

The Greedy Spanner is Existentially Optimal

[Extended Abstract]

Arnold Filtser*

Department of Computer Science
Ben-Gurion University of the Negev
Be'er Sheva, 8410501, Israel
arnoldf@cs.bgu.ac.il

Shay Solomon

School of Computer Science
Tel Aviv University
Tel Aviv 69978, Israel
solo.shay@gmail.com

ABSTRACT

The greedy spanner is arguably the simplest and most well-studied spanner construction. Experimental results demonstrate that it is at least as good as any other spanner construction, in terms of both the size and weight parameters. However, a rigorous proof for this statement has remained elusive.

In this work we fill in the theoretical gap via a surprisingly simple observation: The greedy spanner is *existentially optimal* (or existentially near-optimal) for several important graph families. Focusing on the weight parameter, the state-of-the-art spanner constructions for both general graphs (due to Chechik and Wulff-Nilsen [SODA'16]) and doubling metrics (due to Gottlieb [FOCS'15]) are complex. Plugging our observation on these results, we conclude that the greedy spanner achieves near-optimal weight guarantees for both general graphs and doubling metrics, thus resolving two longstanding conjectures in the area.

Further, we observe that approximate-greedy algorithms are existentially near-optimal as well. Consequently, we provide an $O(n \log n)$ -time construction of $(1+\epsilon)$ -spanners for doubling metrics with constant lightness and degree. Our construction improves Gottlieb's construction, whose runtime is $O(n \log^2 n)$ and whose number of edges and degree are unbounded, and remarkably, it matches the state-of-the-art Euclidean result (due to Gudmundsson et al. [SICOMP'02]) in all the involved parameters (up to dependencies on ϵ and the dimension).

Keywords

Doubling metric, graph spanner, greedy algorithm, greedy spanner, light spanner.

1. INTRODUCTION

1.1 Graph Spanners

*Partially supported by the Lynn and William Frankel Center for Computer Sciences, ISF grant No. (523/12) and by the European Union Seventh Framework Programme (FP7/2007-2013) under grant agreement n° 303809.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions.acm.org.

PODC'16, July 25-28, 2016, Chicago, IL, USA

© 2016 ACM. ISBN 978-1-4503-3964-3/16/07...\$15.00

DOI: <http://dx.doi.org/10.1145/2933057.2933114>

Given a (connected and undirected) n -vertex m -edge graph $G = (V, E, w)$ with positive edge weights and a parameter $t \geq 1$, a subgraph $H = (V, E', w)$ of G ($E' \subseteq E$) is called a t -spanner for G if for all $u, v \in V$, $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$. (Here $\delta_G(u, v)$ and $\delta_H(u, v)$ denote the distances between u and v in the graphs G and H , respectively.) The parameter t is called the *stretch* of H . Spanners constitute a fundamental graph structure, and have been extensively and intensively studied since they were introduced [48, 49].

In many practical applications one is required to construct a spanner that satisfies a number of useful properties, while preserving a small stretch. First, the spanner H should have a small number of edges. Second, its *weight* $w(H) = \sum_{e \in E} w(e)$ should be close to the weight of a minimum spanning tree (MST) of the graph G . We henceforth refer to the normalized notion of weight $\Psi(H) = \frac{w(H)}{w(\text{MST}(G))}$, which is called *lightness*; a *light* spanner is one with small lightness. Third, its *degree* $\Delta(H)$, defined as the maximum number of edges incident on a vertex, should be small.

Light and sparse spanners are particularly useful for efficient broadcast protocols in the message-passing model of distributed computing [3, 4], where efficiency is measured with respect to both the total communication cost (corresponding to the spanner's size and weight) and the speed of message delivery at all destinations (corresponding to the spanner's stretch). Additional applications of such spanners in distributed systems include network synchronization and computing global functions [2, 49, 3, 4, 47]. Light and sparse spanners were also found useful for various data gathering and dissemination tasks in overlay networks [8, 61, 44], in wireless and sensor networks [62, 6, 56], for VLSI circuit design [16, 17, 18, 55], for routing [63, 49, 50, 60], to compute distance oracles and labels [46, 59, 53], and to compute almost shortest paths [15, 54, 22, 25, 29]. Low degree spanners are also very useful in many of these applications. For example, the degree of the spanner is what determines local memory constraints when using spanners to construct network synchronizers and efficient broadcast protocols. In compact routing schemes, the use of low degree spanners enables the routing tables to be of small size. More generally, viewing vertices as processors, in many applications the degree of a processor represents its *load*, hence a low degree spanner guarantees that the load on all the processors in the network will be low.

The *greedy spanner* by Althöfer et al. [1] is arguably the simplest and most well-studied spanner construction. Althöfer et al. showed that for every weighted n -vertex graph $G = (V, E, w)$ and an integer parameter $k \geq 1$, the *greedy algorithm* (see Algorithm 1) constructs a $(2k - 1)$ -spanner with $O(n^{1+1/k})$ edges; assuming Erdős' girth conjecture [26], this size bound is asymptotically tight. Althöfer et al. also showed that the lightness of the greedy spanner is $O(n/k)$. Chandra et al. [12] improved the light-

ness bound, and showed that the greedy spanner for stretch parameter $t = (2k - 1) \cdot (1 + \epsilon)$ (here $k > 1$, $\epsilon > 0$) has lightness $O(k \cdot n^{1/k} \cdot (1/\epsilon)^{1+1/k})$. Two decades later, Elkin, Neiman and the second author [23] improved the analysis of [12] and showed that the greedy $(2k - 1) \cdot (1 + \epsilon)$ -spanner has lightness $O(n^{1/k} \cdot (1 + k / (\epsilon^{1+1/k} \log k)))$. Very recently Chechik and Wulff-Nilsen [13] improved the lightness bound all the way to $O(n^{1/k} (1/\epsilon)^{3+2/k})$. Assuming Erdős’ girth conjecture [26] and ignoring dependencies on ϵ , the bound of [13] on the lightness is asymptotically tight, thus resolving a major open question in this area. However, the result of Chechik and Wulff-Nilsen [13] is not due to a refined analysis of the greedy spanner. Instead, they devised a different construction, which is far more complex, and bounded the lightness of their own construction. The following question was left open.

QUESTION 1. *Is the lightness analysis of [23] for the greedy spanner optimal, or can one refine it to derive a stronger bound? In particular, is the spanner of [13] lighter than the greedy spanner?*

1.2 Spanners for Euclidean and Doubling Metrics.

Consider a set P of n points in \mathbb{R}^d , $d \geq 2$, and a stretch parameter $t \geq 1$. A graph $G = (P, E, w)$ in which the weight $w(p, q)$ of each edge $e = (p, q) \in E$ is equal to the Euclidean distance $\|p - q\|$ between p and q is called a *Euclidean graph*. We say that the Euclidean graph G is a t -spanner for P (or equivalently, for the corresponding Euclidean metric $(P, \|\cdot\|)$) if for every pair $p, q \in P$ of distinct points, there exists a path $\Pi(p, q)$ in G between p and q whose weight (i.e., the sum of all edge weights in it) is at most $t \cdot \|p - q\|$. The path $\Pi(p, q)$ is said to be a t -spanner path between p and q . For Euclidean metrics, one usually focuses on the regime $t = 1 + \epsilon$, for $\epsilon > 0$ being an arbitrarily small parameter. Euclidean spanners were introduced by Chew [14], and were subject to intensive ongoing research efforts since then. We refer to the book “Geometric Spanner Networks” [45], which is devoted almost exclusively to Euclidean spanners and their numerous applications. As with general graphs, it is important to devise Euclidean spanners that achieve small size, lightness and degree.

