

The 4/3 Additive Spanner Exponent Is Tight

AMIR ABOUD and GREG BODWIN, Stanford University

A spanner is a sparse subgraph that approximately preserves the pairwise distances of the original graph. It is well known that there is a smooth tradeoff between the sparsity of a spanner and the quality of its approximation, so long as distance error is measured *multiplicatively*. A central open question in the field is to prove or disprove whether such a tradeoff exists also in the regime of *additive* error. That is, is it true that for all $\varepsilon > 0$, there is a constant k_ε such that every graph has a spanner on $O(n^{1+\varepsilon})$ edges that preserves its pairwise distances up to $+k_\varepsilon$? Previous lower bounds are consistent with a positive resolution to this question, while previous upper bounds exhibit the beginning of a tradeoff curve: All graphs have $+2$ spanners on $O(n^{3/2})$ edges, $+4$ spanners on $\tilde{O}(n^{7/5})$ edges, and $+6$ spanners on $O(n^{4/3})$ edges. However, progress has mysteriously halted at the $n^{4/3}$ bound, and despite significant effort from the community, the question has remained open for all $0 < \varepsilon < 1/3$.

Our main result is a surprising negative resolution of the open question, even in a highly generalized setting. We show a new information theoretic *incompressibility* bound: There is no function that compresses graphs into $O(n^{4/3-\varepsilon})$ bits so distance information can be recovered within $+n^{o(1)}$ error. As a special case of our theorem, we get a *tight* lower bound on the sparsity of additive spanners: the $+6$ spanner on $O(n^{4/3})$ edges cannot be improved in the exponent, even if any *subpolynomial* amount of additive error is allowed. Our theorem implies new lower bounds for related objects as well; for example, the 20-year-old $+4$ emulator on $O(n^{4/3})$ edges also cannot be improved in the exponent unless the error allowance is polynomial.

Central to our construction is a new type of graph product, which we call the *Obstacle Product*. Intuitively, it takes two graphs G, H and produces a new graph $G \otimes H$ whose shortest paths structure looks locally like H but globally like G .

CCS Concepts: • **Mathematics of computing** → **Extremal graph theory**; *Paths and connectivity problems*;

Additional Key Words and Phrases: Graph theory, shortest paths, spanners, lower bounds, graph sketching

ACM Reference format:

Amir Abboud and Greg Bodwin. 2017. The 4/3 Additive Spanner Exponent Is Tight. *J. ACM* 64, 4, Article 28 (September 2017), 20 pages.
<https://doi.org/10.1145/3088511>

1 INTRODUCTION

A *spanner* of a graph is a sparse subgraph that approximately preserves the distances of the original graph. Spanners were introduced by Peleg and Schäffer (1989) after they naturally arose in different contexts in the late 1980s (Awerbuch 1985; Peleg and Ullman 1989; Althöfer et al. 1993). Today, spanners are an indispensable tool in many well-studied fields, such as graph compression,

A preliminary version of this article appeared in STOC 2016.

This work was supported by NSF Grants CCF-1417238, CCF-1528078 and CCF-1514339, and BSF Grant BSF:2012338. This work was done (in part) while the authors were visiting the Simons Institute for the Theory of Computing.

Authors' addresses: A. Abboud and G. Bodwin, Gates Computer Science Building, 353 Serra Mall, Stanford, CA 94305; emails: {abboud, gbodwin}@cs.stanford.edu.

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.

© 2017 ACM 0004-5411/2017/09-ART28 \$15.00

<https://doi.org/10.1145/3088511>

synchronization in distributed networks, routing schemes, approximation algorithms for all pairs shortest paths, and more.

The early results on the topic were mostly about multiplicative t -spanners, which are subgraphs that preserve all distances up to a multiplicative constant t . A landmark upper bound result states that every graph has a multiplicative $(2t - 1)$ -spanner on $O(n^{1+1/t})$ edges for all positive integers t (Althöfer et al. 1993). A well-known lower bound argument shows that this tradeoff is tight under the popular Girth Conjecture of Erdős (1964). Graphs on $\Omega(n^{1+1/t})$ edges with shortest cycles length $2t + 2$ (which exist under the conjecture) cannot be sparsified at all without increasing a distance of a removed edge from 1 to $2t + 1$. Therefore, the picture for multiplicative spanners is essentially complete.

A $+k$ additive spanner of a graph G is a subgraph H that preserves distances up to an additive constant k , that is, for any two nodes u, v in G , we have $\text{dist}_G(u, v) \leq \text{dist}_H(u, v) + k$. In most contexts, additive error is more practically appealing than multiplicative error: A detour that adds two minutes to a trip is preferable to one that triples its length. At first, this seemed hard to achieve; nontrivial additive spanners were only known for special types of graphs (Liestman and Shermer 1991, 1993). However, the seminal article of Aingworth et al. (1996, 1999) showed the following surprise: All graphs have $+2$ additive spanners on just $\tilde{O}(n^{3/2})$ edges (see also Dor et al. (1996), Elkin and Peleg (2004), Roditty et al. (2005), Thorup and Zwick (2006), Knudsen (2014), and Baswana et al. (2010)). In other words, the multiplicative 3-spanner shares a tight sparsity bound with the additive $+2$ spanner. This result caused significant optimism in the community regarding the existence of good additive spanners, which quickly became the central objects of study: Could it be that graphs can be sparsified to near-linear sizes while incurring only a constant additive distance error?

The 20 years that followed witnessed a high throughput of results on additive spanners (see the recent Encyclopedia of Algorithms entry by Chechik (2014)), and while this puzzle has been labelled

- “major” (Thorup and Zwick 2006; Coppersmith and Elkin 2006; Chechik 2013, 2014; Bodwin and Vassilevska Williams 2015, 2016)
- “main” (Baswana et al. 2005, 2010)
- “fascinating” (Woodruff 2006)
- “the chief open question in the field” (Pettie 2009)

and has been raised repeatedly, it has remained inexplicably wide open.

OPEN QUESTION 1. Prove or disprove: For all $\varepsilon > 0$, there is a constant k_ε such that all graphs have a $+k_\varepsilon$ additive spanner on $O(n^{1+\varepsilon})$ edges.

Open Question 1 has been generally regarded with optimism. This was motivated by some interesting new positive results: Chechik (2013) showed that all graphs have $+4$ spanners on $\tilde{O}(n^{7/5})$ edges, and Baswana et al. (2005, 2010) showed that all graphs have $+6$ spanners on $O(n^{4/3})$ edges (see Woodruff (2010) and Knudsen (2014) for followup work). This was also motivated by a lack of negative results: The Girth Conjecture implies an $\Omega(n^{1+1/k})$ lower bound for $+(2k - 2)$ spanners, and the only further progress came from Woodruff (2006), who proved this lower bound unconditionally. These bounds are fully compatible with a positive resolution of Open Question 1. Perhaps the most compelling argument for optimism, though, is the intuitive one: How could it be that spanners enjoy improved sparsity from a $+2$, $+4$, or $+6$ error allowance, but then suddenly the trail goes cold and no further tradeoffs are possible? Despite all the evidence, the size upper bounds have not improved beyond $n^{4/3}$ since the mid-2000s, with no hint as to why.

Meanwhile, considerable research effort has been spent on several promising directions of attack, with the explicit stated goal of making progress on this problem. Elkin and Peleg (2004)

showed that there are near-linear size spanners with constant additive error in addition to a $(1 + \epsilon)$ multiplicative error (see also Elkin (2005), Thorup and Zwick (2006), and Baswana et al. (2010) for more work on these *mixed spanners*). Interesting spanners with non-constant additive error have been discovered (Bollobás et al. 2003; Thorup and Zwick 2006; Pettie 2009; Baswana et al. 2010; Chechik 2013; Bodwin and Vassilevska Williams 2015, 2016). Another example is the study of *pairwise distance preservers* and *pairwise spanners*, where we only require that distances between a small set of node pairs be (approximately) preserved (Bollobás et al. 2003; Coppersmith and Elkin 2006; Pettie 2009; Cygan et al. 2013; Kavitha and Varma 2013; Parter 2014; Kavitha 2015; Bodwin and Vassilevska Williams 2016; Abboud and Bodwin 2016; Bodwin 2017). A lot of progress has also been done on additive spanners for special graph classes; see Chepoi et al. (2005), Dragan et al. (2011), Dragan et al. (2012), Chepoi et al. (2012), and Dragan and Abu-Ata (2014) and the references therein. While these fields have steadily progressed, none have proven sufficient to explain or beat the $n^{4/3}$ bound, and so the puzzle of Open Question 1 has endured.

1.1 Our Results

Our first result is a negative resolution of Open Question 1, and a negative resolution of the even stronger question of the tightness of the $n^{4/3}$ threshold. We prove:

THEOREM 1.1. *For all $\epsilon > 0$, there is a $\delta > 0$ and an infinite family of n node graphs $G = (V, E)$ such that for any subgraph $H = (V, E')$ with $|E'| = O(n^{4/3-\epsilon})$, there exist nodes $u, v \in V$ with*

$$\text{dist}_H(u, v) = \text{dist}_G(u, v) + \Omega(n^\delta).$$

Thus, the edge sparsity of $n^{4/3}$ cannot be improved, even for additive error $n^{o(1)}$, and the +6 spanner of Baswana et al. on $O(n^{4/3})$ edges is essentially the sparsest possible constant error additive spanner.

Spanners are merely one (very appealing) form of graph compression. There are many other well-studied methods of compression that do not restrict the compressed form to be a subgraph of the original graph G ; examples include Distance Oracles, Sketches, Labeling Schemes, Metric Embeddings, Emulators, and so on. For example, an *emulator* is a sparse weighted graph H (not necessarily a subgraph of G) that approximately preserves the distances of G (Dor et al. 1996; Thorup and Zwick 2006; Woodruff 2006; Bodwin and Vassilevska Williams 2015).

