

New Separations and Reductions for Directed Preservers and Hopsets

Gary Hoppenworth*

Yinzhan Xu[†]

Zixuan Xu[‡]

Abstract

We study distance preservers, hopsets, and shortcut sets in n -node, m -edge directed graphs, and show improved bounds and new reductions for various settings of these problems. Our first set of results is about exact and approximate distance preservers. We give the following bounds on the size of directed distance preservers with p demand pairs:

- $\tilde{O}(n^{5/6}p^{2/3} + n)$ edges for exact distance preservers in unweighted graphs, and
- $\Omega(n^{2/3}p^{2/3})$ edges for approximate distance preservers with any given finite stretch, in graphs with arbitrary aspect ratio.

Additionally, we establish a new directed-to-undirected reduction for exact distance preservers. We show that if undirected distance preservers have size $O(n^\lambda p^\mu + n)$ for constants $\lambda, \mu > 0$, then directed distance preservers have size $O\left(n^{\frac{1}{2-\lambda}} p^{\frac{1+\mu-\lambda}{2-\lambda}} + n^{1/2}p + n\right)$. As a consequence of the reduction, if state-of-the-art upper bounds for undirected preservers can be improved for some $p \leq n$, then so can state-of-the-art upper bounds for directed preservers.

Our second set of results is about directed hopsets and shortcut sets. For hopsets in directed graphs, we prove that the hopbound is:

- $\Omega(n^{2/9})$ for $O(m)$ -size shortcut sets, improving the previous $\Omega(n^{1/5})$ bound [Vassilevska Williams, Xu and Xu, SODA 2024],
- $\Omega(n^{2/7})$ for $O(m)$ -size exact hopsets in unweighted graphs, improving the previous $\Omega(n^{1/4})$ bound [Bodwin and Hoppenworth, FOCS 2023], and
- $\Omega(n^{1/2})$ for $O(n)$ -size approximate hopsets with any given finite stretch, in graphs with arbitrary aspect ratio. This result establishes a separation between this setting and $O(n)$ -size approximate hopsets for graphs with polynomial aspect ratio, which are known to have an $\tilde{O}(n^{1/3})$ upper bound [Bernstein and Wein, SODA 2023].

*University of Michigan. garytho@umich.edu. Supported by NSF:AF 2153680.

[†]Massachusetts Institute of Technology. xyzhan@mit.edu. Supported by NSF Grant CCF-2330048, a Simons Investigator Award and HDR TRIPODS Phase II grant 2217058.

[‡]Massachusetts Institute of Technology. zixuanxu@mit.edu.

Contents

1	Introduction	1
1.1	Preservers	1
1.2	Hopsets	4
2	Technical Overview	7
2.1	Exact Preserver Reductions and Upper Bounds	7
2.2	Approximate Hopset and Preserver Lower Bounds	9
2.3	$O(m)$ -size Shortcut Set and Hopset Lower Bounds	11
3	Preliminaries	12
4	Directed to Undirected Preserver Reduction	13
4.1	Directed to DAG Reduction	14
4.2	DAG to Undirected Reduction	18
4.3	Finishing the Reduction	20
5	Exact Distance Preserver Upper Bound	22
5.1	Preliminary Definitions	22
5.2	Ingredient 1: the Dense Low-Diameter Cluster Lemma	23
5.3	Ingredient 2: Sourcewise Distance Preservers	25
5.4	Main Construction	26
5.5	Proving the Dense Low-Diameter Cluster Lemma	28
6	Directed Approximate Preserver Lower Bound	33
7	Directed Approximate Hopset Lower Bound	34
8	$O(m)$-size Shortcut Set Lower Bound	38
9	$O(m)$-size Exact Hopset Lower Bound in Unweighted Graphs	40
A	Proof of Assumptions 3 and 4 from Section 4	47
B	Proof of Lemma 5.7 from Section 5	48

1 Introduction

In network design, two commonly studied categories of problems are graph augmentation problems, where the goal is to augment the graph so that it gains some useful property, and graph sparsification problems, where the goal is to sparsify the graph while preserving some of its features. In this work, we study two fundamental problems from these two categories: hopsets and distance preservers.

The hopset problem is a typical graph augmentation problem. We are given a graph¹, and we want to add a small number of additional edges to the graph so that all pairs of vertices (s, t) have an approximate shortest $s \rightsquigarrow t$ path with few edges. Another popular graph augmentation object, the shortcut set (first introduced by [Tho92]), is actually a special case of hopsets where we only care about reachability. Hopsets and shortcut sets have a wide range of applications in areas such as parallel computing and distributed computing (e.g. [KS97, HKN14, EN16, FN18, JLS19, KS21, EN19, Fin19, ASZ20, BGWN20, CFR20, EM21]).

On the other hand, the (approximate) distance preserver problem is a graph sparsification problem. Here, we are given an input graph G and a set of demand pairs $P \subseteq V(G) \times V(G)$, and we need to find a sparse subgraph H of G that preserves all distances between demand pairs in G up to an approximation factor α . Approximate distance preservers are a generalization of another well-studied graph sparsification object, reachability preservers. Distance preservers have applications to problems such as distance oracles [EP16, ES23] and spanners [BCE05, BV15, AB16, AB17, ABP17, EFN17, GK17, BV21].

Despite being graph augmentation and graph sparsification problems respectively, hopsets and distance preservers are known to be related in some ways. For instance, distance preserver lower bounds imply lower bounds for hopsets [KP22a, BHT23], and small-stretch hopsets can be translated into small-stretch distance preservers in a black-box way [KP22a].

In this paper, we develop better understandings of hopsets and distance preservers in directed graphs. We discuss our results and compare them with prior work in Sections 1.1 and 1.2.

1.1 Preservers

The problem of constructing sparse distance preservers is an important basic combinatorial problem with connections to incidence geometry [CE06, KSSY16], distance oracles [EP16, ES23], and fast graph algorithms [Alo02, CGMW18, ES23]. We now formally define the more general object of approximate distance preservers, or approximate preservers for short.

Definition 1.1 (Approximate Preservers). Given an n -node graph G and a set of demand pairs $P \subseteq V(G) \times V(G)$ of size $|P| = p$, a subgraph H of G is an α -approximate distance preserver of G, P if

$$\text{dist}_H(s, t) \leq \alpha \cdot \text{dist}_G(s, t)$$

for all $(s, t) \in P$.

There are three settings of approximate preservers that have received special attention in prior work:

- $\alpha = 1$ (also known as *Exact Distance Preservers*.)
- $\alpha = 1 + \varepsilon$ (also known as $(1 + \varepsilon)$ -*Approximate Preservers*.)
- $\alpha = \infty$ (also known as *Reachability Preservers*.)

¹We assume graphs to be directed unless otherwise specified.

Setting	Upper Bound	Lower Bound	Beats Consistency?
$\alpha = 1$ undirected	$O(\min\{n^{1/2}p + n, np^{1/2}\})$ [CE06]	$\Omega(n^{2/3}p^{2/3} + n)$ [CE06]	X
$\alpha = 1$ undirected unweighted	$O\left(\min\left\{n^{2/3}p^{2/3} + np^{1/3}, \frac{n^2}{\text{RS}(n)}\right\}\right)$ [BV21, Bod17]	$\Omega(n^{2/(d+1)}p^{(d-1)/d})$ for integer $d \geq 1$ [CE06]	✓
$\alpha = 1$ directed	$O(\min\{n^{2/3}p + n, np^{1/2}\})$ [CE06, Bod21]	$\Omega(n^{2/3}p^{2/3} + n)$ [CE06]	X
$\alpha = 1 + \varepsilon$ undirected	$\tilde{O}_\varepsilon((n + p) \cdot n^{o(1)})$ [KP22a, ES23]	$\Omega(n + p)$	✓
$\alpha = 1 + \varepsilon$ directed	$\tilde{O}_\varepsilon(n^{2/3}p^{2/3} + np^{1/3})$ [KP22a, BW23]	$\Omega(n^{2/(d+1)}p^{(d-1)/d})$ for integer $d \geq 1$ [CE06]	✓
$\alpha = \infty$ directed	$O(n^{3/4}p^{1/2} + n^{5/8}p^{11/16} + n)$ [CE06, BHT23]	$\Omega(n^{2/(d+1)}p^{(d-1)/d})$ for integer $d \geq 1$ [CE06]	✓

Table 1: Some previous bounds for various settings of approximate preservers on n -node graphs with p demand pairs. We say existing upper bounds for approximate preservers beat consistency if they imply a separation with the bounds for consistent tiebreaking schemes for some $p \in [1, n^2]$. In the second row, $\text{RS}(n)$ denotes the Rusza-Szemerédi function from extremal graph theory.

Exact distance preservers were first formally introduced in the seminal work of Coppersmith and Elkin [CE06]. In their work, Coppersmith and Elkin proved upper bounds and lower bounds for exact distance preservers in several different settings. As can be seen in Table 1, many of the bounds proved in [CE06] are still state-of-the-art.

A fundamental tool used in the upper bounds of [CE06] is *consistent tiebreaking schemes*. Given an n -node graph G , a tiebreaking scheme roughly corresponds to a collection of shortest paths Π in G . Informally, we say that tiebreaking scheme Π is *consistent* if no two shortest paths $\pi_1, \pi_2 \in \Pi$ intersect, split apart, and then intersect later again (see Definition 3.2 for a formal definition). Upper bounds for the number of distinct edges in consistent tiebreaking schemes imply upper bounds for exact distance preservers, since we may assume without loss of generality that shortest paths in the input graph are consistent. A line of work initiated by [CE06] has established tight bounds on the size of consistent tiebreaking schemes.

Theorem 1.2 (Consistent Tiebreaking Schemes). *Consistent tiebreaking schemes on n nodes and p paths have*

- $\Theta(\min(n^{1/2}p + n, np^{1/2}))$ edges in undirected graphs [CE06, BV21], and
- $\Theta(\min(n^{2/3}p + n, np^{1/2}))$ edges in directed graphs [CE06, Bod17, BHT23].

As we mentioned earlier, upper bounds on the number of edges in consistent tiebreaking schemes yield upper bounds for exact distance preservers, which in turn imply upper bounds for α -approximate preservers for all $\alpha > 1$. Recently, there has been a tremendous amount of progress in constructing approximate preservers with sparsity *better* than consistent tiebreaking schemes. We summarize this progress below.

- Exact distance preservers of size $O(n^{2/3}p^{2/3} + np^{1/3})$ edges in undirected, unweighted graphs were constructed in [BV21].
- Almost-optimal $(1 + \varepsilon)$ -preservers in undirected graphs were constructed in [KP22a, ES23]. Approximate preservers of size $\tilde{O}_\varepsilon(n^{2/3}p^{2/3} + np^{1/3})$ edges for directed graphs were constructed in [KP22a, BW23].
- Reachability preservers of size $O(n^{3/4}p^{1/2} + n^{2-\sqrt{2}+o(1)}p^{1/\sqrt{2}} + n)$ edges were constructed in [BHT23, BL24], and *online* reachability preservers of size $O(n^{0.72}p^{0.56} + n^{0.6}p^{0.7} + n)$ were constructed in [BL24].

All of these preserver constructions are significantly sparser than consistent tiebreaking schemes in some regime of p . However, as we can see in Table 1, the state-of-the-art bounds for exact distance preservers in weighted (un)directed graphs do not achieve sparsity better than consistent tiebreaking. This motivates the following question.

Open Question 1. *Is there a separation between exact distance preservers and consistent tiebreaking schemes?*

In addition to asking when we can separate approximate preservers from consistent tiebreaking schemes, we can also ask whether we can separate different settings of α -approximate preservers from each other. For example, the $\Omega(n^{2/3}p^{2/3})$ lower bound for exact distance preservers in weighted graphs due to [CE06], along with the recent progress on reachability preserver upper bounds due to [BHT23], imply a separation between α -approximate preservers when $\alpha = 1$ and α -approximate preservers when $\alpha = \infty$. However, the following questions remained:

Open Question 2. *Is there a separation between approximate preservers with finite stretch and approximate preservers with infinite stretch (reachability preservers)? Is there a separation between exact preservers and approximate preservers with finite stretch?*

Towards answering Open Question 1 and Open Question 2, we prove the following results for distance preservers:

Theorem 1.3 (Main Result 1). *For n -node graphs and p demand pairs, the size of the distance preserver is:*

1. (Theorem 5.1) $\tilde{O}(n^{5/6}p^{2/3} + n)$ edges, for exact distance preservers in unweighted graphs, and
2. (Theorem 6.1) $\Omega(n^{2/3}p^{2/3})$ edges, for any given finite stretch.

Remark. At first glance, the lower bound in Theorem 1.3 Item 2 appears to nearly tightly match the $\tilde{O}_\varepsilon(n^{2/3}p^{2/3} + np^{1/3})$ upper bound for $(1 + \varepsilon)$ -approximate preservers due to [KP22a, BW23]. However, the upper bounds of [KP22a, BW23] hold only for graphs with $2^{\tilde{O}(1)}$ aspect ratio, while our lower bound graph has $\exp(n)$ aspect ratio.

Theorem 1.3 has the following consequences:

- Theorem 1.3 Item 1 answers Open Question 1 in the affirmative in the natural graph class of unweighted directed graphs. Specifically, the upper bound of Theorem 1.3 Item 1 beats consistent tiebreaking in the regime of $n^{1/2} \leq p \leq n$.

- Theorem 1.3 Item 2 partially answers Open Question 2 in the affirmative for general graphs, establishing a separation between approximate preservers with finite stretch and reachability preservers. This can be verified by comparing the bound in Theorem 1.3 Item 2 to the upper bound for $\alpha = \infty$ in Table 1.

Additionally, we establish a new directed-to-undirected reduction for exact distance preservers. Let $\text{UDP}(n, p)$ denote the best upper bound for undirected distance preservers on n vertices and p demand pairs, and let $\text{DDP}(n, p)$ denotes the same for directed distance preservers.

Theorem 1.4 (Main Result 2). *If $\text{UDP}(n, p) = O(n^\lambda p^\mu + n)$ for constants $\lambda, \mu > 0$, then*

$$\text{DDP}(n, p) = O\left(n^{\frac{1}{2-\lambda}} p^{\frac{1+\mu-\lambda}{2-\lambda}} + n^{1/2} p + n\right).$$

This reduction has the property that if we plug the state-of-the-art $O(n^{1/2} p + n)$ upper bound for undirected distance preservers into the reduction, we recover the state-of-the-art $O(n^{2/3} p + n)$ upper bound for directed distance preservers. More generally, we obtain the following consequence of the reduction.

Corollary 1.5. *Consistent tiebreaking is optimal for directed distance preservers when $p \leq n^{2/3}$ only if consistent tiebreaking is optimal for undirected distance preservers when $p \leq n$.*

Our reduction makes significant progress on Open Question 1. If consistent tiebreaking is optimal for directed distance preservers for $p \leq n^{2/3}$, then consistent tiebreaking is optimal for undirected distance preservers for $p \leq n$. Conversely, if consistent tiebreaking is suboptimal in undirected graphs for some $p \leq n$, then consistent tiebreaking is suboptimal in directed graphs for some $p \leq n^{2/3}$.

Along the way to proving Theorem 1.4, we also obtain the following interesting algorithmic result (Corollary 4.7 in the main body), which does not seem to appear in the literature to the best of our knowledge, though it could be folklore.

Corollary 1.6. *All-pairs shortest paths (APSP) on n -node, m -edge directed acyclic graphs can be reduced to APSP on n -node, m -edge undirected graphs in $O(m)$ time.*

A related result is the asymptotic equivalence between all-pairs shortest paths in directed dense graphs and all-pairs shortest paths in undirected dense graphs as shown by Vassilevska Williams and Williams [VW18]; however, such an equivalence is not known for sparse graphs.

1.2 Hopsets

A fundamental problem in algorithm design is how to compute (approximate) shortest paths and reachability in directed graphs efficiently. In many algorithmic settings where we want to compute shortest paths, it is useful to assume that shortest paths in the input graph contain few edges. However, in general this assumption does not hold. Instead, we can try adding a small set of additional edges to the input graph and hope that these edges reduce the lengths of shortest paths. This is the idea behind *hopsets*, which we now formally define.

Definition 1.7 (Approximate Hopsets). Given a graph G , an α -approximate hopset with hopbound β is a set of additional edges H such that:

- Every edge $(u, v) \in H$ has weight $w(u, v) = \text{dist}_G(u, v)$.

- For every pair of nodes (s, t) in the transitive closure of G , there exists an $s \rightsquigarrow t$ path π in $G \cup H$ with at most $|\pi| \leq \beta$ edges and weight at most $w(\pi) \leq \alpha \cdot \text{dist}_G(s, t)$.

There are three settings of approximate hopsets that have received special attention in prior work:

- $\alpha = 1$ (also known as *Exact Hopsets*.)
- $\alpha = 1 + \varepsilon$ (also known as $(1 + \varepsilon)$ -*Approximate Hopsets*.)
- $\alpha = \infty$ (also known as *Shortcut Sets*.)

Setting	Size	Upper Bound	Lower Bound	Beats Folklore Sampling?
$\alpha = 1$	$O(n)$	$O(\sqrt{n})$ [UY91, BRR10]	$\tilde{\Omega}(\sqrt{n})$ [BH22]	$\times$
$\alpha = 1 + \varepsilon$	$O_\varepsilon(n)$	$\tilde{O}(n^{1/3})$ [BW23]	$\Omega(n^{1/4})$ [BH22, VXX24]	$\checkmark$
$\alpha = \infty$	$O(n)$	$\tilde{O}(n^{1/3})$ [KP22b]	$\Omega(n^{1/4})$ [BH22, VXX24]	$\checkmark$
$\alpha = \infty$	$O(m)$	$\tilde{O}((n^2/m)^{1/3})$ [KP22b]	$\Omega(n^{1/5})$ [VXX24]	$\checkmark$

Table 2: Some previous bounds on various settings of α -approximate hopsets in directed graphs. Note that the upper bound for $\alpha = 1 + \varepsilon$ holds specifically for graphs with aspect ratio $2^{\tilde{O}(1)}$.