The *doubling dimension* of a metric space (M, δ) is the smallest value ddim such that every ball B in the metric space can be covered by at most 2^{ddim} balls of half the radius of B . This notion generalizes the Euclidean dimension, since the doubling dimension of the Euclidean space \mathbb{R}^d is $\Theta(d)$. A metric space is called *doubling* if its doubling dimension is constant. Spanners for doubling metrics were also subject of intensive research [30, 10, 9, 40, 52, 32, 34, 57, 24, 11, 58]. The basic line of work in this context is to generalize the known Euclidean spanner results for arbitrary doubling metrics.

Das et al. [19] showed that, in low-dimensional Euclidean metrics, the greedy $(1 + \epsilon)$ -spanner has constant degree (and so $O(n)$ edges) and constant lightness. In n -point doubling metrics, the greedy $(1 + \epsilon)$ -spanner has $O(n)$ edges and lightness $O(\log n)$ [57]. As for the degree, there exist n -point metric spaces with doubling dimension 1 for which the greedy spanner has a degree of $n - 1$ [41, 57]. It has been a major open question to determine whether any doubling metric admits a $(1 + \epsilon)$ -spanner with sub-logarithmic lightness. Recently Gottlieb [31] answered this question in the affirmative by devising such a spanner construction with constant lightness. Again, this result is not due to a refined analysis of the greedy spanner. Instead, Gottlieb devised a different construction, which is far more complex, and bounded the lightness of his own construction. The following question was left open.

QUESTION 2. *Is the lightness analysis of [57] for the greedy spanner optimal, or can one refine it to derive a stronger bound? In particular, is the spanner of [31] lighter than the greedy spanner?*

The high runtime of the greedy spanner is a major drawback. The state-of-the-art implementation of the greedy spanner in both Euclidean and doubling metrics requires time $O(n^2 \log n)$ [7]. Building on [19], Das and Narasimhan [20] devised a much faster algorithm that follows the greedy approach. The runtime of their “approximate-greedy” algorithm is $O(n \log^2 n)$, yet its degree and lightness are both bounded by constants (as with the greedy spanner). Gudmundsson et al. [38] improved the result of [20], implementing the approximate-greedy algorithm within time $O(n \log n)$. For doubling metrics, however, the only spanner construction with sub-logarithmic lightness is that of [31]; the runtime of Gottlieb’s construction is $O(n \log^2 n)$ rather than $O(n \log n)$, and the size and degree of his construction are unbounded. Hence, there is a big gap in this context between Euclidean and doubling metrics, leading to the following question.

QUESTION 3. *Can one compute $(1 + \epsilon)$ -spanners with constant lightness in doubling metrics within time $O(n \log n)$? Furthermore, can one extend the state-of-the-art Euclidean result of [38] to arbitrary doubling metrics?*

There have been numerous experimental studies on Euclidean spanners. (See [28, 27], and the references therein.) The conclusion emerging from these experiments is that the greedy Euclidean spanner outperforms the other popular Euclidean spanner constructions, with respect to the size and lightness bounds. (Specifically, the greedy spanner was found to be 10 times sparser and 30 times lighter than any other examined spanner.) It is reasonable to assume that a similar situation occurs in arbitrary doubling metrics.

1.3 Our Contribution

In this work we fill in the theoretical gap by making three important observations.

1. Our first observation is surprisingly simple: The greedy spanner is *existentially optimal* with respect to both the size and the lightness, for any graph family that is closed under edge removal. Applying this observation to the family of general weighted graphs, we conclude that the greedy spanner is just as light as the spanner of [13], thus answering Question 1.
2. The first observation does not hold for doubling metrics. Our second observation is that the greedy spanner is existentially *near-optimal* with respect to both the size and the lightness, for the family of doubling metrics. In particular, it is just as light as the spanner of [31], thus answering Question 2.
3. Our third observation concerns the optimality of the approximate-greedy algorithm of [20, 38] in doubling metrics, and is more intricate than the first two observations. Informally, it states that the approximate-greedy spanner with stretch parameter t is existentially near-optimal with respect to the lightness, for the family of doubling metrics, but when compared to spanners with a slightly smaller stretch parameter $t' < t$. This enables us to conclude that the lightness of the approximate-greedy spanner is close to that of [31]. In this way we manage to extend the state-of-the-art Euclidean result of [38] to arbitrary doubling metrics, thus answering Question 3.¹

¹The $O(n \log n)$ runtime bound of [38] holds in the traditional algebraic computation-tree model with the added power of indirect addressing. Our result applies with respect to the same computation model.




Figure 1: The graph H in the figure is the Petersen graph on 10 vertices, with girth 5 and 15 edges. All edges of H have weight 1, and are colored black. The red dashed edges are the edges of the star S of weight $1 + \epsilon$. The greedy 3-spanner for the graph G obtained as the union of the black and red edges in the figure includes all 15 edges of H , whereas the optimal 3-spanner for G consists of the 9 edges of S .

To clarify the meaning of *existential optimality*, suppose that the greedy spanner is existentially optimal with respect to the lightness, for some graph family \mathcal{F} . This does not imply that for any graph $G \in \mathcal{F}$, the lightness of the greedy spanner for G is bounded by the optimal lightness of any spanner for G . It simply means that for any graph $G \in \mathcal{F}$, there *exists* a graph $G' \in \mathcal{F}$, such that the lightness of the greedy spanner for G is bounded by the optimal lightness of any spanner for G' . Put in other words, for a graph $G \in \mathcal{F}$, let $l(G) = l_t(G)$ denote the optimal lightness of any t -spanner for G , for an arbitrary stretch parameter $t \geq 1$, and let $l(\mathcal{F}) = \max\{l(G) \mid G \in \mathcal{F}\}$ denote the maximum value $l(G)$ over all graphs G in \mathcal{F} . Then, for any graph $G \in \mathcal{F}$, although the lightness of the greedy t -spanner for G may well exceed $l(G)$, it must be upper bounded by $l(\mathcal{F})$. For example, let \mathcal{F} be the family of general weighted graphs on n vertices, and let H be an n -vertex dense graph of high girth, namely, with girth $t + 2$ and $n^{1+\Theta(1/t)}$ edges, where all edge weights are 1. Also, let S be a star on the same vertex set as H rooted at an arbitrary vertex, so that all edges of S that belong to H have weight 1 and all edges of S that do not belong to H have weight $1 + \epsilon$. Finally, let G be the graph containing all edges of H and all edges of S with weight $1 + \epsilon$. Note that the greedy t -spanner for G includes all $n^{1+\Theta(1/t)}$ edges of the high girth graph H , whereas the optimal t -spanner (assuming $t \geq 2 + 2\epsilon$) consists of the edges of the star S , hence is much sparser and lighter. (See Figure 1 for an illustration.) This example, however, does not contradict the existential optimality of the greedy spanner: Although the lightness of the greedy t -spanner for G exceeds $l(G)$, it can be shown that it is equal to $l(H)$, which, in turn, is bounded by $l(\mathcal{F})$.