This leads to the following information theoretic question, which is an extremely relaxed version of Open Question 1: Does there exist an algorithm for compressing a metric defined by an unweighted undirected graph into $O(n^{4/3-\epsilon})$ bits, for some $\epsilon > 0$, such that approximate distances can be recovered from the compressed bitstring within constant additive error? We further generalize Theorem 1.1 to show that our graph family is strongly *incompressible* and provide a negative answer to this question as well.

THEOREM 1.2. *For all $\epsilon > 0$, there exists a $\delta > 0$ such that there is no mapping ψ from graphs G on n nodes to bitstrings of length $O(n^{4/3-\epsilon})$ such that the distances of G can always be recovered from the bitstring $\psi(G)$ within $+n^\delta$ error.*

An interesting consequence of Theorem 1.2 is that the 20-year-old +4 additive emulator on $O(n^{4/3})$ edges of Dor et al. (1996) cannot be improved in the exponent even with any $n^{o(1)}$ additive error.

Finally, our resolution of Open Question 1 shows that polynomial additive distance error must be suffered to obtain near-linear size compression of graphs. Given this, it is natural to wonder *how much* polynomial error is necessary to obtain compression in this regime. We show:

THEOREM 1.3. *For all $\varepsilon > 0$, there exists a $\delta > 0$ such that there is no mapping ψ from graphs G on n nodes to bitstrings of length $O(n^{1+\delta})$ such that the distances of G can always be recovered from the bitstring $\psi(G)$ within $+n^{1/22-\varepsilon}$ error.*

The current best method for compression into $n^{1+o(1)}$ bits obtains $+\tilde{O}(n^{1/3})$ error (Bodwin and Vassilevska Williams 2015). We leave it as an open question to close this gap.

2 THE CONSTRUCTION

The goal of this section is to explicitly construct the graphs mentioned in our theorems. First, we offer a technical overview that gives the intuition behind our construction, which will help demystify some steps in our proof. We will also highlight the novelty of our approach over previous lower bound constructions.

2.1 Technical Overview

Suppose we seek a lower bound against $+k$ spanners (think of k as a small polynomial, like $n^{0.01}$ for now). Our general approach is to start with a graph G with the following special property: There is a set of node pairs P , with each pair at distance k , such that there is a unique shortest path in G between each pair and the edge set of G is precisely the union of these paths. Thus, any edge deletion from G will stretch one of the pairwise distances in P by at least $+1$. We then perform some transformations to G to amplify this error. The main trick here is to pick a family of “obstacle graphs” H and then perform a new kind of replacement product $G \otimes H$, which we call the Obstacle Product (OP). The effect will be that any spanner that deletes too many edges from $G \otimes H$ must increase the distance between one of the (transformed) node pairs in P by at least $+k$.

Think of G as the *outer graph* and of H as the *inner graphs*. The Obstacle Product $G \otimes H$ consists of two steps:

- (1) (*Edge Extension*) Replace every edge in G with a length k path
- (2) (*Inner Graph Replacement*) Replace each *original* node (i.e., not nodes created by the edge extension step) $v \in V(G)$ with an appropriately chosen subgraph H_v from the family of inner graphs.

Finally, for each edge $\{u, v\}$ incident on a node v , one must choose exactly one node in H_v that will serve as the new endpoint of this edge after the inner graph replacement step. We will temporarily skip over the technical detail of how to choose this node. The key outcome of the OP is that the (transformed) shortest paths for the node pairs in P must now wander through many “obstacles” H_v before reaching their final destination.

Our next step is to use a counting argument to say that for a sparse-enough subgraph $F \subseteq G \otimes H$, there exists a (transformed) pair in $\{s, t\} \in P$ such that F is missing at least one of the edges used by the shortest path $\rho_{G \otimes H}(s, t)$ in *every* obstacle graph H_v that this path visits. We can then prove that

$$\text{dist}_F(s, t) \geq \text{dist}_{G \otimes H}(s, t) + k.$$

Our proof of this claim uses two cases. Possibly (1) the shortest paths $\rho_F(s, t)$ and $\rho_{G \otimes H}(s, t)$ follow the same path in the outer graph; that is, they intersect the same set of inner graphs in the same order. In this case, because F is missing an edge used by $\rho_{G \otimes H}(s, t)$ in every inner graph, $\rho_F(s, t)$ must take a $+1$ length detour at every inner graph, for a total error of $+k$. Alternatively, (2) $\rho_F(s, t)$ and $\rho_{G \otimes H}(s, t)$ follow different paths in the outer graph; that is, $\rho_F(s, t)$ intersects a new inner graph. In this case, we argue that $\rho_F(s, t)$ must travel an extra k -length path in the outer graph to reach its final destination, and this is the source of its $+k$ error. See Figures 1 and 2 for a depiction of these two cases.




Fig. 1. Case 1: Perhaps the new shortest path $\rho_F(s, t)$ still passes through the same set of inner graphs as $\rho_{G \otimes H}(s, t)$. In this case, we use a counting argument to show that (for some pair (s, t)) the subgraph F is missing an edge used by $\rho_{G \otimes H}(s, t)$ in every inner graph that it touches. Therefore, $\rho_F(s, t)$ must take a +1 edge detour in every inner graph it touches for a total of $+k$ error.




Fig. 2. Case 2: Perhaps the new shortest path $\rho_F(s, t)$ passes through an inner graph that is *not* touched by the old shortest path $\rho_{G \otimes H}(s, t)$. In this case, we argue that $\rho_F(s, t)$ must travel an extra k -length path to reach its destination, and this is the source of its $+k$ error.

We consider the abstraction of this approach to be the biggest leap in understanding provided by our work. Several important prior constructions can be viewed within this framework: The previous lower bound constructions of Woodruff (2006) and Parter (2014) can both be viewed within the obstacle product framework, with the inner graph as a biclique and the outer graph as a certain type of “layered clique.” A previous lower bound construction by the authors (Abboud and Bodwin 2016) used the same “layered clique” outer graph but allowed for various inner graphs to be substituted in. The construction in this article is the first one that allows for the modular substitution of outer graphs. The ability to choose both inner and outer graphs turns out to be quite powerful: In addition to its central role in this article, this type of product has proved useful in followup work constructing lower bound graphs for various other problems related to sketching graph distances (Bodwin 2017; Abboud et al. 2017).

In this article, flexibility in our choice of outer graph turns out to be powerful enough that we are able to prove our tight spanner lower bounds *while only ever using cliques for our inner graphs*. With this in mind, it will simplify our article greatly to proceed with the restriction that H is a family of cliques. Let us now return to the main construction. With the simplification that all inner graphs are cliques, we will execute the second step of the obstacle product (henceforth, the *clique replacement* step) as follows: Replace each node v with a clique on $\deg(v)$ nodes, and connect each of the $\deg(v)$ edges entering v to a different clique node. Following the argument above, we now have that every subgraph F must keep at least one “clique edge” per pair in P or else it stretches the distance between one of these paths by $+k$. Assuming that no two pairs $p_1, p_2 \in P$ ever use the same clique edge, this implies that any $+(k - 1)$ additive spanner of our graph $G \otimes H$ must keep at least $|P|$ edges in total. The clique edge used by a (transformed) pair $p \in P$ in a clique H_v is determined by the 2-path that the (original) path $\rho_G(p)$ uses to enter and leave the node v . Thus, our lower




Fig. 3. If no two paths enter and leave v in the same way, then after v is replaced with a clique, all of these paths will use a different clique edge.

bound of $|P|$ on the spanner density is realized so long as the shortest paths for every pair of pairs in P is originally 2-path disjoint; that is, $\rho_G(p_1) \cap \rho_G(p_2)$ does not contain any 2-paths in G (see Figure 3).

We now have a lower bound on the number of edges that the spanner must keep; our next step is to obtain a favorable upper bound on the number of nodes in the spanner. The dominant cost here is from the edge extension step, and so the number of nodes in the spanner is roughly equal to $E(G) \cdot k$. We are now able to state the properties that we want G to have:

- (1) G is the union of many shortest paths between a set of node pairs P , with the following properties:
 - Each pair $p \in P$ has distance k
 - Each pair $p \in P$ has a unique shortest path between its endpoints
 - These shortest paths are 2-path disjoint
 Note that we want P to be as large as possible, since we will ultimately obtain a lower bound of $|P|$ on the number of edges in any $+(k-1)$ spanner.
- (2) G has as few edges as possible, because the number of nodes in the spanner (due to the Edge Extension step) will be $k \cdot |E(G)|$.

This completes the technical overview. In the next two subsections, we will describe how to obtain starting graphs G with these properties in full technical detail. After that, we will repeat the details of the obstacle product transformation more formally, and we will fully prove that this series of transformation has the claimed properties.

2.2 Starting Point

Our starting point is the following lemma:

LEMMA 2.1. *For all $\epsilon > 0$, there is a $0 < \delta < \epsilon$, and an infinite family of n node graphs $G = (V, E)$ and pair sets $P \subseteq V \times V$ with the following properties:*

- (1) *For each pair in P , there is a unique shortest path between its endpoints*
- (2) *These paths are edge disjoint*
- (3) *For all $\{s, t\} \in P$, we have $\text{dist}_G(s, t) = \Delta$ for some value $\Delta = \Theta(n^\delta)$*
- (4) *$|P| = \Theta(n^{2-\epsilon})$*

Alon (2001) constructed graphs that prove this lemma in his work on property testing. To make this article self-contained, we have included a full proof in the appendix. The lemma can also be shown by a slight modification of the work of Coppersmith and Elkin (2006) on distance preservers, who proved a version of this lemma with all properties except the third.

An additional property of the graphs from this lemma, implied by the others, is that $|E| = \Omega(n^{2-\varepsilon+\delta})$. This is *too dense* for our purposes; for technical reasons discussed in the overview, we need a sparser object to prove interesting results. The subject of the next lemma is to modify these graphs to reduce the edge count.

