Hierarchical Spanners

Davide Bilò
Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Italy

Luciano Gualà
Department of Enterprise Engineering, "Tor Vergata" University of Rome, Italy

Stefano Leucci
Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Italy

Guido Proietti
Department of Information Engineering, Computer Science and Mathematics, University of L'Aquila, Italy

Alessandro Straziota
Department of Enterprise Engineering, "Tor Vergata" University of Rome, Italy

Abstract

A *hierarchical graph* $\mathcal{G}$ consists of a vertex set $V(\mathcal{G})$ and L pairwise disjoint edge sets $E_1, \dots, E_L$. Such a structure naturally defines a hierarchy of L unweighted graphs, where the ℓ-th graph is $G_\ell = (V, \bigcup_{i=1}^{\ell} E_i)$.

In this paper, we initiate the study of *hierarchical spanners*, namely subgraphs of $\mathcal{G}$ that approximately preserve distances among a given set of pairs of vertices in $V(\mathcal{G})$ at *every* level of the hierarchy. This notion generalizes classical spanners, and thus all known lower bounds extend to this setting; however, it is not clear whether the same size-stretch trade-offs can be achieved.

We investigate this question by devising both upper and lower bounds for hierarchical spanners under various types of stretch and pairs of vertices of interest whose approximate (or exact) distances are to be maintained. On the positive side, a trivial adaptation of the greedy construction yields $(2k-1)$-spanners of size $O(n^{1+1/k})$, matching the classical bounds. However, the non-hierarchical bounds do not extend to the hierarchical case when additive or nearly-additive spanners are considered. For instance, we prove that any β-additive single-pair hierarchical spanner must have size $\Omega\left(n\sqrt{\frac{n}{\beta+1}}\right)$ in the worst case. This bound is tight, as we provide a matching upper bound for every $\beta \geq 0$, which in turn implies a $O(n\sqrt{n})$-size single-pair hierarchical *preserver*.

Finally, we present additional positive results among which a 4-additive all-pairs hierarchical spanner of size $\widetilde{O}(n^{5/3})$, an (essentially tight) single-source hierarchical $(1+\varepsilon)$-spanner of size $\widetilde{O}(n/\varepsilon)$, an all-pairs hierarchical spanner of size $\widetilde{O}\left(n\sqrt{\frac{n}{\varepsilon}}\right)$ achieving stretch $(1+\varepsilon, 2)$, for any constant value of $\varepsilon > 0$, and a subsetwise hierarchical preserver of size $O(n\sqrt{n|S|})$, where $S \subseteq V(\mathcal{G})$.

2012 ACM Subject Classification Theory of computation → Sparsification and spanners; Mathematics of computing → Graph algorithms

Keywords and phrases Graph spanners, Network sparsification, Hierarchical graphs

Digital Object Identifier 10.4230/LIPIcs.ESA.2026.145

Funding *Davide Bilò*: acknowledges the support of the project *SOS-TG: Spanners and Oracles for Static and Temporal Graphs*, funded by Università degli Studi dell'Aquila under the Call for Proposals for Fundamental Research and Early-Career Research Grants – Year 2026.

1 Introduction

Motivated by problems in network sparsification, the design of high-quality graph spanners has received extensive attention. Given parameters $\alpha \geq 1$ and $\beta \geq 0$, a graph G, and a set of vertex pairs of interest P, an (α, β)-spanner H of G (with respect to P) is a subgraph of G that approximately preserves distances between pairs in P. Formally, for every $(u, v) \in P$, H satisfies $d(u, v, H) \leq \alpha d(u, v, G) + \beta$, where $d(\cdot, \cdot, X)$ denotes the shortest-path distance in graph X. The pair (α, β) is referred to as the distortion, or stretch, of the spanner. When $\alpha = 1$, H is said to be a β-additive spanner, while H is called a α-multiplicative spanner, or simply α-spanner, when $\beta = 0$. Spanners with distortion $(1, 0)$ are also called *preservers*.

The quality of a spanner is typically evaluated along two dimensions. On the one hand, one aims to have small *size*, measured as the number of edges. On the other hand, one seeks to guarantee a low distortion. A central goal in the area is to understand the achievable trade-offs between size and distortion for general graphs. For instance, a classical result shows that when P consists of all vertex pairs, every n-vertex graph admits a $(2k - 1, 0)$-spanner with $O(n^{1+1/k})$ edges, which is asymptotically tight under the Erdős girth conjecture [14].[1]

The literature on graph spanners is vast, and a wide range of trade-offs has been established depending on the type of distortion considered, the choice of pairs of interest (e.g., all pairs versus a subset), and additional constraints such as resilience to edge or vertex failures (leading to so-called *fault-tolerant* spanners). We refer the reader to the survey by Bodwin [1] and subsequent works [8, 9, 10, 17] for recent developments.

However, most prior work focuses on the design of spanners for "single-layer" graphs. In this paper, we extend the study of spanners to a more general setting of *hierarchical* (multi-layer) graphs. Specifically, we assume that the edge set of a hierarchical graph is partitioned into L subsets $E_1, \ldots, E_L$, where L is the *depth* of the hierarchy. This naturally defines a hierarchy of graphs, where the graph G_ℓ at level ℓ contains all edges in $E_1 \cup E_2 \cup \cdots \cup E_\ell$.

Hierarchical graphs arise naturally in several practical scenarios. For example, consider a public transportation network in which each edge is associated with a type of subscription required to traverse it; more expensive subscriptions grant access to increasingly larger subsets of edges. Another example comes from wireless networks, where each edge is associated with an activation threshold: an edge is present only if the transmission power at both endpoints exceeds a certain value. In this case, the hierarchical graph represents the hierarchy of networks obtained as a function of a uniform power parameter γ assigned to all nodes.

Given a hierarchical graph $\mathcal{G}$, we define a *hierarchical* (α, β)*-spanner* as a subgraph $\mathcal{H}$ of $\mathcal{G}$ that approximately preserves distances between the pairs of interest at *every* level of the hierarchy, i.e., H_ℓ is an (α, β)-spanner of G_ℓ for every $\ell = 1, \ldots, L$. Note that this notion generalizes the classical spanner problem, since a standard graph can be viewed as a hierarchical graph of depth 1. Consequently, known lower bounds for spanners immediately extend to hierarchical spanners. On the other hand, it is not clear whether the classical size–distortion trade-offs can still be achieved in this more general setting.

1.1 Our results

In this paper we prove several bounds on the size of hierarchical spanners for different choices of pairs of interest and distortions. Some of these upper and lower bounds are summarized in Tables 1 and 2. The following discussion provides a high-level overview with a focus on

[1] Such lower bound has been proven unconditionally for $k = 1$, 2, 3, and 5 [18, 19].

Table 1 Upper bounds for various types of hierarchical spanners. All the bounds hold regardless of the depth L of the hierarchy. The $\widetilde{O}$ notation hides multiplicative factors that are polylogarithmic in n. We also assume that $\varepsilon > 0$ is a constant. Gray rows are special cases obtained by instantiating more general results already listed in the table.

Type	Stretch	Size	Reference
Pairwise	preserver	$O(n\sqrt{n\lvert P\rvert})$	Theorem 5
Pairwise	$1+\varepsilon$	$O(n\lvert P\rvert/\varepsilon)$	Theorem 5
Pairwise	$1+\varepsilon$	$\widetilde{O}(n\sqrt{\lvert P\rvert/\varepsilon})$	Theorem 5
Single-pair	preserver	$O(n\sqrt{n})$	Theorem 5
Single-pair	$1+\varepsilon$	$O(n/\varepsilon)$	Theorem 5
$s \times T$	preserver	$O(n\sqrt{n\lvert T\rvert})$	Theorem 5
Subsetwise	preserver	$\widetilde{O}(n\sqrt{n\lvert S\rvert})$	Theorem 5
Single-source	$1+\varepsilon$	$\widetilde{O}(n/\varepsilon)$	Theorem 6
All-pairs	4-additive	$\widetilde{O}(n^{5/3})$	Corollary 10
All-pairs	$(1+\varepsilon, 2)$	$\widetilde{O}(n\sqrt{n/\varepsilon})$	Corollary 12

the qualitative differences between standard spanners and their hierarchical counterpart. We refer the reader to the technical part of the paper for precise bounds and more fine-grained results, also depending on the depth of the hierarchy.

Single-pair. The simplest case asks to preserve distances between a single pair of vertices, which is trivial in a non-hierarchical setting. We show that $O(n\sqrt{n})$ edges suffice to build a hierarchical preserver. Indeed, tight upper and lower bounds of $\Theta(n\sqrt{\frac{n}{\beta+1}})$ on the size of β-additive hierarchical spanners hold for any $\beta \geq 0$.

All pairs. On the opposite end of the spectrum lies the all-pair setting. For multiplicative stretch, it is well known that every (non-hierarchical) graph admits a $(2k-1)$-spanner with $O(n^{1+1/k})$ edges, obtainable through a simple greedy construction [3].[2] Interestingly, a natural adaptation of this approach also yields hierarchical $(2k-1)$-spanners with the same size.[3] Concerning additive distortions, spanners with subquadratic size are known in the non-hierarchical setting; hence, a natural question is whether subquadratic size is also attainable for hierarchical additive spanners. We answer this question positively, by showing the existence of hierarchical 4-additive spanners with size $\widetilde{O}(n^{5/3})$. Although this bound is worse than that of the non-hierarchical 2-additive spanner of size $O(n^{3/2})$ [2, 15], the existence of hierarchical spanners with size $O(n^{3/2})$ appears to be unlikely even if one allows for any constant additive stretch. Indeed, as discussed above, the simpler single-pair β-additive case already requires $\Omega(n\sqrt{n})$ edges for $\beta = O(1)$. Somewhat surprisingly, allowing a slightly weaker guarantee leads to significantly sparser constructions. In particular, by

[2] Observe that, while the greedy construction holds even in the more general case of weighted (non-hierarchical) graphs, the hierarchical graphs considered in this papers are unweighted.

[3] To construct a $(2k-1)$-spanner of a non-hierarchical graph G, the greedy algorithm starts with a spanning subgraph H of G with no edges, and processes the edges of G in an arbitrary order. Whenever the distance in H between the endpoints of a processed edge exceeds $2k-1$, such edge is added to H. This procedure also works for the hierarchical case once edges are considered in non-decreasing order of the index ℓ of their containing set E_ℓ. Indeed, for each $e \in E_\ell$, the above ordering ensures that $e \in E(H_\ell)$ or that H_ℓ contains a path of length at most $2k-1$ between the endpoints of e.