The meaning of *existential near-optimality* is similar, except that we are allowed to have some *slack*, which may depend on the stretch parameter t as well as on parameters of the graph family of interest \mathcal{F} . As mentioned, in our third observation we compare the lightness of the greedy spanner with a certain stretch parameter t to the optimal lightness of any spanner, but with a slightly smaller stretch parameter t' . This is just one example of how the slack parameter can be used. Another example is to compare the greedy spanner in some graph family \mathcal{F} to the optimal spanner, but with respect to a different (closely related) graph family \mathcal{F}' . In particular, in our second and third observations we compare the lightness

of the greedy spanner in metric spaces of bounded doubling dimension to the optimal lightness of any spanner, but with respect to metric spaces of slightly larger doubling dimension. (See Section 2 for the definition of doubling dimension.) It would be interesting to study additional ways of using the slack parameter, as they may lead to new results in this area.

We remark that light spanners were extensively studied in various graph families such as planar graphs [1, 43], apex graphs [37], bounded pathwidth graphs [36, 42], bounded catwidth graphs [42], bounded genus graphs [35, 37, 21], bounded treewidth graphs [21, 42], and graphs excluding fixed minors [35, 21]. Since all these graph families are closed under edge removal, our first observation implies that the greedy spanner for them is just as good as any other spanner.

1.4 Organization

In Section 2 we present the notation that is used throughout the paper, and summarize some statements from previous work that are most relevant to us. In Section 3 we show that the greedy spanner is existentially optimal for graph families that are closed under edge removal. The basic optimality argument of Section 3 is extended to doubling metrics in Section 4. Finally, in Section 5 we show that the approximate-greedy spanner in doubling metrics is light.

2. PRELIMINARIES

Let $G = (V, E, w)$ be a (connected and undirected) graph with positive edge weights. The weight $w(P)$ of a path P is the sum of all edge weights in it, i.e., $w(P) = \sum_{e \in P} w(e)$. For a pair of vertices $u, v \in V$, let $\delta_G(u, v)$ denote the distance between u and v in G , i.e., the weight of a shortest path between them. We denote by $M_G = (V, \delta_G)$ the (shortest path) metric space *induced* by G ; we will view M_G as a complete weighted graph $(V, \binom{V}{2}, w)$ over the vertex set V , where the weight $w(u, v)$ of an edge (u, v) is given by the graph distance $\delta_G(u, v)$ between its endpoints. A subgraph $H = (V, E', w)$ of G (where $E' \subseteq E$) is called a t -spanner for G if for all $u, v \in V$, $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$. The parameter t is called the *stretch* of the spanner H . If $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$ for all edges $(u, v) \in E$, then it also holds that $\delta_H(u, v) \leq t \cdot \delta_G(u, v)$ for all pairs of vertices $u, v \in V$. Therefore, to bound the stretch of the spanner, one may restrict the attention to the edges of the graph. Let $|H| = |E'|$ denote the *size* of H , and let $w(H) = w(E') = \sum_{e \in E'} w(e)$ denote its *weight*. The *lightness* $\Psi(H)$ of H is the ratio between the weight of H and the weight of an MST for G , i.e., $\Psi(H) = \frac{w(H)}{w(\text{MST}(G))}$. (Throughout the paper all logarithms are in base 2.)

The result of Chechik and Wulff-Nilsen [13] is summarized in the following theorem.

Theorem 1 ([13]) *For every weighted n -vertex graph $G = (V, E, w)$ and parameters $k \geq 1$ and $0 < \epsilon < 1$, there exists a $(2k - 1) \cdot (1 + \epsilon)$ -spanner with $O(n^{1+1/k})$ edges and lightness $O(n^{1/k}(1/\epsilon)^{3+2/k})$. Such a spanner can be constructed in polynomial time.*

2.1 Doubling metrics

The following lemma gives the standard packing property of doubling metrics (see, e.g., [39]).

LEMMA 1. *Let (M, δ) be a metric space with doubling dimension ddim . If $S \subseteq M$ is a subset of points with minimum interpoint distance r that is contained in a ball of radius R , then $|S| \leq \left(\frac{2R}{r}\right)^{O(\text{ddim})}$.*

The following theorem states that any doubling metric admits a constant degree $(1 + \epsilon)$ -spanner.

Theorem 2 ([10, 33]) For any n -point metric space (M, δ) with doubling dimension ddim and parameter $0 < \epsilon < 1/2$, there exists a $(1 + \epsilon)$ -spanner with degree $\epsilon^{-O(\text{ddim})}$. The runtime of this construction is $\epsilon^{-O(\text{ddim})}(n \log n)$.

The result of Gottlieb [31] is summarized in the following theorem.

Theorem 3 ([31]) For any n -point metric space (M, δ) with doubling dimension ddim and parameter $0 < \epsilon < 1/2$, there exists a $(1 + \epsilon)$ -spanner with lightness $(\text{ddim}/\epsilon)^{O(\text{ddim})}$. The runtime of this construction is $(\text{ddim}/\epsilon)^{O(\text{ddim})}(n \log^2 n)$.

2.2 The Greedy Spanner and its Basic Properties

Algorithm 1 Greedy($G = (V, E, w), t$)

```

1:  $H = (V, \emptyset, w)$ .
2: for each edge  $(u, v) \in E$ , in non-decreasing order of weight,
   do
3:   if  $\delta_H(u, v) > t \cdot w(u, v)$  then
4:     Add the edge  $(u, v)$  to  $E(H)$ .
5:   end if
6: end for

```

The greedy spanner algorithm is presented in Algorithm 1. Let $H = (V, E_H, w)$ be the output of an arbitrary execution of the greedy algorithm with stretch parameter t . It is immediate that H has stretch at most t . If the edge weights in the graph are distinct, then H is uniquely defined, but this does not hold in general; nevertheless, by letting H designate an arbitrary such spanner, we may henceforth refer to it as the *greedy t -spanner*. The following observation is immediate (see, e.g., [23, 13]).

OBSERVATION 2. H contains all edges of some MST of G , denoted Z . (Hence Z is also an MST of H .)

3. THE BASIC OPTIMALITY PROOF

In this section we show that the greedy spanner is existentially optimal, with respect to both the size and the lightness, for any graph family that is closed under edge removal.

We start by making the basic observation that the only t -spanner of the greedy t -spanner is itself.

LEMMA 3. Let $G = (V, E, w)$ be any weighted graph, let $t \geq 1$ be any stretch parameter, and let H be the greedy t -spanner of G . If H' is a t -spanner for H , then $H' = H$.

PROOF. Assume for contradiction that H' is a t -spanner for H yet there is an edge $e \in H \setminus H'$. Let P be a shortest path in H' between the endpoints of e . As H' is a t -spanner of H , it holds that $w(P) \leq t \cdot w(e)$. Consider the last edge examined by the greedy algorithm among the edges of P and e , denoted e' . By the description of the greedy algorithm, we have $w(e) \leq w(e')$. Consequently, by the time the greedy algorithm examines edge e' , all the edges of the path $(P \cup e) \setminus e'$ must have already been added to the greedy spanner. (See Figure 2 for an illustration.) This path connects the endpoints of e' , and its weight is given by

$$w(P) - w(e') + w(e) \leq w(P) \leq t \cdot w(e) \leq t \cdot w(e').$$

Hence the greedy algorithm will not add edge e' to H , a contradiction. \square

Equipped with Lemma 3, we now turn to the basic optimality proof.




