On the Size Overhead of Pairwise Spanners

Ofer Neiman ✉
Ben-Gurion University of the Negev, Beer-Sheva, Israel

Idan Shabat ✉
Ben-Gurion University of the Negev, Beer-Sheva, Israel

Abstract

Given an undirected possibly weighted n-vertex graph $G = (V, E)$ and a set $\mathcal{P} \subseteq V^2$ of pairs, a subgraph $S = (V, E')$ is called a $\mathcal{P}$-*pairwise* α-spanner of G, if for every pair $(u, v) \in \mathcal{P}$ we have $d_S(u, v) \leq \alpha \cdot d_G(u, v)$. The parameter α is called the *stretch* of the spanner, and its *size overhead* is define as $\frac{|E'|}{|\mathcal{P}|}$.

A surprising connection was recently discussed between the additive stretch of $(1+\epsilon, \beta)$-spanners, to the hopbound of $(1+\epsilon, \beta)$-hopsets. A long sequence of works showed that if the spanner/hopset has size $\approx n^{1+1/k}$ for some parameter $k \geq 1$, then $\beta \approx \left(\frac{1}{\epsilon}\right)^{\log k}$. In this paper we establish a new connection to the *size overhead* of pairwise spanners. In particular, we show that if $|\mathcal{P}| \approx n^{1+1/k}$, then a $\mathcal{P}$-pairwise $(1+\epsilon)$-spanner must have size at least $\beta \cdot |\mathcal{P}|$ with $\beta \approx \left(\frac{1}{\epsilon}\right)^{\log k}$ (a near matching upper bound was recently shown in [18]). That is, the size overhead of pairwise spanners has similar bounds to the hopbound of hopsets, and to the additive stretch of spanners.

We also extend the connection between pairwise spanners and hopsets to the large stretch regime, by showing nearly matching upper and lower bounds for $\mathcal{P}$-pairwise α-spanners. In particular, we show that if $|\mathcal{P}| \approx n^{1+1/k}$, then the size overhead is $\beta \approx \frac{k}{\alpha}$.

A *source-wise spanner* is a special type of pairwise spanner, for which $\mathcal{P} = A \times V$ for some $A \subseteq V$. A *prioritized* spanner is given also a ranking of the vertices $V = (v_1, \ldots, v_n)$, and is required to provide improved stretch for pairs containing higher ranked vertices. By using a sequence of reductions: from pairwise spanners to source-wise spanners to prioritized spanners, we improve on the state-of-the-art results for source-wise and prioritized spanners. Since our spanners can be equipped with a path-reporting mechanism, we also substantially improve the known bounds for path-reporting prioritized distance oracles. Specifically, we provide a path-reporting distance oracle, with size $O(n \cdot (\log\log n)^2)$, that has a constant stretch for any query that contains a vertex ranked among the first $n^{1-\delta}$ vertices (for any constant $\delta > 0$). Such a result was known before only for non-path-reporting distance oracles.

2012 ACM Subject Classification Theory of computation → Shortest paths

Keywords and phrases Graph Algorithms, Shortest Paths, Spanners

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.83

Related Version *Full Version*: `https://arxiv.org/pdf/2311.13673.pdf`

Funding Partially supported by the Lynn and William Frankel Center for Computer Sciences and ISF grant 970/21.

1 Introduction

Spanners and hopsets are basic graph structures that have been extensively studied, and found numerous applications in graph algorithms, distributed computing, geometric algorithms, and many more. In this work we study pairwise spanners, and prove an intriguing relation between pairwise spanners, general spanners, and hopsets. We then derive several new results on *source-wise* and *prioritized* spanners and distance oracles.

Let $G = (V, E)$ be an undirected n-vertex graph, possibly with nonnegative weights on the edges. An (α, β)-*spanner* is a subgraph $S = (V, E')$ such that for every pair $u, v \in V$

$$d_S(u, v) \leq \alpha \cdot d_G(u, v) + \beta \ , \tag{1}$$

where d_G, d_S stand for the shortest path distances in G, S, respectively. A spanner is called *near-additive* if its multiplicative stretch is $\alpha = 1 + \epsilon$ for some small $\epsilon > 0$.[1] Given a set $\mathcal{P} \subseteq V^2$ of pairs, a $\mathcal{P}$-*pairwise* α-spanner for G has to satisfy Equation (1) (with $\beta = 0$) only for pairs $(u, v) \in \mathcal{P}$. The *size overhead* of a pairwise spanner is defined as $\frac{|E'|}{|\mathcal{P}|}$. A *source-wise spanner*[2] is a special type of pairwise spanner, in which $\mathcal{P} = A \times V$ for some set $A \subseteq V$.

An $(\alpha, \beta)-$*hopset* for a graph $G = (V, E)$ is a set of edges $H \subseteq \binom{V}{2}$ such that for every $u, v \in V$,

$$d^{(\beta)}_{G \cup H}(u, v) \leq \alpha \cdot d_G(u, v) \ .$$

Here, $G \cup H$ denotes the graph G with the additional edges of H, and the weight function $w(x, y) = d_G(x, y)$ for every $\{x, y\} \in H$. The notation $d^{(\beta)}_{G \cup H}(u, v)$ stands for the weight of the shortest $u - v$ path in this graph, among the ones that contain at most β edges. A hopset is called *near-exact* if its stretch is $\alpha = 1 + \epsilon$ for some small $\epsilon > 0$.

1.1 Pairwise Spanners, Near-additive Spanners, and Hopsets

In [16, 21], a connection was discussed between the additive stretch β of near-additive spanners, to the hopbound β of near-exact hopsets. Given an integer parameter $k \geq 1$, that governs the size of the spanner/hopset to be $\approx n^{1+1/k}$, a sequence of works on spanners [17, 29, 26, 2, 13, 12] and on hopsets [8, 6, 19, 23, 14, 20, 15], culminated in achieving $\beta = O\left(\frac{\log k}{\epsilon}\right)^{\log k}$ for both spanners and hopsets. In [2] a lower bound of $\beta \geq \Omega_k\left(\frac{1}{\epsilon}\right)^{\log k - 1}$ was shown. So whenever ϵ is sufficiently small, we have that $\beta \approx \left(\frac{1}{\epsilon}\right)^{\log k}$. As spanners and hopsets are different objects, and β has a very different role for each, this similarity is somewhat surprising (albeit comparable techniques are used for the constructions).

In this paper we establish an additional connection, to pairwise spanners. Here the parameter k governs the number of pairs, $|\mathcal{P}| \approx n^{1+1/k}$,[3] and we show that $\mathcal{P}$-pairwise $(1 + \epsilon)$-spanners must have size at least $\beta \cdot |\mathcal{P}|$ with $\beta \approx \left(\frac{1}{\epsilon}\right)^{\log k}$. So the parameter β, which is the additive stretch for near-exact spanners, and is the hopbound for hopsets, now plays the role of the *size overhead* for pairwise spanners.

An exact version of pairwise spanners (i.e., with $\alpha = 1$) was introduced in [9], where they were called *distance preservers*. The sparsest distance preservers that are currently known are due to [9, 7]. For an n-vertex graph and a set of pairs $\mathcal{P}$, they have size $O\left(\min\{n^{2/3}|\mathcal{P}|, n|\mathcal{P}|^{1/2}\} + n\right)$. So whenever $|\mathcal{P}| \leq n^{2-\delta}$ for some constant $\delta > 0$, the size overhead β is polynomial in n.

When considering near-exact pairwise spanners, following [21], in [18] a $\mathcal{P}$-pairwise $(1 + \epsilon)$-spanner of size $\approx \beta \cdot |\mathcal{P}|$, with $\beta = O\left(\frac{\log k}{\epsilon}\right)^{\log k}$ was shown (where k is such that $|\mathcal{P}| \approx n^{1+1/k}$).

---

[1] Spanners with additive stretch $\beta > 0$ are usually defined on unweighted graphs. A possible variation for weighted graph has also been studied [12, 3], in which β is multiplied by the weight of the heaviest edge on some $u - v$ shortest path.

[2] Source-wise spanners were called *terminal spanners* in [11].

[3] In this paper we focus on $\mathcal{P}$-pairwise spanners with $|\mathcal{P}| \geq n$. Note that for $|\mathcal{P}| < n$ there is a trivial $\Omega(n)$ lower bound on the size.

Our Results

In this paper we show a near matching lower bound for $\mathcal{P}$-pairwise $(1+\epsilon)$-spanners, that $\beta \geq \Omega\left(\frac{1}{\log^2 k\cdot\epsilon}\right)^{\log k-1}$, establishing that the size overhead must be $\beta \approx \left(\frac{1}{\epsilon}\right)^{\log k}$. We derive this lower bound by a delicate adaptation of the techniques of [2] to the case of pairwise spanners.

1.1.1 Larger Stretch Regime

As the stretch grows to be bounded away from 1, the connection between spanners and hopsets diminishes (note that the lower bound of [2] is meaningless for large stretch). For size $\approx n^{1+1/k}$ and constant stretch α, one can obtain (α, β)-spanners and hopsets with $\beta \approx k^{1+1/\log\alpha}$ [5, 24]. However, as the stretch grows to $\alpha = k^\delta$ (for some constant $0 < \delta < 1$), there is a hopset with $\beta \approx k^{1-\delta}$ [5, 24], while spanners must have $\beta = \Omega(k)$.[4] In [24], a lower bound of $\beta = \Omega\left(\frac{k}{\alpha}\right)$ for hopsets was shown.

Our Results

In this work we exhibit a similar tradeoff for pairwise spanners, as is known for hopsets.[5] In particular, we devise a $\mathcal{P}$-pairwise α-spanner with size $\beta \cdot |\mathcal{P}|$ where $\beta \approx \frac{k^{1+1/\log\alpha}}{\alpha}$ (and k is such that $|\mathcal{P}| \approx n^{1+1/k}$). We also show a lower bound $\beta = \Omega\left(\frac{k}{\alpha}\right)$ on the size overhead of such pairwise spanners. Note that this lower bound nearly matches the upper bound for $\alpha \geq k^\delta$ for any constant $\delta > 0$, and in this regime we have $\beta \approx \frac{k}{\alpha}$ for both hopsets and pairwise spanners.

1.2 Source-wise and Prioritized Spanners and Distance Oracles

Source-wise spanners were first studied by [27, 25]. Given an integer parameter $k \geq 1$ and a subset $A \subseteq V$ of sources, they showed a source-wise spanner with stretch $2k-1$ and size $O(kn \cdot |S|^{1/k})$ (the spanners of [27] could also be distance oracles, while [25] had slightly improved stretch $2k-2$ for some pairs). By increasing the stretch to $4k-1$, [11] obtained improved size $O(n + \sqrt{n} \cdot |S|^{1+1/k})$. The current state-of-the-art result is by [21]. For any $0 < \epsilon < 1$, they gave a source-wise spanner with stretch $4k-1+\epsilon$ and size $O(n+|S|^{1+1/k}) \cdot \beta$, where $\beta \approx \left(\frac{\log\log n}{\epsilon}\right)^{\log\log n}$.

Given an undirected possibly weighted graph $G = (V, E)$, a *prioritized* metric structure (such as spanner, hopset, distance oracle) is also given an ordering of the vertices $V = (v_1, \ldots, v_n)$, and is required to provide improved guarantees (e.g., stretch, hopbound, query time, etc.) for higher ranked vertices.

In [10], among other results, prioritized distance oracles for general graphs were shown.[6] For an n-vertex graph with priority ranking $(v_1, \ldots, v_n)$, a distance oracle with size $O(n \log\log n)$, query time $O(1)$, and stretch $O\left(\frac{\log n}{\log n - \log j}\right)$ for any pair containing v_j was shown.[7] Note that the stretch is constant for any pair containing a vertex ranked among

[4] Note that an (α, β)-spanner is also an $(\alpha+\beta, 0)$-spanner. Thus, with size $n^{1+1/k}$ it must have $\alpha + \beta \geq k$ [4].

[5] For stretch $3+\epsilon$ with $0 < \epsilon < 1$, a pairwise spanner of size $\beta \cdot |\mathcal{P}|$ with $\beta \approx k^{\log 3 + O(\log(1/\epsilon))}$ was shown in [18]. However, we are not aware of any result for larger stretch.