Table 2 Lower bounds for single-pair and $s \times T$ hierarchical preserver and β-spanners with $\beta \geq 1$.

Type	Stretch	Size	Reference
Single-pair	β-additive	$\Omega(n\sqrt{\frac{n}{\beta+1}})$	Theorem 15
Single-pair	preserver	$\Omega\left(n\sqrt{n}\right)$	Theorem 15
$s \times T$	preserver	$\Omega(n\|T\|)$	Theorem 13

relaxing the distortion to $(1+\varepsilon, 2)$, for any $\varepsilon \in (0,1]$ of choice, we obtain all-pairs hierarchical spanners with $\widetilde{O}(n\sqrt{n/\varepsilon})$ edges, matching the size achieved by non-hierarchical 2-additive spanners, up to polylogarithmic factors.

Single-source. If distances between a single source vertex and a set T of target vertices are to be preserved, then a simple lower bound of $\Omega(n|T|)$ on the size of any hierarchical preserver holds, which implies a strong $\Omega(n^2)$ lower bound for the single-source case. On the positive side, single-source hierarchical $(1+\varepsilon)$-spanners have size $\widetilde{O}(n/\varepsilon)$ for any $\varepsilon \in (0,1]$, which is tight up to polylogarithmic factors whenever ε is constant.

Pairwise and beyond. In the most general setting, one is only interested in preserving shortest paths between $|P|$ arbitrarily-chosen pairs of vertices. We show that pairwise hierarchical preservers have size $O(n\sqrt{n|P|})$. This resembles the $O(n\sqrt{|P|})$ upper bound on the size of pairwise spanners of non-hierarchical graphs given in [13], since each pair in $(s,t) \in P$ may require the addition of up to $n-1$ shortest paths between s and t when the graph is hierarchical. Our bound can be improved to $\widetilde{O}(n\sqrt{n|S|})$ for the special case in which $P = S \times S$. For better upper bounds on pairwise spanners in the non hierarchical-case, we refer the reader to [1].

Interestingly, if one allows a slightly larger $(1+\varepsilon)$ stretch, for $\varepsilon \in (0,1]$, we can obtain hierarchical pairwise spanners of size $\widetilde{O}(n\sqrt{|P|/\varepsilon})$.

1.2 Other related works

As already discussed in the introduction, the literature on spanners for standard graphs is vast. The generalization addressed in this work is not the only one possible, as several other directions have been explored in the literature. Two are particularly noteworthy: the study of spanners for temporal graphs [5, 6, 11, 12] and spanners for generalized shortest-path metrics, such as Beer distances and group Steiner distances [4, 7].

A closely related problem is investigated in [16], where a hierarchy of graphs is provided as input. In this context, the goal is to construct an efficient data structure (index) capable of answering distance queries between vertex pairs at any given level of the hierarchy. The authors design a 2-hop labeling index whose quality is evaluated experimentally, however no theoretical bounds on its size or query time are provided.

1.3 Paper organization

Our results on pairwise and single-source hierarchical preservers are given in Section 3 and Section 4, respectively. Section 3 also discusses subsetwise and $s \times T$ hierarchical preservers. Some of the above results are then used in Section 5 to derive all-pairs hierarchical spanners. Our lower bounds are given in Section 6, while Section 7 concludes the paper and discusses some problems left open.

2 Preliminary definitions

Given a graph G, we denote by $V(G)$ and $E(G)$ its vertex-set and the set of its edges, respectively. We write $H \subseteq G$ to say that H is a subgraph of G. We denote by $d(s,t,G)$ the hop-distance from $s \in V(G)$ to $t \in V(G)$ in the graph G.

Let $P \subseteq V(G) \times V(G)$ be a set of vertex pairs, and for $\alpha, \beta \in \mathbb{R}$ we say that H is a *pairwise* (α,β)*-spanner* of G w.r.t. P, if H is a spanning subgraph of G such that $d(s,t,H) \leq \alpha \cdot d(s,t,G) + \beta$ for all pairs $(s,t) \in P$. The values α and β are called the *multiplicative* and the *additive stretch* of H, respectively. An $(\alpha,0)$-spanner is also referred to as an α-spanner, and a $(1,\beta)$-spanner is called a β-additive spanner. A 1-spanner is also called a *preserver*. In the following we will consider *single-pair* spanners (i.e., $P = \{(s,t)\}$), $s \times T$ spanners (i.e., $P = \{s\} \times T$, for some $T \subseteq V(G)$), *single-source* spanners (i.e., $P = \{s\} \times V(G)$ spanners), *sourcewise* spanners (i.e., $P = S \times V(G)$, for some set of sources $S \subseteq V(G)$), *subsetwise* spanners (i.e., $P = S \times S$, for some subset $S \subseteq V(G)$), and *all-pairs* spanners (i.e., subsetwise spanners with $S = V(G)$).

A *hierarchical graph* $\mathcal{G} = (V(\mathcal{G}), E_1, E_2, \ldots, E_L)$ consists of a set $V(\mathcal{G})$ of n vertices, and of L pairwise disjoint subsets $E_1, \ldots, E_L$ of edges with endvertices in $V(\mathcal{G})$. We define the set of all edges of $\mathcal{G}$ as $E(\mathcal{G}) = \cup_{i=1}^{L} E_i$, and we let $m = |E(\mathcal{G})|$. Equivalently, one can think of $\mathcal{G}$ as a sequence of graphs $G_0, G_1, \ldots, G_L$, all defined on the same set of vertices $V(\mathcal{G})$, where the set of edges in G_ℓ is $\cup_{i=1}^{\ell} E_i$. Observe that G_0 is the empty graph, and that $G_0 \subseteq G_1 \subseteq G_2 \subseteq \cdots \subseteq G_L$, hence the name hierarchical graph. Throughout the paper, for the sake of simplicity, we assume that G_L is connected.

A hierarchical graph $\mathcal{H} = (V(\mathcal{H}), E'_1, \ldots, E'_L)$ is said to be an *hierarchical subgraph* of $\mathcal{G} = (V(\mathcal{G}), E_1, \ldots, E_L)$ if $E'_\ell \subseteq E_\ell$ for each $\ell = 0, \ldots, L$. A hierarchical (α,β)-spanner $\mathcal{H}$ of $\mathcal{G}$ w.r.t. a set of pairs $P \subseteq V(\mathcal{G}) \times V(\mathcal{G})$ is a *hierarchical subgraph* of $\mathcal{G}$ such that H_ℓ is an (α,β)-spanner of G_ℓ w.r.t. P for every $\ell = 1, \ldots, L$.

3 Pairwise hierarchical spanners and beyond

In this section we describe a simple algorithm to build pairwise hierarchical (α,β)-spanner w.r.t. a set of pairs $P \subseteq V(\mathcal{G}) \times V(\mathcal{G})$. Since the algorithm will be used as a building-block for several other hierarchical spanner constructions, it is useful to introduce a parameter D, and handle the more general case in which the sought pairwise hierarchical (α,β)-spanner $\mathcal{H}$ is only required to approximate *short distances*, i.e., it must hold that $d(s,t,H_\ell) \leq \alpha \cdot d(s,t,G_\ell) + \beta$ for every pair $(s,t) \in P$ and every $\ell \in \{1, \ldots, L\}$ for which $d(s,t,G_\ell) \leq D$.

The algorithm builds a collection Π of paths whose union yields exactly the edge set of the returned spanner $\mathcal{H}$. For each pair $(s,t) \in P$, the algorithm examines all $\ell = 1, \ldots, L$ in order, and greedily adds a shortest path π_ℓ from s to t in G_ℓ if $|\pi_\ell| \leq D$ and no previous path added to Π for this pair guarantees an (α,β)-approximate distance in G_ℓ. The pseudocode is given in Algorithm 1. Moreover, in Appendix A, we discuss how to efficiently compute Π.

Define M as 2 plus the minimum between the following quantities: (i) L, (ii) $\frac{2D}{\beta+1}$, and (iii) $\log_\alpha \frac{D}{\beta+1}$, where (iii) is only part of the minimum when $\alpha > 1$. We first provide a lemma that is useful to bound the total number of paths added to Π by the algorithm.

▶ **Lemma 1.** *For any fixed pair $(s,t) \in P$, the number of shortest paths π_ℓ from s to t added to Π by Algorithm 1 is at most M.*

Proof. Let $\pi^{(k)}, \pi^{(k-1)}, \ldots, \pi^{(0)}$ be the shortest paths π_ℓ added to Π by Algorithm 1 when the pair (s,t) is considered by the outer loop, in this order. In the rest of the proof we assume that k is at least 1 and we upper bound its value.

Algorithm 1 Computes a pairwise hierarchical (α, β)-spanner for short distances of $\mathcal{G}$.

Input : A hierarchical graph $\mathcal{G}$, a set of vertex pairs $P \subseteq V(\mathcal{G}) \times V(\mathcal{G})$, two values $\alpha \geq 1$, $\beta \geq 0$, and a distance threshold D.

Output : A pairwise hierarchical (α, β)-spanner $\mathcal{H}$ for short distances of $\mathcal{G}$ w.r.t. P.

$\Pi \leftarrow \emptyset$;
foreach $(s, t) \in P$ **do**
  $D' \leftarrow +\infty$; `// Length of the shortest path from` s `to` t `added to` Π
  for $\ell = 1, \ldots, L$ **do**
    $\pi_\ell \leftarrow$ a shortest path from s to t in G_ℓ or $\perp$ if no such path exists;
    if $\pi_\ell \neq \perp$ **and** $|\pi_\ell| \leq D$ **and** $D' > \alpha|\pi_\ell| + \beta$ **then**
      $\Pi \leftarrow \Pi \cup \{\pi_\ell\}$;
      $D' \leftarrow |\pi_\ell|$;
$\mathcal{H} \leftarrow$ the hierarchical spanning subgraph of $\mathcal{G}$ containing all edges of the paths in Π;
return $\mathcal{H}$;

The first upper bound is $k < L$ and holds since the algorithm adds at most one such path per graph of the hierarchy.