Figure 2: The path P in H' between the endpoints of edge e is depicted by a dashed line. The path $P \cup e \setminus e'$ between the endpoints of edge e' , all edges of which have been added to H by the time the greedy algorithm examines edge e' , is colored red.

Theorem 4 (Greedy is optimal) Let \mathcal{G} be any family of n -vertex graphs that is closed under edge removal, and let $t = t(n)$ be any stretch parameter. Assume that for every graph $G \in \mathcal{G}$, there exists a t -spanner for G with at most $m(n, t)$ edges and lightness at most $l(n, t)$. Then for every graph $G \in \mathcal{G}$, the greedy t -spanner H of G has at most $m(n, t)$ edges and lightness at most $l(n, t)$.

PROOF. Consider an arbitrary graph G in \mathcal{G} , and let H be the greedy t -spanner of G . Since \mathcal{G} is closed under edge removal and H is a subgraph of G , H belongs to \mathcal{G} . Hence, there exists a t -spanner \mathcal{H} of H with at most $m(n, t)$ edges and lightness at most $l(n, t)$. Lemma 3 implies that $\mathcal{H} = H$, from which the size bound on H immediately follows. The lightness bound is slightly trickier, as the spanner \mathcal{H} is computed on top of the greedy spanner H rather than the original graph G . Nevertheless, Observation 2 implies that G and H have the same MST Z . Since the lightness of \mathcal{H} is at most $l(n, t)$ and Z is an MST for H , it follows that

$$l(n, t) \geq \Psi(\mathcal{H}) = \frac{w(\mathcal{H})}{w(\text{MST}(H))} = \frac{w(\mathcal{H})}{w(Z)},$$

hence $w(\mathcal{H}) \leq l(n, t) \cdot w(Z)$. Using the fact that $\mathcal{H} = H$, we conclude that the lightness of H satisfies

$$\begin{aligned} \Psi(H) &= \frac{w(H)}{w(\text{MST}(G))} = \frac{w(H)}{w(Z)} \\ &= \frac{w(\mathcal{H})}{w(Z)} \leq \frac{l(n, t) \cdot w(Z)}{w(Z)} = l(n, t). \end{aligned}$$

\square

As the family of weighted graphs is closed under edge removal, we can apply Theorem 4 on it. Hence the greedy spanner for general graphs has size and lightness at least as good as in Theorem 1.

COROLLARY 4. For every weighted graph $G = (V, E, w)$ on n vertices and m edges and parameters $k \geq 1$ and $0 < \epsilon < 1$, the greedy $(2k - 1) \cdot (1 + \epsilon)$ -spanner has $O(n^{1+1/k})$ edges and lightness $O(n^{1/k}(1/\epsilon)^{3+2/k})$. (A naive implementation of the greedy algorithm requires $O(mn^{1+1/k})$ time.)

In [5] it was proved that for any parameter $0 < \delta < 1$ and any stretch parameter $t = t(n)$, if every n -vertex weighted graph admits a t -spanner with at most $m(n, t)$ edges and lightness at most $l(n, t)$, then for every such graph there also exists a t/δ -spanner with at most $m(n, t)$ edges and lightness at most $1 + \delta \cdot l(n, t)$.² Applying Theorem 4 again, we derive the following result.

²This reduction appears in the full version of [5], which is currently available via <http://www.cs.bgu.ac.il/~arnoldf/BFN16.pdf>.

COROLLARY 5. *For every weighted n -vertex graph $G = (V, E, w)$ and parameter $0 < \delta < 1$, the greedy $O(\log n/\delta)$ -spanner has $O(n)$ edges and lightness at most $1 + \delta$.*

As mentioned in the introduction, a plethora of graph families that are closed under edge removal were studied extensively in the spanner literature. This includes the families of planar graphs, bounded genus graphs, bounded treewidth graphs, graphs excluding fixed minors, and more. For all these graph families, Theorem 4 shows that the greedy spanner is existentially optimal.

4. THE OPTIMALITY ARGUMENT IN DOUBLING METRICS

The basic optimality argument of Section 3 applies to graph families that are closed under edge removal. Note that metric spaces do not fall into this category. Nevertheless, for metric spaces, the basic optimality argument suffices: On the one hand, the upper bound for general weighted graphs applies to any metric space, and on the other hand, the lower bound due to high girth graphs naturally applies to the induced metric spaces (see, e.g., [1, 51]).

In this section we study the optimality of the greedy spanner for *doubling metrics*. For such metric spaces, one would like to obtain spanners with stretch $1 + \epsilon$, where ϵ is arbitrarily close to 0. We will show that the greedy $(1 + \epsilon)$ -spanner is existentially near-optimal in doubling metrics, with respect to both the size and the lightness.

The next observation and subsequent lemma will be used for proving the lightness optimality.

OBSERVATION 6. *Consider the metric space M_G induced by an arbitrary weighted graph $G = (V, E, w)$. Then any MST of M_G is a spanning tree of G . (Hence there is a common MST for G and M_G , denoted Z .)*

PROOF. Consider an MST Z for M_G , and suppose for contradiction that Z contains an edge e outside G . Since e belongs to $M_G \setminus G$, any path in G between the endpoints of e consists of at least two edges. Consider the (multi) graph obtained from Z by replacing edge e with a shortest path in G between the endpoints of e . It is a spanning subgraph of M_G of weight $w(Z)$, which contains at least $n + 1$ edges (some of which may be multiple edges), and thus at least one cycle. By breaking cycles in this subgraph, we obtain a spanning tree of M_G of weight strictly smaller than $w(Z)$, yielding a contradiction to the weight minimality of Z . \square

LEMMA 7. *Let (M, δ) be any metric space, $t \geq 1$ be some stretch parameter, and H be the greedy t -spanner of M . For every t -spanner H' of the metric space M_H induced by H , we have $w(H) \leq w(H')$.*

PROOF. Let H' be a t -spanner of M_H , and define H'' as the subgraph of H obtained from H' by replacing each edge e of H' with a shortest path in H between the endpoints of e . Clearly, the distances in H'' are no greater than the respective distances in H' . Since H'' is a subgraph of H , it follows that H'' is a t -spanner for H . Lemma 3 implies that $H'' = H$. Finally, noting that $w(H'') \leq w(H')$, we have $w(H) = w(H'') \leq w(H')$. \square

The following lemma will be used for proving the size optimality.

LEMMA 8. *Let (M, δ) be any metric space, $t < 2$ be some stretch parameter, and H be the greedy t -spanner of M . For every t -spanner H' of the metric space M_H induced by H , we have $|H| \leq |H'|$.*

PROOF. For every edge $e' \in H'$, let $P_{e'}$ be a shortest path between the endpoints of e' in H . We say that edge $e' \in H'$ covers all edges of $P_{e'}$, and all these edges are covered by e' . An edge $e' \in H \cap H'$ covers itself.