Hopsets were first studied implicitly in the 90’s by Ullman and Yannakakis [UY91], and defined explicitly by Cohen [Coh00]. In particular, Ullman and Yannakakis gave an implicit construction for exact hopsets using a simple sampling scheme, which has since become folklore (and slightly improved [BRR10]).

Theorem 1.8 (Folklore Sampling [UY91, BRR10]). *Every n -node graph has an $O(n)$ -size hopset with hopbound $O(n^{1/2})$.*

While many algorithmic applications for hopsets have been found since the 90’s, no progress had been made on approximate hopset upper bounds in directed graphs until recently. In 2003, Hesse gave the first polynomial lower bounds for $O(n)$ -size and $O(m)$ -size shortcut sets, proving that indeed polynomial hopbound is necessary for hopsets [Hes03]. Nevertheless, it remained open whether the folklore sampling algorithm could be improved, even in the easiest setting of $\alpha = \infty$.

In a breakthrough result by Kogan and Parter [KP22b], it was shown that there exist $O(n)$ -size shortcut sets (∞ -approximate hopsets) with hopbound $\tilde{O}(n^{1/3})$. Kogan and Parter also proved that there exist $O(n)$ -size $(1 + \varepsilon)$ -approximate hopsets with hopbound $\tilde{O}_\varepsilon(n^{2/5})$. The hopbound for $(1 + \varepsilon)$ -approximate hopsets was subsequently improved to $\tilde{O}_\varepsilon(n^{1/3})$ by Bernstein and Wein [BW23], using a careful analysis of approximate shortest paths in directed graphs.²

The possibility of extending these improvements to exact hopsets was ruled out by [BH22], where it was proved that $O(n)$ -size hopsets have hopbound $\tilde{\Omega}(n^{1/2})$ in the worst case. This lower bound proved that folklore sampling is indeed optimal for exact hopsets. Additionally, this lower

² $\tilde{O}$ hides $\text{polylog}(n)$ factors.

bound, along with the upper bounds of [KP22b, BW23], established a separation between exact hopsets and approximate hopsets.

However, the following question remained:

Open Question 3. *Is there a separation between approximate hopsets with finite stretch and approximate hopsets with infinite stretch (shortcut sets)?*

One of the barriers to making progress on this question is that all known lower bounds for hopsets follow from existing lower bounds for shortcut sets (with the exception of [BH22, KP22a]). Indeed, while there has been a long series of work improving lower bounds for shortcut sets [Hes03, HP21, LVWX22, BH22, VXX24], there has been relatively little progress in proving lower bounds for hopsets that surpass existing lower bounds for shortcut sets. This barrier motivates the following question.

Open Question 4. *Can we prove new lower bounds for hopsets that surpass existing lower bounds for shortcut sets?*

Towards answering these open questions, we show a range of new lower bounds for various settings of hopsets, summarized in Theorem 1.9 below:

Theorem 1.9 (Main Result 3). *For hopsets in directed graphs, the hopbound is:*

1. (Theorem 7.1.) $\Omega(n^{1/2})$ for $O(n)$ -size approximate hopsets with any given finite stretch,
2. (Theorem 8.1.) $\Omega(n^{2/9})$ for $O(m)$ -size hopsets and stretch infinity (shortcut sets), and
3. (Theorem 9.1.) $\Omega(n^{2/7})$, for $O(m)$ -size exact hopsets, in unweighted graphs (this bound holds even for undirected graphs).

Remark. At first glance, the $\Omega(n^{1/2})$ lower bound in Item 1 of Theorem 1.9 for $O(n)$ -size approximate hopsets seems to contradict the $O(n^{1/3})$ upper bound for $\tilde{O}_\varepsilon(n)$ -size $(1+\varepsilon)$ -approximate hopsets due to [BW23]. However, the upper bounds of [BW23] hold only for graphs with $2^{\tilde{O}(1)}$ aspect ratio. Our lower bound graph has $\exp(n)$ aspect ratio.

Theorem 1.9 has the following consequences:

- Theorem 1.9 Item 1 answers Open Question 3 in the affirmative for general graphs. In particular, it shows that for graphs with exponential aspect ratio, the folklore sampling $O(\sqrt{n})$ bound of [UY91, BRR10] is optimal.
- Theorem 1.9 Item 1 also separates approximate hopsets in graphs with polynomial aspect ratio from approximate hopsets in graphs with exponential aspect ratio. This is highly unusual behavior in network design problems. For example, in undirected graphs the state-of-the-art bounds for multiplicative spanners [ADD⁺93], approximate distance preservers [KP22a], and hopsets [EN19, HP19, KP22a] have no dependency on the aspect ratio of the input graph.
- Theorem 1.9 Item 2 improves the previous $\Omega(n^{1/5})$ lower bound for $O(m)$ -size shortcut sets from [VXX24] to $\Omega(n^{2/9})$.
- Theorem 1.9 Item 3 makes progress on Open Question 4. The best known lower bound for $O(m)$ -size shortcut is $\Omega(n^{2/9})$, as shown in Item 2. Our bound in Item 3 beats the current best bound for $O(m)$ -size shortcut, establishing another example where the current

best lower bound for some setting of hopset is higher than that for shortcut set. In fact, our $\Omega(n^{2/7})$ lower bound is even higher than the $\Omega(n^{1/4})$ lower bound for $O(n)$ -size shortcut set [BH22, VXX24]. Improving existing lower bounds for $O(n)$ -size shortcut set beyond $\Omega(n^{1/4})$ is likely to be difficult due to known barriers for improving reachability preserver lower bounds [BHT23].

2 Technical Overview

In this section, we provide intuition and proof overviews for our results.

- In Section 2.1, we discuss several new insights about the (non)existence of cycles in systems of shortest paths in directed graphs, with applications to exact distance preservers and APSP. (This corresponds to Item 1 of Theorem 1.3 and Theorem 1.4.)
- In Section 2.2, we discuss the new weighting scheme we developed to prove new lower bounds for approximate hopsets and preservers (Item 1 of Theorem 1.9 and Item 2 of Theorem 1.3).
- In Section 2.3, we discuss our new lower bound constructions for $O(m)$ -size shortcut sets and $O(m)$ -size exact hopsets in unweighted graphs (Items 2 and 3 of Theorem 1.9).

2.1 Exact Preserver Reductions and Upper Bounds

In this section, we discuss the ideas behind our new reduction from directed distance preservers to undirected distance preservers, as well as our new upper bound for distance preservers in directed unweighted graphs. Both of these results stem from new insights about the existence of acyclic *and cyclic* structures in systems of shortest paths in directed graphs. These insights are summarized in Observation 2.1 and Lemma 2.2, respectively.

Directed to Undirected Reduction. We will briefly sketch the strategy we use to obtain an extremal reduction from directed to undirected distance preservers. Let G be an n -node directed graph, and let P be a set of $|P| = p$ demand pairs in G . Our reduction from directed distance preservers to undirected distance preservers happens in two steps:

1. First, we give an extremal reduction from directed distance preservers to distance preservers in DAGs.
2. Second, we give an extremal reduction from distance preservers in DAGs to distance preservers in undirected graphs.

The key idea behind our first reduction from directed distance preservers to distance preservers in DAGs is summarized in Observation 2.1 and visualized in Figure 1.

Observation 2.1 (DAG Structure (Informal)). *Let Π be a consistent collection of shortest paths in a directed graph G , and let π be a path in Π . Then the paths in $\Pi \setminus \{\pi\}$, restricted to the nodes in path π , form a DAG.*

In Observation 2.1, we define the restriction of a path $\pi' \in \Pi \setminus \pi$ to the nodes in path π as the path $\pi' \cap \pi$, with nodes ordered in the sequence they appear in π' . Then, as we can visualize in Figure 1, Observation 2.1 tells us that the paths in Π , restricted to the nodes in path π , form a DAG with a topological ordering that is the reverse of the order of nodes in π . We note that




Figure 1: This figure depicts an example of a path system $\Pi = \{\pi, \pi', \pi'', \pi'''\}$. The paths in $\Pi \setminus \{\pi\}$, restricted to the nodes in path π , form a DAG.

Observation 2.1 is not strictly true, unless paths in Π are pairwise edge-disjoint. However, edge-disjointness holds without loss of generality for shortest paths in distance preservers, as we discuss at the beginning of Section 4.

Here is an outline of how we can use Observation 2.1 to reduce extremal bounds for directed distance preservers to extremal bounds for distance preservers in DAGs:

- Let H be a minimal distance preserver of directed graph G and demand pairs P .
- If H is dense, then there exists a shortest $s \rightsquigarrow t$ path π in H , for some $(s, t) \in P$, such that π contains many high-degree nodes.
- Let D be the DAG in Observation 2.1 corresponding to the shortest paths in H induced on path π . Then since path π contains many high-degree nodes, DAG D is dense.
- We can apply extremal bounds for distance preservers in DAGs to obtain an upper bound on $|E(D)|$. This will in turn yield an upper bound on $|E(H)|$.

This completes the sketch of our reduction from directed distance preservers to distance preservers in DAGs; the full reduction is given in Section 4.1.

The second step of our reduction is an extremal reduction from distance preservers in DAGs to distance preservers in undirected graphs. This reduction is simple, but to our knowledge previously unknown. Here is a rough sketch of how we can convert shortest paths in a DAG to shortest paths in an undirected graph: Let $v_1, \dots, v_n$ be the topological order of the DAG, and let W be a sufficiently large number (say larger than the sum of all edge weights in the DAG). Then we add $W \cdot (j - i)$ to the weight of any edge (v_i, v_j) in the graph, and make the edge undirected.

A simple analysis of the undirected reweighted graph shows that shortest paths remain the same in this new graph. This reduction immediately implies that the extremal size of distance preservers in DAGs is at most the extremal size of distance preservers in undirected graphs. Additionally, this reduction implies a fast *algorithmic* reduction from APSP in DAGs to APSP in undirected graphs. We give a complete presentation of our reduction and results in Section 4.2.

Exact Preserver Upper Bound. Our exact distance preserver upper bound (Item 2 in Theorem 1.3) for directed unweighted graphs emerges from the following intuition. Let G be an n -node directed, unweighted graph, and let P be a set of $|P| = p$ demand pairs in G . If graph G is a DAG, then we can apply our extremal reduction from distance preservers in DAGs to distance preservers in undirected graphs. This implies that G, P has a distance preserver of size $O(\min\{n^{1/2}p + n, np^{1/2}\})$.

Then if graph G has $|E(G)| = \omega(n^{1/2}p + n)$ edges, we know that G must contain a directed cycle. More generally, if $|E(G)| \gg n^{1/2}p + n$, then G contains *many* short directed cycles. This fact provides the intuition behind the following (informal) lemma.

Lemma 2.2 (Dense Low-Diameter Cluster Lemma (cf. Lemma 5.2)). *Let H be a minimal distance preserver of directed, unweighted graph G and demand pairs P . If H has (sufficiently large) average degree d , then there exists a set $S \subseteq V(G)$ such that*

1. $|S| = \Theta(d)$,
2. $\deg_H(v) = \Omega(d)$ for all $v \in S$, and
3. Set S is strongly connected and has small (weak) diameter in H .

Lemma 2.2 can be understood as a directed generalization of the following standard lemma in undirected graphs:

Lemma 2.3 ([Die00], cf. [BV21]). *Let H be an undirected, unweighted graph with average degree d . Then H has a set S of $|S| = \Theta(d)$ nodes, where $\deg_H(v) = \Omega(d)$ for all $v \in S$ and S has strong diameter 2.*

Note that we cannot hope for a generalization of Lemma 2.3 to general directed graphs because directed graphs can be arbitrarily dense and acyclic (e.g., the directed biclique). We formalize Lemma 2.2 in Lemma 5.2 and we prove it in Section 5.5.

Once we prove the Dense Low-Diameter Cluster Lemma, our distance preserver upper bounds in directed, unweighted graphs follow from a straightforward simulation of the upper bound argument for distance preservers in undirected, unweighted graphs due to [BV21]. The key subroutine of the upper bound in [BV21] is a sourcewise distance preserver construction, using the set S in Lemma 2.3 as the set of source nodes. The authors of [BV21] prove that there exists a particularly sparse sourcewise distance preserver of demand pairs $P \subseteq S \times V(G)$, using a clever application of the pigeonhole principle to shortest path distances. In our upper bound we follow this framework, while using the Dense Low-Diameter Cluster Lemma instead of Lemma 2.3.

2.2 Approximate Hopset and Preserver Lower Bounds

Our lower bounds for approximate hopsets and preservers emerge from the following observation.

Observation 2.4 (Informal). *Existing lower bounds for exact preservers and hopsets in undirected graphs can be converted into lower bounds for α -approximate preservers and hopsets in directed graphs for any finite α .*

We convert lower bounds for exact distance problems in undirected graphs into lower bounds for approximate distance problems in directed graphs using a simple argument. Given an undirected lower bound graph, we do the following:

1. Direct the edges of the lower bound graph in a natural way.
2. Reweight the edges of the lower bound graph using the *Over-Under* weighting scheme.

We now illustrate the Over-Under weighting scheme, by sketching the proof of our directed approximate preserver lower bound. Our application of the Over-Under weighting scheme to directed approximate hopset lower bounds follows a similar argument.

The Over-Under Weighting Scheme for Approximate Preserver Lower Bounds. For exact distance preservers in undirected weighted graphs, the best known size lower bound is $\Omega(n^{2/3}p^{2/3})$ due to a construction by [CE06]. The lower bound graph of [CE06] corresponds to an arrangement of points and lines in the plane with many point-line incidences. Formally, there exists an arrangement of points $\mathcal{P} \subseteq \mathbb{R}^2$ and lines $\mathcal{L}$ in the plane such that:

1. $|\mathcal{P}| = n$ and $|\mathcal{L}| = p$,
2. There are $\Omega(n^{2/3}p^{2/3})$ distinct point-line incidences between points in $\mathcal{P}$ and lines in $\mathcal{L}$,
3. Any two distinct lines $L, L' \in \mathcal{L}$ intersect on at most one point, and
4. Every line $L \in \mathcal{L}$ is a unique shortest path in $\mathbb{R}^2$ under the Euclidean metric.

The set of points $\mathcal{P}$ and set of lines $\mathcal{L}$ can naturally be converted into an n -node graph G embedded in the Euclidean plane and a set of demand pairs P of size $|P| = p$ as follows:

- Let $V(G) = \mathcal{P}$,
- Add edge $(x, y) \in \mathcal{P} \times \mathcal{P}$ to $E(G)$ if points x and y occur consecutively on some line $L \in \mathcal{L}$,
- Assign edge (x, y) weight $w(x, y) = \|x - y\|$, where $\|\cdot\|$ denotes Euclidean distance, and
- Add demand pair $(s, t) \in \mathcal{P} \times \mathcal{P}$ to P if s and t are the first and last points in $\mathcal{P}$ on some line $L \in \mathcal{L}$, respectively.

Using Properties 1-4 of the arrangement of points $\mathcal{P}$ and lines $\mathcal{L}$, it becomes a simple exercise to prove that any exact preserver of G, P has at least $\Omega(n^{2/3}p^{2/3})$ edges, recovering the [CE06] lower bound. However, our goal is to reweight graph G and make it directed, so that any α -approximate preserver of G, P requires $\Omega(n^{2/3}p^{2/3})$ edges. Here is a rough sketch of how we do this:

1. For every undirected edge $(x, y) \in E(G)$, if point x has a smaller first coordinate than point y , then convert edge (x, y) into directed edge $x \rightarrow y$.
2. Let $\mathcal{L} = \{L_1, \dots, L_p\}$ be the lines in $\mathcal{L}$ written by increasing slope. We assign the edges in $E(G)$ corresponding to line L_i weight $(\alpha \cdot n)^i$.

We now give some intuition for why this yields a lower bound against α -approximate preservers. Consider any line $L_i \in \mathcal{L}$, and let $x_1, \dots, x_k$ be the points in $\mathcal{P}$ on L_i ordered by increasing first coordinate. Then $(x_1, x_k) \in P$ is a demand pair in P . We claim that path $\pi_{L_i} = (x_1, \dots, x_k)$ is the unique α -approximate $x_1 \rightsquigarrow x_k$ path in G . This claim is a consequence of the following property of directed paths in G , which can be visualized in Figure 2.

Observation 2.5 (Over-Under Property). *Any $x_1 \rightsquigarrow x_k$ path π such that $\pi \neq \pi_{L_i}$ must contain an edge lying on a line with slope larger than L_i .*

Using the Over-Under Property, it follows that any $x_1 \rightsquigarrow x_k$ path $\pi \neq \pi_{L_i}$ must contain an edge e lying on a line with slope larger than L_i . However, by our exponential weighting scheme,

$$w(e) \geq (\alpha \cdot n)^{i+1} > \alpha \cdot w(\pi_{L_i}).$$

Therefore, any α -approximate preserver of G, P must contain every edge in π_{L_i} . This directly implies an $\Omega(n^{2/3}p^{2/3})$ lower bound for any α -approximate preserver of G, P , completing the sketch of Item 2 of Theorem 1.3.

The Over-Under weighting scheme is also used in our construction for approximate hopset lower bounds, combined with the previous undirected exact hopset lower bound of [BH23] and simplifications from [VXX24].




Figure 2: This figure depicts an example graph constructed by a point-line incidence system. All edges on lines with a larger slope than the line containing x_1 and x_4 are marked as red, wider lines. Except the path (x_1, x_2, x_3, x_4) , all other paths from x_1 to x_4 use red edges, which have weight much larger than the edges on (x_1, x_2, x_3, x_4) .

