q) \leq (1 + 2/\varepsilon)D$, whereas $\mathbf{d}_{G,w}(p, q) \geq 4D/\varepsilon$. Thus $\frac{\mathbf{d}_{G,w}(p, q)}{\mathbf{d}_w(p, q)} \geq \frac{4D/\varepsilon}{(1+2/\varepsilon)D} > 2 - \varepsilon$. \square

In the remainder of this section we will describe a general strategy for computing a $(2 + \varepsilon)$ -spanners for weighted point sets in spaces of bounded doubling dimension that admit a so-called semi-separated pair decomposition [20] of small size. Given the lower bound, the stretch factor is almost optimal in the worst case. We use the strategy to obtain $(2 + \varepsilon)$ -spanners for two cases of weighted point sets: points in \mathbb{R}^d , and points on the boundary of a convex body in \mathbb{R}^d .

The semi-separated pair decomposition. Let (S, \mathbf{d}) be a metric space, where S is a set of n points, and let d be its doubling dimension. We assume that each point of S has a real weight $w(p) \geq 0$. Our spanner construction will be based on a decomposition $\{(A_1, B_1), \dots, (A_m, B_m)\}$ having properties similar to those of the WSPD. As we will see, the number of edges in the spanner for weighted point sets is proportional to $\sum_{i=1}^m (|A_i| + |B_i|)$. Thus, we need a decomposition for which this summation is small. Callahan and Kosaraju [9] have shown that, for the WSPD, this summation can be as large as $\Theta(n^2)$; in other words, we cannot use the WSPD to obtain a non-trivial result. By using a decomposition satisfying a weaker condition, it is possible to make sure the summation is only $O(n \log n)$. This decomposition is the semi-separated pair decomposition, as introduced by Varadarajan [20].

For a real number $s > 0$, two subsets $A, B \subseteq S$ are called *semi-separated with respect to s* , if $\mathbf{d}(A, B) \geq s \cdot \min(\text{diam}_{\mathbf{d}}(A), \text{diam}_{\mathbf{d}}(B))$. A *semi-separated pair decomposition (SSPD)* for the metric space (S, \mathbf{d}) , with respect to s , is defined to be a set $\Psi = \{(A_1, B_1), \dots, (A_m, B_m)\}$ of pairs of non-empty subsets of S , having the same properties as in Definition 1, except that in condition 1, the sets A_i and B_i are semi-separated with respect to s . The quantity $\sum_{i=1}^m (|A_i| + |B_i|)$ is called the *size* of the SSPD.

For the Euclidean distance function in \mathbb{R}^2 , Abam *et al.* [1] showed that an SSPD with $O(n)$ pairs and size $O(n \log n)$ can be computed in $O(n \log n)$ time. The argument of Abam *et al.* [1] can easily be generalized to d -dimensional Euclidean space for $d > 2$; for completeness we give the formal proof in Section 4. The bound on the SSPD that we obtain is optimal: it follows from results by Hansel [14] that, for any set of n points, any SSPD has size $\Omega(n \log n)$; see also Bollobás and Scott [6].

From SSPDs to spanners. Let Ψ be an SSPD for S with respect to some $s > 0$. For each pair $(A, B) \in \Psi$ we will add a set $E(A, B)$ of edges to our spanner such that any two points $a \in A$ and $b \in B$ are connected by a path of length at most $(2 + \frac{3}{s}) \cdot \mathbf{d}_w(a, b)$.

The main idea is quite simple. Assume without loss of generality that $\text{diam}_{\mathbf{d}}(A) \leq \text{diam}_{\mathbf{d}}(B)$. Thus, we have $\mathbf{d}(A, B) \geq s \cdot \text{diam}_{\mathbf{d}}(A)$. Define $\text{center}(A)$ to be a point from A of minimum weight (among the points in A), and let $E_1(A, B) = \{(x, \text{center}(A)) : x \in A \cup B \text{ and } x \neq \text{center}(A)\}$. This provides short connections between the points in A and those in B by going via $\text{center}(A)$: since $\mathbf{d}(A, B) \geq s \cdot \text{diam}_{\mathbf{d}}(A)$, going via $\text{center}(A)$ does not create a large detour in the \mathbf{d} -metric, and since $w(\text{center}(A)) \leq w(a)$ the extra path length caused by $w(\text{center}(A))$ is also limited. In fact, for some pairs of points a, b , the set $E_1(A, B)$ already gives us a path of the required length. The next lemma gives the condition under which this is the case.