For each edge e in $H \setminus H'$, let Q_e be a shortest path between the endpoints of e in H' . Since H' is a t -spanner for M_H , we have $w(Q_e) \leq t \cdot w(e)$. Observe that the edges in $\cup_{e' \in Q_e} P_{e'}$ form a path Π_e in H between the endpoints of e . (It will be shown next that the path Π_e is not simple.) We have

$$w(\Pi_e) \leq \sum_{e' \in Q_e} w(P_{e'}) = \sum_{e' \in Q_e} w(e') = w(Q_e) \leq t \cdot w(e).$$

Next, we argue that the edge e must belong to Π_e . Indeed, otherwise the edges of Π_e contain a simple path in H between the endpoints of e of weight bounded by $t \cdot w(e)$, implying that the heaviest edge among the edges of this path and e would not be added to the greedy t -spanner H . Consequently, at least one edge e' in Q_e must cover e .

We define an injection $f : H \rightarrow H'$ as follows. For each edge $e \in H \cap H'$, $f(e)$ is defined as e ; in this case edge $e = f(e)$ covers itself. For each $e \in H \setminus H'$, $f(e)$ is defined to be an arbitrary edge of Q_e that covers e . To see that f is injective, suppose for contradiction the existence of two distinct edges e_1 and e_2 in H and an edge $e' \in H'$ such that $f(e_1) = f(e_2) = e' \in H'$. It must hold that e_1 and e_2 are in $H \setminus H'$. Assume without loss of generality that $w(e_1) \leq w(e_2)$. Since both e_1 and e_2 are covered by e' , it follows that $w(e') \geq w(e_1) + w(e_2) \geq 2 \cdot w(e_1)$. On the other hand, by the definition of f , the shortest path Q_{e_1} in H' between the endpoints of e_1 contains the edge $e' = f(e_1)$. Hence the weight of a shortest path in H' between the endpoints of e_1 is given by $w(Q_{e_1}) \geq w(e') \geq 2 \cdot w(e_1) > t \cdot w(e_1)$, which contradicts the fact that H' is a t -spanner for H . It follows that f is injective, from which we conclude that $|H| \leq |H'|$. \square

The following observation shows that a small “stretching” of any metric space does not change the doubling dimension of the metric space by much.

OBSERVATION 9. *Let (M, δ) be a metric space with doubling dimension ddim . Let H be a t -spanner of M , for $t \leq 2$. Then the metric space M_H induced by H has doubling dimension at most $2 \cdot \text{ddim}$.*

PROOF. Clearly, any ball of radius r in the “stretched” metric space M_H is contained in the respective ball of the original metric space M . By definition, this ball can be covered by $2^{2 \cdot \text{ddim}}$ balls of radius $\frac{r}{4}$ in M , and so by $2^{2 \cdot \text{ddim}}$ balls of radius $t \cdot \frac{r}{4} \leq \frac{r}{2}$ in the stretched metric space M_H . \square

The existential near-optimality result for doubling metrics is summarized in the following theorem.

Theorem 5 (Greedy is near-optimal in doubling metrics) *Assume that for every n -point metric space (M, δ) with doubling dimension ddim and any stretch parameter $t < 2$, there is a t -spanner with at most $m(n, \text{ddim}, t)$ edges and lightness at most $l(n, \text{ddim}, t)$. Then for every n -point metric space (M, δ) with doubling dimension ddim , the greedy t -spanner has at most $m(n, 2 \cdot \text{ddim}, t)$ edges and lightness at most $l(n, 2 \cdot \text{ddim}, t)$.*

PROOF. Let M be an arbitrary n -point metric space with doubling dimension ddim , let $H = (V, E)$ be the greedy t -spanner for M , and let M_H be the metric space induced by H . By Observation 9, the doubling dimension of M_H is bounded by $2 \cdot \text{ddim}$. Our assumption implies that there exists a t -spanner \mathcal{H} for M_H with at

most $m(n, 2 \cdot \text{ddim}, t)$ edges and lightness at most $l(n, 2 \cdot \text{ddim}, t)$. Lemma 8 implies that $|H| \leq |\mathcal{H}|$, from which the size bound on H immediately follows. As for the lightness bound on H , note that \mathcal{H} is computed on top of M_H rather than the original metric space M . Nevertheless, Observation 2 and Observation 6 imply that M and M_H have the same MST Z . Since the lightness of \mathcal{H} is at most $l(n, 2 \cdot \text{ddim}, t)$, it follows that

$$l(n, 2 \cdot \text{ddim}, t) \geq \Psi(\mathcal{H}) = \frac{w(\mathcal{H})}{w(\text{MST}(M_H))} = \frac{w(\mathcal{H})}{w(Z)},$$

hence $w(\mathcal{H}) \leq l(n, 2 \cdot \text{ddim}, t) \cdot w(Z)$. By Lemma 7, we have $w(H) \leq w(\mathcal{H})$, hence the lightness of H satisfies

$$\begin{aligned} \Psi(H) &= \frac{w(H)}{w(\text{MST}(M))} = \frac{w(H)}{w(Z)} \leq \frac{w(\mathcal{H})}{w(Z)} \\ &\leq \frac{l(n, 2 \cdot \text{ddim}, t) \cdot w(Z)}{w(Z)} = l(n, 2 \cdot \text{ddim}, t). \end{aligned}$$

□

It is known that the greedy $(1 + \epsilon)$ -spanner for n -point doubling metrics has $O(n)$ edges and lightness $O(\log n)$ [57], where the O -notation hides a multiplicative term of $(1/\epsilon)^{O(\text{ddim})}$. Applying Theorem 5 in conjunction with the result of Theorem 3, we reduce the lightness bound of the greedy spanner to constant.

COROLLARY 10. *For every metric space (M, δ) with doubling dimension ddim and any parameter $0 < \epsilon < \frac{1}{2}$, the greedy $(1 + \epsilon)$ -spanner has $n(1/\epsilon)^{O(\text{ddim})}$ edges and lightness $(\text{ddim}/\epsilon)^{O(\text{ddim})}$.*

Remark. Corollary 10 shows that the greedy $(1 + \epsilon)$ -spanner in doubling metrics achieves optimal bounds on the size and the lightness, *disregarding dependencies on ϵ and the doubling dimension*. However, improving these dependencies is a fundamental challenge of practical importance. By Theorem 5, any improvement whatsoever in the dependencies on ϵ and the doubling dimension on either the size or the lightness of *any spanner construction* for doubling metrics – would trigger a similar improvement to the greedy spanner.

5. THE APPROXIMATE-GREEDY SPANNER IN DOUBLING METRICS IS LIGHT

Corollary 10 shows that the greedy $(1 + \epsilon)$ -spanner in doubling metrics achieves near-optimal bounds on the size and the lightness. Nevertheless, this spanner has two major disadvantages. First, as mentioned in the introduction, there exist metric spaces with doubling dimension 1 for which its degree may be unbounded. (This is in contrast to d -dimensional Euclidean metrics, where the greedy $(1 + \epsilon)$ -spanner has degree $\epsilon^{-O(d)}$.) Second, it cannot be constructed within sub-quadratic time in doubling metrics due to a lower bound of [41]. In fact, even in d -dimensional Euclidean metrics, the state-of-the-art implementation of the greedy $(1 + \epsilon)$ -spanner requires time $\epsilon^{-O(d)}(n^2 \log n)$ [7].