2.3 $O(m)$ -size Shortcut Set and Hopset Lower Bounds

Here we give an overview view of our $O(m)$ -size shortcut set and hopset lower bounds.

Shortcut Set Lower Bound. Our construction for an $\Omega(n^{2/9})$ diameter lower bound for $O(m)$ -size shortcut sets is a refinement of a prior shortcut set lower bound due to [VXX24]. The construction of [VXX24] obtains an $\Omega(n^{1/5})$ lower bound for $O(m)$ -size shortcut sets. Their constructed graph is a layered graph with $\Theta(n^{1/5})$ layers, and each layer has $\Theta(n^{4/5})$ vertices. Our construction intuitively can be viewed as a concatenation of $n^{1/5}$ copies of their graphs, where the last layer of the previous copy is identified with the first layer of the next copy. Due to technical reasons, the number of vertices in each layer needs to blow up by a factor of $\Theta((n^{1/5})^3)$. Thus, the total number of vertices is $\Theta(n^{9/5})$, and the number of layers is $\Theta(n^{2/5})$, which intuitively gives an $\Omega(n^{2/9})$ lower bound.

Unweighted Exact Hopset Lower Bound. Our construction for unweighted exact hopset lower bounds is an adaptation of our layered construction for the $\Omega(n^{2/9})$ lower bound for $O(m)$ -size shortcut sets. The high-level intuition for the higher $\Omega(n^{2/7})$ lower bound is the following: in the construction for shortcut sets, we need the critical paths to be the unique paths between their start and end nodes; in hopsets, we only need the critical paths to be the unique *shortest* paths between their start and end nodes. This allows us to create an unlayered version of shortcut set construction, which reduces the total number of vertices in the graph, and leads to a higher lower bound.

3 Preliminaries

In this section, we present the notations and definitions used in our paper.

Notation

- For integers $n \geq 0$, we use $[n]$ to denote $\{1, \dots, n\}$ and $\llbracket n \rrbracket$ to denote $\{0, \dots, n-1\}$. For integers $i, j \geq 0$, where $i < j$, we use $[i, j]$ to denote $\{i, i+1, \dots, j-1, j\}$.
- Let $G = (V, E, w)$ be a weighted graph with weight function $w : E \rightarrow \mathbb{R}$. For vertex $v \in V$, we use $\deg_G(v)$ to denote the number of edges incident to v in G . For $s, t \in V$, we use $w(s, t)$ to denote the edge weight of the edge from s to t , $\text{dist}_G(s, t)$ to denote the distance from s to t in G and $\pi_G(s, t)$ to denote a shortest path from s to t . The subscripts G might be dropped if it is clear from context.
- For $n, p \geq 1$, we use $\text{UDP}(n, p)$ (resp. $\text{DDP}(n, p)$, $\text{DADP}(n, p)$) to denote the best bound on the size of exact distance preservers for n -node undirected (resp. directed, directed acyclic) weighted graphs with p demand pairs.
- For a vertex $v \in V(G)$, we use $\text{outdeg}_G(v)$ to denote the number of outgoing edges incident to v in G , and we use $\text{indeg}_G(v)$ to denote the number of incoming edges incident to v in G . We let $\deg_G(v) = \text{indeg}_G(v) + \text{outdeg}_G(v)$.
- We will frequently use symbol π to denote paths. For a simple path π , $|\pi|$ denotes its number of edges and $w(\pi)$ denotes the total edge weights of edges on π . For two vertices x, y on π where x appears before y , $\pi[x, y]$ or $\pi[x \rightsquigarrow y]$ denotes the subpath of π from x to y .
- Given a path π and a set of nodes S , we use $\pi \cap S$ to denote the subpath of π induced on the nodes of S . Formally, $\pi \cap S$ is the path obtained by listing the nodes in $V(\pi) \cap S$ in the order they appear on path π .
- Given a graph G , a set of demand pairs $P \subseteq V(G) \times V(G)$, and a distance preserver H of G, P , we say that H is a minimal distance preserver if no edge can be deleted from H while ensuring that H remains a distance preserver of G, P . If shortest paths in G are unique shortest paths, then there is a unique minimal distance preserver H of G .

Definitions

Definition 3.1 (Tiebreaking Scheme). Given a graph G and a set of demand pairs $P \subseteq V(G) \times V(G)$, we define a tiebreaking function $\pi(\cdot, \cdot)$ with respect to G, P as follows. For each demand pair $(s, t) \in P$, fix a shortest path $\pi_{s,t}$ in G . Then we define the output of tiebreaking function $\pi(\cdot, \cdot)$ on input (s, t) to be

$$\pi(s, t) = \pi_{s,t}$$

for all $(s, t) \in P$. Additionally, we let $\pi(P)$ denote the image of P under $\pi(\cdot, \cdot)$, i.e., $\pi(P)$ denotes the collection of paths $\{\pi(s, t) \mid (s, t) \in P\}$. Likewise, we let

$$\bigcup_{\pi \in \pi(P)} \pi$$

denote the graph obtained by taking the (non-disjoint) union of the paths in set $\pi(P)$.

Definition 3.2 (Consistent Tiebreaking Scheme). Let $\pi(\cdot, \cdot)$ be a tiebreaking scheme with respect to G, P . Then $\pi(\cdot, \cdot)$ is consistent if for all paths $\pi_1, \pi_2 \in \pi(P)$ and nodes $x, y \in V(G)$ such that x precedes y in both π_1 and π_2 , then we have that $\pi_1[x \rightsquigarrow y] = \pi_2[x \rightsquigarrow y]$.

Definition 3.3 (Branching Events [CE06]). Let G be a directed graph, let P be a set of demand pairs, and let $\pi(\cdot, \cdot)$ be a tiebreaking scheme. Let $H = \cup_{(s,t) \in P} \pi(s, t)$ be the exact distance preserver implied by $\pi(\cdot, \cdot)$.

A branching event is a pair of directed edges $(x, y), (x, z) \in E(H)$ in H , for any $x, y, z \in V(G)$. For each branching event $(x, y), (x, z)$, we associate a pair of paths $\pi_1, \pi_2 \in \pi(P)$ such that $(x, y) \in \pi_1$ and $(x, z) \in \pi_2$. If there is more than one choice of paths π_1, π_2 , then we choose arbitrarily.

We say that a branching event $(x, y), (x, z)$ is *between* paths $\pi_1, \pi_2 \in \pi(P)$ if paths π_1, π_2 are associated with branching event $(x, y), (x, z)$ as in Definition 3.3.

4 Directed to Undirected Preserver Reduction

The goal of this section will be to prove Theorem 1.4 and Corollaries 1.5 and 1.6. We restate Theorem 1.4 below.

Theorem. If $\text{UDP}(n, p) = O(n^\lambda p^\mu + n)$ for constants $\lambda, \mu > 0$, then

$$\text{DDP}(n, p) = O\left(n^{\frac{1}{2-\lambda}} p^{\frac{1+\mu-\lambda}{2-\lambda}} + n^{1/2} p + n\right).$$

We will need several intermediate claims and lemmas before we can prove Theorem 1.4. Let G, P be an n -node directed, weighted graph and a set of $|P| = p$ demand pairs, such that the minimal distance preserver H of G, P has

$$|E(H)| = \Theta(\text{DDP}(n, p))$$

edges. Let $\pi(\cdot, \cdot)$ be a tiebreaking scheme associated with G, P . We will make the following assumptions on G, P, H , and $\pi(\cdot, \cdot)$, which we can assume without loss of generality:

1. Path $\pi(s, t)$ is a unique shortest $s \rightsquigarrow t$ path in G for all $s, t \in V(G)$. In particular, this implies that $\pi(\cdot, \cdot)$ is consistent.
2. Paths in $\pi(P)$ are pairwise edge-disjoint.
3. Every node $v \in V(G)$ has degree at least $\deg_H(v) \geq \frac{|E(H)|}{4n}$ in H .
4. Every path $\pi \in \pi(P)$ has at least $|\pi| \geq \frac{|E(H)|}{4p}$ edges.

Assumptions 1 and 2 hold without loss of generality due to the Independence Lemma for weighted distance preservers (Lemma 33 of [BHT23]). Note that Assumption 1 implies that tiebreaking scheme $\pi(\cdot, \cdot)$ is consistent. Assumptions 3 and 4 follow from standard arguments in this area, so we defer the proof of these assumptions to Appendix A.

Our extremal reduction from $\text{DDP}(n, p)$ to $\text{UDP}(n, p)$ will happen in two steps. First, in Section 4.1, we will reduce DDP to DADP , culminating in the proof of Lemma 4.5. Second, in Section 4.2 we will show that $\text{DADP}(n, p) \leq \text{UDP}(n, p)$ and prove Corollary 1.6. We combine these two reductions in Section 4.3 to complete the proofs of Theorem 1.4 and Corollary 1.5.

4.1 Directed to DAG Reduction

We begin our reduction from DDP to DADP with the following observation about the structure of paths in $\pi(P)$.

Claim 4.1 (Path Intersection Claim). *Fix two distinct shortest paths $\pi_1, \pi_2 \in \pi(P)$. Let*

$$v_1, \dots, v_k \in V(\pi_1) \cap V(\pi_2)$$

be the nodes in $V(\pi_1) \cap V(\pi_2)$, listed in the order they appear in path π_1 . Then nodes in $V(\pi_1) \cap V(\pi_2)$ appear on path π_2 in the reverse order $v_k, \dots, v_1$.

Proof. Suppose for the sake of contradiction that there are two vertices $x, y \in V(\pi_1) \cap V(\pi_2)$, where $x \neq y$, such that node x comes before node y in path π_1 and in path π_2 . Then by the consistency of tiebreaking scheme $\pi(\cdot, \cdot)$, it follows that $\pi_1[x, y] = \pi_2[x, y]$. However, this implies that there exists an edge $e \in E(\pi_1) \cap E(\pi_2)$, contradicting Assumption 2 about the collection of paths $\pi(P)$. $\square$

Now, for any path $\pi^* \in \pi(P)$, we will define an associated collection of acyclic paths Π_{DAG} . This will correspond to the acyclic collection of paths discussed in the technical overview in Observation 2.1.

Definition 4.2 (DAG Path Collection). Fix a path $\pi^* \in \pi(P)$. For every other path $\pi \in \pi(P) \setminus \{\pi^*\}$ we will define a subpath π_{DAG} of π induced on the nodes of π^* . This will yield an acyclic collection of paths $\Pi_{\text{DAG}} = \{\pi_{\text{DAG}} \mid \pi \in \pi(P) \setminus \{\pi^*\}\}$. Given a path $\pi \in \pi(P) \setminus \{\pi^*\}$, we construct π_{DAG} as follows.

- Let $v_1, \dots, v_k \in V(\pi^*) \cap V(\pi)$ be the nodes in $V(\pi^*) \cap V(\pi)$, listed in the order they appear in path π .
- We define π_{DAG} to be

$$\pi_{\text{DAG}} = (v_1, \dots, v_k).$$

- We assign edge $(v_i, v_{i+1}) \in \pi_{\text{DAG}}$ weight $w(v_i, v_{i+1}) = \text{dist}_G(v_i, v_{i+1})$.

We define directed graph D be the graph obtained by taking the (non-disjoint) union of paths in Π_{DAG} , i.e.,

$$D = \bigcup_{\pi_{\text{DAG}} \in \Pi_{\text{DAG}}} \pi_{\text{DAG}}.$$

We now verify that Π_{DAG} is an acyclic collection of shortest paths in graph D .

Claim 4.3. *Fix a path $\pi^* \in \pi(P)$, and let Π_{DAG} and D be as defined in Definition 4.2. Then graph D is a DAG, and π_{DAG} is a unique shortest path in D for all $\pi_{\text{DAG}} \in \Pi_{\text{DAG}}$.*

Proof. Let $<_D$ be the total ordering on the vertices of $V(D)$ that is the reverse of the ordering of the vertices in path π^* . That is, if node x comes before node y in π^* , then $y <_D x$. It is immediate that $<_D$ is a valid topological ordering for directed graph D because every path $\pi_{\text{DAG}} \in \Pi_{\text{DAG}}$ respects $<_D$ by Claim 4.1.

To see that all paths in Π_{DAG} are unique shortest paths in D , observe the following properties of graph D :

- For all $s, t \in V(D)$,

$$\text{dist}_D(s, t) \geq \text{dist}_G(s, t).$$

- For all vertices $s, t \in V(D)$, if $s, t \in V(\pi_{\text{DAG}})$ for some $\pi \in \Pi_{\text{DAG}}$, then

$$\text{dist}_D(s, t) = \text{dist}_G(s, t).$$

These two properties imply that every path $\pi_{\text{DAG}} \in \Pi_{\text{DAG}}$ is a shortest path in D . To see that π_{DAG} is a *unique* shortest path in D , observe that if π_{DAG} is an $s \rightsquigarrow t$ path in D , then path π_{DAG} is the unique (possibly non-contiguous) subpath of path $\pi(s, t) \subseteq G$ in D . Then since $\pi(s, t)$ is a unique shortest path in G , we conclude that π_{DAG} is a unique shortest path in D . $\square$

Our end goal is to upper bound the size $|E(H)|$ of graph H by upper bounding the size $|E(D)|$ of DAG D . As such, we need to prove that $|E(D)|$ grows with $|E(H)|$. We achieve this in a roundabout way with the following claim.

Claim 4.4. *Fix a shortest path $\pi^* \in \pi(P)$, and let Π_{DAG} and D be as defined in Definition 4.2. Let $Q \subseteq P$ be the set of demand pairs (s, t) in P such that $\pi(s, t)$ intersects path π^* at exactly 1 node. Then*

$$\sum_{v \in V(\pi^*)} \deg_H(v) = \Theta(|\pi^*| + |Q| + |E(D)|).$$

Proof. In order to prove Claim 4.4, it will suffice to tightly bound the number of distinct edges incident to nodes in π^* , since

$$\sum_{v \in V(\pi^*)} \deg_H(v) = \Theta(|\{e \in E(H) \mid e = (u, v), \text{ where } u \in \pi^* \text{ or } v \in \pi^*\}|).$$

Let $\Pi \subseteq \pi(P)$ be the set of all paths in $\pi(P)$ in graph H that have a nonempty intersection with path π^* . Observe that every edge incident to a node in π^* in H is contained in a path in Π . We partition the set of paths Π into sets Π_1, Π_2 , and Π_3 , where

$$\Pi_1 = \{\pi^*\}, \quad \Pi_2 = \pi(Q), \quad \text{and} \quad \Pi_3 = \Pi \setminus (\Pi_1 \cup \Pi_2).$$

We let E_1, E_2 , and E_3 denote the set of edges in paths in Π_1, Π_2 , and Π_3 , respectively, that are incident to nodes in π^* . Then since sets Π_1, Π_2 , and Π_3 partition Π , and paths in Π are pairwise edge-disjoint by Assumption 2, we conclude that

$$\sum_{v \in V(\pi^*)} \deg_H(v) = \Theta(|E_1| + |E_2| + |E_3|).$$

We will bound the sizes of E_1, E_2 , and E_3 separately.

- (Size of $|E_1|$.) Set E_1 has size $|E_1| = \Theta(|\pi^*|)$.
- (Size of $|E_2|$.) Every path $\pi \in \pi(Q)$ has either 1 or 2 edges incident to a node in π^* . Then since paths in $\pi(P)$ are pairwise edge-disjoint by Assumption 2, the edges of the paths in $\pi(Q)$ contribute $\Theta(|Q|)$ edges to the sum of the degrees of the nodes in π^* in H .
- (Size of $|E_3|$.) Fix a path $\pi \in \Pi_3$, and let π_{DAG} denote the corresponding path in Π_{DAG} . We observe that since $|V(\pi) \cap V(\pi^*)| > 1$,

$$|\{(u, v) \in E(\pi) \mid u \in V(\pi^*) \text{ or } v \in V(\pi^*)\}| = \Theta(|V(\pi) \cap V(\pi^*)|) = \Theta(|\pi_{\text{DAG}}|).$$

Then we conclude that

$$|E_3| = \sum_{\pi \in \Pi_3} |\{(u, v) \in E(\pi) \mid u \in V(\pi^*) \text{ or } v \in V(\pi^*)\}| = \Theta \left(\sum_{\pi \in \Pi_3} |\pi_{\text{DAG}}| \right) = \Theta(|E(D)|),$$

where the first and last equality follow from the fact that paths in Π are pairwise edge-disjoint and that paths in Π_{DAG} are pairwise edge-disjoint, respectively.

Putting it all together, we conclude that

$$\sum_{v \in V(\pi^*)} \deg_H(v) = \Theta(|E_1| + |E_2| + |E_3|) = \Theta(|\pi^*| + |Q| + |E(D)|),$$

as desired. $\square$

We now have all the ingredients we need to prove our reduction from DDP to DADP.

Lemma 4.5 (Directed to DAG Reduction).

$$\text{DDP}(n, p) \leq O \left(\text{DADP} \left(\frac{\text{DDP}(n, p)}{p}, p \right)^{1/2} \cdot n^{1/2} p^{1/2} + n^{1/2} p + n \right).$$

Proof. Let G, P be an n -node directed, weighted graph and a set of $|P| = p$ demand pairs, such that the minimal distance preserver H of G, P has

$$|E(H)| = \Theta(\text{DDP}(n, p))$$

edges. Let $\pi(\cdot, \cdot)$ be a tiebreaking scheme associated with G, P . We can assume that Assumptions 1-4 discussed at the beginning of Section 4 hold without loss of generality. In particular, these assumptions imply that there exists a path $\pi^* \in \pi(P)$ such that $|\pi^*| = \Theta \left(\frac{|E(H)|}{p} \right)$, since the average length of paths in $\pi(P)$ is $\Theta \left(\frac{|E(H)|}{p} \right)$, and every path in $\pi(P)$ has length at least $\frac{|E(H)|}{4p}$.

