Geometric Spanners for Weighted Point Sets

Mohammad Ali Abam[1,⋆], Mark de Berg[2,⋆⋆], Mohammad Farshi[3,⋆⋆⋆], Joachim Gudmundsson[4], and Michiel Smid[3,⋆⋆⋆]

[1] MADALGO Center, Aarhus University, Denmark
abam@madalgo.au.dk
[2] Department of Computer Science, TU Eindhoven, The Netherlands
m.t.d.berg@tue.nl
[3] School of Computer Science, Carleton University, Ottawa, ON, K1S 5B6, Canada
mfarshi@cg.scs.carleton.ca, michiel@scs.carleton.ca
[4] NICTA, Sydney, Australia
joachim.gudmundsson@nicta.com.au

Abstract. Let $(S, \mathbf{d})$ be a finite metric space, where each element $p \in S$ has a non-negative weight $\mathrm{w}(p)$. We study spanners for the set S with respect to weighted distance function $\mathbf{d}_w$, where $\mathbf{d}_w(p,q)$ is $\mathrm{w}(p) + \mathbf{d}(p,q) + \mathrm{w}(q)$ if $p \neq q$ and 0 otherwise. 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 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.

⋆ MAA was supported by the MADALGO Center for Massive Data Algorithmics, 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 NSERC of Canada.

A. Fiat and P. Sanders (Eds.): ESA 2009, LNCS 5757, pp. 190–202, 2009.

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 $t \geqslant 1$, if for any two points $p, q \in S$ we have $\mathbf{d}_{\mathcal{G}}(p, q) \leqslant t \cdot |pq|$. The smallest t for which $\mathcal{G}$ is a t-spanner is called the *stretch factor* (or *dilation*, or *spanning ratio*) of $\mathcal{G}$. 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 [1] for this and many other results on spanners. Instead of considering points in Euclidean space, one can also consider points in some other metric space. As it turns out, results similar to the Euclidean setting are possible when the so-called *doubling dimension* of the metric space—see footnote 1 on p. 195 for a definition—is bounded by a constant d: in this case there is a $(1 + \varepsilon)$-spanner with $n/\varepsilon^{O(d)}$ edges [2,3].

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 passing through 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 $\mathrm{w}(p)$. We now define a new distance function on S, denoted by $\mathbf{d}_w$, as follows.

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

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 an *additively weighted t-spanner* of S, if for any two points p and q in S we have $\mathbf{d}_{\mathcal{G},w}(p, q) \leqslant t \cdot \mathbf{d}_w(p, q)$. We want to compute an additively weighted t-spanner of 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. (An easy 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 additively weighted spanners 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.* [4] 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 $|pq| - \mathrm{w}(p) - \mathrm{w}(q)$. 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 additively weighted spanners. 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 **d**-metric on the cluster centers, and then connecting each point to its cluster center. We apply our method to obtain, for any $0 < \varepsilon < 1$, additively weighted $(5+\varepsilon)$-spanners 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 an additively weighted $(5+\varepsilon)$-spanner 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 [5,6] 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 $\mathrm{w}(p) \geqslant 0$. We will show that if we can find a good spanner for S in the **d**-metric, we can also find a good additively weighted 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 **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 at most proportional to the weight $\mathrm{w}(p)$ of p. We then show that a t-spanner of the cluster centers in the **d**-metric, while connecting the rest of the points to the center of their clusters, results in an $O(t)$-spanner of S in the $\mathbf{d}_w$-metric.

Clusterings for additively weighted spanners. 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 $k_2 \geqslant 1$. Define a (k_1, k_2)*-clustering of S* to be a partitioning of S into a collection $\{C_1, ..., C_m\}$ of clusters, each with a center denoted by $\mathrm{center}(C_i)$, such that the following three conditions hold:

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

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

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

Let $\{C_1, C_2, \ldots, C_m\}$ be a (k_1, k_2)-clustering of S, and let $c_i = \mathrm{center}(C_i)$. Let $\mathcal{C} = \{c_1, c_2, \ldots, c_m\}$ 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 $E_2 = \{(c_i, p) : 1 \leqslant i \leqslant m \text{ and } p \in C_i \text{ and } p \neq c_i\}$. 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 $\mathbf{d}_{\mathcal{G},w}(p, q) \leqslant t' \cdot \mathbf{d}_w(p, q)$. Let C_i and C_j be the clusters containing p and q, respectively, and consider $c_i = \mathrm{center}(C_i)$ and $c_j = \mathrm{center}(C_j)$. (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,