2.3 Path Packing

Our next move is to prove the following modification of Lemma 2.1:

LEMMA 2.2. *For all $\varepsilon > 0$, there is a $0 < \delta < \varepsilon$, and an infinite family of graphs $G = (V, E)$ and pair sets $P \subseteq V \times V$ with the following properties:*

- (1) *For each pair P in P , there is a unique shortest path between its endpoints*
- (2) *These paths are 2-path disjoint (i.e., no two paths share any pair of consecutive edges)*
- (3) *For all $\{s, t\} \in P$, we have $\text{dist}_G(s, t) = \Delta$ for some value $\Delta = \Theta(n^\delta)$*
- (4) $|P| = \Theta(n^{2-\varepsilon})$
- (5) $|E| = O(n^{3/2})$

There are two important differences between Lemmas 2.1 and 2.2: The paths have become 2-path disjoint rather than edge disjoint, and the edge count has fallen from $\approx n^2$ to $\approx n^{3/2}$. For this reason, we think of this step as *path packing*, as we are essentially packing the same number of shortest paths into many fewer edges while only slightly relaxing the paths' overlap properties.

We will prove this lemma by starting with a graph as described in Lemma 2.1 and performing a transformation to this graph to give it the properties of Lemma 2.2. With this in mind, one should *not* think of the parameters n, ε, δ in Lemma 2.2 as the same as those in Lemma 2.1; despite having the same names, they will change across the transformation.

We will now describe our transformation at a high level. Start with a graph $G = (V, E)$ and pair set $P \subseteq V \times V$ as described in Lemma 2.1. Our goal is to create a new graph $G' = (V', E')$ and pair set $P' \subseteq V' \times V'$ as described in Lemma 2.2. We achieve this by defining G' to be a product of G_1 and G_2 (which are two identical copies of G), and we define P' to be a product of P_1 and P_2 (which are two identical copies of P , although P_1 contains node pairs from G_1 and P_2 contains node pairs from G_2). Intuitively, each new pair $p \in P'$ corresponds to a pair of pairs $(p_1 \in P_1, p_2 \in P_2)$, and a walk between the endpoints of p in G' corresponds to a walk between the endpoints of p_1 in G_1 and simultaneously a walk between the endpoints of p_2 in G_2 . To enable this, each node in G' corresponds to a pair of nodes $(u \in V_{G_1}, v \in V_{G_2})$, and we carefully design the new edge set E' such that an edge in E' corresponds to a step in G_1 or a step in G_2 but not both (e.g., exactly one of the two indices will remain the same across any edge in E'). The shortest paths for P' will be forced to take alternating steps in G_1 and G_2 .

From this, we can argue 2-path disjointness of the shortest paths for P' as follows. Consider a length 2 subpath of the shortest path for some pair $p' \in P'$. These two edges correspond to a step in G_1 and a step in G_2 ; since shortest paths for P_1 in G_1 are completely edge disjoint (and the same for P_2, G_2), these edges *uniquely* determine pairs $p_1 \in P_1$ and $p_2 \in P_2$. Therefore, this information is sufficient to uniquely determine the new pair $p' = (p_1, p_2)$ being considered.

We will now describe this construction in full detail.

- *The nodes:* For each (possibly non-distinct) pair of nodes $u_1, u_2 \in V$ and index $i \in \{1, 2\}$, we add the triple (u_1, u_2, i) as a new node to G' . That is, the node set of our new graph G' will be defined as

$$V' = \{(u_1, u_2, i) \mid u_1, u_2 \in V, i \in \{1, 2\}\}.$$

Semantically, the index i dictates whether we are supposed to take a step in the graph G_1 (represented by the first coordinate) or the graph G_2 (represented by the second coordinate).

– *The edges:* To define the edges of G' we first need to define a *forwards direction* for each edge in E . These directions will be chosen so the shortest path between any pair $\{s, t\} \in P$ can be thought of as a walk that only traverses edges in their forwards direction. We do this by fixing an arbitrary ordering $s \rightsquigarrow t$ for every pair $\{s, t\} \in P$ and then directing the edges of the unique shortest path between this pair from s towards t . In other words, for any edge $\{u, v\}$ on the path, we define the $u \rightarrow v$ direction to be *forward* iff v is closer to t than it is to s . Let \vec{P} be the version of P with an arbitrary ordering imposed on each pair, and let \vec{E} be the version of E with an (ordered) edge $(u, v) \in E$ corresponding to each (unordered) edge $\{u, v\} \in E$ if $u \rightarrow v$ was defined as the forwards direction. A crucial observation here is that the forwards direction is well defined and consistent for all the edges in E , due to the property that every edge in E is on the shortest path for exactly one pair in P .

We are now ready to define the edge set E' of G' . The edges will be defined differently depending on the index value i of the node. For every node $x = (u_1, u_2, 1) \in V'$, we add an edge $\{x, y\}$ to G' if and only if $y = (u'_1, u_2, 2)$ and $(u_1, u'_1) \in \vec{E}$. Additionally, for every node $x = (u_1, u_2, 2) \in V'$, we add an edge $\{x, y\}$ to G' if and only if $y = (u_1, u'_2, 1)$ and $(u_2, u'_2) \in \vec{E}$. Intuitively, this enforces that any interesting shortest path must alternate between taking a forwards step in the first coordinate, and then a forwards step in the second coordinate. More formally, the edge set of our new graph is

$$E' := \left\{ \{(u_1, u_2, 1), (u'_1, u_2, 2)\} \mid (u_1, u'_1) \in \vec{E}, u_2 \in V \right\} \cup \left\{ \{(u_1, u_2, 2), (u_1, u'_2, 1)\} \mid (u_2, u'_2) \in \vec{E}, u_1 \in V \right\}.$$

– *The pair set:* Our new pair set P' will be composed of pairs of old pairs from P , and its definition will rely on the ordering of each pair that we fixed above. For each ordered, possibly non-distinct pair of pairs $(s_1, t_1), (s_2, t_2) \in \vec{P}$, we add the pair $\{(s_1, s_2, 1), (t_1, t_2, 1)\}$ to P' . Intuitively, given pairs $p_1, p_2 \in \vec{P}$, the corresponding new pair dictates that p_1 must be traveled in G_1 and p_2 must be traveled in G_2 . The formal definition of P' is

$$P' := \left\{ \{(s_1, s_2, 1), (t_1, t_2, 1)\} \mid (s_1, t_1), (s_2, t_2) \in \vec{P} \right\}.$$

To prove Lemma 2.2, we will first define a path for each pair in P' that alternates between taking a step in G_1 and a step in G_2 as described above.

Definition 2.3 (The Path ρ). We define

$$\rho(\{(s_1, s_2, 1), (t_1, t_2, 1)\}) := (u_1 = (s_1, s_2, 1), u_2, \dots, u_{k-1}, u_k = (t_1, t_2, 1)),$$

where

- (1) u_i for odd i is equal to $(v_{\lceil i/2 \rceil}^1, v_{\lceil i/2 \rceil}^2, 1)$, and
- (2) u_i for even i is equal to $(v_{i/2+1}^1, v_{i/2}^2, 2)$,

where v_j^1 is the j th node on the unique shortest path between s_1 and t_1 in G , and v_j^2 is the j th node on the unique shortest path between s_2 and t_2 in G .

In other words, $\rho(p')$ is built iteratively by alternatingly changing the first coordinate along the unique shortest path from s_1 to t_1 in G and then changing the second coordinate along the unique

shortest path from s_2 to t_2 in G . To see that the $\rho(p')$ is well defined for all $p' \in P'$ (i.e., the path $\rho(p')$ exists in G'), note that by construction an edge from $(v_{\lfloor i/2 \rfloor}^1, v_{\lfloor i/2 \rfloor}^2, 1)$ to $(v_{\lfloor i/2 \rfloor + 1}^1, v_{\lfloor i/2 \rfloor}^2, 2)$ (for odd i) exists iff the edge $(v_{\lfloor i/2 \rfloor}^1, v_{\lfloor i/2 \rfloor + 1}^1)$ is in \vec{E} , which follows from the fact that $(s_1, t_1) \in \vec{P}$. The other case, where i is even, follows from an entirely symmetric argument.

Also note that the length of $\rho(p')$ is exactly 2Δ for all $p' \in P'$. Our next step is to argue that $\rho(p')$ is the unique shortest path between its endpoints.

CLAIM 1. *For all pairs $p' \in P'$, the path $\rho(p')$ is the unique shortest path between the endpoints of p' .*

PROOF. Let $p' = (s', t')$, where $s' = (s_1, s_2, 1)$ and $t' = (t_1, t_2, 1)$. By construction, any $s' \rightsquigarrow t'$ path $\pi := (s', x_1, \dots, x_k, t')$ corresponds to a $s_1 \rightsquigarrow t_1$ path π_1 in G_1 and a $s_2 \rightsquigarrow t_2$ path π_2 in G_2 , where π_1 and π_2 are given by the sequence of values taken by the first and second indices of the nodes x_i (respectively). Moreover, we have $|\pi| = |\pi_1| + |\pi_2|$.

Note that the path $\rho(p')$ is the unique $s' \rightsquigarrow t'$ path π that corresponds to the unique shortest $s_1 \rightsquigarrow t_1$ path in G_1 and the unique shortest $s_2 \rightsquigarrow t_2$ path in G_2 . Thus, any other $s' \rightsquigarrow t'$ path $\tau(p')$ from s' to t' corresponds to a non-shortest $s_1 \rightsquigarrow t_1$ or $s_2 \rightsquigarrow t_2$ path (without loss of generality, assume the former). We thus have $|\rho(p')| = 2\Delta$, while $|\tau(p')| \geq (\Delta + 1) + \Delta > 2\Delta$. Hence, $\rho(p')$ is the unique shortest $s' \rightsquigarrow t'$ path. \square

Finally, we observe that, from our construction, the paths $\rho(p'), \rho(p'')$ are 2-path disjoint for any distinct $p', p'' \in P'$.