[6] A distance oracle is a data structure that can efficiently report approximate distances. The parameters of interest are usually the query time, the size, and the multiplicative stretch.

[7] Additional results with even smaller size and larger stretch were shown in [10] as well.

the first $n^{1-\delta}$ vertices (for constant $\delta > 0$). However, that distance oracle could only return distances, and could not report paths. In [10] an additional *path-reporting* distance oracle was shown. Given an integer parameter $k \geq 1$, it had size $O(kn^{1+1/k})$, stretch $2\left\lceil\frac{k\log j}{\log n}\right\rceil - 1$ for pairs containing v_j, and query time $O\left(\left\lceil\frac{k\log j}{\log n}\right\rceil\right)$. Note that this oracle size is $\Omega(n\log n)$ for any k, and with size $O(n\log n)$ it has prioritized stretch $2\log j - 1$, which is much worse than the stretch that the previous oracle had for higher ranked vertices.

Our Results

By using a sequence of reductions, from pairwise spanners to source-wise spanners to prioritized spanners, we improve on the state-of-the-art results. In particular, for any integer parameter $k \geq 1$, any $0 < \epsilon < 1$, and a subset $A \subseteq V$, we obtain a source-wise spanner with stretch $4k - 1 + \epsilon$ and size $O(n\log\log n + |S|^{1+1/k} \cdot \beta)$ where $\beta \approx \left(\frac{\log\log n}{\epsilon}\right)^{\log\log n}$. Note that β only multiplies the term $|S|^{1+1/k}$, which could be much smaller than n, while in [21] the size is always at least $\Omega(n \cdot \beta)$.

Using the fact that the constructions of pairwise spanners can also be *path-reporting*, and that our reductions preserve this property, we devise *path-reporting* distance oracles with size $O(n(\log\log n)^2)$, query time $O(1)$, and prioritized stretch $O\left(\frac{\log n}{\log n - \log j - o(\log\log n)^2}\right)$. That is, we nearly achieve the improved parameters of the non-path-reporting oracles of [10] (in particular, we get constant stretch for any pair containing a vertex ranked among the first $n^{1-\delta}$ vertices).

1.3 Technical Overview

1.3.1 Lower Bound for Near-Exact Pairwise Spanners

In [2], a series of lower bounds was shown for graph compression structures, such as near-additive spanners, emulators, distance oracles and hopsets. Specifically, for the first three, [2] proved that for any positive integer κ, any such structure with size $O(n^{1+\frac{1}{2^\kappa-1}})$, that preserves distances d up to $(1+\epsilon)d + \beta$, must have $\beta = \Omega\left(\frac{1}{\epsilon\kappa}\right)^{\kappa-1}$. Almost the same lower bound[8] was proved by [2] for $(1+\epsilon,\beta)$-hopsets of size $O(n^{1+\frac{1}{2^\kappa-1}})$.

All the lower bounds mentioned above were demonstrated in [2] on essentially the same graph. The construction of this graph relied on a *base graph* $\ddot{B}$, that was presented in [1] and had a *layered structure*. That is, the vertices of $\ddot{B}$ are partitioned into $2l + 1$ subsets, such that any edge can only be between vertices of adjacent layers. The first and the last layers of $\ddot{B}$ are called *input ports* and *output ports* respectively. This graph also had a set $\ddot{\mathcal{P}}$ of pairs of input and output ports (u, v) such that there is a unique shortest path $P_{u,v}$ from u to v in $\ddot{B}$, that visits every layer exactly once. Moreover, each edge of $\ddot{B}$ is labeled by a label from a set $\ddot{\mathcal{L}}$, such that the edges of every such path $P_{u,v}$ are alternately labeled by a *unique* pair of labels $a, b \in \ddot{\mathcal{L}}$ (meaning that no other pair $(u', v') \in \ddot{\mathcal{P}}$ has its shortest path $P_{u',v'}$ alternately labeled by the same labels a, b).

On top of the base graph $\ddot{B}$, a graph $\mathcal{H}_\kappa$ was constructed, for every positive integer κ. In fact, for *hopsets*, a slightly different graph was constructed in [2] than for *near-additive spanners* (denoted as $\mathcal{H}_\kappa$ in both cases). The two constructions are recursive, with different

[8] For hopsets of size $O(n^{1+\frac{1}{2^\kappa-1}})$ and stretch $1+\epsilon$, the actual lower bound on β that was proved in [2] was $\beta = \Omega\left(\frac{1}{\epsilon 2^\kappa}\right)^{\kappa-1}$, instead of $\Omega\left(\frac{1}{\epsilon\kappa}\right)^{\kappa-1}$. It was suggested though, that a more careful analysis might achieve the same lower bound as of $(1+\epsilon,\beta)$-spanners.

base cases. Moreover, since the construction for near-additive spanners must be unweighted, edges with large weight in the construction for hopsets, must be replaced by long unweighted paths. In this paper, we construct a graph, still denoted as $\mathcal{H}_\kappa$, that is actually a *combination* of their two constructions. In particular, we use the same base case as for near-additive spanners, which is simply a complete bipartite graph $K_{p,p}$, while using edges of large weight, similarly to the construction of $\mathcal{H}_\kappa$ for hopsets. In addition, we shift the indices of the sequence $\{\mathcal{H}_\kappa\}$ by 1. We describe here the construction of $\mathcal{H}_1$, as the general construction is more involved and appears fully in Section 3.1.

The graph $\mathcal{H}_1$ is the same graph as $\ddot{B}$, where every vertex in its $2l-1$ middle layers is replaced by a copy of the complete bipartite graph $K_{p,p}$, where p is the number of labels in $\ddot{\mathcal{L}}$ (for $\kappa > 1$, these vertices are replaced by copies of $\mathcal{H}_{\kappa-1}$). Each vertex in each side of the copy of $K_{p,p}$ is mapped to a label in $\ddot{\mathcal{L}}$. An original edge of $\ddot{B}$, that had label a, now connects these copies of $K_{p,p}$, by their vertices that correspond to the label a (see Figure 1 in Section 3.1 for a detailed illustration). In addition, these edges are assigned with a large weight of $2l-1$ (for $\kappa > 1$, this weight is $(2l-1)^\kappa$). This means that any shortest path in $\ddot{B}$ between a pair $(u,v) \in \ddot{\mathcal{P}}$ now must pass through $2l-1$ copies of $K_{p,p}$. Other paths that connect u, v, on the other hand, must visit a layer more than once, and therefore their weight is larger than $d_{\mathcal{H}_1}(u,v)$ by at least $2(2l-1)$. Choosing $l \approx \frac{1}{\epsilon}$ (for $\kappa > 1$, $l \approx \left(\frac{1}{\epsilon\kappa}\right)^\kappa$), this means that paths other than the unique shortest path have stretch more then $1 + \frac{2(2l-1)}{d(u,v)} \approx 1 + \frac{2(2l-1)}{2l(2l-1)} \approx 1+\epsilon$. Hence, any $\ddot{\mathcal{P}}$-pairwise $(1+\epsilon)$-spanner must contain the unique shortest paths between any pair $(u,v) \in \ddot{\mathcal{P}}$.

In our proof of a lower bound for pairwise spanners, we utilize an additional property of the graph $\mathcal{H}_1$. Recall that the edges of $P_{u,v}$ (a shortest path in $\ddot{B}$ that connects some $(u,v) \in \ddot{\mathcal{P}}$) are alternately labeled by a unique pair of labels $a, b \in \ddot{\mathcal{L}}$. This means that in $\mathcal{H}_1$, shortest paths that connect different pairs $(u,v), (u',v') \in \ddot{\mathcal{P}}$ cannot share an edge of a copy of $K_{p,p}$. This is due to the fact that if such path goes through the edge (x,y) of a copy of $K_{p,p}$, and x corresponds to a label a, and y to a label b, then it uniquely determines the pair of labels a, b of the path $P_{u,v}$. The result is that any $\ddot{\mathcal{P}}$-pairwise $(1+\epsilon)$-spanner for $\mathcal{H}_1$ must contain a disjoint set of $2l-1$ edges, for every $(u,v) \in \ddot{\mathcal{P}}$. Hence, its size must be at least

$$(2l-1) \cdot |\ddot{\mathcal{P}}| = \Omega\left(\frac{1}{\epsilon} \cdot |\ddot{\mathcal{P}}|\right) .$$

For larger κ's, the number of the edges that these paths do not share grows to $\beta = \Omega_\kappa\left(\frac{1}{\epsilon}\right)^\kappa$, therefore the lower bound for the size is $\beta|\mathcal{P}_\kappa|$, where $\mathcal{P}_\kappa$ is the corresponding set of pairs of the graph $\mathcal{H}_\kappa$. It can be proved that the number of pairs in $\mathcal{P}_\kappa$ is approximately $n^{1+\frac{1}{2^{\kappa+1}-1}}$, where n is the number of vertices in $\mathcal{H}_\kappa$.

1.3.2 Lower Bound for Pairwise Spanners with Large Stretch

Our proof of a lower bound for pairwise spanners with stretch larger than $1+\epsilon$ (typically a constant stretch, or stretch between 2 and k, where k is such that $|\mathcal{P}| \approx n^{1+\frac{1}{k}}$) is demonstrated on an unweighted graph with high number of edges and high girth (the length of the minimal cycle). This kind of graphs was used in [28, 24] to show lower bounds for spanners, distance oracles and hopsets. Specifically, we use a graph $G = (V, E)$ from [22], that is regular, has $\Omega(n^{1+\frac{1}{k}})$ edges, and has girth larger than k.

Given a desired stretch $\alpha > 1$, we consider pairs of vertices in G of distance $\delta = \left\lfloor \frac{k}{\alpha+1} \right\rfloor$, henceforth δ-pairs. The useful property of δ-pairs is that due to the high girth of G, they have a unique shortest path that connects them, while any other path must have stretch more that α. Note the similarity of this property to the property of the unique shortest paths from the lower bound for near-exact pairwise spanners. This property implies that any $\mathcal{P}$-pairwise α-spanner, for any set $\mathcal{P}$ of δ-pairs, must contain all the unique shortest paths that connect the δ-pairs of $\mathcal{P}$. We call such paths the δ-paths that correspond to the δ-pairs in $\mathcal{P}$.

Then, we use the regularity of the graph $G = (V, E)$, as well as the high girth of G, to prove some combinatorial properties of G. We prove that there are many δ-pairs - approximately $n^{1+\frac{\delta}{k}}$ pairs - and that each edge $e \in E$ participates in a large number of corresponding δ-paths - approximately $\delta n^{\frac{\delta-1}{k}}$.

To finally prove our lower bound for pairwise α-spanners, we sample a set $\mathcal{P}$ of $O\left(\frac{n^{1+\frac{1}{k}}}{\delta}\right)$ δ-pairs out of all the possible $n^{1+\frac{\delta}{k}}$ such pairs. This means that the sample probability is $\frac{1}{\delta} n^{-\frac{\delta-1}{k}}$. But since the number of δ-pairs whose corresponding δ-path pass through a specific edge $e \in E$ is approximately $\delta n^{\frac{\delta-1}{k}}$, we expect that a constant fraction of the edges of G still participate in a unique shortest path of the δ-pairs in $\mathcal{P}$. Every $\mathcal{P}$-pairwise α-spanner must contain these edges, and therefore must have size

$$\Omega(|E|) = \Omega(n^{1+\frac{1}{k}}) = \Omega(\delta|\mathcal{P}|) = \Omega\left(\frac{k}{\alpha}|\mathcal{P}|\right) .$$

This proves our lower bound of $\beta = \Omega\left(\frac{k}{\alpha}\right)$.

1.3.3 Upper Bound for Pairwise Spanners with Large Stretch

The state-of-the-art pairwise $(1+\epsilon)$-spanners and pairwise $(3+\epsilon)$-spanners of [18] were achieved using the following *semi-reduction* from hopsets. Suppose that the hopset H has stretch α and hopbound β, for a graph $G = (V, E)$. Given a set of pairs of vertices $\mathcal{P}$, we construct a pairwise spanner S that consists of all the β-edges paths in $G \cup H$ that connect pairs in $\mathcal{P}$ and have stretch α. Of course, these paths include edges of H that are not allowed to be on the final pairwise spanner S (as they don't exist in G). Therefore, an additional step is required, that adds more edges to S, such that every edge $(x, y) \in H$ will have a path in S with weight $w(x, y)$. This way, the distance between every pair $(u, v) \in \mathcal{P}$ is the same distance as in the graph $G \cup H$, which is at most $\alpha \cdot d_G(u, v)$. The size of the pairwise spanner S is $\beta \cdot |\mathcal{P}|$, plus the number of edges that are required to preserve the distances between every $(x, y) \in H$.