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

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 $\mathbf{d}_{\mathcal{G},w}(c_i, c_j) \leqslant k_2 \cdot \mathbf{d}_{\mathcal{G}_1}(c_i, c_j) \leqslant k_2 t \cdot \mathbf{d}(c_i, c_j)$. Finally, we observe that

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

Combing this with our two earlier derivations, we get

$$
\begin{aligned}
\mathbf{d}_{\mathcal{G},w}(p,q) &\leqslant (2+k_1)\cdot \mathrm{w}(p) + \mathbf{d}_{\mathcal{G},w}(c_i,c_j) + (2+k_1)\cdot \mathrm{w}(q)\\
&\leqslant (2+k_1)\cdot \mathrm{w}(p) + k_2 t\cdot \mathbf{d}(c_i,c_j) + (2+k_1)\cdot \mathrm{w}(q)\\
&\leqslant (2+k_1)\cdot \mathrm{w}(p) + k_2 t\cdot (k_1\cdot \mathrm{w}(p) + \mathbf{d}(p,q) + k_1\cdot \mathrm{w}(q)) + (2+k_1)\cdot \mathrm{w}(q)\\
&= (2+k_1+k_1k_2t)\cdot \mathrm{w}(p) + k_2 t\cdot \mathbf{d}(p,q) + (2+k_1+k_1k_2t)\cdot \mathrm{w}(q)\\
&\leqslant \max(2+k_1+k_1k_2t, k_2t)\cdot (\mathrm{w}(p) + \mathbf{d}(p,q) + \mathrm{w}(q))\\
&= \max(2+k_1+k_1k_2t, k_2t)\cdot \mathbf{d}_w(p,q).
\end{aligned}
$$

□

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 $\{C_1, \ldots, C_m\}$ of S.

1. Sort the points of S in nondecreasing order of their weight, and let $p_1, p_2, \ldots, p_n$ 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 $1 \leqslant j \leqslant m$ such that c_j is a $(1+\varepsilon)$-approximate nearest-neighbor of p_i in the set $\mathcal{C}$, in the $\mathbf{d}$-metric.
   (b) If $\mathbf{d}(c_j, p_i) \leqslant k \cdot \mathrm{w}(p_i)$, 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 $\mathcal{C} = \mathcal{C} \cup \{p_i\}$.
4. Return the collection $\{C_1, \ldots, C_m\}$ 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 $\mathrm{w}(c_j) \leqslant \mathrm{w}(p)$ 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 $\mathbf{d}(\text{center}(C_j), p) \leqslant k \cdot \mathrm{w}(p)$. 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 $\mathrm{w}(c) \leqslant \mathrm{w}(c')$. 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 $\mathbf{d}(c_j, c') > k \cdot \mathrm{w}(c')$. Let c'' be the exact nearest-neighbor of c' in $\mathcal{C}$. Then, since $c \in \mathcal{C}$, $\mathbf{d}(c_j, c') \leqslant (1+\varepsilon) \cdot \mathbf{d}(c'', c') \leqslant (1+\varepsilon) \cdot \mathbf{d}(c, c')$. It follows that

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

□

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

Due to space limitation, the proof of the theorem is removed in this version.

Applications: Euclidean spaces and spaces of bounded doubling dimension. Theorem 1 can immediately be used to obtain additively weighted spanners in Euclidean spaces and metric spaces of bounded doubling dimension.[1]

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 an additively weighted $(5+\varepsilon)$-spanner of 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 an additively weighted $(5+\varepsilon)$-spanner of S having $n/\varepsilon^{O(d)}$ edges in $O(n \log n) + n/\varepsilon^{O(d)}$ time.

Proof. Callahan and Kosaraju [7] 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 $E_{\mathrm{sp}}(n) = O(n/\varepsilon^d)$ edges in $T_{\mathrm{sp}}(n) = O(n \log n + n/\varepsilon^d)$ time. Moreover, Arya *et al.* [8] 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 $t = 1 + \varepsilon/10$.

Gottlieb and Roddity [9] 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 $E_{\mathrm{sp}}(n) = n/\varepsilon^{O(d)}$ edges in $T_{\mathrm{sp}}(n) = O(n \log n) + n/\varepsilon^{O(d)}$ 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 $t = 1 + \varepsilon/10$. ⊔⊓

[1] 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 $\{q \in S : \mathbf{d}(p,q) \leqslant R\}$. 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 $R/2$.