Let the collection of paths Π_{DAG} and DAG D be as defined in Definition 4.2, with respect to path π^* . By Assumption 3, we have that

$$\sum_{v \in V(\pi^*)} \deg_H(v) = \Omega \left(\frac{|E(H)|}{p} \cdot \frac{|E(H)|}{n} \right) = \Omega \left(\frac{\text{DDP}(n, p)^2}{np} \right).$$

Applying Claim 4.4, it follows that

$$\Omega \left(\frac{\text{DDP}(n, p)^2}{np} \right) \leq \Theta(|\pi^*| + |Q| + |E(D)|),$$

where $Q \subseteq P$ is the set of demand pairs $(s, t) \in P$ such that $\pi(s, t)$ intersects path π^* at exactly 1 node. We make the following observations:

- $|\pi^*| = \Theta \left(\frac{\text{DDP}(n, p)}{p} \right)$ by our choice of π^* ,
- $|Q| \leq p$ since $Q \subseteq P$, and

- $|E(D)| \leq \text{DADP}(|V(D)|, |\Pi_{\text{DAG}}|)$.

The final observation follows from the fact that each path $\pi_{\text{DAG}} \in \Pi_{\text{DAG}}$ is a unique shortest path between its endpoints in D by Claim 4.3. Consequently, if we define demand pairs

$$P_D = \{(s, t) \in V(D) \times V(D) \mid \exists \pi_{\text{DAG}} \in \Pi_{\text{DAG}} \text{ s.t. } \pi_{\text{DAG}} \text{ is an } s \rightsquigarrow t \text{ path in } D\},$$

then any distance preserver of D, P_D will contain every path in Π_{DAG} and thus $E(D)$. Using our observations and the fact that $|\Pi_{\text{DAG}}| \leq p$ and $|V(D)| \leq |V(\pi^*)| = \Theta(\frac{\text{DDP}(n, p)}{p})$, we obtain the following inequality:

$$\Omega\left(\frac{\text{DDP}(n, p)^2}{np}\right) \leq \Theta(|\pi^*| + |Q| + |E(D)|) \leq \Theta\left(\frac{\text{DDP}(n, p)}{p} + p + \text{DADP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)\right)$$

We now split our analysis into three cases, based on which of the three terms on the right hand side is the largest.

- (Term $\frac{\text{DDP}(n, p)}{p}$ dominates.) In this case, we find that

$$\frac{\text{DDP}(n, p)^2}{np} \leq O\left(\frac{\text{DDP}(n, p)}{p}\right).$$

Rearranging, we conclude that

$$\text{DDP}(n, p) = O(n).$$

- (Term p dominates.) In this case, we find that

$$\frac{\text{DDP}(n, p)^2}{np} \leq O(p).$$

Then multiplying both sides of the inequality by np and taking the square root of both sides, we get

$$\text{DDP}(n, p) = O(n^{1/2}p).$$

- (Term $\text{DADP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)$ dominates.) In this case, our inequality is

$$\frac{\text{DDP}(n, p)^2}{np} \leq O\left(\text{DADP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)\right).$$

Multiplying both sides by np , and taking the square root of both sides, we obtain

$$\text{DDP}(n, p) = O\left(\text{DADP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)^{1/2} \cdot n^{1/2}p^{1/2}\right)$$

Combining our three bounds from our three different cases, we conclude that

$$\text{DDP}(n, p) \leq O\left(\text{DADP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)^{1/2} \cdot n^{1/2}p^{1/2} + n^{1/2}p + n\right),$$

completing the proof. $\square$

4.2 DAG to Undirected Reduction

We now present our reduction to convert shortest paths in DAGs to shortest paths in undirected graphs.

Lemma 4.6 (DAG to Undirected Reduction).

$$\text{DADP}(n, p) \leq \text{UDP}(n, p)$$

Proof. Let $G = (V, E, w)$ be an n -node weighted directed acyclic graph, and let $P \subseteq V(G) \times V(G)$ be a set of $|P| = p$ demand pairs. Let $<_G$ be a topological ordering of DAG G , and let $v_1, \dots, v_n$ be the nodes of G , ordered according to $<_G$. Let W be a sufficiently large number (e.g., chose W to be larger than the sum of all edge weights in G). We may assume without loss of generality that for all $(s, t) \in P$, there exists a unique shortest $s \rightsquigarrow t$ path $\pi(s, t)$ in G , by the Independence Lemma for weighted distance preservers (Lemma 33 of [BHT23]).

We define an undirected graph $G' = (V, E', w')$ as follows. For every directed edge $(v_i, v_j) \in E(G)$, add the undirected edge (v_i, v_j) to $E(G')$, and assign it weight

$$w'(v_i, v_j) := w(v_i, v_j) + W \cdot (j - i).$$

This completes the construction of graph G' .

Fix a pair of nodes $v_i, v_j \in V$ such that v_i can reach v_j in G . Let π denote the unique shortest $v_i \rightsquigarrow v_j$ path in G , and let π' denote a shortest $v_i \rightsquigarrow v_j$ path in G' . We claim that paths π and π' are identical, up to a reweighting of their edges. Formally, we claim if path π has the corresponding node sequence $(x_1, x_2, \dots, x_k)$, then π' has the corresponding node sequence $(x_1, x_2, \dots, x_k)$, as well.

Suppose that path π has the corresponding node sequence $(x_1, x_2, \dots, x_k)$, where $x_1 = v_i$ and $x_k = v_j$. For all $\ell \in [1, k]$, let $q_\ell \in [1, n]$ be the index such that node $x_\ell = v_{q_\ell}$. Notice that since π is a path in DAG G , we must have that $q_\ell < q_{\ell+1}$ for all $\ell \in [1, k-1]$. Then the distance between v_i and v_j in G' is at most

$$\begin{aligned} \text{dist}_{G'}(v_i, v_j) &\leq w'(\pi) = \sum_{i=1}^{k-1} w'(x_i, x_{i+1}) = \sum_{i=1}^{k-1} (w(x_i, x_{i+1}) + W \cdot |q_{i+1} - q_i|) \\ &= \sum_{i=1}^{k-1} (w(x_i, x_{i+1}) + W \cdot (q_{i+1} - q_i)) \\ &= w(\pi) + W \cdot (q_k - q_1) \\ &= \text{dist}_G(s, t) + W \cdot (j - i). \end{aligned}$$

In the above sequence of inequalities, we made use of the fact that since $q_\ell < q_{\ell+1}$ for all $\ell \in [1, k-1]$, it follows that

$$\sum_{i=1}^{k-1} |q_{i+1} - q_i| = \sum_{i=1}^{k-1} (q_{i+1} - q_i) = q_k - q_1 = j - i.$$

Let path π' have the corresponding node sequence $(y_1, y_2, \dots, y_{k'})$, where $y_1 = v_i$ and $y_{k'} = v_j$. For all $\ell \in [1, k']$, let $r_\ell \in [1, n]$ be the index such that node $y_\ell = v_{r_\ell}$. We claim that since π' is a shortest $v_i \rightsquigarrow v_j$ path in G' , we have that $r_\ell < r_{\ell+1}$ for all $\ell \in [1, k'-1]$. Informally speaking, this is equivalent to claiming that path π' respects the total order $<_G$ in undirected graph G' .

Suppose towards contradiction that there exists an $\ell \in [1, k]$ such that $r_\ell > r_{\ell+1}$. Then

$$\sum_{i=1}^{k'-1} |r_{i+1} - r_i| > \sum_{i=1}^{k'-1} (q_{i+1} - q_i) = q_{k'} - q_1 = j - i.$$

This immediately implies that

$$\begin{aligned} w'(\pi') &= \sum_{i=1}^{k'-1} (w(y_i, y_{i+1}) + W \cdot |r_{i+1} - r_i|) \\ &\geq w(\pi') + W \cdot (j - i + 1) \\ &= (w(\pi') + W) + W \cdot (j - i) \\ &> w(\pi) + W \cdot (j - i) \\ &= \text{dist}_{G'}(s, t), \end{aligned}$$

where the final inequality follows the fact that if W is sufficiently large (e.g., larger than the sum of all edge weights in G), then $W > w(\pi)$. This contradicts our assumption that π' is a shortest $v_i \rightsquigarrow v_j$ path in G' , so we conclude that $r_\ell < r_{\ell+1}$ for all $\ell \in [1, k']$.

Then by identical calculations as before, we have that

$$\begin{aligned} w'(\pi') &= \sum_{i=1}^{k'-1} (w(y_i, y_{i+1}) + W \cdot (q_{i+1} - q_i)) \\ &= w(\pi') + W \cdot (q_k - q_1) \\ &= w(\pi') + W \cdot (j - i) \end{aligned}$$

Since π' is a shortest path in G' , then

$$w'(\pi') = w(\pi') + W \cdot (j - i) \leq \text{dist}_{G'}(v_i, v_j) \leq \text{dist}_G(v_i, v_j) + W \cdot (j - i),$$

so $w(\pi') \leq \text{dist}_G(v_i, v_j)$. Since path π is the unique shortest $v_i \rightsquigarrow v_j$ path in G , we conclude that paths π and π' are identical, e.g., $(x_1, \dots, x_k) = (y_1, \dots, y_{k'})$.

Then every path $\pi \in \pi(P)$ is a unique shortest path between its endpoints in undirected graph G' . Informally, this means that the unique shortest paths between demand pairs P are identical in G and G' . As a consequence, any distance preserver of G, P is a distance preserver of G', P , and vice versa. We conclude that $\text{DADP}(n, p) \leq \text{UDP}(n, p)$. $\square$

Our reduction from DADP to UDP also implies an algorithmic reduction from APSP in DAGs to APSP in undirected graphs, as we claimed in Section 1.1 in Corollary 1.6. We now prove Corollary 1.6.

Corollary 4.7 (restatement of Corollary 1.6). *Let G be an n -node, m -edge weighted DAG. Then we can compute all-pairs shortest paths on G in $O(m)$ time, plus the time it takes to compute all-pairs shortest paths on an n -node, m -edge undirected weighted graph.*

Proof. First, we observe that given a directed acyclic graph G , we can construct graph G' from Lemma 4.6 in $O(n + m)$ time. Topological sorting can be done in $O(n + m)$ time. We can reweigh edge (v_i, v_j) by adding $+W \cdot (j - i)$ weight to (v_i, v_j) in $O(1)$ time. Thus constructing undirected, weighted graph G' can be done in $O(n + m)$ time. Let $v_1, \dots, v_n$ be the topological ordering of G from Lemma 4.6.

For all $v_i, v_j \in V(G)$, if

$$\text{dist}_{G'}(v_i, v_j) \geq W \cdot (j - i + 1),$$

then we report that $\text{dist}_G(v_i, v_j) = \infty$. Otherwise, we report that

$$\text{dist}_G(v_i, v_j) = \text{dist}_{G'}(v_i, v_j) - W \cdot (j - i).$$

□

4.3 Finishing the Reduction

We can chain our reductions in Lemmas 4.5 and 4.6 to prove Theorem 1.4.

Proof of Theorem 1.4. Suppose that $\text{UDP}(n, p) = O(n^\lambda p^\mu + n)$. Then we calculate

$$\begin{aligned} \text{DDP}(n, p) &\leq O \left(\text{DADP} \left(\frac{\text{DDP}(n, p)}{p}, p \right)^{1/2} \cdot n^{1/2} p^{1/2} + n^{1/2} p + n \right) \\ &\leq O \left(\text{UDP} \left(\frac{\text{DDP}(n, p)}{p}, p \right)^{1/2} \cdot n^{1/2} p^{1/2} + n^{1/2} p + n \right) \\ &\leq O \left(\left(\left(\frac{\text{DDP}(n, p)}{p} \right)^\lambda p^\mu + \frac{\text{DDP}(n, p)}{p} \right)^{1/2} \cdot n^{1/2} p^{1/2} + n^{1/2} p + n \right) \\ &\leq O \left(\left(\frac{\text{DDP}(n, p)^{\lambda/2}}{p^{\lambda/2}} \cdot p^{\mu/2} + \frac{\text{DDP}(n, p)^{1/2}}{p^{1/2}} \right) \cdot n^{1/2} p^{1/2} + n^{1/2} p + n \right) \\ &\leq O \left(\text{DDP}(n, p)^{\lambda/2} \cdot n^{1/2} p^{\frac{1+\mu-\lambda}{2}} + \text{DDP}(n, p)^{1/2} \cdot n^{1/2} + n^{1/2} p + n \right) \end{aligned}$$

This upper bound has four terms. Suppose the first term $\text{DDP}(n, p)^{\lambda/2} \cdot n^{1/2} p^{\frac{1+\mu-\lambda}{2}}$ dominates. Then

$$\text{DDP}(n, p)^{1-\lambda/2} \leq O \left(n^{1/2} p^{\frac{1+\mu-\lambda}{2}} \right),$$

so in this case we conclude

$$\text{DDP}(n, p) \leq O \left(n^{\frac{1}{2-\lambda}} p^{\frac{1+\mu-\lambda}{2-\lambda}} \right).$$

Likewise, when the second term $\text{DDP}(n, p)^{1/2} \cdot n^{1/2}$ dominates, we get that $\text{DDP}(n, p) = O(n)$. Putting it all together, we conclude that

$$\text{DDP}(n, p) = O \left(n^{\frac{1}{2-\lambda}} p^{\frac{1+\mu-\lambda}{2-\lambda}} + n^{1/2} p + n \right).$$

□

We can use a similar analysis to prove Corollary 1.5. We restate Corollary 1.5 in greater detail below for convenience.

Corollary 4.8 (cf. Corollary 1.5). *Suppose that consistent tiebreaking schemes are optimal for directed distance preservers when $p \leq n^{2/3}$, or equivalently,*

$$\text{DDP}(n, p) = \Theta(\min(n^{2/3} p + n, np^{1/2})) = \Theta(n^{2/3} p + n),$$

when $p \leq n^{2/3}$. Then consistent tiebreaking schemes are optimal for undirected distance preservers when $p \leq n$, or equivalently,

$$\text{UDP}(n, p) = \Theta(\min(n^{1/2}p + n, np^{1/2})) = \Theta(n^{1/2}p + n),$$

when $p \leq n$.

Proof. Suppose that

$$\text{DDP}(n, p) = \Theta(\min(n^{2/3}p + n, np^{1/2})) = \Theta(n^{2/3}p + n) = \Theta(n^{2/3}p),$$

when $n^{1/3} \leq p \leq n^{2/3}$. Now by combining Lemmas 4.5 and 4.6, we get the inequality

$$\text{DDP}(n, p) \leq O\left(\text{UDP}\left(\frac{\text{DDP}(n, p)}{p}, p\right)^{1/2} \cdot n^{1/2}p^{1/2} + n^{1/2}p + n\right).$$

Plugging in $\text{DDP}(n, p) = \Theta(n^{2/3}p)$, we compute that

$$n^{2/3}p \leq O\left(\text{UDP}\left(n^{2/3}, p\right)^{1/2} \cdot n^{1/2}p^{1/2} + n^{1/2}p + n\right).$$

Let $c > 0$ be a universal constant such that

$$n^{2/3}p \leq c \cdot \left(\text{UDP}\left(n^{2/3}, p\right)^{1/2} \cdot n^{1/2}p^{1/2} + n^{1/2}p + n\right),$$

for sufficiently large n . Then for all $10cn^{1/3} \leq p \leq n^{2/3}$,

$$n^{2/3}p \leq 10c \cdot \text{UDP}\left(n^{2/3}, p\right)^{1/2} \cdot n^{1/2}p^{1/2}.$$

Dividing both sides by $10cn^{1/2}p^{1/2}$ and taking the square root, we obtain

$$\text{UDP}(n^{2/3}, p) = \Omega(n^{1/3}p).$$

If we let $N = n^{2/3}$, this is equivalent to

$$\text{UDP}(N, p) = \Omega(N^{1/2}p),$$

for $10c \cdot N^{1/2} \leq p \leq N$. Moreover, it is true unconditionally that $\text{UDP}(N, p) = \Theta(N)$ when $p \leq N^{1/2}$. We conclude that

$$\text{UDP}(N, p) = \Omega(N^{1/2}p + N),$$

for $p \leq N$. Consistent tiebreaking schemes imply that

$$\text{UDP}(N, p) = O(\min(N^{1/2}p + N, Np^{1/2})) = O(N^{1/2}p + N),$$

when $p \leq N$, so we conclude that

$$\text{UDP}(N, p) = \Theta(N^{1/2}p + N) = \Theta(\min(N^{1/2}p + N, Np^{1/2})),$$

when $p \leq N$, as desired. □

5 Exact Distance Preserver Upper Bound

The goal of this section will be to prove the following theorem.

Theorem 5.1. *Every n -node directed unweighted graph G with demand pairs P of size $p = |P|$ admits a distance preserver with*

$$\tilde{O}\left(n^{5/6}p^{2/3} + n\right)$$

edges.

We will need to introduce some preliminary definitions first.

5.1 Preliminary Definitions

Throughout Section 5, we will use G to denote an n -node directed unweighted graph, and we will use $P \subseteq V(G) \times V(G)$ to denote an associated set of demand pairs of size $|P| = p$. Additionally, we will associate with G and P a *consistent* tiebreaking scheme $\pi(\cdot, \cdot)$. Recall that for each demand pair $(s, t) \in P$, our tiebreaking scheme fixes an $s \rightsquigarrow t$ shortest path $\pi(s, t)$. We let $\pi(P)$ denote this collection of shortest paths in G , i.e.,

$$\pi(P) = \{\pi(s, t) \mid (s, t) \in P\}.$$

Throughout this section, we will assume that every edge $e \in E(G)$ is contained in a path in $\pi(P)$. (This assumption is without loss of generality because we can safely remove all other edges in G without increasing distances between demand pairs.)