Lemma 7. *Let $a^* = \text{center}(A)$. Let $b \in B$ be a point such that $w(a^*) \leq w(b) + \mathbf{d}(a^*, b)$. Then, for any $a \in A$, we have $\mathbf{d}_w(a, a^*) + \mathbf{d}_w(a^*, b) \leq (2 + \frac{3}{s}) \cdot \mathbf{d}_w(a, b)$.*

Proof. We have

$$\begin{aligned} \mathbf{d}_w(a, a^*) + \mathbf{d}_w(a^*, b) &= (w(a) + \mathbf{d}(a, a^*) + w(a^*)) + (w(a^*) + \mathbf{d}(a^*, b) + w(b)) \\ &\leq 2 \cdot w(a) + \mathbf{d}(a, a^*) + 2 \cdot (\mathbf{d}(a^*, b) + w(b)) \\ &\leq 2 \cdot w(a) + \mathbf{d}(a, a^*) + 2 \cdot (\mathbf{d}(a^*, a) + \mathbf{d}(a, b) + w(b)) \\ &\leq 2 \cdot (w(a) + \mathbf{d}(a, b) + w(b)) + 3 \cdot \text{diam}_{\mathbf{d}}(A) \\ &\leq 2 \cdot (\mathbf{d}_w(a, b)) + 3 \cdot (\mathbf{d}(a, b)/s) \\ &\leq (2 + \frac{3}{s}) \cdot \mathbf{d}_w(a, b). \quad \square \end{aligned}$$

It remains to establish short paths between the points in A and the points $b \in \overline{B}$, where $\overline{B} = \{b \in B : w(a^*) > w(b) + \mathbf{d}(a^*, b)\}$ with $a^* = \text{center}(A)$. We cannot use any point from A as an intermediate destination on such paths, because the weights of the points from A are too large compared to those in \overline{B} . Hence, we need to go via a point from \overline{B} . However, the diameter of \overline{B} can be large. Therefore we first decompose the set \overline{B} into subsets of small diameter.

The points b in \overline{B} have $\mathbf{d}(a^*, b) < w(a^*)$, so they are contained in a \mathbf{d} -ball C of radius $w(a^*)$. Recall that d is the doubling dimension of (S, \mathbf{d}) . Thus we can cover C by $s^{O(d)}$ balls of radius $w(a^*)/(2s)$. Let C_1, \dots, C_ℓ be such a collection of balls, where $\ell = s^{O(d)}$. We partition \overline{B} into subsets B_1, \dots, B_ℓ in such a way that $B_i \subseteq C_i$ for all $1 \leq i \leq \ell$. For each B_i , let $\text{center}(B_i)$ be a point of minimum weight (among the points in B_i). The next lemma shows that going from any point in A to any point in B_i via $\text{center}(B_i)$ gives us a path of the required length.

Lemma 8. *Let $b_i^* = \text{center}(B_i)$. Then, for two points $a \in A$ and $b \in B_i$ we have $\mathbf{d}_w(a, b_i^*) + \mathbf{d}_w(b_i^*, b) < (2 + \frac{2}{s}) \cdot \mathbf{d}_w(a, b)$.*

Proof. Note that $w(b_i^*) \leq w(b) < w(a^*) \leq w(a)$. Hence, we have

$$\begin{aligned} \mathbf{d}_w(a, b_i^*) + \mathbf{d}_w(b_i^*, b) &= (w(a) + \mathbf{d}(a, b_i^*) + w(b_i^*)) + (w(b_i^*) + \mathbf{d}(b_i^*, b) + w(b)) \\ &< (2 \cdot w(a) + \mathbf{d}(a, b) + \mathbf{d}(b_i^*, b)) + (\mathbf{d}(b_i^*, b) + 2 \cdot w(b)) \\ &\leq 2 \cdot w(a) + \mathbf{d}(a, b) + 2 \cdot \text{diam}_{\mathbf{d}}(B_i) + 2 \cdot w(b) \\ &\leq 2 \cdot w(a) + \mathbf{d}(a, b) + 2 \cdot (w(a^*)/s) + 2 \cdot w(b) \\ &\leq (2 + \frac{2}{s}) \cdot w(a) + \mathbf{d}(a, b) + 2 \cdot w(b) \\ &< (2 + \frac{2}{s}) \cdot \mathbf{d}_w(a, b). \quad \square \end{aligned}$$