Moreover, we must have $|\pi^{(k)}| \leq D$, $|\pi^{(0)}| \geq 1$, and $|\pi^{(i)}| > \alpha|\pi^{(i-1)}| + \beta$ for $i = 1, \ldots, k$, which implies $|\pi^{(i)}| \geq |\pi^{(i-1)}| + \max\{\beta, 1\}$. Then, $D \geq |\pi^{(k)}| > k \cdot \max\{\beta, 1\}$, yielding the second upper bound $k < \frac{D}{\max\{\beta,1\}} \leq \frac{2D}{\beta+1}$.

Finally, whenever $\alpha > 1$, we have $D \geq |\pi^{(k)}| > \alpha^k + \beta \sum_{j=0}^{k-1} \alpha^j \geq \alpha^k + \beta\alpha^{k-1} > \alpha^{k-1}(\beta + 1)$, which provides the third upper bound $k < 1 + \log_\alpha \frac{D}{\beta+1}$. ◀

The next two lemmas show two different upper bounds on the size of the pairwise hierarchical (α, β)-spanner $\mathcal{H}$ computed by the algorithm.

▶ **Lemma 2.** *Algorithm 1 computes a pairwise hierarchical (α, β)-spanner for short distances of size at most $|P|MD$. Moreover, if $\alpha > 1$, the size of the spanner is at most $\frac{\alpha|P|D}{\alpha-1}$.*

Proof. The $|P|MD$ upper bound follows since, by Lemma 1, at most M paths are added for each pair of vertices in P, and that each such path has length at most D by construction.

Consider now the case in which $\alpha > 1$. For a fixed pair of vertices $(s, t) \in P$, the length of the i-th path from s to t added to Π by the algorithm must be at most $\frac{D}{\alpha^{i-1}}$. Therefore, the size of $\mathcal{H}$ is upper bounded by $|P| \sum_{i=1}^{\infty} \frac{D}{\alpha^{i-1}} = \frac{\alpha|P|D}{\alpha-1}$. ◀

▶ **Lemma 3.** *Algorithm 1 computes a pairwise hierarchical (α, β)-spanner for short distances of size $O(n\sqrt{|P|M})$.*

Proof. We consider the paths π_ℓ added to Π in an arbitrary non-decreasing order of ℓ, and we study how a non-hierarchical version $\widetilde{H}$ of $\mathcal{H}$ changes overt time.

Initially $\widetilde{H}$ is the graph containing all vertices in $V(\mathcal{H})$ and no edges, and when we consider some π_ℓ, $\widetilde{H}$ contains all edges of the paths that precede π_ℓ in the ordering.

Let $\Delta = \sqrt{|P|M}$. We say that a vertex in $\widetilde{H}$ has *high degree* if its degree is at least Δ.

For a path π_ℓ, it holds that the the number $\eta(\pi_\ell)$ of vertices $v \in V(\pi_\ell)$ having high degree in $\widetilde{H}$ is at most $\frac{3n}{\Delta}$. Indeed, if more than $\frac{3n}{\Delta}$ such vertices had degree larger than Δ, then there would be some vertex $u \in V(\widetilde{H})$ incident to at least $\left\lceil \frac{(1+\lfloor \frac{3n}{\Delta} \rfloor)\cdot\Delta}{n} \right\rceil \geq 4$ vertices from π_ℓ. This cannot be the case since 2 of such vertices would have distance at least 3 in π_ℓ, but at most 2 in $\widetilde{H}$, contradicting the fact that π_ℓ is a shortest path in G_ℓ (which contains all edges added in $\widetilde{H}$ up to this point).

Consider now the graph $\widetilde{H}$ after all paths π_ℓ have been processed. From the above discussion, we have that the final degree $\deg(v)$ of a vertex in $\widetilde{H}$ is upper bounded by Δ plus the number of edges added the paths in Π when v has high degree.

Then, $\sum_{v\in V(\widetilde{H})} \deg(v) \le n\Delta + 2\sum_{\pi_\ell\in\Pi} \eta(\pi_\ell)$ since a path π_ℓ contains at most 2 edges incident to each of its $\eta(\pi_\ell)$ vertices of high degree. The overall number of edges of $\widetilde{H}$ is upper bounded by:

$$\sum_{v\in V(\widetilde{H})} \deg(v) \le n\Delta + 2\sum_{\pi_\ell\in\Pi} \eta(\pi_\ell) \le n\Delta + |\Pi|\cdot\frac{6n}{\Delta} = O\left(n\Delta + |P|M\frac{n}{\Delta}\right) = O(n\sqrt{|P|M}),$$

where we used $|\Pi| \le |P|M$ since at most M paths π_ℓ are added for each pair in P, as shown by Lemma 1. ◀

We provide an additional technical lemma that will be used to build sparser pairwise hierarchical spanners than those obtained by naïvely using Algorithm 1. More precisely, we will use pivoting techniques to *hit* long paths, thus allowing us to apply Algorithm 1 with more refined values of the parameter D. These techniques have already been successfully applied to design sparse spanners for non-hierarchical graphs.

▶ **Lemma 4.** *Given a collection $\mathcal{Z}$ of subsets of $\{1,\dots,n\}$, where each subset has size at least x and $|\mathcal{Z}|$ is polynomially bounded in n, we can find in polynomial time a subset $R \subseteq \{1,\dots,n\}$ of size $O(\frac{n\log n}{x})$ that hits all subsets in $\mathcal{Z}$, i.e., $R\cap Z \neq \emptyset$ for all $Z \in \mathcal{Z}$.*

Proof. Consider an iterative greedy algorithm that builds a sequence of partial hitting sets R_i while keeping track of number n_i of sets $S \in \mathcal{S}$ that are not hit by R_i. Initially $R_0 = \emptyset$ and $n_0 = |\mathcal{Z}|$. In the generic i-th iteration, the algorithm finds an element $j \in \{1,\dots,n\}$ maximizing the number $c_i(j)$ of sets $Z \in \mathcal{Z}$ such that $j \in S$ and $S \cap R_{i-1} = \emptyset$ and adds it to the next partial hitting set, i.e., it sets $R_i = R_{i-1} \cup \{j\}$ and $n_i = n_{i-1} - c_i(j)$. The algorithm stops as soon as $n_i = 0$ and returns R_i.

In the rest of the proof we show that the number of iterations of the algorithm is at most $1 + \frac{n}{x}\ln|\mathcal{Z}| = O(\frac{n}{\ell}\log n)$, thus simultaneously bounding the running time of the algorithm (notice that each iteration can be performed in polynomial time) and the size of the returned hitting set. At the beginning of the i-th iteration, there are n_{i-1} sets in $\mathcal{Z}$ that are not hit by R_{i-1} and each occurrence of an element j' in any such set contributes 1 to $c_i(j')$. Since each set contains at least x elements, we have $\sum_{j'=1}^{n} c_i(j') \ge n_{i-1}x$ implying that $c_i(j) = \max_{j'=1,\dots,n} c_i(j') \ge \frac{1}{n}\sum_{j'=1}^{n} c_i(j') \ge \frac{n_{i-1}x}{n}$ and hence $n_i = n_{i-1} - c_i(j) \le n_{i-1}\cdot(1-\frac{x}{n}) \le n_0\cdot(1-\frac{x}{n})^i = |\mathcal{Z}|\cdot\left((1-\frac{x}{n})^n\right)^{i/n} \le |\mathcal{Z}|\cdot e^{-xi/n}$. As a consequence, there must be some $i \le 1 + \frac{n}{x}\ln|\mathcal{Z}|$ such that $n_i = 0$. Indeed, for $i > \frac{n}{x}\ln|\mathcal{Z}|$, we have $|\mathcal{Z}|\cdot e^{-xi/n} < 1$. ◀

We are now ready to provide the main results of this section.

▶ **Theorem 5.** *Given a hierarchical graph $\mathcal{G}$ with n vertices and depth L, Algorithm 1 can be used to compute any of the following hierarchical spanners of $\mathcal{G}$ of choice:*

1. *A pairwise hierarchical $(1+\varepsilon)$-spanner w.r.t. $P \subseteq V(\mathcal{G})\times V(\mathcal{G})$, where $\varepsilon \in (0,1]$, of size $O\Big(n\min\{\sqrt{\frac{|P|\log n}{\varepsilon}}, \sqrt{|P|L}, \frac{|P|}{\varepsilon}\}\Big)$;*
2. *A pairwise hierarchical β-additive spanner, with $\beta \ge 0$, of size $O\Big(n\sqrt{|P|\min\{L, \frac{n}{\beta+1}\}}\Big)$ w.r.t. $P \subseteq V(\mathcal{G})\times V(\mathcal{G})$;*
3. *A subsetwise hierarchical preserver of size $O(n\min\{|S|\sqrt{L}, \sqrt{|S|n\log n}\})$ w.r.t. $S \subseteq V(\mathcal{G})$.*

Proof. From Lemma 2 with $D = n$ and $\alpha = 1 + \varepsilon$, the size of a pairwise hierarchical $(1+\varepsilon)$-spanner is at most $\frac{(1+\varepsilon)n}{\varepsilon} = O(\frac{|P|n}{\varepsilon})$. Moreover, from Lemma 3 with $D = n$, we can also upper bound the size as $O(n\sqrt{|P| \min\{L, \frac{\log n}{\varepsilon}\}})$, where we used $\log_{1+\varepsilon} n = \Theta(\frac{\log n}{\varepsilon})$. This proves the first claim.

The second claim follows directly from Lemma 3 once M is replaced with its definition.

In the rest of the proof we focus on the third claim, and we show how to use Algorithm 1 to build a subsetwise hierarchical preserver w.r.t. a subset $S \subseteq V(\mathcal{G})$. We run Algorithm 1 to compute the set of paths Π added to a hierarchical preserver for $P = S \times S$ and $D = n$. Then, we build a collection $\mathcal{Z}$ of subsets of $V(\mathcal{G})$ by considering each path $\pi \in \Pi$ with $|\pi| \geq x - 1$, where $x = \lceil \frac{n \log n}{|S|} \rceil$. For each such π, we let $v_1, \dots, v_{|\pi|+1}$ be the vertices of π, in order, and we add to $\mathcal{Z}$ all sets $V(\pi[v_{ix+1} : v_{(i+1)x}])$ for $i = 0, \dots, \lfloor \frac{|\pi|+1}{x} \rfloor - 1$. We then compute a hitting set R of $\mathcal{Z}$ of size $O(\frac{n \log n}{x}) = O(|S|)$ as shown in Lemma 4.