In [18], preserving the distances between pairs in H is done by relying on the specific properties of the hopset H. Namely, it was observed that the relevant hopset H has small *supporting size* - the minimal number of edges required to preserve the distances in H. We, however, take an alternative approach. Instead of relying on specific properties of the hopset H to preserve it *accurately*, we use a pairwise spanner with low stretch, to preserve the distance in H *approximately*. In fact, we use the very same H-pairwise $(1+\epsilon)$-spanner of [18] to do that.

Considering the process described above, we are only left to choose the hopset H that we use, to achieve a pairwise spanner with relatively large stretch. The advantage of having a large stretch α, is that the hopbound β can be much smaller, and as a result, so is the size overhead of the resulting pairwise spanner. In particular, for pairwise $(1+\epsilon)$-spanners, we know by Section 3.1 that the size overhead must be $\Omega_k\left(\frac{1}{\epsilon}\right)^{\log k}$. But for a larger stretch $O(\alpha)$, we can get a much smaller hopbound, and therefore a much smaller size overhead, of $\beta = \frac{k^{1+\frac{2}{\ln\alpha}}}{\alpha}$. This is done by using the state-of-the-art hopsets with this type of stretch, from [5, 24].

However, the process of directly applying a pairwise spanner with low stretch on a hopset described above, results in a somewhat large size of the pairwise spanner. This is because the size of the hopset, which is roughly $n^{1+\frac{1}{k}}$, is multiplied by the size overhead coefficient β, which is at least poly(k). To reduce this size, and avoid the additional term of poly$(k) \cdot n^{1+\frac{1}{k}}$, we eventually do use certain properties of the hopsets of [24]. Specifically, we prove that their hopsets could be partitioned into three sets H_1, H_2, H_3, such that H_1, H_2 can be efficiently preserved, while H_3 is relatively small. Thus, we can use a pairwise spanner with low stretch only on H_3, instead of using it on the whole hopset H. Then, the coefficient $\beta \approx poly(k)$ multiplies only the size of H_3, which is significantly smaller than $n^{1+\frac{1}{k}}$.

1.3.4 Subset, Source-wise and Prioritized Spanners

Our new results for subset, source-wise and prioritized spanners are achieved via a series of reductions. These reductions implicitly appeared in [18, 10]. The first reduction receives a pairwise spanner and uses it in order to construct a *subset* spanner. A *subset spanner* is a special type of pairwise spanner, for which $\mathcal{P} = A \times A$ for some $A \subseteq V$. The second is a quite simple reduction that turns a subset spanner into a source-wise spanner with almost the same properties. The last reduction uses source-wise spanners to construct prioritized spanners. Using the new state-of-the-art pairwise spanners, we achieve new results for each of these types of spanners.

Besides these new results, we believe that these reductions themselves could be of interest. It is immediate to find backwards reductions, from prioritized spanners to source-wise spanners, and from source-wise spanners to subset spanners. This means that, in a sense, these three[9] types of spanners are equivalently hard to construct. Every new construction of one of these spanners immediately implies new constructions for the others.

Next we shortly describe the three reductions mentioned above.

From pairwise to subset spanners

Let $G = (V, E)$ be an undirected weighted n-vertex graph, and let $A \subseteq V$ be a subset, for which we want to construct a subset spanner. We consider the graph $K = (A, \binom{A}{2})$, where every pair of vertices $u, v \in A$ is connected by an edge of weight $d_G(u, v)$. On the graph K, we apply a known construction of *emulator*. That is, we find a small graph $R = (A, E')$, such that

$$d_G(u, v) = d_K(u, v) \leq d_R(u, v) \leq \alpha_E \cdot d_G(u, v) \ ,$$

for any $u, v \in A$. The parameter α_E is the stretch of the emulator R. For this purpose, we use either the distance oracle of Thorup and Zwick from [28], which can be thought of as an emulator with stretch $\alpha_E = 2k - 1$ and size $O(k|A|^{1+\frac{1}{k}})$, or the emulator from [18], that has stretch $\alpha_E = O(k)$ and size $O(|A|^{1+\frac{1}{k}})$. Here, k is a positive integer parameter for our choice. These two emulators are *path-reporting*, meaning that given $u, v \in A$, they can also efficiently report a path of stretch α_E inside the emulator itself.

Next, we consider the set R as a set of pairs of vertices from G, and apply a pairwise spanner on this set. We use the existing pairwise spanners from [18] that have stretch either $\alpha_P = 1 + \epsilon$ or $\alpha_P = 3 + \epsilon$, and size[10]

$$O(|R| \cdot \beta + n \log \log n) \ .$$

[9] Note that the reduction that constructs a subset spanner, given a pairwise spanner, remains a one-way reduction, in the sense that no efficient reduction in the other direction is known.

[10] In [18], the size of these pairwise spanners was presented as $O(|\mathcal{P}| \cdot \beta + n \log k + n^{1+\frac{1}{k}})$, where $\beta \approx \left(\frac{\log k}{\epsilon}\right)^{\log k}$, or $\beta = k^{O\left(\log \frac{1}{\epsilon}\right)}$, and $\mathcal{P}$ is the set of pairs. We choose here $k = \log n$ since for small

Here, $\beta \approx \left(\frac{\log\log n}{\epsilon}\right)^{\log\log n}$ when $\alpha_P = 1+\epsilon$, and $\beta = (\log n)^{O\left(\log\frac{1}{\epsilon}\right)}$ when $\alpha_P = 3+\epsilon$ (there is also additional dependency on ϵ in the size of the pairwise spanner, in case that $\alpha_P = 1+\epsilon$). These two pairwise spanners are also path-reporting.

Now fix a pair of vertices $(u,v) \in A^2$. In the graph K, the shortest path between u, v is the single edge (u,v), that has weight $d_G(u,v)$, and thus in R there is a path between u, v that has weight $\alpha_E \cdot d_G(u,v)$. Note that this path consists of edges of R, which are not actual edges of the graph G. For that reason, we use our pairwise spanner, to find, for every edge e on this path, a path in G that replaces e and has weight of at most $\alpha_P \cdot w(e)$. The result is a path in G, with stretch $\alpha_E \cdot \alpha_P$. That is, the resulting spanner has a $\alpha_E \cdot \alpha_P$-stretch path for every pair of vertices in A. Hence, this is a subset spanner. Note that since both the emulator R and the pairwise spanner we used were path-reporting, then so is the resulting subset spanner.

From subset to source-wise spanner

Let $G = (V,E)$ be an undirected weighted graph, and let $A \subseteq V$ be a subset. Suppose that S is an A-subset spanner with stretch α. To construct a source-wise spanner for A, we just add to S a shortest path, from every vertex $v \in V$ to its closest vertex $p(v)$ in A. If the shortest paths and the vertices $p(v)$ are chosen in a consistent manner, it is not hard to prove that the added edges form a *forest*. Thus, the increase in the size of the spanner S is by at most $n-1$ edges, and also we can easily navigate from a vertex $v \in V$ to the corresponding $p(v) \in A$ (we simply go towards the root of v's tree).

Given any $v \in V$ and $a \in A$, the spanner contains the $v - p(v)$ shortest path, concatenated with the path in S from $p(v)$ to a. It can be shown that it has at most $2\alpha+1$ stretch. Also, if the subset spanner S is path-reporting, then so is the new source-wise spanner.

From source-wise to prioritized spanner

In the setting of prioritized spanners, we get a priority ranking of the vertices of V: $(v_1, v_2, ..., v_n)$. To construct a prioritized spanner based on source-wise spanners, we consider a sequence of non-decreasing *prefixes* of this list:

$$A_1 = \{v_1, ..., v_{f(1)}\},\ \ A_2 = \{v_1, ..., v_{f(1)}, ..., v_{f(2)}\},\ \cdots\ A_T = \{v_1, ..., v_{f(1)}, ..., v_{f(2)}, ..., v_{f(T)}\}\ ,$$

where f is some non-decreasing function, and we apply a source-wise spanner on each of them. The idea is that in the construction of the source-wise spanners for the first prefixes in the list, we may use a smaller stretch. Then, since these prefixes are small enough, the size of the resulting source-wise spanners will not be too large.

Specifically, our path-reporting prioritized spanner with size $O(n(\log\log n)^2)$ is achieved by using a sequence of prefixes of sizes $n^{\frac{1}{2}}, n^{\frac{3}{4}}, n^{\frac{7}{8}}, ..., n^{1-\frac{1}{2^i}}$. Recall that the size of an A-source-wise spanner with stretch $O(k)$ is $O(|A|^{1+\frac{1}{k}} \cdot \beta + n\log\log n)$, where $\beta = n^{o(1)}$. Note that for $A = A_i$, by choosing $k = 2^i - 1$, the first term is $|A_i|^{1+\frac{1}{k}} \leq n^{\frac{2^i-1}{2^i}\cdot\frac{2^i}{2^i-1}} = n$. Thus, we must choose k a bit larger than $2^i - 1$, in order to obtain an A_i-source-wise spanner with size $O(n\log\log n)$ (this is the sparsest A_i-source-wise spanner we can obtain, since our pairwise/subset/source-wise spanners always have an additive term of $n\log\log n$ in their size). Namely, for cancelling the factor β, we must choose $k = \frac{2^i-1}{1-2^i\cdot o(1)}$.

---

enough $\mathcal{P}$, this is the most sparse that these spanners can get. In our case, the set $\mathcal{P}$ is the emulator R, which is indeed small enough.

In the choice of k above, the $o(1)$ factor in the denominator is actually $\frac{\log \beta}{\log n}$. This means that we cannot make this choice for prefixes A_i with $i > \log\log n - \log\log\beta = (1 - o(1))\log\log n$. For this reason, we cannot choose T such that the prefixes $A_1, ..., A_T$ will cover the entire priority ranking $(v_1, ..., v_n)$, or most of it (to cover $\frac{n}{2}$ vertices, we need $T = \log\log n$). Instead, we choose T to be roughly $\log\log n - \log\log\beta$. The size of the resulting prioritized spanner is $T \cdot O(n \log\log n) = O(n(\log\log n)^2)$, and the covered vertices are the vertices v_j such that $j \leq \frac{n}{\beta}$. For queries that include a covered vertex v_j, the stretch is roughly $O(2^i)$, where i is the minimal integer such that $j \leq n^{1-\frac{1}{2^i}}$. That is, the stretch is approximately $O\left(\frac{\log n}{\log n - \log j}\right)$.

For uncovered vertices (i.e., vertices v_j such that $j > \frac{n}{\beta}$), we simply add a (non-prioritized) spanner, that has stretch $O\left(\frac{\log n}{\log\log\log n}\right)$ for *all* the vertices of G. The size of such spanner may be as low as $O(n(\log\log n)^2)$, using results from [18]. This spanner, as well as all the source-wise spanners we used, are path-reporting, therefore the resulting prioritized spanner is also path-reporting.

We provide another variation of a path-reporting prioritized spanner with reduced stretch, at the cost of increasing the size to $O(n \log n)$. To achieve that, we change the sequence of prefixes we use, to a sequence of prefixes with sizes $n^{\frac{1}{2}}, n^{\frac{2}{3}}, n^{\frac{3}{4}}, ..., n^{1-\frac{1}{i}}$. This sequence grows slower than the previous one. While increasing the size of the resulting prioritized spanner, this enable us to control the stretch of each source-wise spanner more carefully.

1.4 Organization

After some preliminaries in Section 2, we prove our new lower bounds for pairwise spanners in Section 3 (lower bounds for pairwise spanners with stretch $1 + \epsilon$ are in Section 3.1, while Section 3.2 is for higher stretch).

In the full version of this paper, we prove our new upper bounds for pairwise spanners in the high stretch regime, and we show the reductions between pairwise, subset, source-wise and prioritized spanners, which result in new upper bounds for these types of spanners.