Given graph G and demand pairs P , we define the following quantities.

- We define quantity $\ell = \ell(G, P)$ to be

$$\ell(G, P) := \left\lceil \frac{|E(G)|}{p} \right\rceil.$$

Quantity ℓ roughly corresponds to the average number of edges each path in $\pi(P)$ contributes to $|E(G)|$.

- We define quantity $d = d(G, P)$ to be

$$d(G, P) := \left\lceil \frac{|E(G)|}{n} \right\rceil.$$

Quantity d roughly corresponds to the average degree of nodes in G .

We will assume without loss of generality that ℓ and d are sufficiently large. If $\ell \leq c$ or $d \leq c$ for some constant c , then $|E(G)| \leq cp$ or $|E(G)| \leq cn$, respectively. Thus our assumption that $\ell > c$ and $d > c$ for a sufficiently large constant c will add

$$O(\ell p + dn) = O(p + n)$$

additional edges to our final preserver H . With this assumption, it follows that

$$\ell p = \Theta(|E(G)|) = \Theta(dn).$$

We will use the equality repeatedly in our proof. We now define an analogous quantity $\hat{\ell}$ that captures how much paths in $\pi(P)$ overlap.

- We define $\widehat{\ell} = \widehat{\ell}(G, P)$ to be

$$\widehat{\ell} := \frac{\sum_{\pi \in \pi(P)} |\pi|}{p}.$$

Quantity $\widehat{\ell}$ roughly corresponds to the average length of the paths in $\pi(P)$. Note that $\widehat{\ell} \gg \ell$ in general, and $\widehat{\ell} = \Theta(\ell)$ when paths in $\pi(P)$ are pairwise edge-disjoint.

We will make the following assumption about paths in $\pi(P)$. For every path $\pi = \pi(s, t)$, we have that

$$\widehat{\ell}/2 \leq |\pi| \leq 2\widehat{\ell}.$$

This assumption holds without loss of generality by a standard bucketing argument. That is, we can bucket the demand pairs $(s, t) \in P$ into buckets $P_0, \dots, P_{\log n}$ so that

$$P_i = \{(s, t) \in P \mid 2^i \leq \text{dist}_G(s, t) \leq 2^{i+1}\}$$

for $i \in [0, \log n]$. Then we can apply our bound on G, P_i for each $i \in [0, \log n]$. This will worsen our upper bound by at most a logarithmic factor.

5.2 Ingredient 1: the Dense Low-Diameter Cluster Lemma

The first ingredient in our distance preserver construction will be the following lemma, which roughly states that if our distance preserver is sufficiently dense, then it contains a low-diameter subset of vertices containing many high-degree nodes. We defer the proof of this lemma to Section 5.5.

Lemma 5.2 (Dense Low-Diameter Cluster Lemma). *Let G be an n -node directed unweighted graph with associated set of demand pairs $P \subseteq V(G) \times V(G)$ of size $|P| = p$.³ Let $\ell = \ell(G, P)$, $d = d(G, P)$, and $\widehat{\ell} = \widehat{\ell}(G, P)$. If G has average degree d , where $d \geq 10p/n^{1/2}$, then G contains a collection of nodes $S \subseteq V(G)$ such that*

1. $|S| = \Theta(d)$,
2. (Dense.) For all $s \in S$, $\deg_G(s) = \Omega(d)$, and
3. (Low diameter.) For all $s, t \in S$, $\text{dist}_G(s, t) = O\left(\frac{\widehat{\ell}p^2}{nd^2}\right)$.

An important feature of the set of vertices S identified in Lemma 5.2 is that S contains $\Theta(d)$ nodes each of degree $\Omega(d)$ in G . We will see how this property implies that many paths in $\pi(P)$ pass through nodes in S , as we state in the following lemma.

Lemma 5.3. *Let G, P, S and $\ell, d, \widehat{\ell}$ be as described in Lemma 5.2. Let $Q \subseteq P$ be the set of demand pairs $(s, t) \in P$ such that $\pi(s, t)$ contains a node in S . Then*

$$|Q| = \Omega(d^2).$$

³We assume without loss of generality that every edge in $E(G)$ lies on a path in $\pi(P)$ for a consistent tiebreaking scheme $\pi(\cdot, \cdot)$.

Before we can prove Lemma 5.3, we will first need to introduce some new notation and establish several technical claims. Let $\Pi \subseteq \pi(P)$ be a collection of paths in $\pi(P)$, and let $S \subseteq V(G)$ be a set of nodes in G . Let $\Pi[S]$ denote the collection of paths in Π induced on the nodes in set S . Formally, for a path $\pi \in \Pi$, we define the path $\pi[S]$ induced on the set S to be the path obtained by deleting every node in $\pi \setminus S$ from π . Likewise, we let $\Pi[S] = \{\pi[S] \mid \pi \in \Pi\}$. Note that the paths in $\Pi[S]$ do not necessarily correspond to paths in graph G , as $E(\pi[S]) \not\subseteq E(\pi)$, in general.

We quickly observe that if a collection of paths Π is consistent, then the collection of paths $\Pi[S]$ is also consistent.

Claim 5.4. *Let Π be a consistent collection of paths, and let $S \subseteq V(G)$ be a set of nodes. Then the collection of paths $\Pi[S]$ is also consistent.*

We need an additional definition and technical claim before we can prove Lemma 5.3. Given a collection of paths Π and a set of nodes $S \subseteq V(G)$, we define an auxiliary graph $J(S, \Pi)$ and an associated collection of paths Π_J as follows.

Definition 5.5 (Graph J and paths Π_J). Let Π be a collection of paths, and let $S \subseteq V(G)$ be a set of nodes. Then we define a graph $J = J(S, \Pi)$ and associated collection of paths Π_J as follows.

- Let $V(J) = S$,
- Let $\Pi_J = \Pi[S]$, and
- Let the edge set $E(J)$ of J be

$$E(J) = \bigcup_{\pi \in \Pi_J} E(\pi).$$

The definition of graph J and paths Π_J we be useful in the proof of Lemma 5.3, when combined with the following technical claim.

Claim 5.6. *Let G, P, S and $\ell, d, \hat{\ell}$ be as described in Lemma 5.2. Let $Q \subseteq P$ be the set of demand pairs $(s, t) \in P$ such that $\pi(s, t)$ contains a node in S . Let $J = J(S, \pi(Q))$ and $\Pi_J = \pi(Q)[S]$ be the graph and collection of paths specified in Definition 5.5 with respect to set S and paths $\pi(Q)$. Then*

$$|V(J)| = |S| = \Theta(d) \quad \text{and} \quad |E(J)| = \Omega(d^2) - |Q|.$$

Proof. It is immediate from Definition 5.5 that $|V(J)| = |S| = \Theta(d)$. What remains is to prove that $|E(J)| = \Omega(d^2) - |Q|$.

Let $E' \subseteq E(G)$ be the set of edges in G incident to a node in S . Formally, let

$$E' = \{(u, v) \in E(G) \mid u \in S \text{ or } v \in S\}.$$

By Lemma 5.2, $|E'| = \Omega(d) \cdot |S| = \Omega(d^2)$. We can assume without loss of generality that every edge $e \in E'$ lies on a path π in $\pi(Q)$. We will define a subset $E'' \subseteq E'$ of E' as follows.

$$E'' = \{e \in E' \mid \text{there exists } \pi \in \pi(Q) \text{ and } s, t \in V(\pi) \cap S \text{ such that } e \in \pi[s \rightsquigarrow t]\}.$$

Informally, E'' corresponds to the set of edges in E' that lie between two nodes $s, t \in S$ on a path $\pi \in \pi(Q)$. Notice that for every path $\pi \in \pi(Q)$, among the set of edges $E(\pi) \cap E'$, only the first and last edge in $E(\pi) \cap E'$ to appear on path π are not contained in E'' . Consequently,

$$|E(\pi) \cap E''| \geq |E(\pi) \cap E'| - 2.$$

This immediately gives the bound

$$|E''| \geq |E'| - 2|Q| = \Omega(d^2) - 2|Q|.$$

To prove Claim 5.6, we want to translate our lower bound on $|E''|$ into a lower bound on $|E(J)|$. Towards this end, we will define a mapping $\varphi : E'' \mapsto E(J)$ as follows. Let e be an edge in E'' . Since $e \in E''$, there exists a path $\pi \in \pi(Q)$ and nodes $s, t \in V(\pi) \cap S$ such that $e \in \pi[s \rightsquigarrow t]$. Moreover, we can assume without loss of generality that $(s, t) \in E(J)$, by choosing nodes s and t that are closest to edge e on path π . We let $\varphi(e) = (s, t)$. We can define $\varphi(e)$ in this way for all $e \in E''$.

We claim that for every edge $(s, t) \in E(J)$, the number of edges in the preimage of edge (s, t) is at most two, i.e., $|\varphi^{-1}((s, t))| \leq 2$. Fix an edge $(s, t) \in E(J)$. Let $E''_{(s,t)} \subseteq E''$ denote the set of edges e in E'' such that $\varphi(e) = (s, t)$. If $|E''_{(s,t)}| = 0$, then we are done. Otherwise, there must exist a path $\pi^* \in \pi(P)$ and nodes $s, t \in V(\pi^*) \cap S$ such that node s precedes node t in path π^* . Since the collection of paths $\pi(P)$ is consistent, for every path $\pi \in \pi(P)$ such that $s, t \in V(\pi)$ and node s precedes node t , we have that $\pi[s \rightsquigarrow t] = \pi^*[s \rightsquigarrow t]$. Additionally, since $(s, t) \in E(J)$, path $\pi^*[s \rightsquigarrow t]$ is internally vertex-disjoint from set S . Consequently, only the first and last edge of path $\pi^*[s \rightsquigarrow t]$ are contained in E'' . This implies that $|E''_{(s,t)}| \leq 2$, as claimed.

Putting everything together, we find that

$$|E(J)| \geq \frac{1}{2} \cdot |\varphi^{-1}(E(J))| = \frac{1}{2} \cdot |E''| = \Omega(d^2) - |Q|,$$

as claimed. $\square$

With Definition 5.5 and Claim 5.6 in hand, we are now ready to prove Lemma 5.3.

Proof of Lemma 5.3. Let G, P, S and $\ell, d, \widehat{\ell}$ be as described in Lemma 5.2. Let $Q \subseteq P$ be the set of demand pairs $(s, t) \in P$ such that $\pi(s, t)$ contains a node in S . Let $J = J(S, \pi(Q))$ and $\Pi_J = \pi(Q)[S]$ be the graph and collection of paths specified in Definition 5.5 with respect to set S and paths $\pi(Q)$.

Note that by Claim 5.4, the collection of paths Π_J is consistent. Additionally, by Claim 5.6 $|V(J)| = \Theta(d)$ and $|E(J)| \geq \Omega(d^2) - |Q|$. Now if $|Q| \geq \Omega(d^2)$, then we have finished the proof of Lemma 5.3. Otherwise, we may assume that $|E(J)| = \Omega(d^2)$.

Then Π_J is a consistent collection of $|\Pi_J| = |Q|$ paths, over a vertex set S of size $|S| = \Theta(d)$. Moreover, the number of distinct edges in Π_J is at least $|E(J)| = \Omega(d^2)$. By Theorem 1.1 of [CE06], a consistent tiebreaking scheme for x demand pairs over y nodes has at most $O(yx^{1/2})$ distinct edges. Plugging in $x \geq |Q|$ and $y = |S| = \Theta(d)$, this implies that graph J has at most $|E(J)| = O(d|Q|^{1/2})$ distinct edges. Putting everything together,

$$\Omega(d^2) = |E(J)| = O(d|Q|^{1/2}),$$

so $|Q| = \Omega(d^2)$, as claimed. $\square$

5.3 Ingredient 2: Sourcewise Distance Preservers

The second ingredient in our distance preserver construction will be a sourcewise distance preserver upper bound for a collection of sources S with low diameter. This sourcewise preserver is a direct generalization of the sourcewise preserver of [BV21] to the directed setting, and follows from a similar argument.

Lemma 5.7 (cf. Lemma 8 of [BV21]). *Let G be an n -node directed graph, and let $S \subseteq V(G)$ be a set of $|S| = s$ nodes of weak diameter h (i.e., $\text{dist}_G(s, t) \leq h$ for all $s, t \in S$). Let $P \subseteq S \times V$ be a set of $|P| = p$ sourcewise demand pairs in G . Then there exists a distance preserver of G, P of size*

$$O\left((nsp h)^{1/2} + n\right).$$

Proof. We defer the proof of Lemma 5.7 to Appendix B due to its similarities with the proof of Lemma 8 in [BV21]. $\square$

5.4 Main Construction

In the first part of our main construction, we will use Lemmas 5.2 and 5.7 to construct an exact distance preserver H that is sparse when $\hat{\ell}$ is small.

Construction of distance preserver H for small $\hat{\ell}$. Let $\pi(\cdot, \cdot)$ be a consistent tiebreaking scheme of G, P . We may assume without loss of generality that every edge in $E(G)$ is contained in a path in $\pi(P)$. Let φ be a parameter of the construction that we will optimize later. Our choice of φ will satisfy $\varphi \geq 10p/n^{1/2}$. While the average degree of G is at least φ , we will repeat the following process.

By Lemma 5.2 and Lemma 5.3, there exists a set of nodes $S \subseteq V(G)$ of size $|S| = \Theta(\varphi)$ such that:

- The set $Q \subseteq P$ of demand pairs (s, t) such that $\pi(s, t)$ contains a node in S is of size $|Q| = \Omega(\varphi^2)$.
- The set S has weak diameter $O(\hat{\ell} p^2 / (n \varphi^2))$ in G .

We will handle $\Theta(\varphi^2)$ paths in Q using two instances of the sourcewise distance preserver claimed in Lemma 5.7. For each demand pair $(x, y) \in Q$ such that $\pi(x, y)$ contains node $s \in S$, we split demand pair (x, y) into the two demand pairs (x, s) and (s, y) . We let $P_1 \subseteq S \times V$ and $P_2 \subseteq V \times S$ be the two resulting sets of demand pairs. Note that in order to construct a distance preserver of G, Q , it suffices to take the union of a distance preserver of G, P_1 and a distance preserver of G, P_2 . We can apply Lemma 5.7 with parameters $s_1 := |S| = \Theta(\varphi)$, $p_1 := \min(|Q|, \varphi^2) = \Theta(\varphi^2)$, and $h_1 := O(\hat{\ell} p^2 / (n \varphi^2))$ to construct a preserver H_1 of demand pairs Q of size

$$O\left((ns_1 p_1 h_1)^{1/2} + n\right) = O\left(n^{1/2} \varphi^{3/2} \cdot \frac{\hat{\ell}^{1/2} p}{n^{1/2} \varphi} + n\right) = O\left(\hat{\ell}^{1/2} \varphi^{1/2} p + n\right).$$

We add the edges of H_1 to our final distance preserver H of G, P . Since we have handled all demand pairs in Q , we can delete Q from P and repeat the above process. Formally, we update G, P as follows:

- $P \leftarrow P \setminus Q$, and
- $E(G) \leftarrow \bigcup_{(s, t) \in P} E(\pi(s, t))$.

We continue this process until the average degree of G is less than φ .

Since $p_1 = \Theta(\varphi^2)$, we will repeat this process at most $p/p_1 = O(p/\varphi^2)$ times before the average degree of H_1 is less than φ . This phase of the construction contributes at most

$$\frac{p}{p_1} \cdot O\left(\widehat{\ell}^{1/2} \varphi^{1/2} p + n\right) = O\left(\widehat{\ell}^{1/2} \cdot \frac{p^2}{\varphi^{3/2}} + \frac{np}{\varphi^2}\right)$$

edges to H . We finish the construction of H by adding all remaining $O(n\varphi)$ edges in H_1 to H . This yields the following lemma.

Lemma 5.8. *For every $\varphi \geq 10p/n^{1/2}$, there exists a distance preserver H of G, P with size*

$$|E(H)| = \tilde{O}\left(\widehat{\ell}^{1/2} \cdot \frac{p^2}{\varphi^{3/2}} + \frac{np}{\varphi^2} + n\varphi\right).$$

To complement the distance preserver of Lemma 5.8, we will construct a distance preserver that is sparse when $\widehat{\ell}$ is large.

Distance preserver H for large $\widehat{\ell}$. To preserve distances when $\widehat{\ell}$ is large, we will make use of a theorem of [BCE05], which states that we can preserve distances between vertices that are far apart in H using few edges.

Theorem 5.9 (Theorem 2.24 of [BCE05]). *Given a (possibly directed) unweighted graph G , a D -preserver is a subgraph H of G such that for all $s, t \in V(G)$ with $\text{dist}_G(s, t) \geq D$, we have that*

$$\text{dist}_H(s, t) = \text{dist}_G(s, t).$$

There exists a D -preserver H of G with size $O(n^2/\widehat{\ell})$.

As an immediate consequence of Theorem 5.9, we can obtain a distance preserver H of G, P with size $O(n^2/\widehat{\ell})$, since we can assume all shortest paths π in $\pi(P)$ are of length at least $|\pi| \geq \widehat{\ell}/2$ without loss of generality, by our earlier bucketing argument.

Lemma 5.10. *There exists a distance preserver H of G, P with size $O(n^2/\widehat{\ell})$.*

Finishing the proof of Theorem 5.1.