We are now ready to define the set of edges for the pair (A, B) in the SSPD Ψ . Namely, we define $E(A, B) = E_1(A, B) \cup \left(\bigcup_{i=1}^{\ell} E_2(A, B_i) \right)$, where $E_2(A, B_i) = \{(x, \text{center}(B_i)) : x \in A \cup B_i \text{ and } x \neq \text{center}(B_i)\}$. For any two points $a \in A$ and $b \in B$, there exists a path in the graph with edge set $E(A, B)$ of length at most $(2 + \varepsilon) \cdot \mathbf{d}_w(a, b)$. This follows by using Lemmas 7 and 8, and setting $s = \varepsilon/3$. Using that $\ell = s^{O(d)} = 1/\varepsilon^{O(d)}$, we get that the total number of edges in $E(A, B)$ is $|E_1(A, B)| + \sum_{i=1}^{\ell} |E_2(A, B_i)| = |A| + |B| + \sum_{i=1}^{\ell} (|A| + |B_i|) = (1/\varepsilon)^{O(d)} \cdot (|A| + |B|)$. By combining the sets $E(A, B)$ for all pairs $(A, B) \in \Psi$ we get our final spanner. Since, by definition of the SSPD, for any two points $a, b \in S$ there is a pair $(A, B) \in \Psi$ such that $a \in A$ and $b \in B$ (or vice versa), we get the following result.

Lemma 9. *The graph $G = (S, E)$ with $E = \bigcup_{(A, B) \in \Psi} E(A, B)$ is a $(2 + \varepsilon)$ -spanner for the weighted point set S with $(1/\varepsilon)^{O(d)} \cdot \sum_{(A, B) \in \Psi} (|A| + |B|)$ edges.*

As in Section 2, the size of the $(2 + \varepsilon)$ -spanner G and the time to construct it depends on the metric space (S, \mathbf{d}) .

Applications. Let S be a set of n points in \mathbb{R}^d and let $\mathbf{d}(p, q)$ be the Euclidean distance between p and q . Observe that the metric space (S, \mathbf{d}) has doubling dimension $\Theta(d)$. Abam *et al.* [1] have shown that in the plane an SSPD of size $O(s^2 n \log n)$ can be computed in $O(n \log n + s^2 n)$ time, for any $s > 1$. In section 4, we will show that their algorithm in fact also works in higher dimensions; its analysis also goes through, with appropriate changes to the constant factors in certain packing lemmas. This leads to an SSPD of size $O(s^d n \log n)$ that can be computed in $O(n \log n + s^d n)$ time, giving the following result.

Theorem 3. *Given a set S of n points in \mathbb{R}^d , each one having a non-negative weight, and given a real number $0 < \varepsilon < 1$, we can construct a $(2 + \varepsilon)$ -spanner for the weighted point set S having $(n/\varepsilon^{O(d)}) \log n$ edges in $(n/\varepsilon^{O(d)}) \log n$ time.*

We now turn our attention to a set S of points on the boundary of a convex body \mathcal{B} . For any two points p and q of S , let $\mathbf{d}_{\mathcal{B}}(p, q)$ be the geodesic distance between p and q along $\partial \mathcal{B}$. The proof of the following lemma is based on the concept of σ -patches introduced earlier.

Lemma 10. *The metric space $(S, \mathbf{d}_{\mathcal{B}})$ has doubling dimension $\Theta(d)$.*

Let \mathbf{d} denote the Euclidean distance function in \mathbb{R}^d , let $s > 1$ be a real number, and consider an SSPD $\{(A_1, B_1) \dots, (A_m, B_m)\}$ for the metric space (S, \mathbf{d}) , with respect to s , whose size is $O(s^d n \log n)$. We fix a real number σ such that $0 < \sigma < \pi/2$. Let i be an index with $1 \leq i \leq m$. As before, we partition both A_i and B_i into subsets A_i^1, \dots, A_i^ℓ and B_i^1, \dots, B_i^ℓ , respectively, where $\ell = O(1/\sigma^{d-1})$, such that the points in each subset belong to the same σ -patch of $\partial\mathcal{B}$. Now define $\Psi = \{(A_i^j, B_i^k) : 1 \leq i \leq m, 1 \leq j \leq \ell \text{ and } 1 \leq k \leq \ell\}$. The proof of the following lemma is similar to that of Lemma 5.

Lemma 11. *The set Ψ forms an SSPD, with respect to $s \cos \sigma$, for the metric space $(S, \mathbf{d}_{\mathcal{B}})$. The size of this SSPD is $O((s^d/\sigma^{2d-2})n \log n)$.*

We choose $\sigma = \pi/3$, so that $\cos \sigma = 1/2$. We obtain the following result.