2 Preliminaries

Given an undirected weighted graph $G = (V, E)$, we denote by $d_G(u, v)$ the distance between the two vertices $u, v \in V$. When the graph G is clear from the context, we sometimes omit the subscript G and write $d(u, v)$.

When the given graph is weighted, $w(e) > 0$ denotes the weight of the edge e. For every set of edges P (e.g., a tree or a path), we denote $w(P) = \sum_{e \in P} w(e)$. If P is a path, we denote by $|P|$ the length of P, i.e., the number of its edges.

3 Lower Bounds for Pairwise Spanners

3.1 Near-exact Pairwise Spanners

To prove our lower bound for pairwise $(1 + \epsilon)$-spanners, we use almost the same construction of a graph as the one appears in [2]. Our argument that achieves this lower bound, however, is somewhat different. Before we go into the specific details of the construction, we overview the properties of the graph of [2].

The authors of [2] constructed a sequence of graphs $\{\mathcal{H}_\kappa\}_{\kappa=0}^{\infty}$, each with a layered structure. The first and last layers of $\mathcal{H}_\kappa$ serves as *input* and *output* ports (respectively), while the interior layers are made out of many copies of $\mathcal{H}_{\kappa-1}$. Then a relatively large set

$\mathcal{P}_\kappa$ of input-output pairs is defined, such that for every $(u, v) \in \mathcal{P}_\kappa$ there is a unique shortest path in $\mathcal{H}_\kappa$ between u, v, that passes through each layer exactly once. The edges between the layers of $\mathcal{H}_\kappa$ are heavy enough, such that any other path from u to v suffers a large stretch, since it must visit at least two layers more than once (in the unweighted version, the edges between the layers are replaced with long paths).

We now construct a sequence $\{\mathcal{H}_\kappa\}$ with the same properties. The construction is essentially the same as in [2]. However, we fully describe it in details here, because (1) there are slight differences from the original construction, and (2) our lower bound proof refers to the specific details of this graph and the way it was constructed. We do use the *second base graph* from Section 2.2 in [2] as it is (this graph is actually originated in [1]).

The second base graph is an undirected unweighted graph, denoted by $\ddot{B}[p, l]$, where $p, l > 0$ are two positive integer parameters. The vertices of this graph are organized in $2l + 1$ layers, each of them with size p, such that all the edges of the graph are between adjacent layers. The vertices on the first layer of $\ddot{B}[p, l]$ are called *input ports*, and the vertices on the last layer of $\ddot{B}[p, l]$ are called *output ports*. Aside from the graph itself, there is a set of pairs of input-output ports $\ddot{\mathcal{P}}[p, l]$, such that every pair in this set is connected in $\ddot{B}[p, l]$ by a unique shortest path. The size of $\ddot{\mathcal{P}}[p, l]$ is $p^{2-o(1)}$, thus it contains a large portion of all the p^2 possible input-output pairs. Lastly, the graph $\ddot{B}[p, l]$ has labels on its edges, such that the edges on the unique shortest path of every input-output pair in $\ddot{\mathcal{P}}[p, l]$ are labeled alternately by two labels. One can think of these labels as routing directions to get from an input to an output.

Formally, the graph $\ddot{B}[p, l]$ is described in the following lemma. The proof of this lemma is implicit in [2], in which the construction of the graph $\ddot{B}[p, l]$ is described.

► **Lemma 1.** *Let $p, l > 0$ be two integer parameters. There is a function $\xi(p, l) = 2^{O(\sqrt{\log p \log l})}$ which is non-decreasing in the parameter p, and there is an undirected unweighted graph $\ddot{B}[p, l]$, with the following properties.*

1. *The vertices of the graph $\ddot{B}[p, l]$ are partitioned into $2l + 1$ disjoint layers $L_0, L_1, ..., L_{2l}$, each of them of size p, such that every edge of $\ddot{B}[p, l]$ is between vertices that belongs to adjacent layers.*
2. *There is a set of edge-labels $\ddot{\mathcal{L}}[p, l]$ of size $|\ddot{\mathcal{L}}[p, l]| \in \left[\frac{\sqrt{p}}{\xi(\sqrt{p}, l)}, \frac{\sqrt{p}}{2}\right]$, such that for every $i < 2l$, every vertex $x \in L_i$, and every label $a \in \ddot{\mathcal{L}}[p, l]$, there is exactly one edge from x to a vertex $y \in L_{i+1}$, that is labeled by a (the vertex y is different for every label $a \in \ddot{\mathcal{L}}[p, l]$).*
3. *Given a vertex $u \in L_0$ and a pair of labels $(a, b) \in (\ddot{\mathcal{L}}[p, l])^2$, let $P_{u,v}$ be the path that starts at u, and its edges are labeled alternately by the labels a, b (starting by a). Here, $v \in L_{2l}$ is the vertex in which the path $P_{u,v}$ ends, and we denote $v = out(u, a, b)$. Then, $P_{u,v}$ is the unique shortest path in $\ddot{B}[p, l]$. Moreover, for any other $u' \in L_0 \setminus \{u\}$, the path $P_{u',v'}$, that is alternately labeled by the same labels a, b and ends in $v' = out(u', a, b) \in L_{2l}$, is vertex-disjoint from $P_{u,v}$.*

We define the set $\ddot{\mathcal{P}}[p, l] \subseteq L_0 \times L_{2l}$ as

$$\ddot{\mathcal{P}}[p, l] = \{(u, out(u, a, b)) \mid u \in L_0,\ a, b \in \ddot{\mathcal{L}}[p, l]\} . \tag{2}$$

The size of this set is $p \cdot |\ddot{\mathcal{L}}[p, l]|^2$, which is at least $p \cdot \left(\frac{\sqrt{p}}{\xi(\sqrt{p}, l)}\right)^2 = \frac{p^2}{\xi(\sqrt{p}, l)^2}$ and at most $p \cdot \left(\frac{\sqrt{p}}{2}\right)^2 = \frac{p^2}{4}$, by Lemma 1.

The specific details of the construction of the graph $\ddot{B}[p,l]$ appear in [2]. We, however, only use the properties that are described in Lemma 1, and do not need these details for our proof. We now define the sequence of graphs $\{\mathcal{H}_\kappa[p,l]\}_{\kappa=0}^{\infty}$ recursively, where $p,l>0$ are any two integer parameters. For every κ, we also define a corresponding set $\mathcal{P}_\kappa[p,l]$ of pairs of vertices from $\mathcal{H}_\kappa[p,l]$.

The graph $\mathcal{H}_0[p,l]$ is defined to be the complete bipartite graph $K_{p,p}$. The corresponding set of pairs $\mathcal{P}_\kappa[p,l]$ is defined to be all the pairs (u,v), of a vertex u from the left side of $\mathcal{H}_0[p,l]=K_{p,p}$ and a vertex v from its right side.

To construct $\mathcal{H}_\kappa[p,l]$ for $\kappa>0$, we start with the graph $\ddot{B}[p,l]$ from Lemma 1. The vertices of this graph are partitioned into layers $L_0,L_1,...,L_{2l}$, where edges only exist in between adjacent layers. We call the vertices in the first layer L_0 *input ports*, and the vertices in the last layer L_{2l} *output ports*. The rest of the vertices of $\ddot{B}[p,l]$ are called *internal vertices*. The input and output ports also serves as the input/output ports of the graph $\mathcal{H}_\kappa[p,l]$ (in particular, there are p input ports and p output ports in $\mathcal{H}_\kappa[p,l]$). Let $\ddot{\mathcal{L}}[p,l]$ be the set of edge-labels, as described in Lemma 1, and let $\ddot{\mathcal{P}}[p,l]$ be the corresponding set of pairs from Definition 2. Let $p'=|\ddot{\mathcal{L}}[p,l]|$. We fix an arbitrary bijection $\pi:\ddot{\mathcal{L}}[p,l]\to\{1,2,...,p'\}$.

Consider the graph $\mathcal{H}_{\kappa-1}[p',l]$. Using π, we match each input/output port of this graph to a label in $\ddot{\mathcal{L}}[p,l]$. We replace each internal vertex of $\ddot{B}[p,l]$ by a copy of the graph $\mathcal{H}_{\kappa-1}[p',l]$. For a vertex u in $\ddot{B}[p,l]$, denote this copy by $\mathcal{H}^u_{\kappa-1}[p',l]$. Let (u,v) be an edge in $\ddot{B}[p,l]$ with label $a\in\ddot{\mathcal{L}}[p,l]$, such that $u\in L_i$, $v\in L_{i+1}$. In $\mathcal{H}_\kappa[p,l]$, we replace this edge by an edge of weight $(2l-1)^\kappa$ as follows.

- If i is even, the new edge is added from the $\pi(a)$-th *input* port of $\mathcal{H}^u_{\kappa-1}[p',l]$ (or, in case that $i=0$, from u itself) to the $\pi(a)$-th *input* port of $\mathcal{H}^v_{\kappa-1}[p',l]$.
- If i is odd, the new edge is added from the $\pi(a)$-th *output* port of $\mathcal{H}^u_{\kappa-1}[p',l]$ to the $\pi(a)$-th *output* port of $\mathcal{H}^v_{\kappa-1}[p',l]$ (or, in case that $i=2l-1$, to v itself).

In other words, we can imagine that the copies of $\mathcal{H}_{\kappa-1}[p',l]$ are inserted as they are in odd layers, and *reversed* in even layers. This way, input ports are connected to input ports, and output ports are connected to output ports. See Figure 1 for an illustration.

This completes the description of the graph $\mathcal{H}_\kappa[p,l]$. We define the corresponding set $\mathcal{P}_\kappa[p,l]$ as follows. Given some $(u,v)\in\ddot{\mathcal{P}}[p,l]$, let $(a,b)\in(\ddot{\mathcal{L}}[p,l])^2$ be the unique pair of labels such that the $u-v$ shortest path in $\ddot{B}[p,l]$ is labeled alternately with a,b. Henceforth, we call (a,b) the *corresponding labels* to (u,v). Denote by u' the $\pi(a)$-th input port of $\mathcal{H}_{\kappa-1}[p',l]$, and by v' the $\pi(b)$-th output port of $\mathcal{H}_{\kappa-1}[p',l]$. We say that $(u,v)\in\mathcal{P}_\kappa[p,l]$ if and only if $(u',v')\in\mathcal{P}_{\kappa-1}[p',l]$.

The following lemma is our version of Lemma 2.2 from [2]. We prove it here since our construction of $\mathcal{H}_\kappa[p,l],\mathcal{P}_\kappa[p,l]$ is slightly different.