Building on [19, 20], Gudmundsson et al. [38] devised a much faster algorithm that follows the greedy approach, hereafter Algorithm `Approximate-Greedy`. The runtime of this algorithm is $\epsilon^{-O(d)}(n \log n)$, yet the degree and lightness of the approximate-greedy spanner produced by the algorithm are both bounded by $\epsilon^{-O(d)}$, just as with the greedy spanner for Euclidean metrics. The runtime analysis of Algorithm `Approximate-Greedy` [38] does not exploit any properties of Euclidean geometry. Specifically, it relies on the triangle inequality, which applies to arbitrary metric spaces, and on standard packing arguments (cf. Lemma 1), which

apply to arbitrary doubling metrics. Therefore, the runtime of Algorithm `Approximate-Greedy` remains $\epsilon^{-O(\text{ddim})}(n \log n)$ in arbitrary doubling metrics. Moreover, the degree bound of $\epsilon^{-O(d)}$ applies to arbitrary doubling metrics as well. (We refer to Chapter 15 in [45] for an excellent description of this algorithm and its analysis.)

In this section we show that the approximate-greedy spanner of [38] has constant lightness in arbitrary doubling metrics. Consequently, Algorithm `Approximate-Greedy` provides an $O(n \log n)$ -time construction of $(1 + \epsilon)$ -spanners in doubling metrics with lightness and degree both bounded by constants.

5.1 A Rough Sketch of Algorithm `Approximate-Greedy`

In this section we provide a very rough sketch of Algorithm `Approximate-Greedy`, aiming to highlight the high-level ideas behind it. This outline is not required for the analysis that is given in Section 5.2; it is provided here for clarity and completeness.

In metric spaces, the greedy algorithm sorts the $\binom{n}{2}$ interpoint distances and examines the edges by non-decreasing order of weight. For each edge that is examined for inclusion in the spanner, the distance between its endpoints in the current spanner is computed. This is expensive for two reasons: (1) The number of examined interpoint distances is quadratic in n . (2) Computing the *exact* spanner distance between two points is costly.

Suppose we aim for a stretch of $t = 1 + \epsilon$, and let t' be an appropriate parameter satisfying $t' = 1 + O(\epsilon) < t$. (Refer to [38, 45] for the exact constant hiding in the O -notation of $O(\epsilon)$.) Instead of examining all $\binom{n}{2}$ interpoint distances, Algorithm `Approximate-Greedy` computes a bounded degree $\sqrt{t/t'}$ -spanner $G' = (M, E', \delta)$ for the input metric space (M, δ) , and simulates the greedy algorithm with stretch parameter $\sqrt{t \cdot t'}$ only on the edges of G' . The output of the algorithm is a $\sqrt{t \cdot t'}$ -spanner $G = (M, E, \delta)$ for G' , which is a t -spanner for the original metric space (M, δ) by the “transitivity” of spanners. A spanner G' of degree $\epsilon^{-O(\text{ddim})}$ can be constructed in $\epsilon^{-O(\text{ddim})}(n \log n)$ time via Theorem 2. Since the output t -spanner G for (M, δ) is a subgraph of G' , its degree will be at most $\epsilon^{-O(\text{ddim})}$.

The greedy simulation is applied only on the edges of G' that are sufficiently “heavy”. Formally, let D denote the maximum weight of any edge of the bounded degree spanner G' , and let E_0 be the set of *light edges* in E' , namely, of weight at most D/n . As $|E_0| \leq |E'| = O(n)$, we have $w(E_0) = O(D) = O(\text{MST}(M))$. All light edges are taken to the output spanner G , and the greedy simulation is applied only on the edges of $E' \setminus E_0$. (So the output spanner G will contain *all* edges of E_0 and *some* edges of $E' \setminus E_0$.)

As mentioned, computing the *exact* distance between two points is costly; using Dijkstra’s algorithm, it requires $O(n \log n)$ time (see, e.g., Section 2.5 and Corollary 2.5.10 in [45]). Since G' has $O(n)$ edges, the overall runtime will be $O(n^2 \log n)$. To speed up the computation time, Algorithm `Approximate-Greedy` does not compute the exact distance between two points, but rather an approximation of that distance. This is achieved by maintaining a much simpler and coarser *cluster graph* that approximates the original distances, on which the distance queries are performed. More specifically, the algorithm partitions the edge set $E' \setminus E_0$ into $\log_\mu n$ buckets, for an appropriate parameter $1 < \mu = O(\log n)$, such that edge weights within each bucket differ by at most a factor of μ . Then it examines the edges of $E' \setminus E_0$ by going from one bucket to the next, examining edges by non-decreasing order of weight. Whenever all edges of some bucket have been examined, the cluster graph is updated according to the new edges that were added to the spanner. The idea is to periodically make the cluster

graph simpler and coarser, so that the shortest path computations made on it will be fast. The bottom-line is that one does not simulate the greedy algorithm (with stretch parameter $\sqrt{t \cdot t'}$) on the edge set $E' \setminus E_0$, but rather an *approximate version* of it.

5.2 Bounding the Lightness of the Approximate-Greedy Spanner

As mentioned, the runtime of Algorithm `Approximate-Greedy` is $\epsilon^{-O(\text{ddim})}(n \log n)$ in arbitrary doubling metrics. In what follows let $G = (M, E, \delta)$ be the t -spanner for (M, δ) returned by Algorithm `Approximate-Greedy`. Since G is a subgraph of the bounded degree spanner $G' = (M, E', \delta)$, its degree is $\epsilon^{-O(\text{ddim})}$.

It remains to bound the lightness of G . The lightness argument of [38], which relies on previous works [19, 20], is based on rather deep properties from Euclidean geometry, most notably the *leapfrog property*. In particular, this argument does not apply to arbitrary doubling metrics.

Instead, we employ the following lemma, which lies at the heart of the lightness analysis of [38]. While this lemma applies to arbitrary doubling metrics, the way it was used in [38] does not extend to arbitrary doubling metrics. Specifically, it was used in [38] to show that the edge set $E \setminus E_0$ satisfies the leapfrog property. (Recall that E_0 is the set of light edges in $G' = (M, E', \delta)$, all of which are taken to the approximate-greedy spanner $G = (M, E, \delta)$.) In Euclidean metrics, it has been proved [19, 45] that any edge set satisfying the leapfrog property has constant lightness, but this proof does not carry over to arbitrary doubling metrics.

Lemma 11 (Lemma 17 in [38]) *Let $e = (u, v) \in E \setminus E_0$. The weight of the second shortest path between u and v in the approximate-greedy spanner G is greater than $t' \cdot w(e)$. (If there are multiple shortest paths between u and v , then the weight of the second shortest path equals the weight of the shortest path.)*

Remark. The parameter t' in the statement of this lemma depends on the stretch parameters of the spanners G' and G that are constructed by Algorithm `Approximate-Greedy`. Specifically, recall that the output spanner G is a $\sqrt{t \cdot t'}$ -spanner for G' , which is, in turn, a $\sqrt{t/t'}$ -spanner for the input metric space M .

We make the following simple observation.