Claim 5.11. *Together, Lemmas 5.8 and 5.10 imply a distance preserver H of G, P of size*

$$|E(H)| = \tilde{O}\left(\frac{n^{2/3}p^{4/3}}{\varphi} + \frac{np}{\varphi^2} + n\varphi\right).$$

Proof. We will use the construction from Lemma 5.8 when $\widehat{\ell}$ is small, and we will use the construction from Lemma 5.10 when $\widehat{\ell}$ is large. Specifically, we define our threshold t for $\widehat{\ell}$ to be

$$t = \frac{n^{4/3}\varphi}{p^{4/3}}.$$

We calculate the following:

- When $\widehat{\ell} \leq t$,

$$\widehat{\ell}^{1/2} \cdot \frac{p^2}{\varphi^{3/2}} \leq t^{1/2} \cdot \frac{p^2}{\varphi^{3/2}} = O\left(\frac{n^{2/3}p^{4/3}}{\varphi}\right),$$

so the distance preserver H of Lemma 5.8 has size

$$|E(H)| = \tilde{O}\left(\frac{n^{2/3}p^{4/3}}{\varphi} + \frac{np}{\varphi^2} + n\varphi\right),$$

as desired.

- When $\widehat{\ell} > t$, distance preserver H of Lemma 5.10 has size

$$|E(H)| = \tilde{O}(n^2/\widehat{\ell}) \leq \tilde{O}(n^2/t) = \tilde{O}\left(\frac{n^{2/3}p^{4/3}}{\varphi}\right),$$

as desired. □

We are now ready to prove Theorem 5.1.

Proof of Theorem 5.1. The bounds stated in Theorem 5.1 will follow directly from the bounds achieved in Claim 5.11 after parameter balancing. Initially, we will assume that $n^{1/2} \leq p \leq n$.

Straightforward calculations show that the three terms in Claim 5.11 are balanced when

$$\varphi = \frac{p^{2/3}}{n^{1/6}}.$$

Recall that parameter φ must be at least $\varphi \geq 10p/n^{1/2}$ for Lemma 5.2 to go through. This holds for our choice of $\varphi = \frac{p^{2/3}}{n^{1/6}}$ when $p \leq n$. Plugging this value of φ into Claim 5.11, we obtain the upper bound

$$\tilde{O}\left(n^{5/6}p^{2/3} + \frac{n^{4/3}}{p^{1/3}} + n^{5/6}p^{2/3}\right).$$

We observe that $\frac{n^{4/3}}{p^{1/3}} \leq n^{5/6}p^{2/3}$ when $p \geq n^{1/2}$. We conclude that $|E(H)| = \tilde{O}(n^{5/6}p^{2/3})$ when $n^{1/2} \leq p \leq n$. For the settings where $p < n^{1/2}$ or $p > n$, we observe that the existing upper bounds of $\min(n^{2/3}p + n, np^{1/2})$ for directed distance preservers due to [CE06, Bod21] are smaller than our claimed bound of $\tilde{O}(n^{5/6}p^{2/3} + n)$. □

5.5 Proving the Dense Low-Diameter Cluster Lemma

The objective of this section is to prove Lemma 5.2.

Lemma (Dense Low-Diameter Cluster Lemma). Let G be an n -node directed unweighted graph with associated set of demand pairs $P \subseteq V(G) \times V(G)$ of size $|P| = p$. Let $\ell = \ell(G, P)$, $d = d(G, P)$, and $\widehat{\ell} = \widehat{\ell}(G, P)$. If G has average degree d , where $d \geq 10p/n^{1/2}$, then G contains a collection of nodes $S \subseteq V(G)$ such that

1. $|S| = \Theta(d)$,
2. (Dense.) For all $s \in S$, $\deg_G(s) = \Omega(d)$, and

3. (Low diameter.) For all $s, t \in S$, $\text{dist}_G(s, t) = O\left(\frac{\ell p^2}{nd^2}\right)$.

Before proving this lemma, we will need to introduce several new definitions, which will be useful for finding our dense low-diameter cluster. Let G, P, S and $\ell, d, \widehat{\ell}$ be as in the statement of Lemma 5.2.

Recall the definition of a branching event from Definition 3.3. We say that a branching event $(x, y), (x, z)$ is a *high-degree branching event* if $\text{outdeg}_G(x) \geq d/4$. A simple argument will show that G contains many high-degree branching events.

Claim 5.12 ([CE06]). *G contains $\Omega(nd^2)$ high-degree branching events.*

Proof. Since G has average degree d , it follows that

$$\sum_{v \in V(G)} \text{outdegree}_G(v) = nd/2.$$

In particular, we have that

$$\begin{aligned} & \sum_{\{v \in V(G) | \text{outdeg}_G(v) \geq d/4\}} \text{outdegree}_G(v) \\ &= \sum_{v \in V(G)} \text{outdegree}_G(v) - \sum_{\{v \in V(G) | \text{outdeg}_G(v) < d/4\}} \text{outdegree}_G(v) \\ &\geq nd/2 - nd/4 \geq nd/4. \end{aligned}$$

Then the number of high-degree branching events is at least

$$\begin{aligned} & \sum_{\{v \in V(G) | \text{outdeg}_G(v) \geq d/4\}} \text{outdegree}_G^2(v) \\ &\geq \frac{1}{n} \cdot \left(\sum_{\{v \in V(G) | \text{outdeg}_G(v) \geq d/4\}} \text{outdegree}_G(v) \right)^2 \\ &\geq \frac{1}{n} \cdot \frac{n^2 d^2}{16} = \Omega(nd^2), \end{aligned}$$

where the second to last inequality follows from Cauchy-Schwarz. $\square$

We will also introduce an auxiliary graph, called the path intersection graph, which will be useful in our proof of Lemma 5.14.

Definition 5.13 (Path Intersection Graph J). We define the vertex set of J to be

$$V(J) = \pi(P),$$

the set of all shortest paths associated with our tiebreaking scheme $\pi(\cdot, \cdot)$. For every $\pi_1, \pi_2 \in V(J)$, we add the edge (π_1, π_2) to J and we assign this edge weight equal to the number of high-degree branching events between π_1 and π_2 .

Lemma 5.14. *There exists a path $\pi^* \in \pi(P)$ and a collection of paths $\pi(Q)$, where $Q \subseteq P$, satisfying the following properties:*

- Every path $\pi' \in \pi(Q)$ has at least $\Omega(nd^2/p^2)$ high-degree branching events with π^* .

- Let $S \subseteq V(\pi^*)$ be the set of nodes in π^* with outdegree at least $d/4$, i.e.,

$$S = \{v \in V(\pi^*) \mid \text{outdeg}_G(v) \geq d/4\}.$$

The total number of high-degree branching events between path π^* and the paths in $\pi(Q)$ is at least $\Omega(|S|d)$.

Proof. Let b denote the number of high-degree branching events in G . By Claim 5.12, there is a constant $c > 0$ such that G has at least $b \geq cnd^2$ high-degree branching events. Then

$$\sum_{e \in E(J)} w(e) \geq b \geq cnd^2,$$

where $w(e)$ denotes the weight of edge $e \in E(J)$. We will perform the following procedure on J .

1. Delete all edges $e \in E(J)$ with weight

$$w(e) \leq \frac{cnd^2}{4p^2}.$$

2. While there exists a node $\pi \in V(J)$ such that sum of the edge weights incident to π is small, we will delete node π from $V(J)$. Formally, for each node $\pi \in V(J)$, let $S_\pi \subseteq V(\pi)$ be the set of nodes in path $\pi \in \pi(P)$ with outdegree at least $d/4$, i.e.,

$$S_\pi = \{v \in V(\pi) \mid \text{outdeg}_G(v) \geq d/4\}.$$

While there exists a node $\pi \in V(J)$ such that the sum of the edge weights incident to π in J is at most

$$\sum_{\{(\pi, \pi') \in E(J) \mid \pi' \in V(J)\}} w((\pi, \pi')) \leq \frac{|S_\pi|d}{16},$$

we will delete node π from J .

We claim that after the above procedure terminates, graph J is nonempty and still contains a constant fraction of its edge weight. After step 1 of the procedure terminates, J will have total edge weight at least

$$\sum_{e \in E(J)} w(e) \geq b - \frac{cnd^2}{4p^2} \cdot p^2 = b - cnd^2/4 = 3/4 \cdot b,$$

since $|E(J)| = p^2$ and $b \geq cnd^2$. Consequently, a large fraction of the edge weights in J survive step 1 of our procedure. Now we will attempt to argue the same is true for the second step of our procedure.

By the definition of high-degree branching events,

$$b \geq \sum_{\pi \in V(J)} |S_\pi| \cdot \frac{d}{4},$$

since there are at each node $v \in S_\pi$, there are at least $d/4$ high-degree branching events between π and other paths in $V(J)$. Then after step 2 of the procedure terminates, J will have total edge weight at least

$$\sum_{e \in E(J)} w(e) \geq 3/4 \cdot b - \sum_{\pi \in V(J)} \frac{|S_\pi|d}{16} = 3/4 \cdot b - \frac{1}{4} \cdot \left(\sum_{\pi \in V(J)} |S_\pi| \cdot \frac{d}{4} \right) \geq b/2.$$

Consequently, after our procedure terminates, the amount of edge weight remaining in $E(J)$ is at least $b/2$, so $E(J)$ and $V(J)$ must be nonempty.

We have shown that after this procedure terminates, graph J remains nonempty. In particular, we have shown that every surviving node in $V(J)$ is incident to edge weight at least $|S_\pi|d/16$, and every surviving edge in $E(J)$ has weight at least $cdn^2/(4p^2)$. Then we can take π^* to be any surviving node $\pi^* \in V(J)$, and we can take $\pi(Q)$ to be the set of all surviving paths adjacent to π^* in J . The lemma immediately follows. $\square$

We will denote the path described in Lemma 5.14 as π^* , and we will use Q to denote the associated set of demand pairs.

Definition 5.15 (π -close nodes). Fix a path $\pi \in \pi(Q)$, and let

$$B = \{((s_i, t_i), (s_i, u_i))\}_{i \in [1, k]}$$

be the set of all high-degree branching events between π and π^* , written so that

- $(s_i, t_i) \in \pi$ and $(s_i, u_i) \in \pi^*$, and
- If $i < j \in [1, k]$, then node s_i comes before node s_j on path π , and node s_j comes before node s_i on path π^* . Note that nodes s_i and s_j cannot appear in the same order on both π and π^* , by consistency.

We say that a pair of nodes $(x, y) \in V(\pi^*) \times V(\pi^*)$ is π -close if $(x, y) = (s_i, s_{i+1})$ for some $i \in [1, k-1]$.

Claim 5.16. *If nodes x and y are in π^* , and x and y are π_1 -close and π_2 -close for paths $\pi_1, \pi_2 \in \pi(Q)$, then $\pi_1 = \pi_2$.*

Proof. This claim will follow directly from our choice of a *consistent* tiebreaking scheme $\pi(\cdot, \cdot)$ and our definition of branching events. Let x and y be nodes in π^* that are π_1 -close and π_2 -close, for some $\pi_1, \pi_2 \in \pi(Q)$. We may assume wlog that node x comes before node y on path π^* . Since π_1 has branching events with π^* at nodes x and y , by consistency we must have that y comes before x on path π_1 . By an identical argument y comes before x on path π_2 . Then by consistency, it follows that $\pi_1[y, x] = \pi_2[y, x]$.

Let z be the node following y in paths π_1 and π_2 , so that $(y, z) \in \pi$. Let w be the node following y in path π^* . Observe that $(y, z), (y, w)$ is a branching event. Moreover, by the definition of π -close nodes, branching event $(y, z), (y, w)$ is associated with paths π_1, π^* and paths π_2, π^* . However, as stated in Definition 3.3, every branching event is associated with at most one pair of paths. We conclude that $\pi_1 = \pi_2$. $\square$

Definition 5.17 (close nodes). We say that nodes $x, y \in V$ are *close* if

$$\text{dist}_G(x, y) + \text{dist}_G(y, x) = O\left(\frac{\widehat{\ell}p^2}{nd^2}\right).$$

Claim 5.18. *Fix a path $\pi = \pi(s, t) \in \pi(Q)$, and let $C \subseteq V(\pi^*) \times V(\pi^*)$ be the set of all pairs of nodes in π^* that are π -close. Then at least $|C|/2$ node pairs $(x, y) \in C$ are close.*

Proof. Fix a path $\pi \in \pi(Q)$, and let

$$B = \{((s_i, t_i), (u_i, v_i))\}_{i \in [1, k]}$$

be the set of all high-degree branching events between π and π^* , written so that

- $(s_i, t_i) \in \pi$ and $(u_i, v_i) \in \pi^*$, and
- if $i < j \in [1, k]$, then node s_i comes before node s_j on path π , and node s_j comes before node s_i on path π^* .

Recall from Definition 5.15 that the set of all pairs of π -close nodes in π^* is exactly $C = \{(s_i, s_{i+1})\}_{i \in [1, k-1]}$, where $|C| = k-1$. Additionally, as we assumed at the beginning of Section 5, $|\pi| \leq 2\widehat{\ell}$ and $|\pi^*| \leq 2\widehat{\ell}$. In particular, this implies

$$\sum_{i=1}^{k-1} (\text{dist}_G(s_i, s_{i+1}) + \text{dist}_G(s_{i+1}, s_i)) \leq |\pi| + |\pi^*| \leq 4\widehat{\ell}.$$

Let $C_1 \subseteq C$ be the set of all π -close pairs of nodes (s_i, s_{i+1}) with roundtrip distance at least $16\widehat{\ell}/k$, i.e.,

$$C_1 = \{(s_i, s_{i+1}) \mid \text{dist}_G(s_i, s_{i+1}) + \text{dist}_G(s_{i+1}, s_i) \geq 16\widehat{\ell}/k\}.$$

By the previous inequality, we observe that

$$|C_1| \cdot 16\widehat{\ell}/k \leq |\pi| + |\pi^*| \leq 4\widehat{\ell},$$

so $|C_1| \leq k/4$. Let $C_2 = C \setminus C_1$. It follows that $|C_2| \geq |C|/2$. Additionally, by Lemma 5.14, we know that $|C| = k-1 = \Omega(nd^2/p^2)$. Then for all $(s_i, s_{i+1}) \in C_2$,

$$\text{dist}_G(s_i, s_{i+1}) + \text{dist}_G(s_{i+1}, s_i) \leq 16\widehat{\ell}/k = O\left(\frac{\widehat{\ell}p^2}{nd^2}\right).$$

We conclude that all pairs of nodes $(x, y) \in C_2$ are close, establishing the claim. $\square$

Claim 5.19. *Let $S \subseteq V(\pi^*)$ be the set of nodes v in π^* of degree at least $\deg_G(v) \geq d/4$. There are $\Omega(|S|d)$ distinct pairs of nodes in $S \times S$ that are close.*

Proof. Let $\pi(Q) = \{\pi_1, \dots, \pi_k\}$, for some integer k . For $i \in [1, k]$, let $C_i \subseteq V(\pi^*) \times V(\pi^*)$ be the set of all pairs of nodes in π^* that are π_i -close. Note that if $(x, y) \in C_i$, then $\deg_G(x) \geq d/4$ and $\deg_G(y) \geq d/4$ by Definition 5.15, so $(x, y) \in S \times S$. For $i \in [1, k]$, let $C'_i \subseteq C_i$ be the set of all pairs of nodes $(x, y) \in C_i$ that are close. We will need the following statements proved earlier.

- By Claim 5.16, $C_i \cap C_j = \emptyset$ for $i \neq j \in [1, k]$.
- By Claim 5.18, $|C'_i| \geq |C_i|/2$ for $i \in [1, k]$.
- By Lemma 5.14, the total number of high-degree branching events between path π^* and the paths in $\pi(Q)$ is at least $\Omega(|S|d)$. This implies that $\sum_{i \in [1, k]} |C_i| = \Omega(|S|d)$.

Putting it all together, the total number of distinct close pairs in $S \times S$ is at least

$$\left| \bigcup_{i \in [1, k]} C'_i \right| = \sum_{i \in [1, k]} |C'_i| \geq \frac{1}{2} \cdot \sum_{i \in [1, k]} |C_i| = \Omega(|S|d).$$

$\square$

We are now ready to prove Lemma 5.2.

Proof of Lemma 5.2. Let $S \subseteq V(\pi^*)$ be the set of nodes referenced in Claim 5.19. By Claim 5.19, each node $u \in S$ is close to $\Omega(d)$ distinct nodes in S on average. Then there must exist a node $v \in S$ that is close to $\Omega(d)$ distinct nodes in S . Let T be the set of nodes in S close to v . Observe that for all $x, y \in T$,

$$\text{dist}_G(x, y) \leq \text{dist}_G(x, v) + \text{dist}_G(v, y) = O\left(\frac{\widehat{\ell} p^2}{nd^2}\right).$$

Moreover, by our choice of set S in Claim 5.19, $\deg_G(x) = \Omega(d)$ for all $x \in S$. We conclude that set S satisfies all properties described in the statement of Lemma 5.2. $\square$

6 Directed Approximate Preserver Lower Bound

Theorem 6.1. *For any n and $\alpha = \alpha(n)$, there exist n -node directed weighted graphs G with demand pairs P of size $|P| = p$, with $p \in [n^{1/2}, n^2]$, such that any α -approximate preserver on G, P has at least*

$$\Omega(n^{2/3} p^{2/3})$$

edges.

Construction

Let $\mathcal{X}$ and $\mathcal{L}$ be an arrangement of $|\mathcal{X}| = n$ points and $|\mathcal{L}| = p$ distinct lines in $\mathbb{R}^2$ with $\varphi = \Omega(n^{2/3} p^{2/3})$ point-line incidences (e.g., take $\mathcal{X}$ and $\mathcal{L}$ to be the arrangement of points and lines Elekes' construction in [Ele01]). We may assume that every line in $\mathcal{L}$ has finite slope wlog (e.g., by rotating the arrangement of points and lines). Given $\mathcal{X}$ and $\mathcal{L}$, we will construct a directed graph G with associated demand pairs P . We choose the vertices of G to be the points in $\mathcal{X}$, i.e.,

$$V(G) = \mathcal{X}.$$

For each line $L \in \mathcal{L}$, we will add edges to $E(G)$ and demand pairs to P as follows.