▶ **Lemma 2.**

$$|\mathcal{P}_\kappa[p,l]|\geq\frac{p^2}{\xi(\sqrt{p},l)^{2\kappa}}\ .$$

Proof. We prove the lemma by induction on κ. For $\kappa=0$, by definition

$$|\mathcal{P}_0[p,l]|=p^2=\frac{p^2}{\xi(\sqrt{p},l)^{2\cdot 0}}\ .$$

Fix $\kappa>0$, and fix an input port u of $\mathcal{H}_\kappa[p,l]$. Denote by A_u the set of pairs $(u,v)\in\mathcal{P}_\kappa[p,l]$, i.e., the set of pairs in $\mathcal{P}_\kappa[p,l]$ with u as their input port. We show a bijection between the sets A_u and $\mathcal{P}_{\kappa-1}[p',l]$. Let (u,v) be a pair in A_u, and let $(a,b)\in(\ddot{\mathcal{L}}[p,l])^2$ be the

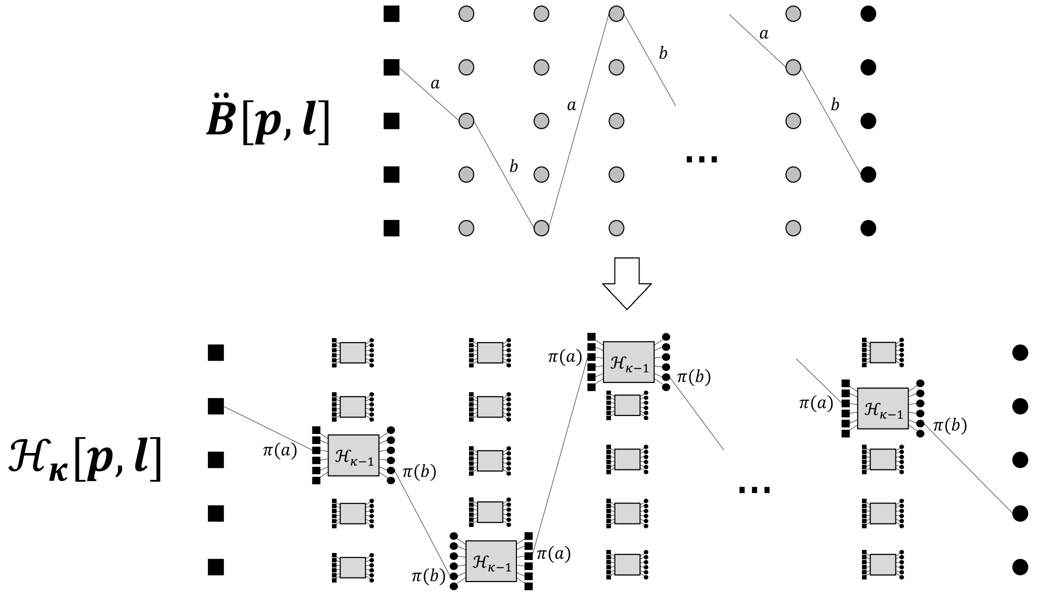

Figure 1 An illustration of the graph $\mathcal{H}_\kappa[p,l]$. The internal vertices of $\ddot{B}[p,l]$ are replaced by copies of $\mathcal{H}_{\kappa-1}[p',l]$, where p' is the number of edges-labels in $\ddot{B}[p,l]$. An edge of $\ddot{B}[p,l]$ that had label a, and is from an even layer to an odd layer, is replaced by an edge that connects the $\pi(a)$-th *input* ports of the corresponding copies of $\mathcal{H}_{\kappa-1}[p',l]$ (input port are represented in the figure by a square shape). If the edge is from an odd layer to an even layer, the same happens for the *output* ports of these copies (represented by circular shape).

corresponding labels to (u,v). By the definition of $\mathcal{P}_\kappa[p,l]$, the $\pi(a)$-th input port u' and the $\pi(b)$-th output port v' of $\mathcal{H}_\kappa[p,l]$ satisfy $(u',v') \in \mathcal{P}_{\kappa-1}[p',l]$. Thus, we map the pair $(u,v) \in A_u$ to the pair $(u',v') \in \mathcal{P}_{\kappa-1}[p',l]$.

To prove that this mapping is a bijection, we now show the inverse mapping. Note that for any $(u',v') \in \mathcal{P}_{\kappa-1}[p',l]$, there are unique labels $a,b \in \ddot{\mathcal{L}}[p,l]$ such that u' is the $\pi(a)$-th input port and v' is the $\pi(b)$-th output port of $\mathcal{H}_{\kappa-1}[p',l]$. This is true since π is a bijection. Let v be the output port of $\mathcal{H}_\kappa[p,l]$, such that $v = out(u,a,b)$ (using the notation $out()$ from Lemma 1). Then, (a,b) are the corresponding labels to (u,v), and since $(u',v') \in \mathcal{P}_{\kappa-1}[p',l]$, we conclude that $(u,v) \in \mathcal{P}_\kappa[p,l]$. That is, $(u,v) \in A_u$.

The two mappings that were described in the two paragraphs above are the inverse of each other, hence our mapping is a bijection, and we conclude that $|A_u| = |\mathcal{P}_{\kappa-1}[p',l]|$. Summing over all the input ports u of $\mathcal{H}_\kappa[p,l]$, and using the induction hypothesis, we get

$$|\mathcal{P}_\kappa[p,l]| = \sum_u |A_u| = \sum_u |\mathcal{P}_{\kappa-1}[p',l]| \geq p \cdot \frac{(p')^2}{\xi(\sqrt{p'},l)^{2\kappa-2}}\ .$$

By Lemma 1, we know that $p' = |\ddot{\mathcal{L}}[p,l]| \geq \frac{\sqrt{p}}{\xi(\sqrt{p},l)}$, and that $\xi(\sqrt{p'},l) \leq \xi(\sqrt{p},l)$, since ξ is a non-decreasing function in the first variable (and $p' \leq \frac{\sqrt{p}}{2} < p$). Hence,

$$|\mathcal{P}_\kappa[p,l]| \geq p \cdot \frac{(p')^2}{\xi(\sqrt{p'},l)^{2\kappa-2}} \geq p \cdot \frac{p}{\xi(\sqrt{p},l)^2 \xi(\sqrt{p'},l)^{2\kappa-2}} \geq \frac{p^2}{\xi(\sqrt{p},l)^{2\kappa}}\ . \qquad \blacktriangleleft$$

Next, we estimate the number of vertices in $\mathcal{H}_\kappa[p,l]$. Denote this number by $n_\kappa[p,l]$.

▶ **Lemma 3.** *For every $\kappa \geq 0$,*

$$\frac{2(2l-1)^\kappa}{\xi(\sqrt{p},l)^{2\kappa}} p^{2-\frac{1}{2^\kappa}} \leq n_\kappa[p,l] \leq 2(2l)^\kappa p^{2-\frac{1}{2^\kappa}} \ .$$

Proof. We again prove the lemma by induction on κ. For $\kappa = 0$, the graph $\mathcal{H}_0[p,l]$ is the complete bipartite graph $K_{p,p}$. Therefore,

$$\frac{2(2l-1)^0}{\xi(\sqrt{p},l)^{2\cdot 0}} p^{2-\frac{1}{2^0}} = 2p = n_0[p,l] = 2p = 2(2l)^0 p^{2-\frac{1}{2^0}} \ .$$

For $\kappa > 0$, recall that $\mathcal{H}_\kappa[p,l]$ was obtained by replacing each of the $(2l-1)p$ vertices in the interior layers of $\ddot{B}[p,l]$, by a copy of $\mathcal{H}_{\kappa-1}[p',l]$. The number of vertices in any of these copies is

$$n_{\kappa-1}[p',l] \leq 2(2l)^{\kappa-1}(p')^{2-\frac{1}{2^{\kappa-1}}} \leq 2(2l)^{\kappa-1}\left(\frac{\sqrt{p}}{2}\right)^{2-\frac{2}{2^\kappa}} < 2(2l)^{\kappa-1}p^{1-\frac{1}{2^\kappa}} \ .$$

On the other hand, this number is also

$$n_{\kappa-1}[p',l] \geq \frac{2(2l-1)^{\kappa-1}}{\xi(\sqrt{p'},l)^{2\kappa-2}}(p')^{2-\frac{1}{2^{\kappa-1}}} \geq \frac{2(2l-1)^{\kappa-1}}{\xi(\sqrt{p'},l)^{2\kappa-2}}\left(\frac{\sqrt{p}}{\xi(\sqrt{p},l)}\right)^{2-\frac{2}{2^\kappa}} \geq \frac{2(2l-1)^{\kappa-1}}{\xi(\sqrt{p},l)^{2\kappa}}p^{1-\frac{1}{2^\kappa}} \ .$$

In these bounds we used the induction hypothesis, the fact that ξ is a non-decreasing function, and the bounds on p' from Lemma 1.

We conclude that the number of vertices in $\mathcal{H}_\kappa[p,l]$ is

$$\begin{aligned} n_\kappa[p,l] &\leq 2p + (2l-1)p \cdot 2(2l)^{\kappa-1}p^{1-\frac{1}{2^\kappa}} \\ &= 2p - 2(2l)^{\kappa-1}p^{2-\frac{1}{2^\kappa}} + 2(2l)^\kappa p^{2-\frac{1}{2^\kappa}} \\ &\overset{\kappa\geq 1}{\leq} 2p - 2\cdot 1 \cdot p^{3/2} + 2(2l)^\kappa p^{2-\frac{1}{2^\kappa}} \leq 2(2l)^\kappa p^{2-\frac{1}{2^\kappa}} \ , \end{aligned}$$

and also

$$n_\kappa[p,l] \geq 2p + (2l-1)p \cdot \frac{2(2l-1)^{\kappa-1}}{\xi(\sqrt{p},l)^{2\kappa}}p^{1-\frac{1}{2^\kappa}} \geq \frac{2(2l-1)^\kappa}{\xi(\sqrt{p},l)^{2\kappa}}p^{2-\frac{1}{2^\kappa}} \ .$$

This completes the inductive proof. ◀

In [2], the authors showed that any $(1+\epsilon,\beta)$-spanner for this graph[11], that has less than $|\mathcal{P}_\kappa[p,l]|$ edges, must have $\beta = \Omega\left(\frac{1}{\epsilon\kappa}\right)^{\kappa-1}$. But note that by Lemma 3, the number of vertices in $\mathcal{H}_\kappa[p,l]$ is $n = n_\kappa[p,l] \approx p^{2-\frac{1}{2^\kappa}}$, while Lemma 2 proves that the number of pairs in $\mathcal{P}_\kappa[p,l]$ is roughly

$$p^2 \approx \left(n^{\frac{1}{2-\frac{1}{2^\kappa}}}\right)^2 = n^{\frac{2^\kappa+1}{2^{\kappa+1}-1}} = n^{1+\frac{1}{2^{\kappa+1}-1}} \ .$$

Thus, the result of [2] means that less than $n^{1+\frac{1}{2^{\kappa+1}-1}}$ edges in a near-additive spanner implies $\beta = \Omega\left(\frac{1}{\epsilon\kappa}\right)^{\kappa-1}$.

[11] For the lower bound for near-additive spanners, one need to use an *unweighted graph*. Thus, [2] actually used a similar graph where the edges between the copies of $\mathcal{H}_{\kappa-1}[p',l]$ are replaced by paths of length $(2l-1)^{\kappa-1}$. This graph has essentially the same properties as the graph $\mathcal{H}_\kappa[p,l]$ described here.

In our case, we will show that any $\mathcal{P}_\kappa[p,l]$-pairwise $(1+\epsilon)$-spanner for $\mathcal{H}_\kappa[p,l]$ must have at least $\beta|\mathcal{P}_\kappa[p,l]|$ edges, for $\beta = \Omega\left(\frac{1}{\epsilon\kappa^2}\right)^\kappa$. Otherwise, the stretch guarantee will not hold for at least one of the pairs in $\mathcal{P}_\kappa[p,l]$.

To achieve this goal, we prove some properties of the shortest paths that connect the pairs in $\mathcal{P}_\kappa[p,l]$. A key notion will be that of a *critical* edge, which also appears in [2].

▶ **Definition 4.** *An edge e of $\mathcal{H}_\kappa[p,l]$ is said to be* critical *if it lies in a copy of $\mathcal{H}_0[p,l]$.*

The following lemma is parallel to Lemma 2.3 in [2].

▶ **Lemma 5.** *The distance between any input and output port of $\mathcal{H}_\kappa[p,l]$ is at least $(2l\kappa+1)(2l-1)^\kappa$.*

Moreover, for every pair $(u,v)\in\mathcal{P}_\kappa[p,l]$, there is a unique shortest path $P_{u,v}$ in $\mathcal{H}_\kappa[p,l]$ that connects u,v, and has weight $w(P_{u,v})=(2l\kappa+1)(2l-1)^\kappa$. This path does not share a critical edge with any other shortest path $P_{u',v'}$, for $(u',v')\in\mathcal{P}_\kappa[p,l]\setminus\{(u,v)\}$. That is, there are no critical edges in $P_{u,v}\cap P_{u',v'}$, for any pair $(u',v')\neq(u,v)$ in $\mathcal{P}_\kappa[p,l]$.

Proof. We prove the Lemma by induction over $\kappa\geq 0$. For $\kappa=0$, the graph $\mathcal{H}_0[p,l]$ is the complete bipartite graph $K_{p,p}$. One of its sides consists of the input ports, and the other consists of the output ports. Thus, the distance between any input port and output port is at least $1=(2l\cdot 0+1)(2l-1)^0$. For every $(u,v)\in\mathcal{P}_0[p,l]$, the edge (u,v) is clearly the unique shortest path between u,v that has weight 1. In addition, this path does not share its only critical edge (u,v) with any other $u'-v'$ shortest path in $\mathcal{H}_0[p,l]=K_{p,p}$, for $(u',v')\in\mathcal{P}_0[p,l]$.