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 ∂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 [7] for the Euclidean metric and by Talwar [3] 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 $\mathbf{d}(A, B) \geqslant s \cdot \max(\text{diam}_{\mathbf{d}}(A), \text{diam}_{\mathbf{d}}(B))$.

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 $\{(A_1, B_1), \ldots, (A_m, B_m)\}$ 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 $1 \leqslant i \leqslant m$ 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 [11], shows how a spanner can be obtained from a WSPD. They prove the lemma for Euclidean spaces, but exactly the same proof applies to any metric space.

Lemma 3. **[11]** *Let $(S, \mathbf{d})$ be a finite metric space and let $t > 1$ be a real number. Furthermore, let $\{(A_1, B_1), \ldots, (A_m, B_m)\}$ be a WSPD for $(S, \mathbf{d})$, with respect to $s = \frac{2(t+1)}{t-1}$ and, for each i with $1 \leqslant i \leqslant m$, 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 $E = \{(a_i, b_i) : 1 \leqslant i \leqslant m\}$ 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 a 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 [2] 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 ∂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. [12].

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

The normal vector of each point of $\partial\mathcal{B}$ at 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 $\{(A_1, B_1), \ldots, (A_m, B_m)\}$ 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 $A_i^1, \ldots, A_i^\ell$ and $B_i^1, \ldots, B_i^\ell$, respectively, where $\ell = O(1/\sigma^{d-1})$. The partitioning is done such that the points in each subset belong to the same σ-patch. Define $\Psi = \{(A_i^j, B_i^k) : 1 \leqslant j \leqslant \ell \text{ and } 1 \leqslant k \leqslant \ell\}$.

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

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

We first show that

$$\mathrm{diam}_{\mathbf{d}}(A_i^j) \geqslant \mathrm{diam}_{\mathbf{d}_{\mathcal{B}}}(A_i^j)\cos\sigma. \tag{2}$$

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

$$\mathrm{diam}_{\mathbf{d}}(B_i^k) \geqslant \mathrm{diam}_{\mathbf{d}_{\mathcal{B}}}(B_i^k)\cos\sigma. \tag{3}$$

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 $\mathbf{d}_{\mathcal{B}}(a,b) \geqslant \mathbf{d}(a,b) \geqslant s \cdot \max(\mathrm{diam}_{\mathbf{d}}(A_i), \mathrm{diam}_{\mathbf{d}}(B_i)) \geqslant s \cdot \max(\mathrm{diam}_{\mathbf{d}}(A_i^j), \mathrm{diam}_{\mathbf{d}}(B_i^k))$. Combining this with (2) and (3), it follows that $\mathbf{d}_{\mathcal{B}}(a,b) \geqslant s\cos\sigma \cdot \max(\mathrm{diam}_{\mathbf{d}_{\mathcal{B}}}(A_i^j), \mathrm{diam}_{\mathbf{d}_{\mathcal{B}}}(B_i^k))$. This proves that (1) holds. ◻

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 an additively weighted $(5+\varepsilon)$-spanner of S having $O(n/\varepsilon^d)$ edges.*

3 An Additively Weighted $(2+\varepsilon)$-Spanner

In each of the applications considered in the previous section, our method generated an additively weighted $(5+\varepsilon)$-spanner. The goal of this section is to see if we can obtain additively weighted spanners 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 additively weighted 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 $\mathbf{d}_w(p,q) \leqslant (1 + 2/\varepsilon)D$, whereas $\mathbf{d}_{\mathcal{G},w}(p,q) \geqslant 4D/\varepsilon$. Thus $\frac{\mathbf{d}_{\mathcal{G},w}(p,q)}{\mathbf{d}_w(p,q)} \geqslant \frac{4D/\varepsilon}{(1+2/\varepsilon)D} > 2 - \varepsilon$. ⊓⊔

In the remainder of this section we will describe a general strategy for computing additively weighted $(2+\varepsilon)$-spanners for spaces of bounded doubling dimension. Given the lower bound, the stretch factor is almost optimal in the worst case. Our method is based on the so-called semi-separated pair decomposition, as introduced by Varadarajan [6]. We use the strategy to obtain additively weighted $(2+\varepsilon)$-spanners for two cases: 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 $\text{w}(p) \geqslant 0$. Our spanner construction will be based on a decomposition $\{(A_1, B_1), \ldots, (A_m, B_m)\}$ having properties similar to those of the WSPD. As we will see, the number of edges in the additively weighted spanner 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 [7] 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 in [6].