CLAIM 2. *For any two distinct pairs $p' \neq p'' \in P'$, the paths $\rho(p'), \rho(p'')$ are 2-path disjoint.*

PROOF. Consider two pairs $p' = \{(s'_1, s'_2, 1), (t'_1, t'_2, 1)\} \in P'$ and $p'' = \{(s''_1, s''_2, 1), (t''_1, t''_2, 1)\} \in P'$ for which there is a 2-path (a, b, c) that is a subpath of both $\rho(p')$ and $\rho(p'')$. We will show that p' and p'' must be the same pair.

By definition of ρ , the subpath (a, b, c) must have one of two forms: either $a = (x_1, y_2, 1), b = (x'_1, y_2, 2), c = (x'_1, y'_2, 1)$ or $a = (x_1, y_2, 2), b = (x_1, y'_2, 1), c = (x'_1, y'_2, 2)$. Assume we are in the first case, and the second case is symmetric. Again, by the definition of ρ , we must have that $(x_1, x'_1) \in \vec{E}$ is on the unique (s'_1, t'_1) shortest path and on the unique (s''_1, t''_1) shortest path in G . Since the shortest paths between pairs in P are edge disjoint in G , this implies that $(s'_1, t'_1) = (s''_1, t''_1)$. Moreover, we have that $(y_2, y'_2) \in \vec{E}$ is on the unique shortest paths for both (s'_2, t'_2) and (s''_2, t''_2) , and so we also have that $(s'_2, t'_2) = (s''_2, t''_2)$. Together, these imply that $p' = p''$. \square

We can now prove Lemma 2.2:

PROOF OF LEMMA 2.2. It is immediate from the construction that $N := |V'| = \Theta(n^2)$ and that $|P'| = |P|^2 = \Theta(n^{4-2\epsilon}) = \Theta(N^{2-\epsilon})$ (thus, the new value of ϵ is the *same* as the old value of ϵ used to create G via Lemma 2.1). Additionally, a loose (but sufficient for our purposes) upper bound on $|E'|$ follows from the observation that each node can have at most $2n$ neighbors and thus $|E'| = O(n^3) = O(N^{3/2})$.

We have demonstrated above that $\rho(p')$ is the unique shortest path for any pair $p' \in P'$ and that $\rho(p'_1)$ and $\rho(p'_2)$ are 2-path disjoint for any $p'_1 \neq p'_2 \in P'$. We additionally have $|\rho(p')| = 2\Delta = \Theta(n^\delta) = \Theta(N^{\delta/2})$ (and so the new value of δ is half the old value of δ used to create G via Lemma 2.1). Since the values of ϵ, δ used to create G satisfied $0 < \delta < \epsilon$, and our new value of ϵ has remained unchanged while our new value of δ has been halved, we then still have $0 < \delta < \epsilon$. \square

2.4 The Obstacle Product

Our final move is to produce an *Obstacle Product* (OP) $G \otimes K$, where G is a graph produced by Lemma 2.2, and K is a clique. For intuition on the obstacle product and why it is useful, we refer

the reader to the overview at the beginning of this section. Its ultimate purpose is to prove our main theorem, which we restate below for convenience:

THEOREM 2.4. *For all $\varepsilon > 0$, there is a $\delta > 0$ and an infinite family of n node graphs $G = (V, E)$ such that for any subgraph $H = (V, E')$ with $|E'| = O(n^{4/3-\varepsilon})$, there exist nodes $u, v \in V$ with*

$$\text{dist}_H(u, v) = \text{dist}_G(u, v) + \Omega(n^\delta).$$

Let us now fix an $\varepsilon > 0$. Our next steps will be to choose an appropriate G, K ; describe how to build the obstacle product $G \otimes K$; and then argue that this is sufficient to prove Theorem 2.4.

Starting Graph. We start by applying Lemma 2.2 with parameter $\alpha = \varepsilon/2$ to obtain a graph $G = (V, E)$ on $|V| = n$ nodes, $|E| = O(n^{3/2})$ edges, and a pair set $P \subseteq V \times V$ of size $|P| = \Theta(n^{2-\alpha})$ such that the distance between any pair $\{s, t\} \in P$ is exactly Δ , for some $\Delta = \Theta(n^\beta)$ and $0 < \beta < \alpha$ (which implies $0 < \beta < \varepsilon/2$), and the shortest paths between pairs in P are 2-path disjoint.

We next take the OP of G with a clique K . This OP consists of the following two transformations:

OP – Edge Extension. First, we replace every edge in G with a path of length $\ell = 3\Delta$.¹

More formally, if the edge $e = \{u, v\} \in E$, then we add the nodes $(e, 1), \dots, (e, \ell)$ to V' and add the edges $\{u, (e, 1)\}, \{v, (e, \ell)\}$, and $\{(e, i), (e, i + 1)\}$ for all $i \in [\ell - 1]$ to E' . The choice of which endpoint $(e, 1)$ or (e, ℓ) we connect to u or v is arbitrary.

Note that the total number of nodes in G' after this transformation is $O(n^{3/2+\beta})$.

OP – Clique Replacement. Next, for each *original* node $v \in V$ (i.e., not nodes introduced by the edge extension step), we replace v with a clique on $\deg_G(v)$ nodes, with each incoming edge connected to a unique node in the new clique.

More formally, let $I(v) \subseteq E$ be the set of edges incident to $v \in V$, and introduce a new node (v, e) for each edge $e \in I(v)$. These new nodes will be connected in a clique: For each $e_i, e_j \in I(v)$, we add the edge $\{(v, e_i), (v, e_j)\}$ to E' . The clique that replaces a node v will be denoted K_v . We will call an edge contained in one of these cliques a *clique edge*. After this replacement, each node (v, e) will still be attached to an endpoint of the path corresponding to the edge e , that is, we will have an edge $\{(v, e), (e, \alpha)\}$, where $\alpha \in [1, \ell]$ depends on the (arbitrary) ordering we chose in the edge-extension step.

The number of nodes added in this clique replacement step is exactly $\sum_{v \in V} \deg_G(v) = 2|E|$, which is $O(n^{3/2})$. An important feature of this step is that we have introduced many new edges to G' ; we will implicitly discuss this in the proof of correctness.

Let $k = \Delta - 1 = \Theta(n^\beta)$. To complete the proof, we will argue that any subgraph of G' with fewer than $O(n^{2-\alpha})$ edges must distort the distances by more than $+k$.

Proof of Correctness. First, we will build a pair set $P' \subseteq V' \times V'$; to prove Theorem 2.4, it will be sufficient to only consider node pairs in P' . For each pair $\{s, t\} \in P$, let the nodes on the unique (s, t) -shortest path in G be $s = u_0 \rightarrow u_1 \rightarrow \dots \rightarrow u_{\Delta-1} \rightarrow u_\Delta = t$, and let the Δ edges on this path be denoted $e_i = \{u_{i-1}, u_i\}$ for all $i \in [\Delta]$ (these definitions use the feature of Lemma 2.2 that all pairs in P are at distance exactly Δ in G). Consider the nodes $s' = (s, e_1), t' = (t, e_\Delta) \in V'$, and add $\{s', t'\}$ as a pair to P' .

We will next reason about the structure of the shortest path between s' and t' in G' . The following (s', t') path will exist by construction: Starting from $s' = (s_1, e_1)$, walk the path in G' that replaced the edge e_1 in G , reaching the node (u_1, e_1) . Then, walk the clique edge $\{(u_1, e_1), (u_1, e_2)\}$,

¹This is overkill: The proof still works if we simply set $\ell = \Delta$. However, a few additional minor technical details must be observed to push this argument through, so we use $\ell = 3\Delta$ here to maintain a bit of simplicity.

which we will denote by $e_1^{s,t}$, since it is the first clique edge on the (s', t') -path. Then, similarly, walk the path that replaced to the edge e_2 and then walk the clique edge $\{(u_2, e_2), (u_2, e_3)\}$, and so on, until we reach t' . This walk will traverse each of the ℓ -length paths that replaced our edges e_i (as well as some clique edges). The i th clique edge crossed will be $e_i^{s,t} = \{(u_i, e_i), (u_i, e_{i+1})\}$, for $i \in [\Delta - 1]$. Thus, this walk will reach the node $t' = (t, e_\Delta)$ after it has taken exactly Δ paths of length ℓ , plus $(\Delta - 1)$ clique edges, that is, $D := \Delta \cdot \ell + (\Delta - 1)$ edges in total. Let $\rho(s', t')$ be the path in G' defined by this walk, and let $C^{s,t}$ be the set of clique edges on this path.

The following claim is an integral part of our proof:

CLAIM 3. *For any two distinct pairs $\{s_1, t_1\} \neq \{s_2, t_2\} \in P$, the corresponding clique edge sets $C^{s_1, t_1}, C^{s_2, t_2}$ are disjoint.*

PROOF. Assume towards a contradiction that $C^{s_1, t_1}, C^{s_2, t_2}$ share a clique edge $\{(u, e_i), (u, e_j)\}$ for some $u \in V$ and $e_i, e_j \in I_G(v)$. This implies that the 2-path in G composed of e_i and e_j is a subpath of the unique shortest (s_1, t_1) -path in G and of the unique shortest path (s_2, t_2) -path. By Lemma 2.2 this can only happen if $\{s_1, t_1\} = \{s_2, t_2\}$, which is a contradiction. \square

Next, we examine the structure of alternate short paths between s and t .