Theorem 4. *Given a convex body \mathcal{B} in \mathbb{R}^d and a set S of n points on the boundary of \mathcal{B} . Assume that each point of S has a non-negative real weight. Let $0 < \varepsilon < 1$ be a real number. We can construct a $(2 + \varepsilon)$ -spanner for the weighted point set S having $(n/\varepsilon^{O(d)}) \log n$ edges in $(n/\varepsilon^{O(d)}) \log n$ time.*

Remark 1. It follows from the proofs of Lemmas 7 and 8 that the graph G has *spanner diameter* 2. That is, for any two points p and q of S , the graph G contains a path between p and q that contains at most two edges and whose \mathbf{d}_w -length is at most $(2 + \varepsilon) \cdot \mathbf{d}_w(p, q)$. If we want to keep this property, then the number of edges in our spanner is worst-case optimal: For any real number $t > 1$, there exists a set S of n points on the real line such that every t -spanner for S having spanner diameter 2 has $\Omega(n \log n)$ edges—see Exercise 12.10 in Narasimhan and Smid [17]. Of course, this then also holds for spanners for weighted point sets. Note that if all weights are equal and very large compared to the \mathbf{d} -diameter of the set, then any 2-spanner of the weighted point set must have spanner diameter 2. (This does not imply, however, that $\Omega(n \log n)$ is a lower bound on the worst-case size of 2-spanners of weighted point sets.)

4 Computing an SSPD in \mathbb{R}^d

Abam et al. [1] have shown how to compute an SSPD of size $O(n \log n)$ in \mathbb{R}^2 . Their result can be generalized to $d > 2$ fairly easily; for completeness we describe the generalization below, repeating the various concepts and definitions where necessary.

To compute an SSPD for a given point set S , we use a bounded aspect ratio (BAR) tree, as introduced by Duncan *et al.* [11]. A BAR-tree for a point set S is a binary space partition (BSP) tree with the following properties:

1. each leaf region contains at most one point from S ,
2. the tree has size $O(n)$,
3. if we go down two levels in the tree then the size of the subtree reduces at least with a factor of $\beta = (d + 1)/(d + 2)$, so its depth is $O(\log_{1/\beta} n) = O(d \log n)$,
4. the region $\mathcal{R}(\nu)$ associated with an (internal or leaf) node ν has aspect ratio at most α for some constant $\alpha > 18d^{3/2}$, that is, there are concentric balls $D_I \subset \mathcal{R}(\nu)$ and $D_O \supset \mathcal{R}(\nu)$ with $\text{radius}(D_O) = \alpha \cdot \text{radius}(D_I)$. Moreover, the complexity of the region $\mathcal{R}(\nu)$ is $O(2^d)$.

Let \mathcal{T} be a BAR-tree on the point set S . For a node ν , we use $\text{pa}(\nu)$ to denote the parent of ν , and we use $S(\nu)$ to denote the subset of points from S that are stored in the leaves of the subtree \mathcal{T}_ν rooted at ν . The *weight* of a node ν is the number of points in $S(\nu)$, and is denoted $|S(\nu)|$. We say that a node ν in \mathcal{T} has *weight class* ℓ , for some integer ℓ , if and only if $|S(\nu)| \leq n/2^\ell$ and $|S(\text{pa}(\nu))| > n/2^\ell$. The weight class of the root is defined to be zero. We denote the collection of nodes of weight class ℓ by $N(\ell)$. Obviously we have $\lceil \log n \rceil$ weight classes. Note that some of the nodes in the tree may not be in any weight class; this can happen when the weight of a node ν is almost the same as the weight of its parent. For example, this happens when $|S(\text{pa}(\nu))| = n/2^\ell$ for some ℓ and $|S(\nu)| = n/2^\ell - 1$. It can also happen that a node belongs to more than one weight class, namely when the weight of a node is much smaller than the weight of its parent. The following lemma is straightforward.

Lemma 12. *Every leaf node is in weight class ℓ_{\max} , where $\ell_{\max} = \lfloor \log n \rfloor$. Furthermore, on any root-to-leaf-path there is exactly one node with weight class ℓ , for any $0 \leq \ell \leq \ell_{\max}$.*

For a node $\nu \in N(\ell)$, we define its ℓ -parent to be the node $\nu' \in N(\ell - 1)$ that is on the path from the root of \mathcal{T} to ν (including ν itself). We denote the ℓ -parent of ν by $\text{pa}(\ell, \nu)$. Observe that ν can be its own ℓ -parent, namely when $\nu \in N(\ell)$ and $\nu \in N(\ell - 1)$. By Lemma 12, if $\nu \in N(\ell)$ then one of its ancestors (possibly itself) must be in weight class $\ell - 1$, so it must have an ℓ -parent. If μ is the ℓ -parent of ν then we call ν a ℓ -child of μ .