For a real number $s > 0$, two subsets $A, B \subseteq S$ are called *semi-separated with respect to* s, if $\mathbf{d}(A, B) \geqslant s \cdot \min(\text{diam}_{\mathbf{d}}(A), \text{diam}_{\mathbf{d}}(B))$. A *semi-separated pair* *decomposition (SSPD)* for the metric space $(S, \mathbf{d})$, with respect to s, is defined to be a set $\Psi = \{(A_1, B_1), \ldots, (A_m, B_m)\}$ 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.

The SSPD was introduced by Varadarajan [6]. For the Euclidean distance function in $\mathbb{R}^2$, Abam *et al.* [5] showed that an SSPD with $O(n)$ pairs and size $O(n \log n)$ can be computed in $O(n \log n)$ time. It is known that for any set of n points, any SSPD has size $\Omega(n \log n)$; see [13,14].

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 $\mathrm{diam}_{\mathbf{d}}(A) \leqslant \mathrm{diam}_{\mathbf{d}}(B)$. Thus, we have $\mathbf{d}(A, B) \geqslant s \cdot \mathrm{diam}_{\mathbf{d}}(A)$. Define center$(A)$ to be a point from A of minimum weight (among the points in A), and let $E_1(A, B) = \{(x, \mathrm{center}(A)) : x \in A \cup B \text{ and } x \neq \mathrm{center}(A)\}$. This provides short connections between the points in A and those in B by going via center(A): since $\mathbf{d}(A, B) \geqslant s \cdot \mathrm{diam}_{\mathbf{d}}(A)$, going via center$(A)$ does not create a large detour in the $\mathbf{d}$-metric, and since $\mathrm{w}(\mathrm{center}(A)) \leqslant \mathrm{w}(a)$ the extra path length caused by $\mathrm{w}(\mathrm{center}(A))$ 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* $c = \mathrm{center}(A)$. *Let* $b \in B$ *be an points such that* $\mathrm{w}(c) \leqslant \mathrm{w}(b) + \mathbf{d}(c, b)$. *Then, for any* $a \in A$, *we have* $\mathbf{d}_w(a, c) + \mathbf{d}_w(c, b) \leqslant \left(2 + \frac{3}{s}\right) \cdot \mathbf{d}_w(a, b)$.

Proof. We have

$$\begin{aligned}
\mathbf{d}_w(a, c) + \mathbf{d}_w(c, b) &= (\mathrm{w}(a) + \mathbf{d}(a, c) + \mathrm{w}(c)) + (\mathrm{w}(c) + \mathbf{d}(c, b) + \mathrm{w}(b)) \\
&\leqslant (2 \cdot \mathrm{w}(a) + \mathrm{diam}_{\mathbf{d}}(A)) + 2 \cdot (\mathbf{d}(c, b) + \mathrm{w}(b)) \\
&\leqslant (2 \cdot \mathrm{w}(a) + \mathrm{diam}_{\mathbf{d}}(A)) + 2 \cdot (\mathbf{d}(c, a) + \mathbf{d}(a, b) + \mathrm{w}(b)) \\
&\leqslant 2 \cdot (\mathrm{w}(a) + \mathbf{d}(a, b) + \mathrm{w}(b)) + 3 \cdot \mathrm{diam}_{\mathbf{d}}(A) \\
&\leqslant 2 \cdot (\mathbf{d}_w(a, b)) + 3 \cdot (\mathbf{d}(a, b)/s) \\
&\leqslant (2 + \tfrac{3}{s}) \cdot \mathbf{d}_w(a, b).
\end{aligned}$$

□

It remains to establish short paths between the points in A and the points $b \in \overline{B}$, where $\overline{B} = \{b \in B : \mathrm{w}(c) > \mathrm{w}(b) + \mathbf{d}(c, b)\}$ with $c = \mathrm{center}(A)$. 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 $\mathbf{d}(c, b) < \mathrm{w}(c)$, so they are contained in a $\mathbf{d}$-ball C of radius $\mathrm{w}(c)$. Recall that d is the doubling dimension of $(S, \mathbf{d})$. Thus we can cover C by $s^{O(d)}$ balls of radius $\mathrm{w}(c)/(2s)$. Let $C_1, \ldots, C_\ell$ be such a collection of balls, where $\ell = s^{O(d)}$. We partition $\overline{B}$ into subsets $B_1, \ldots, B_\ell$ in such a way that $B_i \subseteq C_i$ for all $1 \leqslant i \leqslant \ell$. For each B_i, let 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 center(B_i) gives us a path of the required length.