The returned subsetwise hierarchical preserver $\mathcal{H}$ is the one obtained by running Algorithm 1 once more with $P = (S \cup R) \times (S \cup R)$ and $D = 2x - 1$. From Lemma 3, the size of $\mathcal{H}$ is $O(n\sqrt{|P|M}) = O(n\sqrt{(|S| + |R|)^2 \min\{L, x\}}) = O\left(n \min\{|S|\sqrt{L}, \sqrt{|S| n \log n}\}\right)$.

We now argue that $\mathcal{H}$ is a subsetwise hierarchical preserver w.r.t $S \times S$. Fix some $\ell \in \{1, \dots, L\}$ and a pair $(s,t) \in S \times S$ such that s and t are connected in G_ℓ. Let π be the shortest path from s to t in G_ℓ that is in Π, let $v_1, \dots, v_{|\pi|+1}$ be the vertices of π, in order, and choose an arbitrary $r_i \in V(\pi[v_{ix+1} : v_{(i+1)x}]) \cap R$ for each $i = 0, \dots, k$ with $k = \lfloor \frac{|\pi|+1}{x} \rfloor - 1$. Define $r_0 = v_1 = s$, and $r_{k+1} = v_{|\pi|+1} = t$. For $i = 0, \dots, k$, both r_i and r_{i+1} belong to $S \cup R$ and satisfy $|\pi[r_i : r_{i+1}]| \leq (i+2)x - (ix+1) = 2x - 1$, hence $d(r_i, r_{i+1}, H_\ell) = d(r_i, r_{i+1}, G_\ell)$. The claim follows from $d(s,t,H_\ell) \leq \sum_{i=0}^{k} d(r_i, r_{i+1}, H_\ell) = \sum_{i=0}^{k} d(r_i, r_{i+1}, G_\ell) = d(s,t,G_\ell)$. ◀

Theorem 5 has the following interesting implications. The size of an $s \times T$ hierarchical preservers is $O(n\sqrt{n|T|})$, which becomes $O(n\sqrt{n})$ in the special case of single-pair hierarchical preservers. This latter bound is asymptotically tight, as shown in Section 6. Moreover, single-pair hierarchical $(1+\varepsilon)$-spanners size $O(\frac{n}{\varepsilon})$, which is linear for constant values of ε. We also note that, for $\beta = o(\frac{|S|}{\log n})$, the size of our subsetwise hierarchical preserver is asymptotically smaller than that of the subsetwise hierarchical β-additive spanner.

4 Single-source hierarchical spanners

In this section we show how the iterative pivoting techniques – which have been used to obtain sparse spanners for non-hierarchical graphs – can be adapted to design single-source hierarchical $(1+\varepsilon)$-spanners of size $O(\frac{n \log^4 n}{\varepsilon})$. Throughout the rest of this section, we denote the source vertex as s.

Given $x > 0$, we define the x-suffix of a path π from s to t as the subpath of π induced the last $\min\{|\pi|, \lfloor x \rfloor\}$ edges of π in a traversal from s to t.

Our construction is given by an iterative algorithm that starts with a hierarchical spanning subgraph $\mathcal{H}$ with no edges. The algorithm performs $k = \lceil \log n \rceil$ iterations, and in a generic iteration i, the algorithm augments $\mathcal{H}$ with suitable suffixes of a collection of a hierarchical paths from s to the vertices in a *pivot set* R_{i-1}, and computes the next pivot set R_i. The initial pivot set R_0 is $V(\mathcal{G})$. In details, the i-th iteration invokes Algorithm 1 with $\alpha = (1+\varepsilon)^{\frac{1}{k}}$, to build a collection Π_i of paths from s to R_{i-1} whose union yields an $s \times R_{i-1}$ hierarchical α-spanner of $\mathcal{G}$. Then, we examine each path $\pi \in \Pi_i$ and we add all the edges of its x_i-suffix π' to $\mathcal{H}$, where $x_i = n^{\frac{i}{k}} \log^{1-\frac{i}{k}} n$. If $|\pi| \geq x_i$, we also add the set $V(\pi')$ (of size at least x_i) to a collection $\mathcal{Z}_i$, which is initially empty. Finally, we compute R_i as a hitting set of $\mathcal{Z}_i$ of size $O(\frac{n}{x_i} \log n)$ using Lemma 4.

▶ **Theorem 6.** *For any $\varepsilon \in (0,1]$, the hierarchical graph $\mathcal{H}$ returned by the algorithm is a single-source hierarchical $(1+\varepsilon)$-spanner of $\mathcal{G}$ of size $O(\frac{n\log^4 n}{\varepsilon})$.*

Proof. We start by shoring that the multiplicative stretch of $\mathcal{H}$ is $1+\varepsilon$. More precisely, we prove the following: for every $i \in \{0,1,\ldots,k-1\}$ and for every $\ell \in \{1,\ldots,L\}$, we have $d(s,t,H_\ell) \le \alpha^{k-i} d(s,t,G_\ell)$ for all $t \in R_i$. The claim follows by choosing $i=0$ since $R_0 = V(\mathcal{G})$ and $\alpha = (1+\varepsilon)^{\frac{1}{k}}$.

The proof by reverse induction on i. The base case is $i = k-1$. In the k-th iteration, the algorithm computed the set Π_k of paths of an $s \times R_{k-1}$ hierarchical α-spanner of $\mathcal{G}_\ell$, and all paths in $\pi \in \Pi_k$ are added to $\mathcal{H}$ since $|\pi| \le n = x_k$.

Regarding the inductive step, assume that the claim holds for some $i+1 \le k-1$, and consider i. Given any $\ell \in \{1,\ldots,L\}$ and any $t \in R_i$, the set Π_{i+1} computed during the $(i+1)$-th iteration of the algorithm contains some path π from s to t of G_ℓ such that $|\pi| \le \alpha \cdot d(s,t,G_\ell)$. If $|\pi| \le x_{i+1}$, then $\mathcal{H}$ contains all edges of π, hence $d(s,t,H_\ell) \le |\pi|$ and the claim holds. Otherwise, $|\pi| > x_{i+1}$, and the edges of a x_{i+1}-suffix π' of π are in $\mathcal{H}$. Moreover, the set of vertices $V(\pi')$ is in the collection $\mathcal{Z}_{i+1}$, hence there is some $r \in R_{i+1} \cap V(\pi')$. By inductive hypothesis, $d(s,r,H_\ell) \le \alpha^{k-(i+1)} d(s,r,G_\ell)$. By the triangle inequality, we have:

$$\begin{aligned} d(s,t,H_\ell) &\le d(s,r,H_\ell) + d(r,t,H_\ell) \le \alpha^{k-(i+1)} d(s,r,G_\ell) + |\pi[r:t]| \\ &\le \alpha^{k-(i+1)} |\pi[s:r]| + |\pi[r:t]| \le \alpha^{k-(i+1)} (|\pi[s:r]| + |\pi[r:t]|) \\ &= \alpha^{k-(i+1)} |\pi| \le \alpha^{k-i} d(s,t,G_\ell). \end{aligned}$$

This concludes the proof by induction. We now bound the size of $\mathcal{H}$. For $i = 0,\ldots,n$, the set R_i contains $O(n^{1-\frac{i}{k}} \log^{\frac{i}{k}} n)$ vertices. Indeed, when > 1 we have $|R_i| = O(\frac{n}{x_i}\log n)$ where $x_i = n^{\frac{i}{k}} \log^{1-\frac{i}{k}} n$, while $|R_0| = O(n)$.

Consider a generic iteration $i = 1,\ldots,k$. The number of edges added to $\mathcal{H}$ during such iteration is at most $|\Pi_i| x_i$, where we have $|\Pi_i| \le |R_{i-1}| \log_\alpha n$ as shown by invoking Lemma 1 with $D = n$ and $\beta = 0$. Then, the overall number of edges added to $\mathcal{H}$ during iteration i is at most:

$$O(|R_{i-1}| x_i \log_\alpha n) = O(n^{1-\frac{i-1}{k}} \log^{\frac{i-1}{k}} n \cdot n^{\frac{i}{k}} \log^{1-\frac{i}{k}} n \log_\alpha n) = O(n^{1+\frac{1}{k}} \log^{1-\frac{1}{k}} n \cdot \log_\alpha n).$$

Since there are $k = O(\log n)$ iterations, the overall size of $\mathcal{H}$ is

$$O\left(kn^{1+\frac{1}{k}} \log^{1-\frac{1}{k}} n \cdot \log_\alpha n\right) = O\left(\frac{kn^{1+\frac{1}{k}} \log^2 n}{\log \alpha}\right) = O\left(\frac{k^2 n^{1+\frac{1}{k}} \log^2 n}{\log(1+\varepsilon)}\right) = O\left(\frac{n\log^4 n}{\varepsilon}\right),$$

where we used $n^{\frac{1}{k}} = O(1)$, and $\log(1+\varepsilon) = \Theta(\varepsilon)$ for every $\varepsilon \in (0,1]$. ◀

5 All-pairs hierarchical spanners

In this section we show a general technique to transform a construction of a subsetwise or sourcewise hierarchical spanner into one for all-pairs hierarchical spanners whose resulting size and stretch both depend on the corresponding parameters of the input construction. The application of such techniques to the constructions of Section 3 and Section 4 allows us to obtain hierarchical 4-additive spanners of size $O(n^{\frac{5}{3}} \log^{\frac{1}{3}} n)$ and hierarchical $(1+\varepsilon, 2)$-spanners of size $O(\frac{n\sqrt{n}\log^2 n}{\sqrt{\varepsilon}})$ for any $\varepsilon \in (0,1]$ of choice.

The technique we use is based on a *clustering* procedure, which is another known technique to design sparse spanners for non-hierarchical graphs. However, we have to carefully adapt the technique to let it work to the case in which we have a hierarchy of graphs.