For a node ν in the BAR-tree, the region corresponding to ν is denoted by $\mathcal{R}(\nu)$ and for a region R , we let $\text{diam}(R)$ denote the diameter of the region R . As mentioned before, all nodes in the BAR-tree have bounded aspect ratio, that is, all aspect ratios are bounded by some fixed constant α .

Lemma 13 (Lemma 3.5 from [1]). *If $\mathbf{d}(\mathcal{R}(\nu), \mathcal{R}(\mu)) \geq \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\nu)), \text{diam}(\mathcal{R}(\mu))\}$ then there are two balls $D_\nu \supset \mathcal{R}(\nu)$ and $D_\mu \supset \mathcal{R}(\mu)$ such that*

$$\mathbf{d}(D_\nu, D_\mu) \geq s \cdot \min\{\text{radius}(D_\nu), \text{radius}(D_\mu)\}.$$

Now we construct an SSPD \mathcal{S} of the point set S using the following algorithm.

1. Construct a BAR tree \mathcal{T} on S . Let α be the maximum aspect ratio of the region $\mathcal{R}(\nu)$ for any node $\nu \in \mathcal{T}$. Compute the weight classes of all nodes in \mathcal{T} .
2. For each weight class ℓ with $0 \leq \ell \leq \ell_{\max}$ do the following: add to \mathcal{S} all pairs $(S(\nu), S(\mu))$ such that
 - (i) $\nu, \mu \in N(\ell)$,
 - (ii) ν and μ are leaves or $\mathbf{d}(\mathcal{R}(\nu), \mathcal{R}(\mu)) \geq \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\nu)), \text{diam}(\mathcal{R}(\mu))\}$, and
 - (iii) $\mathbf{d}(\mathcal{R}(\text{pa}(\ell, \nu)), \mathcal{R}(\text{pa}(\ell, \mu))) < \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\text{pa}(\ell, \nu))), \text{diam}(\mathcal{R}(\text{pa}(\ell, \mu)))\}$.

Lemma 14 (Lemma 3.6 from [1]). *\mathcal{S} is an SSPD for S with respect to s .*

To bound the size of the SSPD, we first prove two auxiliary lemmas.

Lemma 15. *A node ν in \mathcal{T} can be an ℓ -parent of at most 2^{2d} nodes in \mathcal{T} .*

Proof. Consider a node $\nu \in N(\ell - 1)$ and let ν' be a node such that $\nu = \text{pa}(\ell, \nu')$. Then ν' is a node in \mathcal{T}_ν (the subtree of \mathcal{T} rooted at ν) in weight class ℓ . Note that no other node than ν' in \mathcal{T}_ν can have ν as its ℓ -parent. Recall that the weight of a node reduces with a factor of β when we go down two levels in a BAR-tree. Since $\nu' \in N(\ell)$, its (normal) parent has weight at least $n/2^\ell$. On the other hand $\nu \in N(\ell - 1)$, so the weight of ν is at most $n/2^{\ell-1}$. Hence, the path between ν and ν' consists of at most $2k$ links, where $\beta^k = 1/2$ and $\beta = (d+1)/(d+2)$. It follows that the total number of nodes in \mathcal{T} that have ν as a ℓ -parent is bounded by 2^{2d} —note that $\log \beta \geq -1/d$. \square

Lemma 16. *Let $\bar{\mathcal{S}}(\ell)$ be the set of all pairs (ν, μ) such that $\nu, \mu \in N(\ell)$ and $\mathbf{d}(\mathcal{R}(\nu), \mathcal{R}(\mu)) < \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\nu)), \text{diam}(\mathcal{R}(\mu))\}$, where $0 \leq \ell \leq \ell_{\max}$. Then $|\bar{\mathcal{S}}(\ell)| = O(\alpha^{2d} d^{d/2} (s+1)^d \cdot 2^\ell)$ and*

$$\sum_{(\nu, \mu) \in \bar{\mathcal{S}}(\ell)} (|S(\nu)| + |S(\mu)|) = O\left(\alpha^{2d} d^{d/2} (s+1)^d \cdot n\right).$$

Proof. We reorder the nodes in the pairs (ν, μ) such that

$$\text{diam}(\mathcal{R}(\nu)) \leq \text{diam}(\mathcal{R}(\mu)).$$

We claim that any node ν appears in a constant number of pairs as the first element of the pair. To show this let (ν, μ) be an arbitrary ordered pair. Let $D_{S(\nu)}$ be the smallest enclosing ball of $S(\nu)$ and let o be its center. Consider the annulus A between the balls D_1 and D_2 with center o and radii $r_1 := ((s+1)\alpha + 1) \cdot \text{radius}(D_{S(\nu)})$ and $r_2 := r_1 + \text{radius}(D_{S(\nu)})$. Note that

$$\text{diam}(\mathcal{R}(\nu))/2 \leq \text{radius}(D_{S(\nu)}) \leq \text{diam}(\mathcal{R}(\nu)).$$