Fix $\kappa>0$. Every path that starts from an input port of $\mathcal{H}_\kappa[p,l]$ and ends in an output port must visit at least $2l-1$ copies of $\mathcal{H}_{\kappa-1}[p',l]$, each one of them replaces a vertex from a different layer of $\ddot{B}[p,l]$. By the induction hypothesis, passing through a copy of $\mathcal{H}_\kappa[p',l]$, from an input port to an output port (or vice versa), requires a path of weight at least $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}$. The edges that connect the different copies are of weight $(2l-1)^\kappa$. Note that any path from the input layer of $\mathcal{H}_\kappa[p,l]$ to its output layer must pass through at least $2l$ of these edges. Overall, such path must have weight of at least

$$2l(2l-1)^\kappa+(2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1}=(2l+2l(\kappa-1)+1)(2l-1)^\kappa=(2l\kappa+1)(2l-1)^\kappa\ .$$

Now fix a pair $(u,v)\in\mathcal{P}_\kappa[p,l]$. Recall that by definition, we know in particular that u,v are vertices of $\ddot{B}[p,l]$, and $(u,v)\in\ddot{\mathcal{P}}[p,l]$. By Lemma 1, in the graph $\ddot{B}[p,l]$ there is a unique $u-v$ shortest path $P=(u=u_0,u_1,u_2,...,u_{2l}=v)$, labeled by some two labels $a,b\in\ddot{\mathcal{L}}[p,l]$. Denote by $H_1,H_2,...,H_{2l-1}$ the copies of $\mathcal{H}_{\kappa-1}[p',l]$ that replaced the vertices $u_1,u_2,...,u_{2l-1}$ in the construction of $\mathcal{H}_\kappa[p,l]$.

Recall that by its construction, the graph $\mathcal{H}_\kappa[p,l]$ contains the following edges. For every *even* $i\in[0,2l-1]$, it contains an edge of weight $(2l-1)^\kappa$ from the $\pi(a)$-th input port of H_i (or from u itself if $i=0$) to the $\pi(a)$-th input port of H_{i+1}. For every *odd* $i\in[0,2l-1]$, it contains an edge of weight $(2l-1)^\kappa$ from the $\pi(b)$-th output port of H_i to the $\pi(b)$-th output port of H_{i+1} (or to v itself if $i=2l-1$). Also, recall that since $(u,v)\in\mathcal{P}_\kappa[p,l]$, it means that $(\pi(a),\pi(b))\in\mathcal{P}_{\kappa-1}[p',l]$ (here, and in the rest of this proof, we identify $\pi(a),\pi(b)$ with the $\pi(a)$-th input port and $\pi(b)$-output port of $\mathcal{H}_{\kappa-1}[p',l]$). Thus, when using the $2l$ edges described above, we can also find paths inside the copies $H_1,...,H_{2l-1}$, each of them with weight $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}$. We conclude that there is a path in $\mathcal{H}_\kappa[p,l]$ with weight

$$2l(2l-1)^\kappa+(2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1}=(2l\kappa+1)(2l-1)^\kappa\ .$$

Denote this path by $P_{u,v}$. Since we already proved that the distance between any input port and any output port is at least $(2l\kappa+1)(2l-1)^\kappa = w(P_{u,v})$, we deduce that $P_{u,v}$ is a shortest path between u, v.

Let $P'_{u,v}$ be a different path than $P_{u,v}$ between u and v in $\mathcal{H}_\kappa[p,l]$. Consider the list of copies of $\mathcal{H}_{\kappa-1}[p',l]$ that $P'_{u,v}$ passes through, by the same order they appear on $P'_{u,v}$. Since $P'_{u,v} \neq P_{u,v}$, there are two cases: either this list is identical to $H_1, H_2, ..., H_{2l-1}$, but for at least one j the path $P'_{u,v}$ passes through H_j using a different path than $P_{u,v}$, or this list is not identical to $H_1, H_2, ..., H_{2l-1}$.

In the first case, by the induction hypothesis, the path that $P'_{u,v}$ uses inside H_j has weight strictly larger than $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}$. The path inside the other copies has weight of at least $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}$, again by the induction hypothesis. Together with the $2l$ edges with weight $(2l-1)^\kappa$ that connect these copies, we get that the weight of $P'_{u,v}$ is strictly more than

$$2l(2l-1)^\kappa + (2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} = (2l\kappa+1)(2l-1)^\kappa \ .$$

In the second case, we "translate" the path $P'_{u,v}$ into a path Q in $\ddot{B}[p,l]$, by replacing each copy of $\mathcal{H}_{\kappa-1}[p',l]$ it passes through by the corresponding vertex of $\ddot{B}[p,l]$. The path Q is different from the path $P = (u, u_1, u_2, ..., u_{2l-1}, v)$. Since the latter is the unique $u-v$ shortest path in $\ddot{B}[p,l]$, the path Q must pass through a layer of $\ddot{B}[p,l]$ more than once (otherwise its weight would be equal to the weight of P). That is, Q passes through at least $2l+1$ internal vertices of $\ddot{B}[p,l]$ and contains at least $2l+2$ edges. This means that the path $P'_{u,v}$ contains at least $2l+2$ edges of weight $(2l-1)^\kappa$. Also, note that $P'_{u,v}$ (like any other input-output path in $\mathcal{H}_\kappa[p,l]$) must pass through at least $2l-1$ copies of $\mathcal{H}_{\kappa-1}[p',l]$. By the induction hypothesis, the weight of $P'_{u,v}$ is at least

$$(2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} + (2l+2)(2l-1)^\kappa > (2l\kappa+1)(2l-1)^\kappa \ .$$

In conclusion, the path $P_{u,v}$ has length $(2l\kappa+1)(2l-1)^\kappa$, while any other $u-v$ path in $\mathcal{H}_\kappa[p,l]$ has a larger weight. Thus, $P_{u,v}$ is a unique shortest path between u, v in $\mathcal{H}_\kappa[p,l]$.

To complete the proof, we have to show that if $P_{u,v}$ and $P_{u',v'}$ share a critical edge, for some $(u,v),(u',v') \in \mathcal{P}_\kappa[p,l]$, then it must be that $(u,v) = (u',v')$. Let $(u,v),(u',v')$ be such two pairs. Since their paths share a critical edge, they must pass through the same copy of $\mathcal{H}_{\kappa-1}[p',l]$. Denote this copy by H. We saw that the path $P_{u,v}$ is originated in a path P in the graph $\ddot{B}[p,l]$, which is the unique shortest path between u, v that satisfies $(u,v) \in \ddot{\mathcal{P}}[p,l]$, and that the edges of P are alternately labeled by some $a, b \in \ddot{\mathcal{L}}[p,l]$. Moreover, $(\pi(a),\pi(b)) \in \mathcal{P}_{\kappa-1}[p',l]$, because $(u,v) \in \ddot{\mathcal{P}}[p,l]$ (recall that we still identify the numbers $\pi(a), \pi(b)$ with input and output ports of $\mathcal{H}_{\kappa-1}[p',l]$). Symmetrically, the path P' and the labels a', b' that correspond to the path $P_{u',v'}$ in $\ddot{B}[p,l]$ satisfy $(\pi(a'),\pi(b')) \in \mathcal{P}_{\kappa-1}[p',l]$. Note that the unique shortest paths between $\pi(a), \pi(b)$ and between $\pi(a'), \pi(b')$ still share a critical edge. Thus, by the induction hypothesis, $(\pi(a),\pi(b)) = (\pi(a'),\pi(b'))$, or equivalently $(a,b) = (a',b')$.

Notice that the two paths P, P' pass through the same vertex in $\ddot{B}[p,l]$, because $P_{u,v}, P_{u',v'}$ pass through the same copy H. We also know that they are alternately labeled by the same to labels $(a,b) = (a',b')$. By Lemma 1, it must be that $P = P'$, and in particular $(u,v) = (u',v')$, otherwise they cannot have the same pair of labels a, b. This completes the inductive proof. ◀

We are now ready to prove our main theorem.

▶ **Theorem 6.** *For infinitely many integers $n > 0$, and for any integer $1 < \kappa \le \log\log n$ and real $0 < \epsilon < \frac{1}{12\kappa}$, there is an n-vertex graph $G = (V, E)$ and a set of pairs $\mathcal{P} \subseteq V^2$ with size at least $n^{1+\frac{1}{2^{\kappa+1}-1}-o(1)}$, such that any $\mathcal{P}$-pairwise $(1+\epsilon)$-spanner for G must have at least $\beta \cdot |\mathcal{P}|$ edges, where $\beta = \Omega\left(\frac{1}{\epsilon\kappa^2}\right)^{\kappa}$.*

Proof. Fix κ and $0 < \epsilon \le \frac{1}{12\kappa}$. Let $G, \mathcal{P}$ be $\mathcal{H}_\kappa[p,l], \mathcal{P}_\kappa[p,l]$, for $l = \left\lfloor \frac{1}{6\epsilon\kappa} \right\rfloor \ge 1$ and some arbitrary p. Denote $b = \left\lfloor \frac{2l-1}{\kappa} \right\rfloor + 1$. We will show that any $\mathcal{P}$-pairwise $(1+\epsilon)$-spanner for G must have size of at least $b^\kappa|\mathcal{P}|$. This proves the theorem for

$$\beta = b^\kappa \ge \left(\frac{2l-1}{\kappa}\right)^\kappa \ge \left(\frac{\frac{1}{3\epsilon\kappa} - 3}{\kappa}\right)^\kappa \ge \left(\frac{\frac{1}{3\epsilon\kappa} - \frac{3}{12\epsilon\kappa}}{\kappa}\right)^\kappa = \left(\frac{1}{12\epsilon\kappa^2}\right)^\kappa .$$

For the size of $\mathcal{P}$, recall that the number of vertices in G is $n \le 2(2l)^\kappa p^{2-\frac{1}{2^\kappa}}$, by Lemma 3. Thus,

$$p^2 \ge \left(\frac{n}{2(2l)^\kappa}\right)^{\frac{2}{2-\frac{1}{2^\kappa}}} \ge n^{1+\frac{1}{2^{\kappa+1}-1}} \cdot \frac{(2l)^{-2\kappa}}{4} .$$

Therefore, by Lemma 2, the size of $\mathcal{P}$ is at least

$$\frac{p^2}{\xi(\sqrt{p}, l)^{2\kappa}} \ge n^{1+\frac{1}{2^{\kappa+1}-1}} \cdot \frac{(2l)^{-2\kappa}}{4 \cdot 2^{O\left(\kappa\sqrt{\log p \log l}\right)}} = n^{1+\frac{1}{2^{\kappa+1}-1}} \cdot 2^{-2\kappa\log(2l)-2-O\left(\kappa\sqrt{\log n \log l}\right)}$$
$$= n^{1+\frac{1}{2^{\kappa+1}-1}-o(1)} .$$

Here we used Lemma 1 to bound $\xi(\sqrt{p}, l)$, we used the fact that $\kappa \le \log\log n$, and we used the fact that $p \le n$ (this is trivial, by the construction of $\mathcal{H}_\kappa[p,l]$).

Let S be a subset of the edges of $G = \mathcal{H}_\kappa[p,l]$, with $|S| < b^\kappa|\mathcal{P}| = b^\kappa|\mathcal{P}_\kappa[p,l]|$. By Lemma 5, for every pair $(u,v) \in \mathcal{P}$, the unique $u-v$ shortest path in G has a disjoint set of critical edges that it goes through. Therefore, there must be some $(u,v) \in \mathcal{P}$ such that S contains less than b^κ of its critical edges.

We prove by induction on $\kappa \ge 0$ that in the graph $\mathcal{H}_\kappa[p,l]$, if a pair $(u,v) \in \mathcal{P}_\kappa[p,l]$ has less than b^κ of its critical edges in S, then

$$d_S(u,v) \ge (2l\kappa+1)(2l-1)^\kappa + 2(2l-b)^\kappa .$$

For $\kappa = 0$, a pair $(u,v) \in \mathcal{P}_0[p,l]$ that has less than $b^0 = 1$ of its critical edges in S, means a pair such that $(u,v) \notin S$. Since the graph $\mathcal{H}_0[p,l] = K_{p,p}$ is bipartite, $d_S(u,v) \ge 3 = (2l \cdot 0 + 1)(2l-1)^0 + 2(2l-b)^0$.