CLAIM 4. *Let $\{s, t\}$ be a pair in P , let $\{s', t'\}$ be the corresponding pair in P' , and let $\rho'(s', t')$ be a (not necessarily shortest) (s', t') path of length less than $D + \Delta$. Let $\mathcal{K} = (K_s, K_{v_1}, \dots, K_{v_{\Delta-1}}, K_t)$ be the sequence of cliques created in the Clique Replacement step that intersect $\rho(s', t')$, and let $\mathcal{K}' = (K_s, K_{v'_1}, \dots, K_{v'_{x-1}}, K_t)$ be the sequence of cliques that intersect $\rho'(s', t')$. Then $\mathcal{K} = \mathcal{K}'$.*

PROOF. Note that, by construction, the node sequence $(s, v_1, \dots, v_{\Delta-1}, t)$ is the unique shortest (s, t) path in G . Similarly, note that the node sequence $(s, v'_1, \dots, v'_{x-1}, t)$ is an (s, t) path in G . If these node sequences are identical, then the claim holds. Assume towards a contradiction that the node sequences differ. Then the latter sequence must be at least one element longer than the former sequence (i.e., $x > \Delta$). Both $\rho(s', t')$ and $\rho'(s', t')$ must travel a path of length ℓ between any two cliques, and so the number of path edges traveled by $\rho'(s', t')$ is at least $\ell = 3\Delta$ more than the number of path edges traveled by $\rho(s', t')$. Meanwhile, $\rho(s', t')$ walks exactly Δ clique edges, while $\rho'(s', t')$ walks exactly x clique edges. Since $x \geq \Delta$, this implies that the total length of $\rho'(s', t')$ is at least $+3\Delta$ longer than $\rho(s', t')$. This is a contradiction, and so it must be the case that the original node sequences are equal, and so $\mathcal{K} = \mathcal{K}'$. \square

This lets us prove the next claim, which is the key to lower bounding the number of edges in any additive $+k$ spanner (recall that we chose $k = \Delta - 1$).

CLAIM 5. *For any pair $\{s, t\} \in P$, any path of length less than $D + k$ from s' to t' must use at least one of the clique edges $C^{s,t}$.*

PROOF. Let $\rho'(s', t')$ be a path from s' to t' of length less than $D + k$. By Claim 4, we know that $\rho'(s', t')$ must pass through exactly the same cliques in the same order as $\rho(s', t')$, which implies that it must walk exactly the same set of paths introduced in the Edge Extension step of the OP. This implies that for all $i \in [\Delta]$, $\rho'(s', t')$ contains a path from (u_i, e_i) to (u_i, e_{i+1}) as a subpath. Let the length of this subpath be d_i for some $d_i \geq 1$, and note that $d_i = 1$ if and only if $\rho'(s', t')$ uses the clique edge $e_i^{s,t} \in C^{s,t}$. A direct calculation shows that the length of $\rho'(s', t')$ is equal to

$$\ell(\Delta - 1) + \sum_{i=1}^{\Delta-1} d_i = D + \sum_{i=1}^{\Delta-1} (d_i - 1).$$

If $d_i > 1$ for all i , then this is equal to at least $D + (\Delta - 1) = D + k$, which contradicts the assumption that $\rho'(s', t')$ has length less than $D + k$. Thus, $d_i = 1$ for some i , and so $\rho'(s', t')$ uses the clique edge $\{(u_i, e_i), (u_i, e_{i+1})\}$. \square

We can now prove our main theorem from a direct counting argument. Recall that our graph G' has $N = O(n^{3/2+\beta})$ nodes and that our pair set P' is of size $|P'| = \Theta(n^{2-\alpha})$. Assume towards a contradiction that there is a subgraph $H = (V', E_H)$ of G' on $|E_H| = O(N^{4/3-\epsilon}) = O(n^{(3/2+\beta)(4/3-\epsilon)}) = O(n^{2+\frac{4}{3}\beta-\frac{3}{2}\epsilon-\epsilon\beta})$ edges in which for all pairs $\{s', t'\} \in P'$ the distance in H is no more than the distance in G' plus k , that is, at most $D + k$. By Claim 5, for each such pair $\{s', t'\} \in P'$ at least one of the clique edges in $C^{s', t'}$ must exist in H , and from Claim 3 we know that all these edges are distinct. Therefore, E_H must contain at least one distinct edge per pair in P' , for a total of $|E_H| = \Omega(n^{2-\alpha})$ edges. These two bounds on $|E_H|$ imply that

$$\begin{aligned} 2 - \alpha &\leq 2 + \frac{4}{3}\beta - \frac{3}{2}\epsilon - \epsilon\beta \\ \left(\frac{3}{2} + \beta\right)\epsilon &\leq \alpha + \frac{4}{3}\beta. \end{aligned}$$

We additionally know that $0 < \beta < \alpha \leq \epsilon/2$. We complete the proof by arguing that there is no possible setting of α, β that makes these inequalities hold simultaneously, and so we have a contradiction and the subgraph H cannot exist. To see this, we will substitute $\beta = 0$ into the left-hand side of the inequality (giving the left-hand side of the inequality its lowest conceivable value) and simultaneously substitute $\alpha = \beta = \frac{\epsilon}{2}$ into the right-hand side of the inequality (giving the right-hand side of the inequality its highest conceivable value). We obtain:

$$\begin{aligned} \left(\frac{3}{2} + 0\right)\epsilon &\leq \frac{\epsilon}{2} + \frac{4}{3} \cdot \frac{\epsilon}{2} \\ \frac{3}{2} \cdot \epsilon &\leq \frac{7}{6} \cdot \epsilon, \end{aligned}$$

which contradicts the fact that $\epsilon > 0$. This implies Theorem 2.4 with parameter δ set to β (since $k = \Omega(n^\beta)$).

2.5 Strong Incompressibility

Finally, we provide some additional observations about our construction to prove our strong incompressibility results. We will prove:

THEOREM 2.5. *For all $\epsilon > 0$, there exists a $\delta > 0$ such that there is no mapping ψ from graphs G on n nodes to bitstrings of length $O(n^{4/3-\epsilon})$ such that the distances of G can always be recovered from the bitstring $\psi(G)$ within $+n^\delta$ error.*

Our argument proceeds along the lines of the well-known proof that the Girth Conjecture implies that graphs cannot be compressed into $o(n^{1+1/t})$ bits without incurring a multiplicative error of $(2t - 1)$ (Thorup and Zwick 2005), which is based on the works of Matoušek (1996) and Bourgain (1985).

For any subset of pairs $T \subseteq P'$, we define $G_T = (V, E_T)$ to be the subgraph of G' obtained by removing all clique edges $C^{s', t'}$ corresponding to all pairs $\{s', t'\}$ in T . Note that there are $2^{|P'|}$ such subgraphs G_T .

CLAIM 6. *For all $T \subseteq P'$ and $\{s', t'\} \in P' \setminus T$, we have that $\text{dist}_{G_T}(s', t') \leq D$.*

PROOF. The only edges that are missing from G_T are clique edges corresponding to pairs in T . By Claim 3, none of these clique edges are in $C^{s', t'}$ for any pair $\{s', t'\} \notin T$. Therefore, the length D path $\rho(s', t')$ remains in G_T . \square

CLAIM 7. *For all $T \subseteq P'$ and $\{s', t'\} \in T$ we have that $\text{dist}_{G_T}(s', t') > D + k$.*

PROOF. Any path from s' to t' in G_T cannot use any of the clique edges in $C^{s,t}$. By Claim 5, the length of such path cannot be $\leq D + k$. \square

The above two claims show that there exists a set \mathcal{G} of $2^{|P'|} = 2^{\Theta(N^{4/3-\epsilon'})}$ graphs (which are G_T for each possible pair subset $T \subseteq P'$) on the same node set V' of size N , as well as a pair set $P' \subseteq V' \times V'$ of size $|P'| = \Theta(n^{4/3-\epsilon'})$ (for some $\epsilon' > 0$), such that the following condition is true:

For any two graphs $G_1, G_2 \in \mathcal{G}$ there is a pair of nodes $\{s', t'\} \in P'$ for which

$$|\text{dist}_{G_1}(s', t') - \text{dist}_{G_2}(s', t')| \geq k,$$

where $k = \Omega(N^\delta)$. This means that for any compression function ψ that takes a graph G on n nodes and produces a bitstring of length less than $|P|$, there exist two different graphs $G_1, G_2 \in \mathcal{G}$ that map to the same bitstring. Thus, when the distance between s' and t' is recovered from this bitstring, it will be at least $\pm\Omega(n^\delta)$ from the correct value for either G_1 or G_2 . Therefore, no such compression function can output bitstrings of length $O(n^{4/3-\epsilon'})$ such that distances can be recovered within $+O(n^\delta)$ error.