Since

$$\mathbf{d}(\mathcal{R}(\nu), \mathcal{R}(\mu)) < \frac{(s+1)\alpha}{2} \cdot \text{diam}(\mathcal{R}(\nu))$$

the region $\mathcal{R}(\mu)$ intersects D_1 —see Fig. 1. Now we have two cases:

Case 1: The region $\mathcal{R}(\mu)$ lies partially outside D_2 . Using the packing lemma by Duncan *et al.* [11, Lemma 3.2] this can happen for $O\left((\alpha\sqrt{d})^d (r_1/(r_2 - r_1)^{d-1})\right) = O(\alpha^{2d-1} d^{d/2} (s+1)^{d-1})$ regions.

Case 2: In this case the region $\mathcal{R}(\mu)$ lies inside D_2 . Because the aspect ratio of the region $\mathcal{R}(\mu)$ is at most α , there are two d -balls D_I and D_O such that $D_I \subset \mathcal{R}(\mu) \subset D_O$ and $\text{volume}(D_O) \leq \alpha^d \cdot \text{volume}(D_I)$, where $\text{volume}(A)$ denotes the volume of the region A . Note that volume of a d -ball with radius r is $C_d r^d$, where

$$C_d = \begin{cases} \frac{(2\pi)^{d/2}}{2 \times 4 \times \dots \times d} & \text{if } d \text{ is even} \\ \frac{(2\pi)^{(d-1)/2}}{1 \times 3 \times \dots \times d} & \text{if } d \text{ is odd.} \end{cases}$$

Therefore

$$\begin{aligned} \text{volume}(\mathcal{R}(\mu)) &\geq \text{volume}(D_I) \\ &\geq \frac{1}{\alpha^d} \cdot \text{volume}(D_O) \\ &= \frac{1}{\alpha^d} \cdot C_d (\text{radius}(D_O))^d \\ &\geq \frac{1}{(2\alpha)^d} \cdot C_d (\text{diam}(\mathcal{R}(\mu)))^d \\ &\geq \frac{1}{(2\alpha)^d} \cdot C_d (\text{diam}(\mathcal{R}(\nu)))^d \\ &\geq \frac{C_d r_2^d}{(2\alpha)^d \cdot ((s+1)\alpha + 2)^d}. \end{aligned}$$

On the other hand, the volume of D_2 is $C_d r_2^d$, which means we can have at most $O(\alpha^{2d}(s+1)^d)$ such regions.

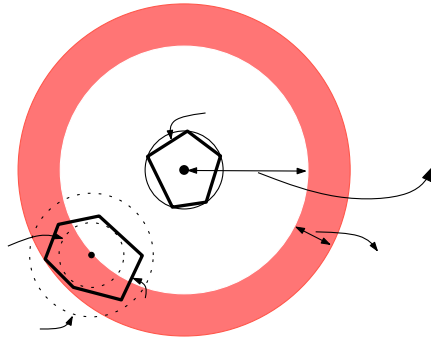


Fig. 1. Illustrating the proof of Lemma 16 (in 2D).

Hence in total we can have $O(\alpha^{2d} d^{d/2} (s+1)^d)$ pairs that have ν as the first element. Since $|N(\ell)| = O(2^\ell)$, we can have $O(2^\ell)$ nodes as the first element of the pair so $|\overline{S}(\ell)| = O(\alpha^{2d} d^{d/2} (s+1)^d \cdot 2^\ell)$. The lemma follows since $|S(\nu)| \leq n/2^\ell$ for each $\nu \in N(\ell)$. \square

Corollary 3. *The number of pairs in the SSPD \mathcal{S} generated by the construction algorithm is $O((4\alpha)^{2d} d^{d/2} (s+1)^d \cdot n)$.*

Proof. By the construction algorithm, if $(S(\nu), S(\mu)) \in \mathcal{S}$ and $\nu, \mu \in N(\ell)$ then $(\text{pa}(\ell, \nu), \text{pa}(\ell, \mu)) \in \overline{\mathcal{S}}(\ell - 1)$. By combining this with Lemma 15, we conclude that the number of pairs in \mathcal{S} is bounded by $O(\sum_{\ell=0}^{\log n} 4^{2d} \cdot |\overline{\mathcal{S}}(\ell)|)$. Using Lemma 16 we have

$$\begin{aligned} |\mathcal{S}| &= O\left(\sum_{\ell=0}^{\log n} 4^{2d} \cdot |\overline{\mathcal{S}}(\ell)|\right) \\ &= O\left(\sum_{\ell=0}^{\log n} (4\alpha)^{2d} d^{d/2} (s+1)^d \cdot 2^\ell\right) \\ &= O\left((4\alpha)^{2d} d^{d/2} (s+1)^d \cdot n\right). \end{aligned}$$

□

Now we are finally ready to bound the size of \mathcal{S} .