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

The proof of the lemma is similar to the proof of Lemma 7 and is removed because of the space limitation.

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 an additively weighted* $(2 + \varepsilon)$*-spanner for* S *with* $(1/\varepsilon)^{O(d)} \cdot \sum_{(A,B) \in \Psi}(|A| + |B|)$ *edges.*

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.* [5] 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 $s > 1$. 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 an additively weighted* $(2+\varepsilon)$*-spanner of* 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 ∂B. The proof of the following lemma is based on the concept of σ-patches introduced earlier, and removed due to space limitation.

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 $\{(A_1, B_1) \ldots, (A_m, B_m)\}$ 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 $1 \leqslant i \leqslant m$. As before, we partition both A_i and B_i into subsets $A_i^1, \ldots, A_i^\ell$ and $B_i^1, \ldots, B_i^\ell$, 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 $\Psi = \{(A_i^j, B_i^k) : 1 \leqslant j \leqslant \ell, 1 \leqslant k \leqslant \ell\}$. 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 an additively weighted $(2+\varepsilon)$-spanner of 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 metric space $(S, \mathbf{d})$ 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 [1]. Of course, this then also holds for additively weighted spanners. Note that if all weights are equal and very large compared to the $\mathbf{d}$-diameter of the set, then any additively weighted 2-spanner 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 additively weighted 2-spanners.)

References

1. Narasimhan, G., Smid, M.: Geometric Spanner Networks. Cambridge University Press, Cambridge (2007)
2. Har-Peled, S., Mendel, M.: Fast construction of nets in low-dimensional metrics and their applications. SIAM J. on Computing 35, 1148–1184 (2006)
3. Talwar, K.: Bypassing the embedding: algorithms for low dimensional metrics. In: STOC 2004, pp. 281–290 (2004)
4. Bose, P., Carmi, P., Couture, M.: Spanners of additively weighted point sets. In: Gudmundsson, J. (ed.) SWAT 2008. LNCS, vol. 5124, pp. 367–377. Springer, Heidelberg (2008)
5. Abam, M.A., de Berg, M., Farshi, M., Gudmundsson, J.: Region-fault tolerant geometric spanners. In: SODA 2007, pp. 1–10 (2007)
6. Varadarajan, K.R.: A divide-and-conquer algorithm for min-cost perfect matching in the plane. In: FOCS 1998, pp. 320–331 (1998)
7. Callahan, P.B., Kosaraju, S.R.: A decomposition of multidimensional point sets with applications to k-nearest-neighbors and n-body potential fields. J. of the ACM 42, 67–90 (1995)
8. Arya, S., Mount, D.M., Netanyahu, N.S., Silverman, R., Wu, A.: An optimal algorithm for approximate nearest neighbor searching in fixed dimensions. J. of the ACM 45, 891–923 (1998)

9. Gottlieb, L.A., Roditty, L.: An optimal dynamic spanner for doubling metric spaces. In: Halperin, D., Mehlhorn, K. (eds.) ESA 2008. LNCS, vol. 5193, pp. 478–489. Springer, Heidelberg (2008)
10. Cole, R., Gottlieb, L.A.: Searching dynamic point sets in spaces with bounded doubling dimension. In: STOC 2006, pp. 574–583 (2006)
11. Callahan, P.B., Kosaraju, S.R.: Faster algorithms for some geometric graph problems in higher dimensions. In: SODA 1993, pp. 291–300 (1993)
12. Agarwal, P.K., Har-Peled, S., Sharir, M., Varadarajan, K.R.: Approximate shortest paths on a convex polytope in three dimensions. J. of the ACM 44, 567–584 (1997)
13. Hansel, G.: Nombre minimal de contacts de fermeture nécessaires pour réaliser une fonction booléenne symétrique de n variables. Comptes Rendus de l'Académie des Sciences 258, 6037–6040 (1964)
14. Bollobás, B., Scott, A.D.: On separating systems. European J. of Combinatorics 28, 1068–1071 (2007)