Our clustering procedure is parameterized by an integer Δ. The procedure returns an ordered collection $\mathcal{C} = \langle C_1, C_2, \dots \rangle$ of *clusters*. Each cluster C_i is a triple $\langle V_i, F_i, c_i \rangle$ where $V_i \subseteq V(\mathcal{G})$, $F_i \subseteq E(\mathcal{G})$ is a set of edges each having at least one endvertex in V_i, and $c_i \in V_i$ is the *center* of cluster C_i. The sets V_i will be pairwise disjoint, and F_i will contain all edges (c_i, v) for $v \in V_i \setminus \{c_i\}$.

We build our collection of clusters by examining the edges of $\mathcal{G}$ in some fixed order $e_1, e_2, \dots$ that places all edges of E_ℓ before those in $E_{\ell+1}$, for every $\ell = 1, \dots, L-1$, where edges in the same E_ℓ are ordered arbitrarily. During the process, we keep track of an auxiliary (non-hierarchical) graph $\widetilde{G}$ containing all *unclustered* vertices of $\mathcal{G}$, and all edges, each connecting two unclustered vertices, among those that have already been examined by the procedure. Initially, $\widetilde{G}$ contains all vertices of $\mathcal{G}$ and no edges. Whenever an edge (u, v) is examined, we add (u, v) to $\widetilde{G}$ if both u and v are vertices in $V(\widetilde{G})$. If the edge is added, we additionally check whether at least one endvertex $x \in \{u, v\}$ has degree at least Δ in $\widetilde{G}$ and, in such a case, we create the next cluster $C_i = \langle V_i, F_i, c_i \rangle$ of the sequence, where $c_i = x$, V_i is the closed neighborhood of x in $\widetilde{G}$, and F_i contains all edges in $\widetilde{G}$ that have at least one endvertex in V_i. Finally, we delete V_i and F_i from $\widetilde{G}$.

Let $\mathcal{H}_\mathcal{C}$ be the hierarchical spanning subgraph of $\mathcal{G}$ containing all edges of $\mathcal{G}$ whose both endvertices are unclustered (i.e., exactly all remaining edges in $\widetilde{G}$ at the end of the clustering procedure), plus all clustered edges (i.e., all edges in $\cup_{i=1}^{|\mathcal{C}|} F_i$). Hence, if an edge $e \in E(\mathcal{G})$ is missing from $E(\mathcal{H}_\mathcal{C})$, then at least one endvertex of e must belong to some cluster (while the converse is not necessarily true). The following lemma shows that the $\mathcal{H}_\mathcal{C}$ is sparse and also provides a bound on the total number of clusters returned by the clustering procedure.

▶ **Lemma 7.** *It holds that $|\mathcal{C}| \leq \frac{n}{\Delta}$ and that $|E(\mathcal{H}_\mathcal{C})| = O(n\Delta)$.*

Proof. Since sets V_i of different clusters C_i are pairwise disjoint, and each set contains at least Δ vertices by construction, the procedure creates at most $\frac{n}{\Delta}$ clusters.

Consider then the set F_i of edges of some generic cluster C_i. Cluster C_i was created following the addition fo some edge e_j. Therefore, immediately before the addition of e_j to $\widetilde{G}$, each vertex in V_i had degree smaller than Δ. Hence, the number of edges in F_i is at most $|V_i|\Delta + 1$. Hence, $\sum_{i=1}^{|\mathcal{C}|} |F_i| \leq n\Delta + |\mathcal{C}| = O(n\Delta)$. Finally, each vertex in $\widetilde{G}$ at the end of the clustering procedure has degree at most Δ, hence the number of edges with at least one unclustered endpoint is at most $n\Delta$. ◀

The next lemma shows a key property that is guaranteed by our clustering procedure for all the edges of $\mathcal{G}$ that are not added to $\mathcal{H}_\mathcal{C}$.

▶ **Lemma 8.** *Let $(u, v) \in E_\ell$ be an edge that is not contained in $\mathcal{H}_\mathcal{C}$. There exists a cluster $C_i \in \mathcal{C}$, such that $V_i \cap \{u, v\} \neq \emptyset$ and $F_i \subseteq E(G_\ell)$.*

Proof. Observe that all edges that are inserted in $\widetilde{G}$ at some point during the clustering procedure belong to $\mathcal{H}_\mathcal{C}$. Indeed, they are either added to some set F_j of some cluster C_j, or they remain in $\widetilde{G}$ until the end of the procedure (and all such remaining edges are part of $\mathcal{H}_\mathcal{C}$ by its definition).

Then, (u, v) is never inserted into $\widetilde{G}$, which implies that at least one of u and v belongs to some cluster. Let C_i be the cluster having $V_i \cap \{u, v\} \neq \emptyset$ that minimizes i. Since (u, v) is never inserted into $\widetilde{G}$, cluster C_i must have been built before the clustering procedure examined (u, v). By the order in which edges are examined, each edge in F_i belongs to some $E_{\ell'}$ with $\ell' \leq \ell$, hence $F_i \subseteq E(G_\ell)$, as claimed. ◀

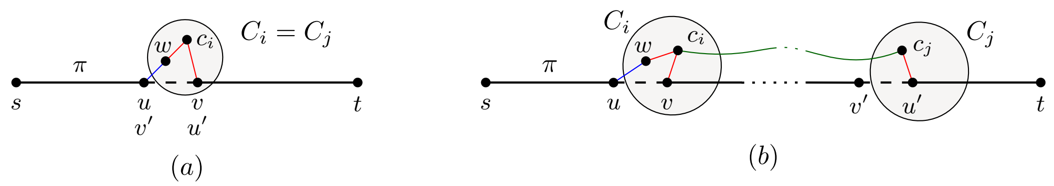

Figure 1 A qualitative representation of the proof of Theorem 9. Red edges are in the sets F_i and F_j corresponding to clusters C_i and C_j. The cross edges added during the second augmentation step are shown in blue. Finally, the edges of a shortest path from c_i to c_j in $\mathcal{H}_S$ (added during the first augmentation step) are shown in green.

5.1 All-pairs hierarchical spanners from subsetwise hierarchical spanners

Let $\mathcal{C}$ be the ordered collection of clusters computed by the clustering procedure and let $S = \{c_1, c_2, \dots, c_{|\mathcal{C}|}\}$ be the cluster centers. In this section we assume the existence of an algorithm that constructs a subsetwise hierarchical (α, β)-spanner $\mathcal{H}_S$ of $\mathcal{G}$ w.r.t. S of size at most $f(n, |S|)$, where n is the number of vertices of the input hierarchical graph and f is non-decreasing in both n and $|S|$. We show that we can use such a construction to obtain a hierarchical $(\alpha, 2\alpha + \beta + 2)$-spanner $\mathcal{H}$ of size $O(f(n, \frac{n}{\Delta}) + n\Delta + \frac{n^2}{\Delta})$.

We build our hierarchical spanner $\mathcal{H}$ by augmenting the hierarchical spanning subgraph $\mathcal{H}_\mathcal{C}$ computed by the clustering procedure with two types of edges:

Spanner edges: which are all the edges of $\mathcal{H}_S$;

Cross edges: For each cluster $C_i \in \mathcal{C}$, and for each vertex $u \in V(\mathcal{G}) \setminus V_i$, we add the edge $e_j = (u, w)$ with $w \in V_i$ that minimizes j w.r.t. the ordering of the edges considered by the clustering procedure, if any.

▶ **Theorem 9.** *$\mathcal{H}$ is a hierarchical $(\alpha, 2\alpha + \beta + 2)$-spanner of size $O(f(n, \frac{n}{\Delta}) + n\Delta + \frac{n^2}{\Delta})$.*

Proof. By Lemma 7, $\mathcal{H}_\mathcal{C}$ has $O(n\Delta)$ edges, and $|S| = |\mathcal{C}| \le \frac{n}{\Delta}$. Hence the number of cross edges is at most $n|\mathcal{C}| \le \frac{n^2}{\Delta}$. Since the size of $\mathcal{H}_S$ is $f(n, |S|) \le f(n, \frac{n}{\Delta})$, this proves the claim on the size of $\mathcal{H}$.

We now argue that $\mathcal{H}$ is a hierarchical $(\alpha, 2\alpha + \beta + 2)$-spanner. Let $s, t \in V(\mathcal{G})$ be two distinct vertices, and consider a shortest path π from s to t in G_ℓ, for some $\ell \in \{1, \dots, L\}$, such that π is not entirely contained in $\mathcal{H}$. Let $e = (u, v)$ (resp. $e' = (v', u')$) be the first (resp. last) edge of π (when traversed from s to t) that is not contained in H_ℓ, where u is closer to s than v (resp. u' is closer to t than v'). Observe that e and e' might coincide. By Lemma 8, some endvertex x of e (resp. x' of e') is contained in the set V_i of some cluster C_i with center c_i and $F_i \subseteq E(G_\ell)$ (resp. in the set V_j of some cluster C_j with center c_j and $F_j \subseteq E(G_\ell)$). See Figure 1. Observe that, for any vertex $w \in V_i$, we must have $d(w, c_i, H_\ell) \le 1$ since F_i either $w = c_i$ or $(w, c_i) \in F_i \subseteq E(H_\ell)$. Similarly, $d(w, c_j, H_\ell) \le 1$ for all $w \in V_j$.

We now argue that $d(s, c_i, H_\ell) \le d(s, x, G_\ell) + 1$. The proof depends on whether $x = u$ or $x = v$. If $x = u$, then $d(s, c_i, H_\ell) \le |\pi[s : u]| + d(u, c_i, H_\ell) \le d(s, u, G_\ell) + 1$. Consider then the case $x = v$. Since $(u, v) \in E(G_\ell)$, our augmentation step must have added some cross edge (u, w) with $w \in V_i$ which precedes (u, v) in the ordering, thus implying that $(u, w) \in E(H_\ell)$. Then: $d(s, c_i, H_\ell) \le d(s, u, H_\ell) + 1 + d(w, c_i, H_\ell) = d(s, v, H_\ell) + d(w, c_i, H_\ell) \le d(s, v, H_\ell) + 1$.

A symmetric argument shows that $d(c_j, t, H_\ell) \le (x', t, G_\ell) + 1$.