Lemma 17. *For the SSPD \mathcal{S} generated by the construction algorithm we have*

$$\sum_{(\nu, \mu) \in \mathcal{S}} (|S(\nu)| + |S(\mu)|) = O\left((4\alpha)^{2d} d^{d/2+1} (s+1)^d \cdot n \log n\right).$$

Proof. Since the number of weight classes in \mathcal{T} is $O(\log n)$ it suffices to prove that for every fixed ℓ it holds that:

$$\sum_{\substack{(\nu, \mu) \in \mathcal{S} \\ \nu, \mu \in N(\ell)}} (|S(\nu)| + |S(\mu)|) = O((4\alpha)^{2d} d^{d/2} (s+1)^d n). \quad (4)$$

Obviously $|S(\nu)| \leq |S(\text{pa}(\ell, \nu))|$ for each node ν , so we can bound (4) by

$$\sum_{\substack{(\nu, \mu) \in \mathcal{S} \\ \nu, \mu \in N(\ell)}} (|S(\text{pa}(\ell, \nu))| + |S(\text{pa}(\ell, \mu))|). \quad (5)$$

From the algorithm we know that

$$\mathbf{d}(\mathcal{R}(\text{pa}(\ell, \nu)), \mathcal{R}(\text{pa}(\ell, \mu))) < \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\text{pa}(\ell, \nu))), \text{diam}(\mathcal{R}(\text{pa}(\ell, \mu)))\}.$$

Furthermore, by Lemma 15 each node can be an ℓ -parent of 2^{2d} nodes. Hence, (5) can be bounded by

$$\sum_{(\nu, \mu) \in \overline{\mathcal{S}}(\ell-1)} 4^{2d} \cdot O(|S(\nu)| + |S(\mu)|), \quad (6)$$

where $\overline{\mathcal{S}}(\ell - 1)$ is the set of all pair (ν, μ) such that $\nu, \mu \in N(\ell - 1)$ and

$$\mathbf{d}(\mathcal{R}(\nu), \mathcal{R}(\mu)) < \frac{(s+1)\alpha}{2} \cdot \min\{\text{diam}(\mathcal{R}(\nu)), \text{diam}(\mathcal{R}(\mu))\}.$$

According to Lemma 16 summation (6) is $O((4\alpha)^{2d} d^{d/2+1} (s+1)^d \cdot n)$, which completes the proof of the lemma. □

Lemma 18. *The SSPD of a set S of n points w.r.t. a constant s can be computed in $O((4\alpha)^{2d} d^{d/2} (s+1)^d n + d2^d n \log n)$ time.*

Proof. The proof is the same as the proof of Lemma 3.10 from [1] where the dependency on the dimension follows from plugging in Corollary 3 and the fact that the BAR tree \mathcal{T} requires $O(d2^d n \log n)$ time to compute [11]. □

The following theorem summarizes the results on the SSPD construction.

Theorem 5. *Given a set S of n points in \mathbb{R}^d and $s > 0$ we can compute an SSPD with respect to s of size $O((4\alpha)^{2d} d^{d/2} (s+1)^d n)$ and weight $O((4\alpha)^{2d} d^{d/2+1} s^d n \log n)$ in time $O((4\alpha)^{2d} d^{d/2} s^d n + d^2 n \log n)$, where α is a constant greater than $18d\sqrt{d}$.*

5 Conclusions

We have studied spanners for vertex-weighted point sets in a metric space, where the length of an edge connecting two points is the length of the edges in the underlying metric space plus the weights of two incident points. We presented a general method to compute a $(5 + \varepsilon)$ -spanner of linear size for this setting, provided that we can compute a linear-size $(1 + \varepsilon)$ -spanner in the underlying metric space. We also presented a method to compute a $(2 + \varepsilon)$ -spanner; this spanner, however, has size $O(n \log n)$. We furthermore showed that in general one cannot obtain a spanner with a subquadratic number of edges whose dilation is smaller than 2. Hence, the main open problem is whether there exists a linear-size spanner with dilation close to 2.