- Let $x_1, \dots, x_k$ denote the points lying on L , ordered so that the first coordinates of points are strictly increasing.
- Add directed edge (x_i, x_{i+1}) to $E(G)$ for $i \in [1, k-1]$.
- Add demand pair (x_1, x_k) to P .

What remains is to assign weights to the edges of G . Let $L_1, \dots, L_p$ denote the lines in $\mathcal{L}$, ordered by increasing slope (breaking ties arbitrarily). If edge $(x, y) \in E(G)$ satisfies $x, y \in L_i$, then we assign edge (x, y) weight

$$w((x, y)) = (2\alpha n)^i.$$

Weight $w((x, y))$ is well-defined because distinct lines intersect on at most one point.

Analysis

Claim 6.2. *Fix a line $L \in \mathcal{L}$. Let $x_1, \dots, x_k$ denote the points lying on L , ordered by increasing first coordinate. Then $(x_1, \dots, x_k)$ is the unique α -approximate shortest path between x_1 and x_k in G .*

Proof. Let $L = L_i$ for some $i \in [1, p]$. Let $\pi = (x_1, \dots, x_k)$ denote the (x_1, x_k) -path in G traveling along line L . Observe that

$$w(\pi) = \sum_{j=1}^{k-1} w((x_j, x_{j+1})) = k \cdot (2\alpha n)^i \leq 2^i \alpha^i n^{i+1}.$$

Let π' be an (x_1, x_k) -path such that $\pi' \neq P$. Let $(x, y) \in E(G)$ be the first edge on path π' such that $(x, y) \in E(\pi') \setminus E(\pi)$. If $x, y \in L_j$ for some $j > i$, then

$$w(\pi') \geq w((x, y)) = (2\alpha)^j n^j \geq (2\alpha)^{i+1} n^{i+1} > \alpha \cdot w(\pi),$$

so π' is not an α -approximate shortest path.

Otherwise, $x, y \in L_j$ for some $j < i$. Consider the line L' passing through points y and x_k . Since $x, y \in L_j$ for some $j < i$, the slope of line L' is strictly greater than the slope of line L_i . Additionally, by construction, the first coordinates of the nodes on path π' are strictly increasing. Then the subpath $\pi'[y, x_k]$ of π' must contain an edge $(x', y') \in E(G)$ such that $x', y' \in L_{j'}$ for some $j' > i$. We conclude that π' is not an α -approximate shortest path. $\square$

Claim 6.3. *Any α -approximate preserver on G, P has at least $\varphi - p = \Omega(n^{2/3} p^{2/3})$ edges.*

Proof. By Claim 6.2, every edge $(x, y) \in E(G)$ lies on a unique α -approximate shortest path between a demand pair $(s, t) \in P$. Then any α -approximate preserver of G must contain every edge in $E(G)$, and must be of size at least

$$|E(G)| = \sum_{L \in \mathcal{L}} |\{x \in \mathcal{X} \mid x \in L\}| - 1 = \varphi - p.$$

$\square$

7 Directed Approximate Hopset Lower Bound

Theorem 7.1. *For any n and $\alpha = \alpha(n)$, there exists an n -node directed weighted graph G such that any $O(n)$ -size α -approximate hopset H of G reduces the hopbound of $G \cup H$ to $\Omega(n^{1/2})$.*

Our directed approximate hopset lower bound construction is based on the undirected exact hopset lower bound of [BH23], but will use a weighting scheme similar to our approximate preserver lower bound in Section 6. We will assume wlog that n is a power of 4.

Vertex Set V . Our vertex set V will correspond to integer points in $\mathbb{Z}^2$. Specifically, we take V to be

$$V = [\sqrt{n}/2] \times [2\sqrt{n}].$$

(Recall that $\sqrt{n}$ is a power of 2, since n is a power of 4.) For each $i \in [\sqrt{n}/2]$, let $L_i = \{(i, y) \in \mathbb{Z}^2 \mid (i, y) \in V\}$ denote the nodes in V with first coordinate i .

Edge Set E . For each node $(x, y) \in V$, we add the directed edges

$$((x, y), (x + 1, y)) \text{ and } ((x, y), (x + 1, y + 1))$$

to E , if they are contained in $V \times V$. We will now assign edge weights to edges in E .

Our weighting scheme will rely on a permutation q on $\llbracket \sqrt{n}/2 \rrbracket$ that was used in [VXX24]. For any integer $i \in \llbracket \sqrt{n}/2 \rrbracket$, q_i is defined as the integer whose $(\log(\sqrt{n}/2))$ -bit binary representation is the reversal of the $(\log(\sqrt{n}/2))$ -bit binary representation of i . For instance $q_0 = 0$ and $q_1 = \sqrt{n}/4$ (recall $\sqrt{n}$ is a power of 2).

For each edge e in $L_i \times L_{i+1}$, $i \in [\sqrt{n} - 1]$:

- If $e = ((x, y), (x + 1, y))$ for some $(x, y) \in L_i$, then assign edge e weight

$$w(e) = 1.$$

- Otherwise, $e = ((x, y), (x + 1, y + 1))$ for some $(x, y) \in L_i$. Assign edge e weight

$$w(e) = (2\alpha n)^{q_i}.$$

Critical Paths P . We will define a collection of directed paths in G with certain properties that will guarantee that G implies an α -approximate hopset lower bound.

For each node $v \in [1] \times [\sqrt{n}]$ and each $d \in [\sqrt{n}]$, we iteratively define a critical path π as follows. The starting node of path π is v . If path π contains its i th node v_i in layer L_i , then we define its $(i + 1)$ th node v_{i+1} in layer L_{i+1} as follows:

- If $d \leq q_i$, then let

$$v_{i+1} = v_i + (1, 0),$$

if $v_{i+1} \in V$ (in fact, this condition always holds).

- If $d > q_i$, then let

$$v_{i+1} = v_i + (1, 1),$$

if $v_{i+1} \in V$ (in fact, this condition always holds).

The path π begins at node $v_1 = v \in L_1$, and terminates at node $v_{\sqrt{n}/2} \in L_{\sqrt{n}/2}$. We add π to our collection of paths P .

Lemma 7.2. *The above graph has the following properties:*

1. $|V| = n$, $|E| = \Theta(n)$, and $|P| = \Theta(n)$.
2. Every path $\pi \in P$ contains exactly one node in each layer L_i , $i \in [\sqrt{n}/2]$.
3. Every path $\pi \in P$ is the unique α -approximate shortest path between its endpoints.
4. For any k distinct paths $\pi_1, \dots, \pi_k \in P$, where $k \in [|P|]$, the intersection $\cap_{i=1}^k \pi_i$ has size at most

$$\left| \bigcap_{i=1}^k \pi_i \right| \leq \frac{2\sqrt{n}}{k}.$$

Proof. We prove each of these items below.

1. The size bounds on $|V|$ and $|E|$ follow directly from the construction. In our construction of critical paths P , we add $\Theta(n)$ paths to P , but it is not immediate that these paths are distinct. Note that if two paths have distinct parameters $d_1 < d_2$ respectively, they will use different edges between L_j and L_{j+1} , for $q_j = d_1$. If $d_1 = d_2$, then the two paths begin at different vertices on L_1 . Thus, $|P| = \Theta(n)$.
2. Each path $\pi \in P$ begins at a vertex $v_1 = (1, y) \in [1] \times [\sqrt{n}]$. Each edge in π corresponds to a vector of form $(1, 0)$ or $(1, 1)$ in G . Then the i th node v_i of π (if it exists) will satisfy

$$v_i \in \{i\} \times [y, y + i] \subseteq \{i\} \times [\sqrt{n}, \sqrt{n} + i] \subseteq V,$$

for $i \in [\sqrt{n}/2]$. We conclude that π contains a node in each layer L_i , $i \in [1, \sqrt{n}/2]$. Moreover, π can contain at most one node in each layer L_i because every edge in G is a directed edge in $L_i \times L_{i+1}$.

3. Fix a path $\pi \in P$ with start node $s \in L_1$ and end node $t \in L_{\sqrt{n}/2}$. Let $d \in [\sqrt{n}]$ be the parameter used to construct π . Consider an $s \rightsquigarrow t$ path π' , where $\pi' \neq \pi$. Suppose that π' contains an edge $e' \in L_i \times L_{i+1}$ corresponding to vector $(1, 1)$, and $d \leq q_i$. Then

$$w(\pi') \geq w(e') = (2\alpha n)^{q_i} \geq (2\alpha n)^d \geq 2\alpha \cdot (n \cdot (2\alpha n)^{d-1}) \geq 2\alpha \cdot w(\pi),$$

where the final inequality follows from the fact that every edge in π has weight at most $(2\alpha n)^{d-1}$. Then path π' is not an α -approximate shortest $s \rightsquigarrow t$ path.

This implies that any α -approximate shortest $s \rightsquigarrow t$ path π' takes an edge $e' \in L_i \times L_{i+1}$ corresponding to vector $(1, 1)$, only if path π takes an edge $e \in L_i \times L_{i+1}$ corresponding to vector $(1, 1)$. Moreover, since π and π' are both $s \rightsquigarrow t$ paths, they must take exactly the same number of edges corresponding to vector $(1, 1)$. We conclude that path π' takes an edge $e' \in L_i \times L_{i+1}$ corresponding to vector $(1, 1)$ if and only if path π takes an edge $e \in L_i \times L_{i+1}$ corresponding to vector $(1, 1)$. Then $\pi = \pi'$, completing the proof.

4. Let $d_1, \dots, d_k$ be the parameters used to construct $\pi_1, \dots, \pi_k$. First, observe that two distinct paths sharing the same parameter d cannot intersect, so all values $d_1, \dots, d_k$ are distinct. wlog, we assume $d_1 \leq d_k - k + 1$. Second, by Item 3, all paths $\pi_1, \dots, \pi_k$ are unique shortest paths, so $\cap_{i=1}^k \pi_i$ must be a contiguous path.

Suppose for some $i \in [\sqrt{n}/2 - 1]$, $d_1 \leq q_i < d_k$, then π_1 and π_k cannot use the same edge between L_i and L_{i+1} , implying $\cap_{i=1}^k \pi_i$ also does not contain such an edge. The interval $[d_1, d_k]$ must contain an interval of the form $\mathcal{I} = [a \cdot 2^b, (a+1) \cdot 2^b]$ (a dyadic interval) where $2^b \geq k/4$. By construction of the permutation q , for all values of i whose last $(\log((\sqrt{n}/2)) - b)$ -bit binary representation is the reversal of a , q_i falls in the interval $\mathcal{I}$, and hence $d_1 \leq q_i < d_k$. Regardless where $\cap_{i=1}^k \pi_i$ begins, the next such i is within $2^{\log((\sqrt{n}/2)) - b} \leq \frac{2\sqrt{n}}{k}$ layers away, at which point $\cap_{i=1}^k \pi_i$ must end. Therefore, $|\cap_{i=1}^k \pi_i| \leq \frac{2\sqrt{n}}{k}$.

□

Using the hopset lower bound argument of [BH23], we can use Lemma 7.2 to argue that any $0.01n$ -size α -approximate hopset H of G has hopbound $\Omega(\sqrt{n})$. In order to extend this lower bound to any hopset H of size $|H| \leq cn$ for an arbitrarily large constant $c > 0$, we will need to modify our construction of G and P , as specified in the following lemma.

Lemma 7.3. *For any $n, c > 0$, there exists a directed weighted graph $G' = (V', E')$ and associated collection of paths P' satisfying the following properties:*

1. $|V'| = \Theta(n/c)$, and $|P'| = \Theta(n)$.
2. Every path $\pi \in P'$ is of length $\Theta(n/c)$.
3. Every path $\pi \in P'$ is the unique α -approximate shortest path between its endpoints.
4. For any k distinct paths $\pi_1, \dots, \pi_k \in P$, where $k \in [|P|]$, the intersection $\cap_{i=1}^k \pi_i$ has size at most

$$\left| \bigcap_{i=1}^k \pi_i \right| \leq \frac{2\sqrt{n}}{ck}.$$

Proof. We will first construct G', P' and then analyze them.

Construction. We will choose V' to be a subset of V . Specifically, we define V' to contain every ci -th layer L_{ci} of G for $i \in [\lfloor \sqrt{n}/(2c) \rfloor]$, i.e.,

$$V' = L_c \cup L_{2c} \cup \dots \cup L_{\lfloor \sqrt{n}/(2c) \rfloor \cdot c}.$$

For each path $\pi \in P$ and each index $i \in [\lfloor \sqrt{n}/(2c) \rfloor - 1]$, we will add an edge to E' as follows.

- Note that π contains exactly one node in layer L_{ci} and exactly one node in layer $L_{c(i+1)}$ by Property 2 of Lemma 7.2. Let $u = \pi \cap L_{ci}$, let $v = \pi \cap L_{c(i+1)}$.
- Add edge (u, v) to E' , and assign (u, v) weight equal to the distance in G from u to v , i.e.,

$$w((u, v)) = \text{dist}_G(u, v).$$

For each path $\pi \in P$, we define a path π' in G' by removing all nodes in $\pi \setminus V'$ from π and taking π' to be the resulting path. Path π' is well-defined in G' by our construction of E' .

Analysis. Items 1 to 3 follow directly from Lemma 7.2 and the observation that for all $u, v \in V'$, $\text{dist}_{G'}(u, v) = \text{dist}_G(u, v)$. To prove Item 4, note that if k paths $\pi'_1, \dots, \pi'_k \in P'$ have an intersection of size $|\cap_i \pi'_i| > \frac{2\sqrt{n}}{ck}$, then there exist k paths $\pi_1, \dots, \pi_k \in P$ with an intersection of size $|\cap_i \pi_i| > \frac{2\sqrt{n}}{ck}$, contradicting Item 4 of Lemma 7.2. $\square$

We are now ready to finish the proof of Theorem 7.1.

Proof of Theorem 7.1. Let graph $G' = (V', E')$ and paths P' be constructed with parameters $n, c > 0$. We will show that for any hopset H of size $|H| \leq n = O(c|V'|)$, $G' \cup H$ has hopbound $\Omega(\sqrt{n}/c)$, completing the proof. We use $\text{hopdist}_{G'}(u, v)$ to denote the α -approximate hopdistance from u to v in G' , and for each $\pi \in P'$, we let s_π and t_π denote its start and end node, respectively.

Similar to the proof of [BH23], we define a potential function with respect to any hopset H (with a slight abuse of notation, we use $G' \cup H$ to denote $(V', E' \cup H)$ in the following):

$$\Phi(H) = \sum_{\pi \in P'} \text{hopdist}_{G' \cup H}(s_\pi, t_\pi).$$

Initially, $\Phi(\emptyset) = |P'| \cdot \Theta(\sqrt{n}/c) = \Theta(n^{3/2}/c)$, since every path $\pi \in P'$ is a unique α -approximate shortest path between its endpoints in G' and has hopbound $|\pi| = \Theta(\sqrt{n}/c)$. We will now upper bound the quantity $\Phi(H) - \Phi(H \cup \{(u, v)\})$ for all $(u, v) \in V' \times V'$ in the transitive closure of

G' . This quantity corresponds to the decrease in potential after adding one additional edge to our hopset.

Observe that the hopset edge (u, v) only decreases the hoplength of a path $\pi \in P$ if $u, v \in \pi$. Let $\pi_1, \dots, \pi_k \in P'$ be the paths in P' containing nodes u and v . Then $\text{hopdist}_{G'}(u, v) \leq \frac{2\sqrt{n}}{ck}$, by Item 4 of Lemma 7.3. It follows that

$$\Phi(H) - \Phi(H \cup \{(u, v)\}) \leq k \cdot \frac{2\sqrt{n}}{ck} = O(\sqrt{n}/c).$$

Then for any hopset H of size $|H| \leq n$, it follows that

$$\Phi(H) \geq \Phi(\emptyset) - n \cdot O(\sqrt{n}/c) = \Omega(n^{3/2}/c).$$

Then by an averaging argument, there is a path $\pi \in P'$ such that the α -approximate hopdistance between s_π and t_π in $G' \cup H$ is $\Phi(H)/|P'| = \Omega(\sqrt{n}/c)$. We conclude that the α -approximate hopbound of $G' \cup H$ is $\Omega(\sqrt{n}/c)$ as desired. $\square$

8 $O(m)$ -size Shortcut Set Lower Bound

Theorem 8.1. *There exist n -node directed graphs G with $|E(G)| = m$ edges, such that any $O(m)$ -size shortcut set H of G reduces the diameter of $G \cup H$ to $\Omega(n^{2/9})$.*

Our construction is parameterized by two parameters r and c , where r is a parameter controlling the size of the graph already used in [VXX24] and c intuitively denotes the number of copies of their construction we use. We assume r is a power of 2 wlog.

Vertex Set V . Our graph has $2cr + 1$ layers, indexed by integers between 0 and $2cr$. Each layer contains vertices from $[4cr] \times [4cr] \times [4cr^2]$, which can be viewed as points in the 3D grid. For $i \in \llbracket 2cr + 1 \rrbracket$ and $\mathbf{x} \in [4cr] \times [4cr] \times [4cr^2]$, we also index the vertex with coordinates $\mathbf{x}$ in the i -th layer by $(i, \mathbf{x})$.