Fix $\kappa > 0$, and let $P'_{u,v}$ be a $u-v$ shortest path in S, for a pair $(u,v) \in \mathcal{P}_\kappa[p,l]$ that has less than b^κ of its critical edges in S. Consider the path Q in $\ddot{B}[p,l]$ that is obtained by replacing each copy of $\mathcal{H}_{\kappa-1}[p',l]$ that $P'_{u,v}$ passes through by the corresponding vertex of $\ddot{B}[p,l]$. If Q is not the unique shortest path P between u,v in $\ddot{B}[p,l]$, then the path Q must pass through a layer of $\ddot{B}[p,l]$ more than once (otherwise its weight would be equal to the weight of P). That is, Q passes through at least $2l+1$ internal vertices of $\ddot{B}[p,l]$ and contains at least $2l+2$ edges. This means that the path $P'_{u,v}$ contains at least $2l+2$ of weight $(2l-1)^\kappa$. Note that $P'_{u,v}$ must pass through at least $2l-1$ copies of $\mathcal{H}_{\kappa-1}[p',l]$, just to get from u to v. By Lemma 5, the weight of $P'_{u,v}$ is at least

$$\begin{aligned}
& (2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} + (2l+2)(2l-1)^\kappa \\
= \;& (2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} + 2l(2l-1)^\kappa + 2(2l-1)^\kappa \\
= \;& (2l\kappa+1)(2l-1)^\kappa + 2(2l-1)^\kappa \ge (2l\kappa+1)(2l-1)^\kappa + 2(2l-b)^\kappa .
\end{aligned}$$

If the path Q equals to P, the unique $u-v$ shortest path in $\ddot{B}[p,l]$, then it passes through exactly $2l-1$ internal vertices of $\ddot{B}[p,l]$. Hence, $P'_{u,v}$ passes through exactly $2l-1$ of copies of $\mathcal{H}_{\kappa-1}[p',l]$. We would like to know how many of them contain less than $b^{\kappa-1}$ critical edges of (u,v). Note that there cannot be less than $2l-b$ such copies: otherwise the number of copies with *at least* $b^{\kappa-1}$ critical edges of (u,v) is more than $2l-1-(2l-b)=b-1$, i.e., at least b, so there are at least $b\cdot b^{\kappa-1}=b^\kappa$ critical edges of (u,v) in S, in contradiction. Therefore, the number t of copies in which there are less than $b^{\kappa-1}$ critical edges of (u,v) is at least $2l-b$. In these copies, $P'_{u,v}$ suffers a weight of at least $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}+2(2l-b)^{\kappa-1}$, by the induction hypothesis. In the other $2l-1-t$ copies, $P'_{u,v}$ must have a weight of at least $(2l(\kappa-1)+1)(2l-1)^{\kappa-1}$, by Lemma 5. Together with the $2l$ edges that connect these copies and have weight of $(2l-1)^\kappa$, we get

$$\begin{aligned}
d_S(u,v) &= w(P'_{u,v}) = 2l(2l-1)^\kappa + (2l-1-t)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} \\
&\quad +t\left[(2l(\kappa-1)+1)(2l-1)^{\kappa-1}+2(2l-b)^{\kappa-1}\right] \\
&= 2l(2l-1)^\kappa + (2l-1)(2l(\kappa-1)+1)(2l-1)^{\kappa-1} + t\cdot 2(2l-b)^{\kappa-1} \\
&= (2l\kappa+1)(2l-1)^\kappa + t\cdot 2(2l-b)^{\kappa-1} \\
&\geq (2l\kappa+1)(2l-1)^\kappa + (2l-b)\cdot 2(2l-b)^{\kappa-1} \\
&= (2l\kappa+1)(2l-1)^\kappa + 2(2l-b)^\kappa \ .
\end{aligned}$$

This completes the inductive proof. It shows that there is a pair $(u,v)\in\mathcal{P}_\kappa[p,l]=\mathcal{P}$ with

$$\begin{aligned}
d_S(u,v) &\geq (2l\kappa+1)(2l-1)^\kappa + 2(2l-b)^\kappa \\
&= (2l\kappa+1)(2l-1)^\kappa\left(1+\frac{2(2l-b)^\kappa}{(2l\kappa+1)(2l-1)^\kappa}\right) \\
&= d_G(u,v)\cdot\left(1+\frac{2}{(2l\kappa+1)}\left(\frac{2l-b}{2l-1}\right)^\kappa\right) \\
&= d_G(u,v)\cdot\left(1+\frac{2}{(2l\kappa+1)}\left(1-\frac{b-1}{2l-1}\right)^\kappa\right) \\
&\geq d_G(u,v)\cdot\left(1+\frac{2}{(2l\kappa+1)}\left(1-\frac{1}{\kappa}\right)^\kappa\right) \overset{\kappa>1}{\geq} d_G(u,v)\cdot\left(1+\frac{2}{(2l\kappa+1)}\cdot\frac{1}{4}\right) \\
&> d_G(u,v)\cdot\left(1+\frac{2}{3l\kappa}\cdot\frac{1}{4}\right) = d_G(u,v)\cdot\left(1+\frac{1}{6l\kappa}\right) \geq d_G(u,v)\cdot(1+\epsilon) \ ,
\end{aligned}$$

Where in the last step we used the definition of $l=\left\lfloor\frac{1}{6\epsilon\kappa}\right\rfloor$. Thus, S cannot be a $\mathcal{P}_\kappa[p,l]$-pairwise $(1+\epsilon)$-spanner. In other words, any $\mathcal{P}_\kappa[p,l]$-pairwise $(1+\epsilon)$-spanner must have size of at least $\beta|\mathcal{P}|$, where $\beta\geq\left(\frac{1}{12\epsilon\kappa^2}\right)^\kappa$. ◀

▶ Remark 7. Notice that the proof of Theorem 6 also works for every subset of the pairs $\mathcal{P}_\kappa[p,l]$. Therefore, the phrasing of Theorem 6 may be strengthen as follows.

For infinitely many integers $n>0$, and for any integer $1<\kappa\leq\log\log n$ and real $0<\epsilon<\frac{1}{12\kappa}$, there is an n-vertex graph $G=(V,E)$ and a number $Q=n^{1+\frac{1}{2^{\kappa+1}-1}-o(1)}$, such that for every integer $q\leq Q$, there is a set of pairs $\mathcal{P}\subseteq V^2$ with size q, such that any $\mathcal{P}$-pairwise $(1+\epsilon)$-spanner for G must have at least $\beta\cdot|\mathcal{P}|$ edges, where $\beta=\Omega\left(\frac{1}{\epsilon\kappa^2}\right)^\kappa$.

3.2 Large Stretch

The lower bound for pairwise spanners with large stretch is achieved using a graph with high girth and large number of edges. This properties of a graph were used for proving lower bounds for distance oracles and spanners (see [28]) and also for hopsets (see [24]). In particular, we use the same graph that was used in [24], which was introduced by Lubotzky, Phillips and Sarnak in [22]. This graph has the additional convenient property of being *regular*, besides its high girth and large number of edges. Its exact properties are described in the following theorem.

▶ **Theorem 8** ([22]). *For infinitely many integers $n \in \mathbb{N}$, and for every integer $k \geq 1$, there exists a $(p+1)$-regular graph $G = (V, E)$ with $|V| = n$ and girth at least $\frac{4}{3}k(1 - o(1))$, where $p = D \cdot n^{\frac{1}{k}}$, for some universal constant D.*

Fix $\alpha, k \geq 1$ such that $k \geq \alpha + 1$, and a large enough $n \in \mathbb{N}$ as in Theorem 8. Let $G = (V, E)$ be the corresponding $(p+1)$-regular graph from Theorem 8. The girth of G is at least $\frac{4}{3}k(1 - o(1))$, thus larger than k. Denote $\delta = \left\lfloor \frac{k}{\alpha+1} \right\rfloor$. Define the following set of pairs $\mathcal{P}_0 \subseteq V^2$.

$$\mathcal{P}_0 = \{(u, v) \in V^2 \mid d_G(u, v) = \delta\} \ .$$

▶ **Lemma 9.** *For every $(u, v) \in \mathcal{P}_0$, there is a unique shortest path $P_{u,v}$ between u, v in G. Furthermore, for every tour $P'_{u,v}$ between u, v, that has length $|P'_{u,v}| \leq \alpha\delta$, we have $P_{u,v} \subseteq P'_{u,v}$.*

Proof. Let $P_{u,v}$ be some $u-v$ shortest path in G, and let $P'_{u,v}$ be a $u-v$ tour with $|P'_{u,v}| \leq \alpha\delta$. If $P_{u,v} \not\subseteq P'_{u,v}$, then the union $P_{u,v} \cup P'_{u,v}$ contains a cycle. This cycle is of length at most

$$|P_{u,v}| + |P'_{u,v}| \leq \delta + \alpha\delta = (\alpha + 1)\delta \leq k \ ,$$

by the definition of $\delta = \left\lfloor \frac{k}{\alpha+1} \right\rfloor$. This is a contradiction to the fact that the girth of G is larger than k. Hence, $P_{u,v} \subseteq P'_{u,v}$. In case $P'_{u,v}$ is also a $u - v$ shortest path, i.e., $|P'_{u,v}| = \delta$, then $P_{u,v} \subseteq P'_{u,v}$ implies $P_{u,v} = P'_{u,v}$. That is, $P_{u,v}$ is the unique shortest path between u, v in G. ◀

Henceforth, we use the notations from Lemma 9, that is, we denote by $P_{u,v}$ the $u - v$ shortest path in G.

▶ **Lemma 10.** *Let $\mathcal{P} \subseteq \mathcal{P}_0$ be some subset, and suppose that S is a $\mathcal{P}$-pairwise α-spanner for G. Then,*

$$\bigcup_{(u,v)\in\mathcal{P}} P_{u,v} \subseteq S \ .$$

Proof. Fix some $(u, v) \in \mathcal{P} \subseteq \mathcal{P}_0$. Since S has stretch α for every pair in $\mathcal{P}$, we know that there is a $u - v$ path $P'_{u,v} \subseteq S$ with $|P'_{u,v}| \leq \alpha|P_{u,v}| = \alpha\delta$. By Lemma 9, the $u - v$ shortest path satisfies $P_{u,v} \subseteq P'_{u,v} \subseteq S$. In conclusion, $P_{u,v} \subseteq S$ for every $(u, v) \in \mathcal{P}$, thus $\bigcup_{(u,v)\in\mathcal{P}} P_{u,v} \subseteq S$. ◀

The following lemma describes several combinatorial properties of the graph G.

▶ **Lemma 11.** *The number of edges in $G = (V, E)$ is $|E| = \frac{n(p+1)}{2}$. The number of paths in $\mathcal{P}_0$ is $\frac{n(p+1)p^{\delta-1}}{2}$. For every edge $e \in E$, there are $\delta p^{\delta-1}$ pairs $(u, v) \in \mathcal{P}_0$ such that $e \in P_{u,v}$.*

Proof. The number of edges $|E|$ is half the sum of the degrees in G. Since G is $(p+1)$-regular, we get $|E| = \frac{n(p+1)}{2}$.

Now fix some $u \in V$, and consider its BFS tree up to distance δ. The root u has $p + 1$ children in this tree, and every other internal vertex has a set of p children, disjoint from the children set of any other vertex in this tree. This is true since otherwise there would be a cycle of length at most $2\delta \le (\alpha+1)\delta \le k$, in contradiction to the girth of G being larger than k. Thus, the number of vertices v in the δ-th layer of this tree, is $(p+1)p^{\delta-1}$. That is, the number of $v \in V$ such that $d_G(u,v) = \delta$ is $(p+1)p^{\delta-1}$. Hence,

$$\sum_{u\in V} |\{v \in \ | d_G(u,v) = \delta\}| = \sum_{u\in V}(p+1)p^{\delta-1} = n(p+1)p^{\delta-1} \ .$$

In this sum, each pair $(u,v) \in \mathcal{P}_0$ is counted exactly twice, therefore $|\mathcal{P}_0| = \frac{n(p+1)p^{\delta-1}}{2}$.