If $e = e'$, then $C_i = C_j$, $c_i = c_j$, and $x = x'$ (see Figure 1 (a)), therefore we can use the fact that $x = x'$ lies on the shortest path π in G_ℓ to write:

$$d(s, t, H_\ell) \le d(s, c_i, H_\ell) + d(c_j, t, H_\ell) \le d(s, x, G_\ell) + d(x', t, G\ell) + 2 = d(s, t, G_\ell) + 2.$$

If $e \neq e'$, then x is encountered no later than x' when π is traversed from s to t, hence $d(s, x, G_\ell) + d(x, x', G_\ell) + d(x', t, G_\ell) = d(s, t, G_\ell)$. Since our augmentation step adds all edges of a $S \times S$ hierarchical (α, β)-spanner $\mathcal{H}_S$ of $\mathcal{G}$ to $\mathcal{H}$, where S is the set of cluster centers, it holds that $d(c_i, c_j, H_\ell) \leq \alpha d(c_i, c_j, G_\ell) + \beta$, and we have (see Figure 1 (b)):

$$\begin{aligned} d(s, t, H_\ell) &\leq d(s, c_i, H_\ell) + d(c_i, c_j, H_\ell) + d(c_j, t, H_\ell) \\ &\leq d(s, x, G_\ell) + \alpha d(c_i, c_j, G_\ell) + \beta + d(x', t, G_\ell) + 2 \\ &\leq d(s, x, G_\ell) + \alpha\big(d(c_i, x, G_\ell) + d(x, x', G_\ell) + d(x', c_j, G_\ell)\big) + \beta + d(x', t, G_\ell) + 2 \\ &= d(s, x, G_\ell) + \alpha d(x, x', G_\ell) + d(x', t, G_\ell) + 2\alpha + \beta + 2 \\ &\leq \alpha d(s, t, G_\ell) + 2\alpha + \beta + 2. \end{aligned}$$ ◀

▶ **Corollary 10.** *It is possible to construct hierarchical 4-additive spanners of size $O(n^{5/3} \log^{1/3} n)$.*

Proof. We set $\Delta = n^{2/3} \log^{\frac{1}{3}} n$ and combine Theorem 9 with the construction of Theorem 5 (3) providing a subsetwise hierarchical preserver of size at most $f(n, |S|) = O(n\sqrt{|S| n \log n})$. ◀

5.2 All pairs hierarchical spanners from sourcewise hierarchical spanners

Let $\mathcal{C}$ be the ordered collection of clusters computed by the clustering procedure and let $S = \{c_1, c_2, \dots, c_{|\mathcal{C}|}\}$ be the cluster centers. In this section we assume the existence of an algorithm that constructs a sourcewise hierarchical (α, β)-spanner $\mathcal{H}_S$ of $\mathcal{G}$ w.r.t. S of size at most $f(n, |S|)$, where n is the number of vertices of the input hierarchical graph and f is non-decreasing in both n and $|S|$. We show that we can use such a construction to obtain a hierarchical $(\alpha, \alpha + \beta + 1)$-spanner $\mathcal{H}$ of size $O(f(n, \frac{n}{\Delta}) + n\Delta + \frac{n^2}{\Delta})$.

This is done by augmenting the hierarchical spanning subgraph $\mathcal{H}_\mathcal{C}$ resulting from our clustering procedure with two types of edges:

Spanner edges: which are all edges of $\mathcal{H}_S$;

Cross edges: For each cluster $C_i \in \mathcal{C}$, and for each vertex $u \in V(\mathcal{G}) \setminus V_i$, we add the edge $e_j = (u, w)$ with $w \in V_i$ that minimizes j w.r.t. the ordering of the edges considered by the clustering procedure, if any.

▶ **Theorem 11.** *$\mathcal{H}$ is a hierarchical $(\alpha, \alpha + \beta + 1)$-spanner of size $O(f(n, \frac{n}{\Delta}) + n\Delta + \frac{n^2}{\Delta})$.*

Proof. By Lemma 7, $\mathcal{H}_\mathcal{C}$ has $O(n\Delta)$ edges, and $|S| = |\mathcal{C}| \leq \frac{n}{\Delta}$. Hence the number of cross edges is at most $n|\mathcal{C}| \leq \frac{n^2}{\Delta}$. Since the size of $\mathcal{H}_S$ is $f(n, |S|) \leq f(n, \frac{n}{\Delta})$, this proves the claim on the size of $\mathcal{H}$.

Let $s, t \in V(\mathcal{G})$ be two distinct vertices, and consider a shortest path π from s to t in G_ℓ, for some $\ell \in \{1, \dots, L\}$, such that π is not entirely contained in $\mathcal{H}$. Let $e = (u, v)$ be the first edge of π (when traversed from s to t) that is not contained in H_ℓ. By Lemma 8, there is an endvertex x of e that is contained in the set V_i of some cluster C_i with center c_i and $F_i \subseteq E(G_\ell)$. As argued in the proof of Theorem 9, we have $d(s, c_i, H_\ell) \leq d(s, x, G_\ell) + 1$.

Since $\mathcal{H}$ contains an $S \times V(\mathcal{G})$ hierarchical (α, β)-spanner of $\mathcal{G}$, where S is the set of cluster centers, it holds that $d(c_i, t, H_\ell) \leq \alpha d_G(s, t, G_\ell) + \beta$. Then:

$$\begin{aligned} d(s, t, H_\ell) &\leq d(s, c_i, H_\ell) + d(c_i, t, H_\ell) \leq d(s, x, G_\ell) + \alpha d(c_i, t, G_\ell) + \beta + 1 \\ &\leq d(s, x, G_\ell) + \alpha(d(x, t, G_\ell) + 1) + \beta + 1 \\ &= d(s, x, G_\ell) + \alpha d(x, t, G_\ell) + \alpha + \beta + 1 \\ &\leq \alpha d(s, t, G_\ell) + \alpha + \beta + 1. \end{aligned}$$ ◀

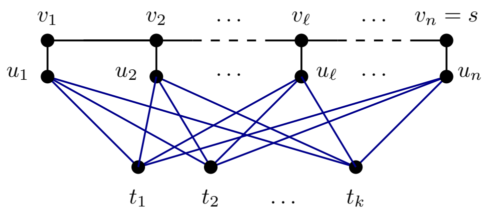

Figure 2 Lower bound construction for an $s \times T$ hierarchical preserver $\mathcal{H}$ with $T = \{t_1, \ldots, t_k\}$. The blue edges form a complete bipartite graph, between the vertices in $\{u_1, \ldots, u_n\}$ and T. The graph H_ℓ must contain all edges in $\{u_\ell\} \times T$.

▶ **Corollary 12.** *Given any $\varepsilon \in (0, 1]$, it is possible to construct hierarchical $(1+\varepsilon, 2)$-spanners of size $O\left(\frac{n\sqrt{n}\log^2 n}{\sqrt{\varepsilon}}\right)$.*

Proof. Let $\varepsilon' = \frac{\varepsilon}{2}$. Given a set S of vertices, we can take the union of the single-source hierarchical $(1+\varepsilon')$-spanner of Theorem 6 from each source $s \in S$, in order to obtain a sourcewise hierarchical $(1+\varepsilon')$-spanner of size at most $f(n, |S|) = O\left(\frac{|S|n\log^4 n}{\varepsilon}\right)$.

We choose $\Delta = \sqrt{\frac{n}{\varepsilon}}\log^2 n$ and combine Theorem 11 with the above subsetwise construction. This results in a hierarchical $(1+\varepsilon', 2+\varepsilon')$-spanner $\mathcal{H}$ having size $O\left(\frac{n\sqrt{n}\log^2 n}{\sqrt{\varepsilon}}\right)$.

To see that $\mathcal{H}$ is a hierarchical $(1+\varepsilon, 2)$-spanner of $\mathcal{G}$, consider any distinct pair of vertices $s, t \in V(\mathcal{G})$, and any $\ell \in 1, \ldots, \ell$. It holds that:

$$\begin{aligned} d(s,t,H_\ell) &\le (1+\varepsilon')d(s,t,G_\ell) + 2 + \varepsilon' = (1+\varepsilon)d(s,t,G_\ell) + 2 + \varepsilon' - \varepsilon' d(s,t,G_\ell) \\ &\le (1+\varepsilon)d(s,t,G_\ell) + 2. \end{aligned}$$ ◀

6 Lower Bounds

This section is devoted to proving lower bounds on the size of hierarchical spanners. We start with a simple construction, depicted in Figure 2, which yields the lower bound stated in the following theorem.

▶ **Theorem 13.** *Given n and k with $1 \le k \le n$, there exists a hierarchical graph $\mathcal{G}$ with $\Theta(n)$ vertices, a vertex $s \in V(\mathcal{G})$, and subset T of k vertices from $V(\mathcal{G})$, such that every $s \times T$ hierarchical preserver of $\mathcal{G}$ has size $\Omega(n|T|)$.*

Proof. The vertices of $\mathcal{G}$ are those in the union of three disjoint sets $\{v_1, \ldots, v_n\}$, $U = \{u_1, \ldots, u_n\}$, and $T = \{t_1, \ldots, t_k\}$. We choose $v_n = s$. Let π be the path from v_1 to v_n traversing all vertices v_i in order, and let $B = \{(u, t) \mid u \in U, t \in T\}$ be the edges of the complete bipartite graph with sides U and T. The depth of G is $L = n$. We define $E_1 = E(\pi) \cup \{(v_1, u_1)\} \cup B$, and $E_\ell = \{(v_\ell, u_\ell)\}$ for $\ell = 2, \ldots, L$. This concludes the description of our graph $\mathcal{G}$. See Figure 2 for an example.

The claim follows by observing that any $\{s\} \times T$ hierarchical preserver $\mathcal{H}$ of $\mathcal{G}$ must include all $n|T|$ edges in B. Indeed, for any choice of $t \in T$ and $\ell = 1, \ldots, L$, the unique shortest path from s to t in G_ℓ traverses the sequence of vertices $s = v_n, v_{n-1}, \ldots, v_{\ell+1}, v_\ell, u_\ell, t$, in this order, implying that $(u_\ell, t) \in E(H_\ell)$. ◀

Next, we prove a lower bound of $\Omega\left(n\sqrt{\frac{n}{\beta+1}}\right)$ on the size of any single-pair hierarchical β-additive spanner. The special case $\beta = 0$ provides a lower bound of $\Omega(n\sqrt{n})$ on the size of single-pair hierarchical preserves.