Recently Abam and Har-Peled [2] showed that any n -point metric space with doubling dimension dim admits an SSPD of expected size $O(n/\varepsilon^{dim})$ such that the total expected weight of the pairs is $(n/\varepsilon^{O(dim)} \log n)$, which can be computed in $(n/\varepsilon^{O(dim)} \log^2 n)$ expected time. By Lemma 9 this immediately implies that any metric space of doubling dimension dim admits a $(2 + \varepsilon)$ -spanner of weighted point sets with $(n/\varepsilon^{O(dim)} \log n)$ edges.

References

1. Mohammad Ali Abam, Mark de Berg, Mohammad Farshi, and Joachim Gudmundsson. Region-fault tolerant geometric spanners. *Discrete and Computational Geometry*, 41:556–582, 2009.
2. Mohammad Ali Abam and Sariel Har-Peled. New constructions of sspds and their applications. In *SCG'10: Proceedings of the 26th Annual ACM Symposium on Computational Geometry*, pages 192–200, 2010.
3. Pankaj K. Agarwal, Sariel Har-Peled, Micha Sharir, and Kasturi R. Varadarajan. Approximate shortest paths on a convex polytope in three dimensions. *Journal of the ACM*, 44:567–584, 1997.
4. Sunil Arya, David M. Mount, N. S. Netanyahu, R. Silverman, and A. Wu. An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. *Journal of the ACM*, 45:891–923, 1998.
5. Patrice Assouad. Plongements lipschitziens dans r^n . *Bulletin de la Société Mathématique de France*, 111(4):429–448, 1983.
6. Béla Bollobás and Alex Scott. On separating systems. *European Journal of Combinatorics*, 28(4):1068–1071, 2007.
7. Prosenjit Bose, Paz Carmi, and Mathieu Courture. Spanners of additively weighted point sets. In *SWAT'08: Proceedings of the 11th Scandinavian Workshop on Algorithm Theory*, volume 5124 of *Lecture Notes in Computer Science*, pages 367–377. Springer, 2008.
8. Paul B. Callahan and S. Rao Kosaraju. Faster algorithms for some geometric graph problems in higher dimensions. In *SODA '93: Proceedings of the 4th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 291–300, Philadelphia, PA, USA, 1993. Society for Industrial and Applied Mathematics.
9. Paul B. Callahan and S. Rao Kosaraju. A decomposition of multidimensional point sets with applications to k -nearest-neighbors and n -body potential fields. *Journal of the ACM*, 42:67–90, 1995.
10. Richard Cole and Lee-Ad Gottlieb. Searching dynamic point sets in spaces with bounded doubling dimension. In *STOC'06*, pages 574–583, 2006.
11. Christian A. Duncan, Michael T. Goodrich, and Stephen Kobourov. Balanced aspect ratio trees: Combining the advances of k -d trees and octrees. *Journal of Algorithms*, 38:303–333, 2001.
12. Lee-Ad Gottlieb and Liam Roditty. An optimal dynamic spanner for doubling metric spaces. In *Proceedings of the 16th Annual European Symposium on Algorithms*, volume 5193 of *Lecture Notes in Computer Science*, pages 478–489. Springer-Verlag, 2008.
13. Anupam Gupta, Robert Krauthgamer, and James R. Lee. Bounded geometries, fractals, and low-distortion embeddings.

14. G. Hansel. Nombre minimal de contacts de fermeture nécessaires pour réaliser une fonction booléenne symétrique de n variables. *C. R. Acad. Sci. Paris*, 258:6037–6040, 1964. Russian transl., Kibern. Sb. (Nov. Ser.) 5 (1968), 47-52.
15. Sariel Har-Peled and Manor Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM Journal on Computing*, 35(5):1148–1184, 2006.
16. Juha Heinonen. *Lectures on analysis on metric spaces*. Universitext. Springer-Verlag, New York, 2001.
17. Giri Narasimhan and Michiel Smid. *Geometric spanner networks*. Cambridge University Press, 2007.
18. Liam Roditty. Fully dynamic geometric spanners. In *SCG'07: Proceedings of the 23rd Annual ACM Symposium on Computational Geometry*, pages 373–380, 2007.
19. Kunal Talwar. Bypassing the embedding: algorithms for low dimensional metrics. In *STOC'04: Proceedings of the 36th Annual ACM Symposium on Theory of Computing*, pages 281–290, New York, NY, USA, 2004. ACM.
20. Kasturi R. Varadarajan. A divide-and-conquer algorithm for min-cost perfect matching in the plane. In *FOCS'98: Proceedings of the 39th Annual IEEE Symposium on Foundations of Computer Science*, pages 320–331, 1998.