The proof of the third property is very similar. Fix some edge $e = (v_1, v_2) \in E$. For every integer $i \in [0, \delta-1]$, consider the BFS tree T_1^i of v_1 up to distance i. Symmetrically, T_2^i denotes the BFS tree of v_2 up to distance i. As before, the children sets of the vertices in $T_1^i \cup T_2^{\delta-1-i}$ are disjoint - otherwise there would be a cycle of length at most

$$\max\{i+1+\delta-1-i, 2i, 2(\delta-1-i)\} < 2\delta \le k \ ,$$

in contradiction. Hence, the number of vertices in the i-th layer of T_1^i is p^i (note that v_1 has p children in this tree), and the number of vertices in the $(\delta-1-i)$-th layer of $T_2^{\delta-1-i}$ is $p^{\delta-1-i}$. For every pair $(u,v) \in \mathcal{P}_0$ such that $e \in P_{u,v}$, the path $P_{u,v}$ has one end in the i-th layer of T_1^i and the other end in the $(\delta-1-i)$-th layer of $T_2^{\delta-1-i}$, for some $i \in [0, \delta-1]$. See Figure 2 for an illustration. Thus, the number of such pairs is

$$\sum_{i=0}^{\delta-1} p^i \cdot p^{\delta-1-i} = \delta p^{\delta-1} \ . \qquad ◀$$

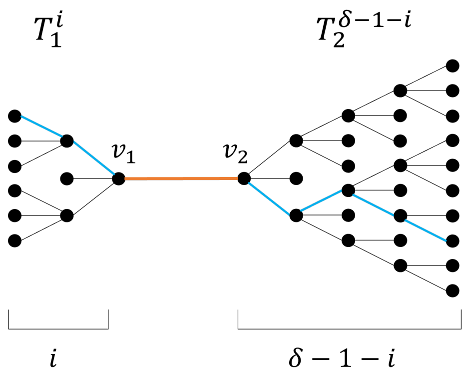

Figure 2 Given an edge $e = (v_1, v_2)$ (colored orange in the figure), we consider the BFS trees T_1^i and $T_2^{\delta-1-i}$ rooted at v_1 and v_2 respectively, up to distance i and $\delta-1-i$ respectively. There are no cycles within these two trees, because of the girth guarantee. By regularity, we know that each vertex in these trees, except the leaves, has exactly p children. Every path of length δ that passes through e, such as the blue path in the figure, is determined by a leaf of T_1^i and a leaf of $T_2^{\delta-1-i}$.

We are now ready for the main theorem of this section.

▶ **Theorem 12.** *For infinitely many integers $n > 0$, and for any real $\alpha \ge 1$ and integer k such that $\alpha+1 \le k \le \log n$, there is an n-vertex graph $G = (V,E)$ and a set of pairs $\mathcal{P} \subseteq V^2$ with size $\Theta\left(\frac{\alpha}{k} \cdot n^{1+\frac{1}{k}}\right)$, such that any $\mathcal{P}$-pairwise α-spanner for G must have at least $\Omega(n^{1+\frac{1}{k}})$ edges, that is, the size overhead is $\beta = \Omega\left(\frac{k}{\alpha}\right)$.*

Proof. By Theorem 8, for infinitely many integers $n > 0$, there is a $(p+1)$-regular graph $G = (V, E)$ with n vertices and girth larger than k, where $p = \Theta(n^{\frac{1}{k}})$. We use the same notations for δ and $\mathcal{P}_0$ as in the beginning of this section.

Let $\mathcal{P} \subseteq \mathcal{P}_0$ be a subset that is formed by sampling each pair in $\mathcal{P}_0$ independently with probability $\frac{1}{\delta p^{\delta-1}}$. The expected number of pairs in $\mathcal{P}$ is

$$\frac{|\mathcal{P}_0|}{\delta p^{\delta-1}} = \frac{n(p+1)p^{\delta-1}}{2\delta p^{\delta-1}} = \frac{n(p+1)}{2\delta} = \frac{|E|}{\delta} ,$$

by Lemma 11. Moreover, by Chernoff bound,

$$\Pr\left[\left||\mathcal{P}| - \frac{|E|}{\delta}\right| > \frac{|E|}{2\delta}\right] \leq 2e^{-\frac{|E|}{12\delta}} = 2e^{-\frac{n(p+1)}{24\delta}} \leq 2e^{-\frac{n}{24\log n}} \leq 2e^{-2} , \tag{3}$$

for large enough n, where we used the fact that $\delta \leq k \leq \log n$.

We say that a pair $(u, v) \in \mathcal{P}_0$ *covers* an edge $e \in E$ if $e \in P_{u,v}$. For an edge $e \in E$, the number of pairs in $\mathcal{P}_0$ that cover e is $\delta p^{\delta-1}$, by Lemma 11. Therefore, the probability that none of the pairs that cover e are in $\mathcal{P}$ is $(1 - \frac{1}{\delta p^{\delta-1}})^{\delta p^{\delta-1}} \leq e^{-1}$. Hence, If we denote by $E' \subseteq E$ the set of edges that are not covered by any $(u, v) \in \mathcal{P}$, then $\mathbb{E}[|E'|] \leq |E| \cdot e^{-1}$. By Markov's inequality,

$$\Pr\left[|E'| > \frac{2}{e}|E|\right] \leq \frac{1}{2} \tag{4}$$

Now, by the union bound, the probability that either $\left||\mathcal{P}| - \frac{|E|}{\delta}\right| > \frac{|E|}{2\delta}$, or $|E'| > \frac{2}{e}|E|$, is at most $2e^{-2} + \frac{1}{2} < 1$, using Inequalities (3) and (4). Therefore, there is a way to choose the subset $\mathcal{P} \subseteq \mathcal{P}_0$, such that the number of edges in E that are not covered by any $(u, v) \in \mathcal{P}$ is at most $\frac{2}{e}|E|$, and such that $\frac{|E|}{2\delta} \leq |\mathcal{P}| \leq \frac{3|E|}{2\delta}$. In particular,

$$|\mathcal{P}| = \Theta\left(\frac{|E|}{\delta}\right) = \Theta\left(\frac{n(p+1)}{2\delta}\right) = \Theta\left(\frac{n^{1+\frac{1}{k}}}{\delta}\right) = \Theta\left(\frac{\alpha}{k} \cdot n^{1+\frac{1}{k}}\right) .$$

For this choice of $\mathcal{P}$, the number of edges $e \in E$ that satisfy $e \in P_{u,v}$ for some $(u, v) \in \mathcal{P}$ is at least $\left(1 - \frac{2}{e}\right)|E|$. Notice that these are exactly the edges in $\bigcup_{(u,v)\in\mathcal{P}} P_{u,v}$. By Lemma 10, any $\mathcal{P}$-pairwise α-spanner for G must contain this set, and therefore must have size at least

$$\left(1 - \frac{2}{e}\right)|E| \geq \frac{1}{4}|E| \geq \frac{1}{4} \cdot \frac{2\delta}{3}|\mathcal{P}| = \frac{\delta}{6}|\mathcal{P}| .$$

This proves the theorem for $\beta = \frac{\delta}{6} = \Omega\left(\frac{k}{\alpha}\right)$. ◀

References

1. Amir Abboud and Greg Bodwin. The 4/3 additive spanner exponent is tight. *Journal of the ACM (JACM)*, 64(4):1–20, 2017.
2. Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. *SIAM Journal on Computing*, 47(6):2203–2236, 2018.
3. Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Stephen Kobourov, and Richard Spence. Weighted additive spanners. In *Graph-Theoretic Concepts in Computer Science: 46th International Workshop, WG 2020, Leeds, UK, June 24–26, 2020, Revised Selected Papers 46*, pages 401–413. Springer, 2020.
4. I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9:81–100, 1993.

5 Uri Ben-Levy and Merav Parter. New (α, β) spanners and hopsets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1695–1714. SIAM, 2020.

6 Aaron Bernstein. Fully dynamic (2+ ε) approximate all-pairs shortest paths with fast query and close to linear update time. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 693–702. IEEE, 2009.

7 Greg Bodwin. New results on linear size distance preservers. *SIAM J. Comput.*, 50(2):662–673, 2021. doi:10.1137/19M123662X.

8 Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM*, 47(1):132–166, 2000. doi:10.1145/331605.331610.

9 D. Coppersmith and M. Elkin. Sparse source-wise and pair-wise distance preservers. In *SODA: ACM-SIAM Symposium on Discrete Algorithms*, pages 660–669, 2005.

10 Michael Elkin, Arnold Filtser, and Ofer Neiman. Prioritized metric structures and embedding. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 489–498, 2015.

11 Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. *Theor. Comput. Sci.*, 697:1–36, 2017. doi:10.1016/j.tcs.2017.06.021.

12 Michael Elkin, Yuval Gitlitz, and Ofer Neiman. Almost shortest paths with near-additive error in weighted graphs. In Artur Czumaj and Qin Xin, editors, *18th Scandinavian Symposium and Workshops on Algorithm Theory, SWAT 2022, June 27-29, 2022, Tórshavn, Faroe Islands*, volume 227 of *LIPIcs*, pages 23:1–23:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.SWAT.2022.23.

13 Michael Elkin and Ofer Neiman. Efficient algorithms for constructing very sparse spanners and emulators. *ACM Transactions on Algorithms (TALG)*, 15(1):1–29, 2018.

14 Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. *SIAM Journal on Computing*, 48(4):1436–1480, 2019.

15 Michael Elkin and Ofer Neiman. Linear-size hopsets with small hopbound, and constant-hopbound hopsets in rnc. In *The 31st ACM Symposium on Parallelism in Algorithms and Architectures*, pages 333–341, 2019.

16 Michael Elkin and Ofer Neiman. Centralized, parallel, and distributed multi-source shortest paths via hopsets and rectangular matrix multiplication. In Petra Berenbrink and Benjamin Monmege, editors, *39th International Symposium on Theoretical Aspects of Computer Science, STACS 2022, March 15-18, 2022, Marseille, France (Virtual Conference)*, volume 219 of *LIPIcs*, pages 27:1–27:22. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.STACS.2022.27.

17 Michael Elkin and David Peleg. (1+epsilon, beta)-spanner constructions for general graphs. *SIAM J. Comput.*, 33(3):608–631, 2004. doi:10.1137/S0097539701393384.

18 Michael Elkin and Idan Shabat. Path-reporting distance oracles with near-logarithmic stretch and linear size. *CoRR*, abs/2304.04445, 2023. doi:10.48550/arXiv.2304.04445.

19 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. *Journal of the ACM (JACM)*, 65(6):1–40, 2018.

20 Shang-En Huang and Seth Pettie. Thorup-zwick emulators are universally optimal hopsets. *Information Processing Letters*, 142, April 2017. doi:10.1016/j.ipl.2018.10.001.

21 Shimon Kogan and Merav Parter. Having hope in hops: New spanners, preservers and lower bounds for hopsets. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 766–777. IEEE, 2022. doi:10.1109/FOCS54457.2022.00078.

22 Alexander Lubotzky, Ralph Phillips, and Peter Sarnak. Ramanujan graphs. *Combinatorica*, 8(3):261–277, 1988.

23 Gary L. Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets. In *Proceedings of the 27th ACM Symposium on Parallelism in Algorithms and Architectures*, SPAA '15, pages 192–201, New York, NY, USA, 2015. ACM. doi:10.1145/2755573.2755574.

24 Ofer Neiman and Idan Shabat. A unified framework for hopsets. In *30th Annual European Symposium on Algorithms (ESA 2022)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2022.

25 Merav Parter. Bypassing erdős' girth conjecture: Hybrid stretch and sourcewise spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 608–619. Springer, 2014.

26 Seth Pettie. Distributed algorithms for ultrasparse spanners and linear size skeletons. In *Proceedings of the twenty-seventh ACM symposium on Principles of distributed computing*, pages 253–262, 2008.

27 Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 261–272. Springer, 2005.

28 M. Thorup and U. Zwick. Approximate distance oracles. In *Proc. of the 33rd ACM Symp. on Theory of Computing*, pages 183–192, 2001.

29 M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. In *Proc. of Symp. on Discr. Algorithms*, pages 802–809, 2006.