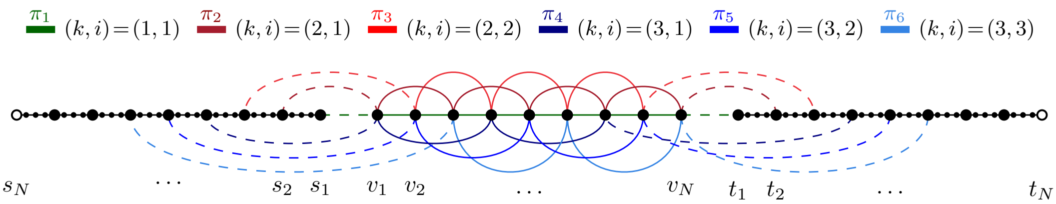

Figure 3 The hierarchical graph $\mathcal{X}$ for $N = 9$. Unnamed vertices are depicted with smaller circles. The paths π_ℓ for different values of ℓ are shown with different colors. For each path π_ℓ, its first and last edges are dashed, while the remaining edges are solid.

We start by defining an auxiliary hierarchical graph $\mathcal{X} = (V(\mathcal{X}), F_1, \dots, F_L)$ which is parameterized by an integer $N \geq 1$ and has the following properties: $\mathcal{X}$ has $9N - 6$ vertices, a depth of $L = \Theta(N)$, two distinguished vertices s_N, t_N, and contains L pairwise edge-disjoint paths $\pi_1, \dots, \pi_L$, each of length at least $\lfloor\sqrt{N}\rfloor$, such that π_ℓ is a subpath of all shortest paths from s' to t' in X_ℓ.

We now describe how $\mathcal{X}$ is built (see Figure 3 for an example with $N = 9$). The set of vertices $V(\mathcal{X})$ contains N vertices $s_1, \dots, s_N$, N vertices $t_1, \dots, t_N$, N vertices $v_1, \dots, v_N$ plus $6N - 6$ additional unnamed vertices. Let π' a path from s_1 to s_N obtained by choosing s_i as its $(1 + 4(i - 1))$-th vertex and traversing 3 unnamed vertices between s_i and s_{i+1}. Let π'' be the edges of a similarly defined path from t_1 to t_N, which is vertex-disjoint from π'.

Consider all integers $k \in \{1, \dots, \lfloor\sqrt{N}\rfloor\}$, which we refer to as *jump lengths*. For each jump length, we further consider all integers *offsets* $i \in \{1, \dots, k\}$.

We choose $L = \frac{1}{2}\lfloor\sqrt{N}\rfloor \cdot (\lfloor\sqrt{N}\rfloor + 1)$. The path π_ℓ has s_ℓ and t_ℓ as its endvertices and is defined by considering the unique pair (k, i) of a jump length k and a corresponding offset i such that $\ell = \frac{k(k-1)}{2} + i$. The internal vertices of π_ℓ are all v_{i+hk} for $h = 0, 1, \dots, \lfloor\frac{N-i}{k}\rfloor$, in increasing order of their index. Observe that $|\pi_\ell| = \lfloor\frac{N-i}{k}\rfloor + 2$. We choose $F_1 = E(\pi') \cup E(\pi'') \cup E(\pi_1)$, and $F_\ell = E(\pi_\ell)$ for $\ell = 2, \dots, L$.

The following lemma immediately implies that all edges in the paths π_ℓ must be contained in any single-pair hierarchical preserver of $\mathcal{X}$ w.r.t. the pair (s_N, t_N).

▶ **Lemma 14.** *For $\ell = 1, \dots, L$, π_ℓ is a subpath of all shortest paths from s_N to t_N in X_ℓ. Moreover, all paths π_ℓ are pairwise edge-disjoint and have a length of at least $\sqrt{N}$.*

Proof. Since the π_ℓ are pairwise edge-disjoint by construction, and their length is lower bounded by $\left\lfloor\frac{N-\lfloor\sqrt{N}\rfloor}{\lfloor\sqrt{N}\rfloor}\right\rfloor + 2 \geq \sqrt{N}$, we fix an $\ell \in \{1, \dots, L\}$ and show π_ℓ is in all shortest paths from s_N to t_N in X_ℓ. Let (k, i) be the pair of a jump length k and a corresponding offset i such that $\ell = \frac{k(k-1)}{2} + i$.

Let $\widetilde{\pi}$ be any shortest path from s_N to t_N in X_ℓ. Let s_j the last vertex from $\{s_1, \dots, s_N\}$ encountered when traversing $\widetilde{\pi}$ from s_N to t_N. Similarly, let t_q be the first vertex from $\{t_1, \dots, t_N\}$ encountered when traversing $\widetilde{\pi}[s_j : t_N]$ from s_j to t_N.

By the choice of s_j and t_q, all internal vertices of $\widetilde{\pi}[s_j : t_q]$ belong to $\{v_1, \dots, v_N\}$. Moreover, all edges (v_a, v_b) with both endpoints in $\{v_1, \dots, v_N\}$ in X_ℓ satisfy $|a - b| \leq k$.

We fist handle the case in which $s_j = s_\ell$ and $t_q = t_\ell$. This implies that the first and last edges of $\widetilde{\pi}[s_\ell : t_\ell]$ coincide with the first and last edges of π_ℓ, respectively, i.e., they are (s_ℓ, v_i) and (v_{i+hk}, t_ℓ) for $h = \lfloor\frac{N-i}{k}\rfloor$. Therefore, $|\widetilde{\pi}[s_\ell, t_\ell]| = 2 + |\widetilde{\pi}[v_i : v_{i+hk}]| \geq 2 + \frac{i+hk-i}{k} = 2 + h = |\pi_\ell|$ and equality is only attainable if all edges (v_a, v_b) of $\widetilde{\pi}[v_i : v_{i+hk}]$ have $|a - b| = k$. Hence, $\widetilde{\pi}[s_\ell : t_\ell]$ coincides with π_ℓ.

We now argue that the complementary case cannot happen. Assume towards a contradiction that at least of the following two conditions holds: (i) $s_j \neq s_\ell$ or (ii) $t_q \neq t_\ell$.

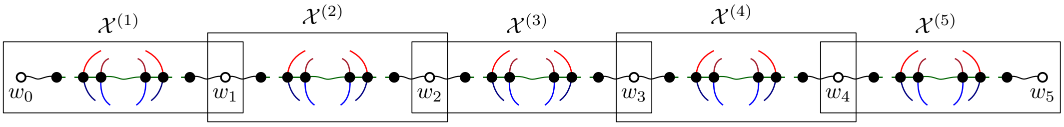

Figure 4 A qualitative representation of the graph $\mathcal{G}$ obtained in our lower-bound construction for single-pair hierarchical β-additive spanners when $\beta = 2$.

Let $\pi^* = \pi'[s_N : s_\ell] \circ \pi_\ell \circ \pi''[t_\ell : t_N]$ and observe that this is a valid path from s_N to t_N in X_ℓ. We have $|\pi^*| = 4(N-\ell) + \lfloor \frac{N-i}{k} \rfloor + 2 + 4(N-\ell) \le 8(N-\ell) + \frac{N}{k} + 2$.

By construction, we must have $j < \ell$ or $q < \ell$ (possibly both) and $|\widetilde{\pi}| = |\widetilde{\pi}[s_N : s_j]| + |\widetilde{\pi}[s_j : t_q]| + |\widetilde{\pi}[t_q : t_N]| \ge 4(N-\ell) + |\widetilde{\pi}[s_j : t_q]| + 4(N-\ell) + 4 = 8(N-\ell) + 4 + |\widetilde{\pi}[s_j : t_q]|$. Moreover, the first (resp. last) internal vertex of $\widetilde{\pi}[s_j : t_q]$ can only be in $\{v_1, \dots, v_k\}$ (resp. $\{v_{N-k+1}, \dots, v_N\}$). It follows that $|\widetilde{\pi}[s_j : t_q]| \ge 1 + \lceil \frac{N-k+1-k}{k} \rceil + 1 \ge \frac{N}{k}$. Hence $|\widetilde{\pi}| \ge 8(N-\ell) + 4 + \frac{N}{k} \ge |\pi^*| + 2$, providing the sought contradiction. ◀

By concatenating multiple copies of $\mathcal{X}$, as shown in Figure 4, to obtain our lower bound on the size of single-pair hierarchical β-additive spanners.

▶ **Theorem 15.** *Given n, and $\beta \ge 0$, there are hierarchical graphs $\mathcal{G}$ of $\Theta(n)$ vertices, such that every single-pair hierarchical β-additive spanner of $\mathcal{G}$ (w.r.t. a suitable pair of vertices) has size $\Omega\left(n\sqrt{\frac{n}{\beta+1}}\right)$.*

Proof. Since a trivial lower bound of $\Omega(n)$ holds regardless of the values of β, in the rest of the proof we assume $\beta = o(n)$,

The lower bound hierarchical graph $\mathcal{G}$ is constructed by starting with $2\beta+1$ copies $\mathcal{X}^{(1)}, \mathcal{X}^{(2)}, \dots \mathcal{X}^{(2\beta+1)}$ of the auxiliary graph $\mathcal{X}$ with $N = \left\lceil \frac{n}{2\beta+1} \right\rceil$. Then, for $i = 1, \dots, \beta$, we identifying vertex t_N of the i-th copy $\mathcal{X}^{(i)}$ with vertex s_N of $\mathcal{X}^{(i+1)}$, and we name the resulting vertex w_i. As corner cases, we let w_0 be vertex s_N in $\mathcal{X}^{(1)}$ and $w_{2\beta+1}$ be vertex t_N in $\mathcal{X}^{(2\beta+1)}$. The of vertices whose distances are to be preserved is $(w_0, w_{2\beta+1})$. See Figure 4 for a qualitative example with $\beta = 2$.

Let $\mathcal{H}$ be a single-pair hierarchical β-additive spanner of $\mathcal{G}$, and let $\mathcal{H}^{(i)}$ be the hierarchical subgraph of $\mathcal{H}$ containing all vertices of $V(\mathcal{X}^{(i)})$ and all edges in $E(\mathcal{H}) \cap E(\mathcal{X}^{(i)})$. We argue that, for any $\ell \in \{1, \dots, L\}$, at least $\beta + 1$ of the $2\beta+1$ copies of $\mathcal{X}$ in $\mathcal{G}$ must contain all the edges of (the corresponding copy of) path π_ℓ.