OBSERVATION 12. *Let H be an arbitrary weighted graph, and let t be any stretch parameter. For any t -spanner H' of H , $w(\text{MST}(H')) \leq t \cdot w(\text{MST}(H))$.*

PROOF. Consider an MST Z for H . Replace each edge of Z by a t -spanner path in H' between the endpoints of that edge, and then break cycles. The resulting structure Z' is a spanning tree of H' , hence $w(\text{MST}(H')) \leq w(Z')$, and we have $w(\text{MST}(H')) \leq w(Z') \leq t \cdot w(Z) = t \cdot w(\text{MST}(H))$. \square

The following lemma bounds the lightness of G . Its proof is based on the somewhat surprising observation that the lightness of the t -spanner G produced by Algorithm `Approximate-Greedy` is existentially near-optimal with respect to stretch parameter $t' < t$ (rather than t). We remark that G is not a greedy spanner, but rather an approximate-greedy spanner, and it is inherently different than the greedy t -spanner and the greedy t' -spanner. In particular, its weight may be larger than the weights of both these greedy spanners. Nevertheless, our existential near-optimality argument suffices to derive the required lightness bound.

LEMMA 13. *The lightness of G is $(\frac{\text{ddim}}{t'-1})^{O(\text{ddim})}$.*

PROOF. Recall that $G = (M, E, \delta)$ is a t -spanner for M , where $t = 1 + \epsilon$, $\epsilon < 1$, and let M_G be the metric space induced by G . By Observation 9, the doubling dimension of M_G is bounded by $2 \cdot \text{ddim}$. Let H' be the t' -spanner of M_G with lightness $(\text{ddim}/(t'-1))^{O(\text{ddim})}$ that is guaranteed by Theorem 3, where $t' = 1 + O(\epsilon) < t$ is the parameter appearing in the statement of Lemma 11, which is optimized as part of Algorithm `Approximate-Greedy`. As in the proof of Lemma 7, we transform H' into a t' -spanner H'' of G of weight at most $w(H')$. By Observation 6 and Observation 12, the MST weights for all graphs M, G, M_G, H' and H'' are the same, up to a factor of $t \cdot t' = O(1)$.

We argue that every edge $e \in E \setminus E_0$ belongs to H'' . Suppose for contradiction that there is an edge $e \in E \setminus E_0$ that does not belong to H'' . Let P be a shortest path between the endpoints of e in H'' . Since H'' is a t' -spanner of G , we have $w(P) \leq t' \cdot w(e)$. Note that this path is contained in G . Since $e \in G$ and M is a metric space, the weight of the second shortest path between u and v is at most $w(P) \leq t' \cdot w(e)$. On the other hand, By Lemma 11, the weight of this path is greater than $t' \cdot w(e)$, a contradiction. It follows that

$$\begin{aligned} w(G) &= w(E \setminus E_0) + w(E_0) \\ &\leq w(H'') + w(E_0) \\ &\leq w(H') + w(E_0) \\ &= \left(\frac{\text{ddim}}{t'-1}\right)^{O(\text{ddim})} \cdot w(\text{MST}(M)). \end{aligned}$$

\square

Setting $t = 1 + \epsilon$ and $t' = 1 + c \cdot \epsilon$ (for an appropriate constant c ; see [38, 45]), we conclude:

THEOREM 6. *For any metric space (M, δ) with doubling dimension ddim and parameter $0 < \epsilon < \frac{1}{2}$, Algorithm `Approximate-Greedy` returns a $(1+\epsilon)$ -spanner with lightness $(\frac{\text{ddim}}{\epsilon})^{O(\text{ddim})}$ and degree $\epsilon^{-O(\text{ddim})}$. The runtime of Algorithm `Approximate-Greedy` is $\epsilon^{-O(\text{ddim})}(n \log n)$.*

Remark. Theorem 6 should be compared to Theorem 3 due to [31]. Both constructions achieve the same lightness bound, but the degree and number of edges in the spanner construction of [31] are unbounded. Moreover, the runtime of the construction of [31] is $(\text{ddim}/\epsilon)^{O(\text{ddim})}(n \log^2 n)$, whereas that of Theorem 6 is $\epsilon^{-O(\text{ddim})}(n \log n)$. By combining the light spanner H_1 of [31] with a bounded degree spanner H_2 , one can obtain a spanner with constant degree and lightness. Specifically, such a spanner \mathcal{H} is obtained by replacing each edge of H_1 with a shortest (or approximately shortest) path in H_2 between the endpoints of that edge. The lightness of the resulting spanner \mathcal{H} will not exceed that of H_1 by much, whereas the degree bound will follow from that of H_2 . There is a major problem with this approach: The runtime needed for computing spanner \mathcal{H} may be very high. Indeed, although there are efficient ways to estimate the weight of an approximately shortest path in H_2 between two points, we must *compute* the corresponding path in H_2 . In particular, to achieve the degree bound of H_2 , one may not use edges outside H_2 . Moreover, even regardless of this computation time, such a path may contain many edges that already belong to the gradually growing spanner \mathcal{H} . Deciding which edges of this path should be added to \mathcal{H} may be very costly by itself.

Acknowledgements

We are grateful to Michael Elkin and Ofer Neiman for fruitful discussions.