3 LOWER BOUNDS FOR LINEAR SIZE GRAPH COMPRESSION

We next calculate our lower bound on the amount of (polynomial) additive error in distances that must be suffered to compress a graph into near-linear space. We show:

THEOREM 3.1. *For all $\epsilon > 0$, there exists a $\delta > 0$ such that there is no mapping ψ from graphs G on n nodes to bitstrings of length $O(n^{1+\delta})$ such that the distances of G can always be recovered from the bitstring $\psi(G)$ within $+n^{1/22-\epsilon}$ error.*

We begin with the following graphs:

THEOREM 3.2 (COPPERSMITH AND ELKIN 2006).² *For any n and $p = O(n^{3/2})$, there is an infinite family of n node graphs $G = (V, E)$ on $\Theta(n^{3/5}p^{3/5})$ edges, as well as pair sets $P \subseteq V \times V$ of size $|P| = p(n)$, with the following properties:*

- (1) *For each pair in P , there is a unique shortest path between its endpoints*
- (2) *These paths are edge disjoint*
- (3) *The edge set of G is precisely the union of these paths*
- (4) *For all $\{s, t\} \in P$, we have $\text{dist}_G(s, t) = \Delta$ for some value $\Delta = \Theta(n^{3/5}p^{-2/5})$*

This construction can be found in Section 4 of Coppersmith and Elkin (2006) (the statement is also in Section 1.2). Note that their construction does not have property 4 above. However, it is easy to enforce this property using the following folklore technique. Given an instance $G = (V, E), P$ satisfying the first three properties of Theorem 3.2 but not necessarily the fourth, let $\Delta := \frac{|E|}{2|P|}$ be half the average length of a path. For each pair $(s, t) \in P$, split its unique shortest path π into as many edge-disjoint subpaths as possible of length exactly Δ each. Remove (s, t) from P and add the endpoints of each such subpath as a new pair to P . If Δ does not divide $|\pi|$ and so there is a “remainder” subpath at the end of length less than Δ , then we discard this subpath and delete all of its edges from G . At the end of this process, it is clear that all four properties in Theorem 3.2 hold for the new instance $G' = (V, E'), P'$. Additionally, note that (1) we delete at most $|P| \cdot \Delta = \frac{|E|}{2}$ edges from G to produce G' , so the density of G' has only changed by a constant factor, and (2)

²Coppersmith and Elkin prove a host of lower bounds, parameterized by an integer $d \geq 2$. To help alleviate some tedious algebra, we have instantiated their construction here with $d = 3$, which optimizes our lower bound ($d = 4$ optimizes equally well).

we have $|P| = \frac{|E|}{\Delta}$ and $|P'| = \frac{|E'|}{\Delta}$, and since $|E'| = \Theta(|E|)$ the size of P' has also only changed by a constant factor.

Unlike Lemma 2.1, the construction in Alon (2001) or the one given in the appendix do not suffice; one must use the construction in Coppersmith and Elkin (2006) to obtain this particular tradeoff between the number of pairs p and the density of the graph.

We start with a graph $G = (V, E)$ drawn from Theorem 3.2, choosing $p = n^{11/9+\delta}$ and so G has $\Theta(n^{4/3+3\delta/5})$ edges. Our construction proceeds exactly as before; we feed this graph through the path packing step, the edge extension step, and the clique replacement step. Looking forward, we will reserve n for the number of nodes in the original graph G and N to describe the number of nodes in G after all transformations have been applied.

By the exact same argument for incompressibility given earlier in this article, we have:

LEMMA 3.3. *There is no mapping ψ from graphs G on N nodes to bitstrings of length $o(p^2)$ such that the distances of G can always be recovered from the bitstring $\psi(G)$ within $+o(\Delta)$ error.*

We now simply calculate these parameters to obtain Theorem 3.1.

- We have $p = n^{11/9+\delta}$
- We have $\Delta = \Theta(n^{3/5}p^{-2/5}) = \Theta(n^{3/5}(n^{11/9+\delta})^{-2/5}) = \Theta(n^{1/9-2\delta/5})$.
- Originally, we have n nodes and $\Theta(n^{4/3+3\delta/5})$ edges in G . After the path packing step, we have n^2 nodes and $\Theta(n^{7/3+3\delta/5})$ edges in G . The edge extension step then introduces

$$\Delta \cdot \Theta(n^{7/3+3\delta/5}) = \Theta\left(n^{1/9-2\delta/5} \cdot n^{7/3+3\delta/5}\right) = \Theta\left(n^{22/9+\delta/5}\right)$$

new nodes to G , while the clique replacement step introduces $\Theta(n^{7/3+3\delta/5})$ new nodes to G (i.e., one node for each edge in G prior to the edge extension step). Note that $n^{22/9+\delta/5}$ dominates $n^{7/3+3\delta/5}$ for sufficiently small $\delta > 0$. Thus, if we choose δ small enough, then we have $N = \Theta(n^{22/9+\delta/5})$.

– Note: We then have

$$\Delta = \Theta\left(N^{(1/9-2\delta/5)/(22/9+\delta/5)}\right) = \Theta\left(N^{1/22-O(\delta)}\right)$$

and

$$p^2 = \Theta\left(N^{(11/9+\delta)/(22/9+\delta/5)}\right)^2 = \Theta\left(N^{(22/9+2\delta)/(22/9+\delta/5)}\right) = \Theta\left(N^{1+81\delta/100-O(\delta^2)}\right).$$

We now prove Theorem 3.1 by directly substituting these computed values into Lemma 3.3. Given some $\varepsilon > 0$ (in the context of Theorem 3.1), we first choose $\delta > 0$ (controlling the size of the pair set in the graph G from Theorem 3.2) small enough that

$$\Delta > N^{1/22-\varepsilon/2}.$$

This is possible because $\Delta = \Theta(N^{1/22-O(\delta)})$. Hence, the fact that distances cannot be recovered within $+o(\Delta)$ error (from Lemma 3.3) implies that distances cannot be recovered within $+N^{1/22-\varepsilon}$ error, as desired.

Meanwhile, the compression threshold in Lemma 3.3 is

$$o(p^2) = o\left(N^{1+81\delta/100-O(\delta^2)}\right).$$

Thus, if we choose δ sufficiently small and choose some constant $\delta' = 4\delta/5 > 0$, then one cannot compress graphs into bitstrings of size $O(n^{1+\delta'})$ and still recover distances within the error bound $+N^{1/22-\epsilon}$ given above. This completes the proof of Theorem 3.1.

4 CONCLUSION

Our work shows that graphs cannot be compressed into $O(n^{4/3-\epsilon})$ bits while maintaining distances within a constant (or even subpolynomial) additive error. We thus implicitly show that the 4/3 exponent is tight for additive spanners and emulators and unbeatable for any other notion of graph compression where distance error is measured additively. We consider it particularly interesting that the analogous statement is false for multiplicative error; that is, many graph compression schemes are capable of obtaining near-linear sized compression with only a constant multiplicative distance error.

Our main technical contribution is the Obstacle Product, which seems to be a powerful framework for generating families of graphs that are hard to efficiently sketch while approximately maintaining distances. Below, we mention some open questions from the realm of spanners that are left open after this work and for which our framework could be applicable.

Completing the Picture of Constant Error Additive Spanners. The only remaining gap (in the exponent) in our understanding of constant error additive spanners is the +4 error case. The lower bounds allow for the possibility of +4 spanners on $O(n^{4/3})$ edges, while upper bounds have only realized $\tilde{O}(n^{7/5})$ (Chechik 2013). Note that this gap does not exist for emulators. Closing this gap is perhaps the most natural and intriguing open question that remains.

Our work also leaves $n^{o(1)}$ gaps between upper and lower bounds; that is, our lower bound holds against spanners on $O(n^{4/3-\epsilon})$ edges, but it is still conceivable that an upper bound of $O(n^{4/3-o(1)})$ edges could be achieved for a small-enough $o(1)$ factor. Hence, the following question remains open: Do all graphs have + k spanners on $O(\frac{n^{4/3}}{\log n})$ edges or maybe even $O(\frac{n^{4/3}}{\log^k n})$? It would be interesting to close this $n^{o(1)}$ gap.

Other Kinds of Spanners. We have shown that $+\Omega(n^\delta)$ additive error is required for spanners on $O(n^{4/3-\epsilon})$ edges. It is then natural to ask precisely how much polynomial error is required for sparse spanners; for example, what is the smallest constant δ such that all graphs have $+O(n^\delta)$ spanners of nearly linear size? The current best upper bound on δ is $\frac{3}{7}$, due to Bodwin and Vassilevska Williams (2016). For general compression (specifically emulators), the upper bound on δ is $\frac{1}{3}$, due to Bodwin and Vassilevska Williams (2015). Finding the exact tradeoff between the error parameter n^δ and the sparsity $n^{4/3-\epsilon}$ is still open. Related notions are *sublinear error* spanners, where if the original distance is d it must be preserved up to $d + o(d)$, and (α, β) *mixed* spanners, where even $(1 + \alpha) \cdot d + \beta$ is allowed. The framework described in this article was adapted by the authors and Pettie (Abboud et al. 2017) to obtain nearly-tight lower bounds in these settings as well, although some notable gaps still remain here to be closed. Finally, it would be interesting to understand the exact tradeoff for *pairwise* spanners (Cygan et al. 2013; Coppersmith and Elkin 2006; Kavitha and Varma 2013; Kavitha 2015). Both of our main results hold even if only the distances within a pair set of size $\Theta(n^{4/3})$ need to be preserved. When only smaller pair sets are considered, it is known that one can obtain +6 spanners on $O(n^{4/3-\epsilon})$ edges (Kavitha 2015), although the exact nature of the tradeoff between pair set size, error, and sparsity remains open.

APPENDIX

A GRAPHS WITH MANY LONG DISJOINT SHORTEST PATHS

In this section, we provide a proof of Lemma 2.1 which is the starting point for our lower-bound constructions. One way to prove the Lemma is by simple modifications to the constructions of Coppersmith and Elkin (2006) for *pairwise distance preserver* lower bounds. Their construction relied on lemmas from discrete geometry regarding the size of the convex hull of points in a radius r ball in the d -dimensional integer lattice. We will provide a simpler proof of this lemma (our proof is less efficient in the tradeoff between ε and δ but will suffice for our purposes). Our proof relies on generalizations of the well-known arithmetic progression free sets that have found countless applications in combinatorics and theoretical computer science.

k-Average-Free Sets. We say that a set of integers $S \subseteq [N]$ is *k-Average-Free* if, for any collection of k integers that are not all equal, $x_1, \dots, x_k \in S$, we have

$$\sum_{i=1}^k x_i \neq k \cdot \bar{x}$$

for all $\bar{x} \in S$. In other words, S does not contain the arithmetic mean of any k of its elements, except for the trivial case when all of these elements are equal. An adaptation of a construction by Behrend (1946) proves the existence of k -average-free sets with surprisingly large density.