Since each w_i with $1 \le i \le 2\beta$ is a cut-vertex of H_ℓ that separates the vertices in $\{w_0, w_1, \dots, w_{w-1}\}$ from those in $\{w_{i+1}, w_{i+2}, \dots, w_{2\beta+1}\}$, we have that any shortest path from w_0 to $w_{2\beta+1}$ in G_ℓ must traverse $w_1, w_2, \dots, w_\beta$ in this order. It follows that the same holds in H_ℓ, and we have $d(w_0, w_{2\beta+1}, H_\ell) = \sum_{i=1}^{2\beta+1} d(w_{i-1}, w_i, H_\ell) = \sum_{i=1}^{2\beta+1} d(w_{i-1}, w_i, H_\ell^{(i)})$. By Lemma 14, the edges of π_ℓ belong to all shortest paths from s_N to t_N in X_ℓ, hence if $H_\ell^{(i)}$ does not contain all (the copies of the) edges of π_ℓ, we must have $d(w_{i-1}, w_i, H_\ell^{(i)}) \ge d(w_{i-1}, w_i, X_\ell^{(i)}) + 1$. This implies that the edges of at most β of the $2\beta+1$ copies of π_ℓ are not entirely contained in $\mathcal{H}$.

Since there are at least $\beta+1$ (edge disjoint) copies of each path π_ℓ, each of which has length at least $\sqrt{N} = \Omega\left(\sqrt{\frac{n}{\beta+1}}\right)$, and all such paths are pairwise edge-disjoint, the size of $\mathcal{H}$ is $\Omega\left(L(\beta+1)\sqrt{\frac{n}{\beta+1}}\right) = \Omega\left(n\sqrt{\frac{n}{\beta+1}}\right)$, where we used $L = \Theta(N) = \Theta(\frac{n}{\beta+1})$. ◀

7 Conclusions and open problems

In this paper, we initiated the study of graph spanners for hierarchical graphs, a framework that generalizes the standard non-hierarchical setting. Our results reveal a complex relationship between hierarchical and non-hierarchical spanners. On the one hand, for the all-pairs case and multiplicative distortions, the established upper bounds for standard graphs can be effectively extended to the hierarchical setting. On the other hand, the situation changes significantly for additive distortions, where hierarchical spanners are subject to a lower bound of $\Omega(n\sqrt{n})$ even in the single-pair case. Many of the bounds provided in this work are not tight and deserve a further investigation. Two concrete open problems that we deem particularly significant are (i) devising a lower bound of $\Omega(n^{\frac{3}{2}+\varepsilon})$, for some constant $\varepsilon > 0$, on the size of all-pairs hierarchical 2-additive spanners, and (ii) pinpointing (asymptotically) tight bounds for hierarchical preservers for the single-source multi-destination ($s \times T$) case.

References

1 Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020. doi:10.1016/j.cosrev.2020.100253.

2 Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28(4):1167–1181, 1999. doi:10.1137/S0097539796303421.

3 Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete Comput. Geom.*, 9(1):81–100, December 1993. doi:10.1007/BF02189308.

4 Joyce Bacic, Saeed Mehrabi, and Michiel Smid. Shortest beer path queries in outerplanar graphs. *Algorithmica*, 85(6):1679–1705, 2023. doi:10.1007/S00453-022-01045-4.

5 Davide Bilò, Gianlorenzo D'Angelo, Luciano Gualà, Stefano Leucci, and Mirko Rossi. Sparse temporal spanners with low stretch. In Shiri Chechik, Gonzalo Navarro, Eva Rotenberg, and Grzegorz Herman, editors, *30th Annual European Symposium on Algorithms, ESA 2022, Berlin/Potsdam, Germany, September 5-9, 2022*, LIPIcs, pages 19:1–19:16. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2022. doi:10.4230/LIPIcs.ESA.2022.19.

6 Davide Bilò, Gianlorenzo D'Angelo, Luciano Gualà, Stefano Leucci, and Mirko Rossi. Blackout-tolerant temporal spanners. *J. Comput. Syst. Sci.*, 141:103495, 2024. doi:10.1016/J.JCSS.2023.103495.

7 Davide Bilò, Luciano Gualà, Stefano Leucci, and Alessandro Straziota. Graph spanners for group steiner distances. In Timothy M. Chan, Johannes Fischer, John Iacono, and Grzegorz Herman, editors, *32nd Annual European Symposium on Algorithms, ESA 2024, Royal Holloway, London, United Kingdom, September 2-4, 2024*, LIPIcs, pages 25:1–25:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2024. doi:10.4230/LIPIcs.ESA.2024.25.

8 Greg Bodwin, Michael Dinitz, Ama Koranteng, and Lily Wang. Light edge fault tolerant graph spanners. In Keren Censor-Hillel, Fabrizio Grandoni, Joël Ouaknine, and Gabriele Puppis, editors, *52nd International Colloquium on Automata, Languages, and Programming, ICALP 2025, Aarhus, Denmark, July 8-11, 2025*, LIPIcs, pages 32:1–32:20. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.ICALP.2025.32.

9 Greg Bodwin, Bernhard Haeupler, and Merav Parter. Fault-tolerant spanners against bounded-degree edge failures: Linearly more faults, almost for free. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms, SODA 2024, Alexandria, VA, USA, January 7-10, 2024*, pages 2609–2642. SIAM, 2024. doi:10.1137/1.9781611977912.93.

10 Greg Bodwin and Gary Hoppenworth. New additive spanner lower bounds by an unlayered obstacle product. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 778–788. IEEE, 2022. doi:10.1109/FOCS54457.2022.00079.

11 Daniele Carnevale, Arnaud Casteigts, and Timothée Corsini. Dismountability in temporal cliques revisited. In Kitty Meeks and Christian Scheideler, editors, *4th Symposium on Algorithmic Foundations of Dynamic Networks, SAND 2025, Liverpool, UK, June 9-11, 2025*, LIPIcs, pages 6:1–6:18. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2025. doi:10.4230/LIPIcs.SAND.2025.6.

12 Arnaud Casteigts, Joseph G. Peters, and Jason Schoeters. Temporal cliques admit sparse spanners. In Christel Baier, Ioannis Chatzigiannakis, Paola Flocchini, and Stefano Leonardi, editors, *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019, Patras, Greece, July 9-12, 2019*, LIPIcs, pages 134:1–134:14. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019. doi:10.4230/LIPIcs.ICALP.2019.134.

13 Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discret. Math.*, 20(2):463–501, 2006. doi:10.1137/050630696.

14 Paul Erdös. On some extremal problems in graph theory. *Israel Journal of Mathematics*, 3:113–116, 1965.

15 Mathias Bæk Tejs Knudsen. Additive spanners and distance oracles in quadratic time. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, Warsaw, Poland, July 10-14, 2017*, LIPIcs, pages 64:1–64:12. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2017. doi:10.4230/LIPIcs.ICALP.2017.64.

16 You Peng, Zhuo Ma, Wenjie Zhang, Xuemin Lin, Ying Zhang, and Xiaoshuang Chen. Efficiently answering quality constrained shortest distance queries in large graphs. In *2023 IEEE 39th International Conference on Data Engineering (ICDE)*, pages 856–868, 2023. doi:10.1109/ICDE55515.2023.00071.

17 Zihan Tan and Tianyi Zhang. Almost-optimal sublinear additive spanners. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 281–294. ACM, 2023. doi:10.1145/3564246.3585125.

18 Jacques Tits. Sur la trialité et certains groupes qui s'en déduisent. *Publications Mathématiques de l'Institut des Hautes Études Scientifiques*, 2:14–60, 1959.

19 R. Wenger. Extremal graphs with no c4,s, or c10,s. *J. Comb. Theory Ser. B*, 52(1):113–116, May 1991. doi:10.1016/0095-8956(91)90097-4.

A An efficient implementation of the algorithm of Section 3

A naïve implementation of Algorithm 1 requires time $O(|P|Lm)$. To implement Algorithm 1 efficiently, we observe that each pair $(s,t) \in P$ can be handled independently. Notice that once one knows all *critical* values $\ell \in \{1,\dots,L\}$ for which the shortest path π_ℓ from s to t is added to Π, for all pairs (s,t), it is easy to build $\mathcal{H}$ in time $O(|\Pi|m)$. Moreover, we can find all pairs for which at least one critical value exists in time $O(m\min\{n,|P|\})$ by computing their pairwise distances in G_L.

We now fix a pair $(s,t) \in P$ with at least one critical value, and we show how to find all critical values for such pair. Define $\phi(\ell) = d(s,t,G_\ell)$, and observe that ϕ is monotonically non-increasing (where some values might be $+\infty$). A binary search can be used to find the smallest index ℓ for which $\phi(\ell) \le D$, which is exactly the smallest critical value. Moreover, given any critical value ℓ, a further binary search can be used to find the next smallest critical value $\ell' > \ell$, or to detect that no such value exists. Since each binary search requires time $O(m\log L)$ and causes the addition of a path in Π, the overall time spent to handle all pairs in P is $O(m|\Pi|\log L + m\min\{n,|P|\}) = O(m|P|M\log L)$.

For hierarchical α-spanners with constant $\alpha > 1$, the above time complexity becomes $O(m|P|\log^2 n)$, i.e., it is almost-linear whenever $|P| = O(1)$, as in the single-pair case.

A faster implementation can be provided for preservers. Indeed, for a fixed pair (s,t), the largest critical value ℓ can be found by solving the *minimum bottleneck shortest path* (MBSP) problem from s to t in the (unweighted) graph G_ℓ, where each edge e is associated to a cost c corresponding to the set E_c containing e.[4] Similarly, given a critical value $\ell > 2$, one can find the previous largest critical value $\ell' < \ell$ by solving the MBSP problem on $G_{\ell-1}$. This can be repeated until the length of the found shortest path exceeds D or the last found critical value is 1. Since MBSP can be solved in time $O(m)$ by a simple modification of the breadth-first-search algorithm, we have that the overall running time to handle all pairs in P is $O(|\Pi|m + m \min\{n, |P|\}) = O(m|P|M)$. In the case of single-pair preservers, the above time complexity becomes $O(mn)$.

[4] Given an unweighted graph in which each edge has an assigned cost, the bottleneck shortest path from s to t is a shortest path (w.r.t. hop-lengths) that minimizes the maximum cost of a traversed edge.