6. REFERENCES

- [1] I. Althöfer, G. Das, D. P. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.
- [2] B. Awerbuch. Communication-time trade-offs in network synchronization. In *Proc. of 4th PODC*, pages 272–276, 1985.
- [3] B. Awerbuch, A. Baratz, and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proc. of 9th PODC*, pages 177–187, 1990.
- [4] B. Awerbuch, A. Baratz, and D. Peleg. Efficient broadcast and light-weight spanners. *Manuscript*, 1991.
- [5] Y. Bartal, A. Filtser, and O. Neiman. On notions of distortion and an almost minimum spanning tree with constant average distortion. In *Proc. of 27th SODA*, pages 873–882, 2016.
- [6] Y. Ben-Shimol, A. Dvir, and M. Segal. SPLAST: a novel approach for multicasting in mobile wireless ad hoc networks. In *Proc. of 15th PIMRC*, pages 1011–1015, 2004.
- [7] P. Bose, P. Carmi, M. Farshi, A. Maheshwari, and M. H. M. Smid. Computing the greedy spanner in near-quadratic time. *Algorithmica*, 58(3):711–729, 2010.
- [8] R. Braynard, D. Kostic, A. Rodriguez, J. Chase, and A. Vahdat. Opus: an overlay peer utility service. In *Prof. of 5th OPENARCH*, 2002.
- [9] H. T.-H. Chan and A. Gupta. Small hop-diameter sparse spanners for doubling metrics. In *Proc. of 17th SODA*, pages 70–78, 2006.
- [10] H. T.-H. Chan, A. Gupta, B. M. Maggs, and S. Zhou. On hierarchical routing in doubling metrics. In *Proc. of 16th SODA*, pages 762–771, 2005.
- [11] T.-H. H. Chan, M. Li, L. Ning, and S. Solomon. New doubling spanners: Better and simpler. In *Proc. of 40th ICALP*, pages 315–327, 2013.
- [12] B. Chandra, G. Das, G. Narasimhan, and J. Soares. New sparseness results on graph spanners. In *Proc. of 8th SOCG*, pages 192–201, 1992.
- [13] S. Chechik and C. Wulff-Nilsen. Near-optimal light spanners. In *Proc. of 27th SODA*, pages 883–892, 2016.
- [14] L. P. Chew. There is a planar graph almost as good as the complete graph. In *Proc. of 2nd SOCG*, pages 169–177, 1986.
- [15] E. Cohen. Fast algorithms for constructing t -spanners and paths with stretch t . In *Proc. of 34th FOCS*, pages 648–658, 1993.
- [16] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Performance-driven global routing for cell based ics. In *Proc. of 9th ICCD*, pages 170–173, 1991.
- [17] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good algorithms for performance-driven global routing. In *Proc. of 5th ISCAS*, pages 2240–2243, 1992.
- [18] J. Cong, A. B. Kahng, G. Robins, M. Sarrafzadeh, and C. K. Wong. Provably good performance-driven global routing. *IEEE Trans. on CAD of Integrated Circuits and Sys.*, 11(6):739–752, 1992.
- [19] G. Das, P. J. Heffernan, and G. Narasimhan. Optimally sparse spanners in 3-dimensional Euclidean space. In *Proc. of 9th SOCG*, pages 53–62, 1993.
- [20] G. Das and G. Narasimhan. A fast algorithm for constructing sparse Euclidean spanners. *Int. J. Comput. Geometry Appl.*, 7(4):297–315, 1997.
- [21] E. D. Demaine, M. Hajiaghayi, and B. Mohar. Approximation algorithms via contraction decomposition. *Combinatorica*, 30(5):533–552, 2010.
- [22] M. Elkin. Computing almost shortest paths. *ACM Transactions on Algorithms*, 1(2):283–323, 2005.
- [23] M. Elkin, O. Neiman, and S. Solomon. Light spanners. In *Proc. of 41th ICALP*, pages 442–452, 2014.
- [24] M. Elkin and S. Solomon. Fast constructions of light-weight spanners for general graphs. In *Proc. of 24th SODA*, pages 513–525, 2013.
- [25] M. Elkin and J. Zhang. Efficient algorithms for constructing $(1+\epsilon, \beta)$ -spanners in the distributed and streaming models. *Distributed Computing*, 18(5):375–385, 2006.
- [26] P. Erdős. Extremal problems in graph theory. In *Proc. of Sympos. Smolenice*, pages 29–36, 1964.
- [27] M. Farshi. *A theoretical and experimental study of geometric networks*. PhD thesis, 2008.
- [28] M. Farshi and J. Gudmundsson. Experimental study of geometric t -spanners. In *Proc. of 13th ESA*, pages 556–567, 2005.
- [29] J. Feigenbaum, S. Kannan, A. McGregor, S. Suri, and J. Zhang. Graph distances in the streaming model: the value of space. In *Proc. of 16th SODA*, pages 745–754, 2005.
- [30] J. Gao, L. J. Guibas, and A. Nguyen. Deformable spanners and applications. In *Proc. of 20th SoCG*, pages 190–199, 2004.
- [31] L. Gottlieb. A light metric spanner. In *Proc. of 56th FOCS*, pages 759–772, 2015.
- [32] L. Gottlieb and L. Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proc. of 19th SODA*, pages 591–600, 2008.
- [33] L. Gottlieb and L. Roditty. Improved algorithms for fully dynamic geometric spanners and geometric routing. In *Proc. of 19th SODA*, pages 591–600, 2008.
- [34] L. Gottlieb and L. Roditty. An optimal dynamic spanner for doubling metric spaces. In *Proc. of 16th ESA*, pages 478–489, 2008.
- [35] M. Grigni. Approximate TSP in graphs with forbidden minors. In *Proc. of 27th ICALP*, pages 869–877, 2000.
- [36] M. Grigni and H. Hung. Light spanners in bounded pathwidth graphs. In *Proc. of 37th MFCS*, pages 467–477, 2012.
- [37] M. Grigni and P. Sissokho. Light spanners and approximate TSP in weighted graphs with forbidden minors. In *Proc. of 13th SODA*, pages 852–857, 2002.
- [38] J. Gudmundsson, C. Levcopoulos, and G. Narasimhan. Fast greedy algorithms for constructing sparse geometric spanners. *SIAM J. Comput.*, 31(5):1479–1500, 2002.
- [39] A. Gupta, R. Krauthgamer, and J. R. Lee. Bounded geometries, fractals, and low-distortion embeddings. In *Proc. of 44th FOCS*, pages 534–543, 2003.
- [40] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [41] S. Har-Peled and M. Mendel. Fast construction of nets in low-dimensional metrics and their applications. *SIAM J. Comput.*, 35(5):1148–1184, 2006.
- [42] H. Hung. Light spanner and monotone tree. *CoRR*, abs/1207.3807, 2012.
- [43] P. N. Klein. A linear-time approximation scheme for planar weighted TSP. In *Proc. of 46th FOCS*, pages 647–657, 2005.

- [44] D. Kostic and A. Vahdat. Latency versus cost optimizations in hierarchical overlay networks. *Technical report, Duke University*, (CS-2001-04), 2002.
- [45] G. Narasimhan and M. Smid. *Geometric Spanner Networks*. Cambridge University Press, 2007.
- [46] D. Peleg. Proximity-preserving labeling schemes and their applications. In *Proc. of 25th WG*, pages 30–41, 1999.
- [47] D. Peleg. *Distributed Computing: A Locality-Sensitive Approach*. SIAM, Philadelphia, PA, 2000.
- [48] D. Peleg and A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.
- [49] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.
- [50] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.
- [51] Y. Rabinovich and R. Raz. Lower bounds on the distortion of embedding finite metric spaces in graphs. *Discrete & Computational Geometry*, 19(1):79–94, 1998.
- [52] L. Roditty. Fully dynamic geometric spanners. In *Proc. of 23rd SoCG*, pages 373–380, 2007.
- [53] L. Roditty, M. Thorup, and U. Zwick. Deterministic constructions of approximate distance oracles and spanners. In *Proc. of 32nd ICALP*, pages 261–272, 2005.
- [54] L. Roditty and U. Zwick. On dynamic shortest paths problems. In *Proc. of 32nd ESA*, pages 580–591, 2004.
- [55] F. S. Salman, J. Cheriyan, R. Ravi, and S. Subramanian. Approximating the single-sink link-installation problem in network design. *SIAM Journal on Optimization*, 11(3):595–610, 2001.
- [56] H. Shpungin and M. Segal. Near optimal multicriteria spanner constructions in wireless ad-hoc networks. *IEEE/ACM Transactions on Networking*, 18(6):1963–1976, 2010.
- [57] M. H. M. Smid. The weak gap property in metric spaces of bounded doubling dimension. In *Proc. of Efficient Algorithms*, pages 275–289, 2009.
- [58] S. Solomon. From hierarchical partitions to hierarchical covers: optimal fault-tolerant spanners for doubling metrics. In *Proc. of 46th STOC*, pages 363–372, 2014.
- [59] M. Thorup and U. Zwick. Approximate distance oracles. In *Proc. of 33rd STOC*, pages 183–192, 2001.
- [60] M. Thorup and U. Zwick. Compact routing schemes. In *Proc. of 13th SPAA*, pages 1–10, 2001.
- [61] J. Vogel, J. Widmer, D. Farin, M. Mauve, and W. Effelsberg. Priority-based distribution trees for application-level multicast. In *Proc. of 2nd NETGAMES*, pages 148–157, 2003.
- [62] P. von Rickenbach and R. Wattenhofer. Gathering correlated data in sensor networks. In *Proc. of DIALM-POMC*, pages 60–66, 2004.
- [63] B. Y. Wu, K. Chao, and C. Y. Tang. Light graphs with small routing cost. *Networks*, 39(3):130–138, 2002.