LEMMA A.1. *For all $\varepsilon > 0$, there is a $\delta > 0$ such that for infinitely many integers N there exist k -average-free sets of size $\Theta(N^{1-\varepsilon})$, with $k = \Omega(N^\delta)$.*

PROOF. Fix $\varepsilon > 0$, and let $d = \lceil 3/\varepsilon \rceil \in \mathbb{N}$ and $\delta = 1/(2d^2) > 0$. Let p be a parameter that can take any value in \mathbb{N} ; we will construct a k -average-free set $A \subseteq [N]$, where k, N , and $|A|$ are functions of p .

Let $X = [p]^d$ be the set of all d -dimensional vectors whose coordinates are integers in $[p]$. For any $r \in [dp^2]$, let

$$X_r := \{x \in X \mid \|x\|_2^2 = r\},$$

that is, X_r is the set of vectors for which the square of the Euclidean norm is r . Observe that for all $x \in X$ we have

$$1 \leq \|x\|_2^2 \leq dp^2,$$

and so the sets X_1, \dots, X_{dp^2} form a partition of the set X . Let $V = X_r$ be the largest of these sets. Since $|X| = p^d$, we have that

$$|V| \geq \frac{p^d}{dp^2} = \frac{p^{d-2}}{d}.$$

We now claim that the set V is the analogue of a vector-valued “ k -average-free set” for any $k \in \mathbb{N}$, in the sense that the average of any set of vectors (not all equal) in V is not an element of V . This follows simply from the fact that all vectors in V have the same Euclidean norm, and so they form a strictly convex set in \mathbb{R}^d . By definition of convexity, for any $x_1, \dots, x_k, \bar{x} \in V$, we have

$$\left\| \sum_{i=1}^k x_i \right\|_2 < \|k\bar{x}\|_2$$

whenever x_1, \dots, x_k, \bar{x} are not all equal; thus, $\sum_{i=1}^k x_i \neq k\bar{x}$, and so V is a k -average-free set of vectors.

Our next step is to convert the vector set V into an integer set A in a way that preserves summation. We remark that this step is not strictly necessary; that is, it would be possible to prove Lemma 2.1 using only a vector-valued average free set. However, it will simplify the proofs

that follow to consider integers rather than vectors. Our transformation will be defined by the function $f : [p]^d \rightarrow [N]$, where $N = ((k+1)p)^d$, in which for any vector $v = (v_1, \dots, v_d) \in [p]^d$, we interpret v_1, \dots, v_d as the digits of a number in base $(k+1)p$. More formally, we define

$$f(v) := v_1 \cdot q^0 + v_2 \cdot q^1 + \dots + v_d \cdot q^{d-1},$$

where $q = (k+1)p$. Note that this sum cannot exceed q^d , and so we have $f(v) \in [N]$. An important feature of our transformation is that there will be no carry bits that transfer between the “coordinates” when summing k integers that correspond to vectors, and so f distributes over k -length sums of vectors in V . In other words, for any set of k vectors $v^1, \dots, v^k \in V$, we have

$$\sum_{i=1}^k f(v^i) = \left(\sum_{i=1}^k v_1^i \right) q^0 + \dots + \left(\sum_{i=1}^k v_d^i \right) q^{d-1}.$$

Note that $\sum_{i=1}^k v_j^i \leq \sum_{i=1}^k p \leq kp$, and so we have

$$\left(\sum_{i=1}^k v_1^i \right) q^0 + \dots + \left(\sum_{i=1}^k v_d^i \right) q^{d-1} = f\left(\sum_{i=1}^k v^i \right).$$

And thus we have $\sum_{i=1}^k f(v^i) = f(\sum_{i=1}^k v^i)$, and so because V is a k -average-free set of vectors, it follows that the corresponding integer set $A = f(V)$ is a k -average-free set of integers.

The final step in our proof is to verify that the set A has size $\Theta(N^{1-\varepsilon})$ when $k = \Omega(N^\delta)$, as claimed. The parameters will work out correctly when we set $k = \lfloor p^{\frac{\delta d}{1-\delta d}} \rfloor - 1$. Note that we now have

$$N = ((k+1)p)^d = \Theta\left(\left(p^{\frac{1}{1-\delta d}}\right)^d\right) = \Theta(p^{d/(1-\delta d)})$$

and so $k = \Omega(N^\delta)$. Last, we must show that $|A| = \Omega(N^{1-\varepsilon})$. Recall that $|A| \geq \frac{p^{d-2}}{d}$ and that $d = \lceil 3/\varepsilon \rceil$ is a constant, which implies that $|A| = \Omega(p^{d-2})$. We must then show that

$$p^{d-2} = \Omega(N^{1-\varepsilon}) = \Omega((p^{d/(1-\delta d)})^{1-\varepsilon}) = \Omega(p^{(1-\varepsilon) \cdot d/(1-\delta d)})$$

or, equivalently, that

$$\begin{aligned} d-2 &\geq (1-\varepsilon) \cdot \frac{d}{1-\delta d} \\ (d-2)(1-\delta d) &\geq (1-\varepsilon) \cdot d \\ \varepsilon d - 2 + (2\delta d - \delta d^2) &\geq 0. \end{aligned}$$

Substituting in our previous choices of $d \geq 3/\varepsilon$ and $\delta = \frac{1}{2d^2}$, we can verify that this inequality holds:

$$\varepsilon d - 2 + (2\delta d - \delta d^2) \geq 3 - 2 + (0 - 1/2) > 0.$$

This shows that $|A| = \Omega(N^{1-\varepsilon})$. To prove the claim that $|A| = \Theta(N^{1-\varepsilon})$, simply note that we may discard elements from A at will without destroying its average freeness property. \square

From k -Average-Free Sets to Graphs. We are now ready to prove Lemma 2.1, which we repeat here for convenience. The key idea is to let each node represent an integer in a careful way such that the k -average-free property translates into the uniqueness of certain shortest paths.

LEMMA A.2. *For all $\varepsilon > 0$, there is a $0 < \delta < \varepsilon$, and an infinite family of graphs $G = (V, E)$ and pair sets $P \subseteq V \times V$ with the following properties:*

- (1) For each pair in P , there is a unique shortest path between its endpoints
- (2) These paths are edge disjoint
- (3) For all $\{s, t\} \in P$, we have $\text{dist}_G(s, t) = \Delta$ for some value $\Delta = \Theta(n^\delta)$
- (4) $|P| = \Theta(n^{2-\epsilon})$

PROOF. Fix $\epsilon > 0$ and apply Lemma A.1 with parameter $\alpha = \epsilon/2 > 0$ to obtain an infinite sequence of integers N and k -average-free sets $A \subseteq [N]$ of size $|A| = \Omega(N^{1-\alpha})$, in which $k = \Omega(N^\beta)$ for some $\beta > 0$. We will assume that $0 < \beta$ is no more than $\epsilon/10$ (otherwise, we may decrease β without destroying the above properties). Consider the infinite sequence of integers defined by $n = N \cdot k^2 = \Theta(N^{1+2\beta})$, and we will describe how to construct a graph $G = (V, E)$ on n nodes that will be in our family, using the corresponding k -average-free set $A \subseteq [N]$.

The Nodes. For each integer $x \in [(k+1) \cdot N]$ and “layer index” $j \in \{0, \dots, k\}$, we define a node (x, j) and add it to G . In other words, we define

$$V := \{(x, j) \mid x \in [(k+1)N], j \in \{0, \dots, k\}\}.$$

The Edges. The edge set of G will be defined using the set A . Our graph will be layered such that all edges have the form $\{(x, j), (y, j+1)\}$. For all $j \in \{0, \dots, k-1\}$ and for each integer $a \in A$, we add an edge $\{(x, j), (y, j+1)\}$ to E for all integers $x, y \in [kN]$ such that $x + a = y$. That is:

$$E := \{(x, j), (y, j+1)\} \mid j \in \{0, \dots, k-1\} \exists a \in A x + a = y\}.$$

The Pair Set. Next, we define our pair set $P \subseteq V \times V$. For every integer $x \in [N]$ and integer $a \in A$, we add the pair $\{s, t\}$ to P , where $s = (x, 0)$ and $t = (x + k \cdot a, k)$. An immediate observation is that $|P| = N \cdot |A| = \Omega(N^{2-\alpha})$. Formally,

$$P := \{(x, 0), (x + k \cdot a, k)\} \mid x \in [N], a \in A\}.$$

We now turn to reasoning about the shortest paths between the pairs in P . First, note that for any pair $s = (x, 0)$ and $t = (x + k \cdot a, k)$ in P , the edge set

$$E^{s,t} := \{(x + (j-1) \cdot a, j-1), (x + j \cdot a, j)\} \mid j \in [k]\}$$

forms a path of length k from s to t . We will refer to this path as $\rho(s, t)$. Since our graph is layered, any (s, t) path must contain at least k edges, and so $\rho(s, t)$ is a shortest path between s and t . Additionally, any other path $\rho'(s, t)$ of equal length must use an edge set of the form

$$E'^{s,t} = \{(y_{j-1}, j-1), (y_{j-1} + a_j, j)\} \mid j \in [k], \{y_j\} \in [kN], \{a_j\} \in A\}.$$

Since these edges form a path from s to t , we deduce that $y_0 = x$ and $y_k = x + k \cdot a$. Combining this with the equations for these k edges, we obtain

$$k \cdot a = \sum_{j=0}^{k-1} a_j.$$

Since A is a k -average free set, this implies that $a = a_0 = \dots = a_{k-1}$, and so $\rho'(s, t) = \rho(s, t)$. Therefore, $\rho(s, t)$ is the unique shortest path between s and t , with length k .

Next, we show that these shortest paths are edge disjoint. Any pair is determined by a starting point $(x, 0)$ for some $x \in [N]$ and an element $a \in A$. For any such pair, all the edges on the shortest path have the form $\{(c, j), (c + a, j+1)\}$. Therefore, two pairs that disagree on the element a cannot share any edge on the shortest path. Moreover, if two pairs agree on the element $a \in A$ but disagree on the starting point (i.e., one starts at $(x, 0)$ and the other starts at $(y, 0)$), then the corresponding shortest paths cannot even share a node: All nodes on the first path have the form $(x + j \cdot a, j)$, while nodes on the second path have the form $(y + j \cdot a, j)$. Therefore, the only way two paths can share a node is if $a + j \cdot a = b + j \cdot a$ and so $a = b$.

To conclude the proof, we bound the parameters of our construction. Since $\alpha = \varepsilon/2$ and $\beta < \varepsilon/10$, we have that

$$|P| = \Omega(N^{2-\alpha}) = \Omega\left(n^{\frac{2-\alpha}{1+2\beta}}\right) \geq \Omega\left(n^{\frac{2-\varepsilon/2}{1+2\beta}}\right) = \Omega\left(n^{2-\frac{9\varepsilon}{2\varepsilon+10}}\right) \geq \Omega(n^{2-\varepsilon}),$$

where last inequality follows from the fact that $0 < \varepsilon \leq 1$. In the statement of the lemma we actually require that $|P| = \Theta(n^{2-\varepsilon})$ as opposed to $\Omega(n^{2-\varepsilon})$, but this follows because picking any appropriately-sized subset of P , which will still satisfy all the other properties. Finally, note that the distances are $k = \Theta(N^\beta)$, which is $\Theta(n^\delta)$ for $\delta = \frac{\beta}{1+2\beta}$. Finally, this setting of δ straightforwardly implies that $0 < \delta < \beta < \varepsilon$, so the condition $0 < \delta < \varepsilon$ holds. \square

ACKNOWLEDGMENTS

We are grateful to Virginia Vassilevska Williams for her sustained advice throughout this project and to Seth Pettie and Ryan Williams for advice on an earlier draft of this article that has considerably improved its presentation. We thank Noga Alon and Uri Zwick for helpful comments. We also thank an anonymous reviewer for useful corrections and writing suggestions.

REFERENCES

- Amir Abboud and Greg Bodwin. 2016. Error amplification for pairwise spanner lower bounds. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*. 841–854.
- Amir Abboud, Greg Bodwin, and Seth Pettie. 2017. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA'17)*. 568–576.
- Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. 1999. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.* 28 (1999), 1167–1181.
- Donald Aingworth, Chandra Chekuri, and Rajeev Motwani. 1996. Fast estimation of diameter and shortest paths (without matrix multiplication). In *Proceedings of the 7th ACM-SIAM Symposium on Discrete Algorithms (SODA'96)*. 547–553.
- Noga Alon. 2001. Testing subgraphs in large graphs. In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science (FOCS'01)*. 434–441.
- I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. 1993. On sparse spanners of weighted graphs. *Discr. Comput. Geom.* 9 (1993), 81–100.
- Baruch Awerbuch. 1985. Complexity of network synchronization. *J. ACM* 32, (1985), 804–823.
- Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. 2005. New constructions of (α, β) -spanners and purely additive spanners. In *Proceedings of the 16th ACM-SIAM Symposium on Discrete Algorithms (SODA'05)*. 672–681.
- Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. 2010. Additive spanners and (α, β) -spanners. *ACM Trans. Algor.* 7, 1 (2010), 5. DOI: <http://dx.doi.org/10.1145/1868237.1868242>
- Felix A. Behrend. 1946. On sets of integers which contain no three terms in arithmetical progression. *Proc. Natl. Acad. Sci. U.S.A.* 32, 12 (1946), 331.
- Greg Bodwin. 2017. Linear size distance preservers. In *Proceedings of the the 28th ACM-SIAM Symposium on Discrete Algorithms (SODA'17)*. 600–615.
- Greg Bodwin and Virginia Vassilevska Williams. 2015. Very sparse additive spanners and emulators. In *Proceedings of the 6th International Conference on Innovations in Theoretical Computer Science (ITCS'15)*. 377–382.
- Greg Bodwin and Virginia Vassilevska Williams. 2016. Better distance preservers and additive spanners. In *Proceedings of the 27th ACM-SIAM Symposium on Discrete Algorithms (SODA'16)*.
- Béla Bollobás, Don Coppersmith, and Michael Elkin. 2003. Sparse distance preservers and additive spanners. In *Proceedings of the 14th ACM-SIAM Symposium on Discrete Algorithms (SODA'03)*. 414–423.
- Jean Bourgain. 1985. On lipschitz embedding of finite metric spaces in hilbert space. *Israel J. Math.* 52, 1–2 (1985), 46–52.
- Shiri Chechik. 2013. New additive spanners. In *Proceedings of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA'13)*. 498–512.
- Shiri Chechik. 2014. Additive spanners. In *Encyclopedia of Algorithms*. Springer, New York.
- Victor Chepoi, Feodor F. Dragan, Bertrand Estellon, Michel Habib, Yann Vaxès, and Yang Xiang. 2012. Additive spanners and distance and routing labeling schemes for hyperbolic graphs. *Algorithmica* 62, 3–4 (2012), 713–732. <http://dx.doi.org/10.1007/s00453-010-9478-x>
- Victor Chepoi, Feodor F. Dragan, and Chenyu Yan. 2005. Additive sparse spanners for graphs with bounded length of largest induced cycle. *Theor. Comput. Sci.* 347, 1–2 (2005), 54–75. <http://dx.doi.org/10.1016/j.tcs.2005.05.017>

- Don Coppersmith and Michael Elkin. 2006. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discr. Math.* 20, 2 (2006), 463–501.
- Marek Cygan, Fabrizio Grandoni, and Telikepalli Kavitha. 2013. On pairwise spanners. In *Proceedings of the 30th Symposium on Theoretical Aspects of Computer Science (STACS'13)*. 209–220.
- Dorit Dor, Shay Halperin, and Uri Zwick. 1996. All pairs almost shortest paths. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science (FOCS'96)*. 452–461.
- Feodor F. Dragan and Muad Abu-Ata. 2014. Collective additive tree spanners of bounded tree-breadth graphs with generalizations and consequences. *Theor. Comput. Sci.* 547 (2014), 1–17. <http://dx.doi.org/10.1016/j.tcs.2014.06.007>
- Feodor F. Dragan, Derek G. Corneil, Ekkehard Köhler, and Yang Xiang. 2012. Collective additive tree spanners for circle graphs and polygonal graphs. *Discr. Appl. Math.* 160, 12 (2012), 1717–1729. <http://dx.doi.org/10.1016/j.dam.2012.03.036>
- Feodor F. Dragan, Fedor V. Fomin, and Petr A. Golovach. 2011. Spanners in sparse graphs. *J. Comput. Syst. Sci.* 77, 6 (2011), 1108–1119. <http://dx.doi.org/10.1016/j.jcss.2010.10.002>
- Michael Elkin. 2005. Computing almost shortest paths. *ACM Trans. Algor.* 1, 2 (2005), 283–323.
- Michael Elkin and David Peleg. 2004. $(1 + \epsilon, \beta)$ -spanner constructions for general graphs. *SIAM J. Comput.* 33, 3 (2004), 608–631.
- Paul Erdős. 1964. Extremal problems in graph theory. *Theory of Graphs and Its Applications* (1964), 29–36.
- Telikepalli Kavitha. 2015. New pairwise spanners. In *Proceedings of the 32nd Symposium on Theoretical Aspects of Computer Science (STACS'15)*. 513–526.
- Telikepalli Kavitha and Nithin Varma. 2013. Small stretch pairwise spanners. In *Proceedings of the 40th International Colloquium on Automata, Languages, and Programming (ICALP'13)*. 601–612.
- Mathias Bæk Tejs Knudsen. 2014. Additive spanners: A simple construction. In *Proceedings of the 14th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT'14)*. 277–281.
- Arthur L. Liestman and Thomas C. Shermer. 1991. Additive spanners for hypercubes. *Parallel Process. Lett.* 1, 1 (1991), 35–42.
- Arthur L. Liestman and Thomas C. Shermer. 1993. Additive graph spanners. *Networks* 23, 4 (1993), 343–363.
- Jiří Matoušek. 1996. On the distortion required for embedding finite metric spaces into normed spaces. *Israel J. Math.* 93, 1 (1996), 333–344.
- Merav Parter. 2014. Bypassing erdős' girth conjecture: Hybrid stretch and sourcewise spanners. In *Proceedings of the 41st International Colloquium on Automata, Languages, and Programming (ICALP'14)*. 608–619.
- David Peleg and Alejandro Schaffer. 1989. Graph spanners. *J. Graph Theory* 13 (1989), 99–116.
- David Peleg and Jeffrey D. Ullman. 1989. An optimal synchronizer for the hypercube. *SIAM J. Comput.* 18 (1989), 740–747.
- Seth Pettie. 2009. Low distortion spanners. *ACM Trans. Algor.* 6, 1 (2009). DOI : <http://dx.doi.org/10.1145/1644015.1644022>
- Liam Roditty, Mikkel Thorup, and Uri Zwick. 2005. Deterministic constructions of approximate distance oracles and spanners. In *Proceedings of the 32nd International Colloquium on Automata, Languages, and Programming (ICALP'05)*. 261–272.
- Mikkel Thorup and Uri Zwick. 2005. Approximate distance oracles. *J. ACM* 52, 1 (2005), 1–24.
- Mikkel Thorup and Uri Zwick. 2006. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th ACM-SIAM Symposium on Discrete Algorithms (SODA'06)*. 802–809.
- David P. Woodruff. 2006. Lower bounds for additive spanners, emulators, and more. In *Proceedings of the 47th IEEE Symposium on Foundations of Computer Science (FOCS'06)*. 389–398.
- David P. Woodruff. 2010. Additive spanners in nearly quadratic time. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming (ICALP'10)*. 463–474.

Received November 2016; revised April 2017; accepted April 2017