Edge Set E . Similar to Section 7, we also use a permutation q on $\llbracket r \rrbracket$ that was used in [VXX24]. For any integer $i \in \llbracket r \rrbracket$, q_i is defined as the integer whose $(\log r)$ -bit binary representation is the reversal of the $(\log r)$ -bit binary representation of i . For instance $q_0 = 0$ and $q_1 = r/2$ (recall r is a power of 2). For every even $i \in \llbracket 2cr \rrbracket$, $\mathbf{w}_i = (1, 0, q_{(i/2) \bmod r})$; for every odd $i \in \llbracket 2cr \rrbracket$, $\mathbf{w}_i = (0, 1, q_{((i-1)/2) \bmod r})$. Note that the sequence of vectors $\{\mathbf{w}_i\}_i$ has period $2r$.

Then for every $i \in \llbracket 2cr \rrbracket$, and every point $\mathbf{x} \in [4cr] \times [4cr] \times [4cr^2]$, we add an edge from $(i, \mathbf{x})$ to $(i + 1, \mathbf{x})$, and an edge from $(i, \mathbf{x})$ to $(i + 1, \mathbf{x} + \mathbf{w}_i)$ if $\mathbf{x} + \mathbf{w}_i$ stays inside the grid.

Critical Paths P . We define a set of paths in the graph, which has certain desired properties. Eventually, we will argue that at least one of these paths will have long distances between its endpoints even if we add $O(m)$ shortcut edges to the graph.

For every $\mathbf{x} \in [2cr] \times [2cr] \times [2cr^2]$ and every $(d_1, d_2) \in [r] \times [r]$, we define a critical path as follows. The start node of the path is $(0, \mathbf{x})$, and the path travels from layer 0 to layer $2cr$. Between layer i and layer $i + 1$, there are two potential edges for this path to travel, one corresponding to $\mathbf{w}_i$, and one corresponding to $\mathbf{0} = (0, 0, 0)$. The path takes the edge corresponding to $\mathbf{w}_i$ if $\mathbf{w}_i \in (\{1\} \times \{0\} \times \llbracket d_1 \rrbracket) \cup (\{0\} \times \{1\} \times \llbracket d_2 \rrbracket)$ (i.e., if $\mathbf{w}_i$ is in the form of $(1, 0, z)$, then we require $z < d_1$; otherwise, $\mathbf{w}_i$ is in the form of $(0, 1, z)$, and we require $z < d_2$), otherwise, the path takes the edge corresponding to $\mathbf{0}$. It is not difficult to verify that the critical paths defined this way do not fall out of the grid.

Lemma 8.2. *The above graph has the following properties:*

1. *The number of vertices $|V| = \Theta(c^4 r^5)$, the number of edges $|E| = \Theta(c^4 r^5)$, the number of critical paths $|P| = \Theta(c^3 r^6)$.*
2. *Every path $\pi \in P$ is the unique path from its start node to its end node.*
3. *For any path σ in the graph of length $g \leq 2r$, there are $O((r/g)^2)$ critical paths in P containing σ .*

Proof. We prove each of the items below.

1. These bounds are elementary.
2. Fix an arbitrary critical path π from $(0, \mathbf{x})$ to $(2cr, \mathbf{y})$, associated with $(d_1, d_2) \in [r] \times [r]$. Each path from $(0, \mathbf{x})$ to $(2cr, \mathbf{y})$ corresponds to a unique subset S of $\llbracket 2cr \rrbracket$, i.e., if $i \in S$, the path uses the edge corresponding to $\mathbf{w}_i$ between layer i and layer $i+1$ and otherwise it uses the edge corresponding to $\mathbf{0}$. Thus, it suffices to argue that there is a unique subset S of $\llbracket 2cr \rrbracket$ so that $\sum_{i \in S} \mathbf{w}_i = \mathbf{y} - \mathbf{x}$.

We have c copies of the vector $(1, 0, z)$ for $z \in \llbracket r \rrbracket$, and c copies of the vector $(0, 1, z)$ for $z \in \llbracket r \rrbracket$. By construction, $\mathbf{y} - \mathbf{x} = (cd_1, cd_2, c \cdot \frac{d_1(d_1-1)}{2} + c \cdot \frac{d_2(d_2-1)}{2})$. In order for any subset of $\{\mathbf{w}_i\}_i$ to be equal to $\mathbf{y} - \mathbf{x}$, we must pick exactly cd_1 vectors of the form $(1, 0, z)$, and cd_2 vectors of the form $(0, 1, z)$. Also, $c \cdot \frac{d_1(d_1-1)}{2} + c \cdot \frac{d_2(d_2-1)}{2}$ is the smallest sum of obtainable by picking such a subset of vectors, and it is simple to see that no other subset can achieve the same sum.

3. First of all, it is simple to see that two critical paths with the same associated values of (d_1, d_2) do not share any vertices. Therefore, it suffices to bound the number of distinct (d_1, d_2) so that there is a critical path associated with (d_1, d_2) that uses σ . We additionally assume $g \geq 2$, as otherwise, the claimed bound is simply $O(r^2)$, which is the total number of distinct (d_1, d_2) .

Now consider a path π associated with value (d_1, d_2) and another path π' associated with value (d'_1, d'_2) where $|d_1 - d'_1| > 16r/g$. We will show that π and π' cannot both contain σ as a subpath. First, because σ has length g , it must travel at least $(g-1)/2$ layers where the corresponding vector w_i has the form $(1, 0, z)$ for $z \in \llbracket r \rrbracket$, where z can take value $q_{j \bmod r}$ for j in an interval $\mathcal{I}$ of length at least $(g-1)/2$. This interval must contain a dyadic interval $\mathcal{I}'$ of length at least $(g-1)/8$ (dyadic intervals are intervals of the form $[a \cdot 2^k, (a+1) \cdot 2^k)$ for nonnegative integers a, k). Say $\mathcal{I}' = [a \cdot 2^k, (a+1) \cdot 2^k)$ for some nonnegative integers a, k . Then $\{q_{i \bmod r}\}_{i \in \mathcal{I}'}$ contains all numbers from $\llbracket r \rrbracket$ that is congruent to $\hat{a}$ modulo $\frac{r}{2^k}$, where $\hat{a}$ is the $\log(\frac{r}{2^k})$ -bit reversal of a . Now, suppose for the sake of contradiction that $|d_1 - d'_1| > 16r/g$. Then because $2^k = |\mathcal{I}'| \geq (g-1)/8 \geq g/16$, we must have $|d_1 - d'_1| > \frac{r}{2^k}$. Then there must exist some $z = q_i$ for some $i \in \mathcal{I}'$ such that $z \in [\min\{d_1, d'_1\}, \max\{d_1, d'_1\})$. By construction, exactly one of π and π' uses the edge corresponding to $\mathbf{w}_i = (1, 0, z)$ and the other uses the edge corresponding to $\mathbf{0}$, so they cannot both contain σ as a subpath. This is a contradiction, so we must have $|d_1 - d'_1| \leq 16r/g$.

Similarly, we can show that $|d_2 - d'_2| \leq 16r/g$. Thus, if a path π' with value (d'_1, d'_2) needs to contain σ together with the path π with value (d_1, d_2) , we must have $|d_1 - d'_1| \leq 16r/g$ and $|d_2 - d'_2| \leq 16r/g$, so the number of such distinct values (d'_1, d'_2) is $O((r/g)^2)$. By previous discussions, this imply that the number of critical paths containing σ is $O((r/g)^2)$.

□

Proof of Theorem 8.1. We consider the graph $G = (V, E)$ with the critical paths P constructed above.

We let E_Δ be the set of shortcut edges in G that connects every pair of vertices (u, v) where u can reach v and u and v are at most Δ layers away for some parameter Δ . For any critical path $\pi \in P$, we use s_π to denote its start node and t_π to denote its end node.

We define a potential function with respect to any shortcut set H :

$$\Phi(H) := \sum_{\pi \in P} \text{dist}_{G \cup E_\Delta \cup H}(s_\pi, t_\pi).$$

Initially, we clearly have $\Phi(\emptyset) = |P| \cdot cr/\Delta = \Theta(c^4 r^7 / \Delta)$.

Now we upper bound $\Phi(H) - \Phi(H \cup \{e\})$, i.e., the potential drop after we add one additional edge to the shortcut set. There are several cases to consider:

- e connects two vertices that are at most Δ layers apart. In this case, $E \cup E_\Delta \cup H = E \cup E_\Delta \cup H \cup \{e\}$, as $e \in E_\Delta$, so $\Phi(H) - \Phi(H \cup \{e\}) = 0$.
- e connects two vertices u, v that are $g > \Delta$ layers apart. Consider any path $\pi \in P$ that contains both u and v . It is not difficult to see that $\text{dist}_{G \cup E_\Delta \cup H}(s_\pi, t_\pi) - \text{dist}_{G \cup E_\Delta \cup H \cup \{e\}}(s_\pi, t_\pi) \leq g/\Delta$. As there is a unique path in G from s_π to t_π by Lemma 8.2 Item 2, there must also be a unique path σ in G from u to v . Thus, by Lemma 8.2 Item 3, there are $O(1 + (r/g)^2)$ critical paths in P containing σ , which are exactly the critical paths that contain both u and v . Therefore, the total potential drop in this case is

$$O(1 + (r/g)^2) \cdot \frac{g}{\Delta} = O\left(\frac{g}{\Delta} + \frac{r^2}{g\Delta}\right),$$

which is always upper bounded by

$$O\left(\frac{cr}{\Delta} + \frac{r^2}{\Delta^2}\right)$$

for $g \in (\Delta, 2cr]$.

Therefore, if $|H| = \varepsilon_1 \min\{c^3 r^6, c^4 r^5 \Delta\}$ for some sufficiently small constant $\varepsilon_1 > 0$, total potential drop $\Phi(\emptyset) - \Phi(H)$ would be bounded by $\varepsilon_2 c^4 r^7 / \Delta$ for some small constant ε_2 . Furthermore, the parameter ε_2 can be made arbitrarily small depending on ε_1 , and when it is sufficiently small, we will have $\Phi(H) = \Theta(c^4 r^7 / \Delta)$. In this case, by averaging, there must exist a path $\pi \in P$ where $\text{dist}_{G \cup H}(s_\pi, t_\pi) \geq \text{dist}_{G \cup E_\Delta \cup H}(s_\pi, t_\pi) = \Theta(cr/\Delta)$, i.e., the diameter of the graph is at least $\Theta(cr/\Delta)$ after adding a shortcut set of size $|H| = \varepsilon_1 \min\{c^3 r^6, c^4 r^5 \Delta\}$.

By setting $c = r/B$ for some large constant B , we get that $|V|, |E| = \Theta(r^9/B^4)$ by Lemma 8.2 Item 1, and the number of edges allowed in the shortcut set is $\varepsilon_1 \min\{r^9/B^3, r^9\Delta/B^4\}$, which can be any constant factor larger than $|E|$ by setting B and Δ big enough. The diameter bound will be $\Theta(r^2/B\Delta)$, which is $\Theta(|V|^{2/9})$ for constant values of B and Δ .

□

9 $O(m)$ -size Exact Hopset Lower Bound in Unweighted Graphs

The following construction is similar to the construction in Section 8, and the main difference is that the construction in Section 8 is a layered graph, while the construction in this section is

an unlayered graph. In comparison, we are able to utilize an unlayered graph construction in this section because it suffices to ensure unique shortest paths have good structures in a construction for hopset lower bound, while the construction in Section 8 requires unique paths to have good structures. The unlayered construction complicates the analysis, but leads to an improved lower bound.

Theorem 9.1. *There exist n -node undirected unweighted graphs G with $|E(G)| = m$ edges, such that any $O(m)$ -size exact hopset H of G reduces the hopbound of $G \cup H$ to $\Omega(n^{2/7})$.*

Similar as before, there are two parameters r and c , and we assume r is a power of 2.

Vertex Set V . The important vertices in the graph are indexed by $\llbracket 2cr + 1 \rrbracket \times [4cr] \times [3cr^2]$. We will say that a vertex (i, j, k) is on the i -th layer. There might also be other (unimportant) vertices defined in the following.

Edge Set E . Recall q is a permutation on $\llbracket r \rrbracket$. For $i \in \llbracket 2cr \rrbracket, j \in [4cr], k \in [3cr^2]$, we add the following edges (if some edges fall out of the grid, then we do not add them):

- An edge between (i, j, k) and $(i + 1, j, k)$;
- A chain of two edges between (i, j, k) and $(i + 1, j, k + q_{(i/2) \bmod r})$ if i is even (this creates a new vertex).
- A chain of two edges between (i, j, k) and $(i + 1, j + 1, k + q_{((i-1)/2) \bmod r})$ if i is odd (this creates a new vertex).

Critical Paths P . Each critical path is associated with a start node $(0, j, k)$ for some $j \in [cr]$ and $k \in [cr^2]$, together with a pair of values $(d_1, d_2) \in [r] \times [r]$ where $r/2 < d_2 < d_1$. When at a vertex (i, j', k') for some j', k' and even i , the path will travel to $(i + 1, j', k' + q_{(i/2) \bmod r})$ (using a chain of two edges) if $q_{(i/2) \bmod r} \geq d_1$; otherwise, it will travel to $(i + 1, j', k')$ (using one edge). When at a vertex (i, j', k') for some j', k' and odd i , the path will travel to $(i + 1, j' + 1, k' + q_{((i-1)/2) \bmod r})$ (using a chain of two edges) if $q_{((i-1)/2) \bmod r} \geq d_2$; otherwise, it will travel to $(i + 1, j', k')$ (using one edge). The path stops when it reaches some vertex on the $(2cr)$ -th layer. It is not difficult to check that the above rules uniquely determine the path.

Lemma 9.2. *The above graph has the following properties:*

1. *The number of vertices $|V| = \Theta(c^3 r^4)$, the number of edges $|E| = \Theta(c^3 r^4)$, the number of critical paths $|P| = \Theta(c^2 r^5)$.*
2. *Every path $\pi \in P$ is the unique shortest path from its start node to its end node.*
3. *For any shortest path σ in the graph of length $g \leq 2r$, there are $O((r/g)^2)$ critical paths in P containing σ .*

Proof. We consider the three items separately:

- These bounds follow straightforwardly from the construction.

- Fix any $\pi \in P$ with start node $s_\pi = (0, j, k)$, end node t_π and associated values (d_1, d_2) . We will argue that there is no other path from s_π to t_π using fewer or same number of edges.

First, if a path only uses one of the two edges in the chains of two edges we added, and then travels back, then it can never be a shortest path. Thus, for the purpose of this proof, we can view the chains of two edges as single edges with cost 2.

For any path travelling from s_π to t_π and for any fixed $i \in \llbracket 2cr \rrbracket$, the number of times it crosses from layer i (the set of vertices whose first coordinate equal i) to layer $i + 1$ must be exactly one more than the number of times it crosses from layer $i + 1$ back to layer i . As a result, we can pair up each edge going from layer i to layer $i + 1$ with an edge going from layer $i + 1$ to layer i , so that after the pairing up, we are left with one edge going from layer i to layer $i + 1$. Each of these set of edges (either a pair of two edges or a single edge left at the end) changes the coordinate of the current vertex, i.e., add a vector to the coordinate of the vertex, and we consider the average (over the total costs of edges in the set) dot product between the vector and $\delta = (1, 1 - \frac{d_2 - 0.5}{d_1 - 0.5}, \frac{1}{d_1 - 0.5})$.

1. A single forward edge:

- (1, 0, 0): the dot product between it and δ is 1.
- (1, 0, x): the dot product between it and δ is $1 + \frac{x}{d_1 - 0.5}$, and the average per cost is $(1 + \frac{x}{d_1 - 0.5})/2$, which is greater than 1 when $x \geq d_1$, and is less than 1 when $x < d_1$.
- (1, 1, y): the dot product between it and δ is $2 - \frac{d_2 - 0.5}{d_1 - 0.5} + \frac{y}{d_1 - 0.5}$, and the average per cost is $(2 + \frac{y - d_2 + 0.5}{d_1 - 0.5})/2$, which is greater than 1 when $y \geq d_2$, and is less than 1 when $y < d_2$.

2. One arbitrary forward edge and one arbitrary backward edge:

- these two edges are the negations of each other: in this case, the sum of the two vectors corresponding to the two edges is 0, and thus the average dot product between it and δ per cost is 0, which is less than 1.
- (1, 0, 0) – (1, 0, x): in this case, the total cost of the two edges is 3, so the average dot product between it and δ per cost is $(-\frac{x}{d_1 - 0.5})/3 < 0 < 1$.
- (1, 0, x) – (1, 0, 0): in this case, the total cost of the two edges is 3, so the average dot product between it and δ per cost is $(\frac{x}{d_1 - 0.5})/3$. Recall $d_1 > r/2$, so $d_1 - 0.5 \geq r/2$, which implies $(\frac{x}{d_1 - 0.5})/3 \leq (\frac{r}{r/2})/3 < 1$.
- (1, 0, 0) – (1, 1, y): in this case, the total cost of the two edges is 3, so the average dot product between it and δ per cost is $(-1 + \frac{d_2 - 0.5}{d_1 - 0.5} - \frac{y}{d_1 - 0.5})/3 < 0 < 1$ (recall $d_2 < d_1$).
- (1, 1, y) – (1, 0, 0): in this case, the total cost of the two edges is 3, so the average dot product between it and δ per cost is $(\frac{y + d_2 - 0.5}{d_1 - 0.5})/3$. Recall $d_1 > r/2$ so $d_1 - 0.5 \geq r/2$ and $d_2 < d_1$, so $(\frac{y + d_2 - 0.5}{d_1 - 0.5})/3 < (\frac{y}{d_1 - 0.5} + 1)/3 \leq (\frac{r}{r/2} + 1)/3 \leq 1$.

Summarizing the above, the only cases where the average dot product is greater than 1 is (1, 0, x) for $x \geq d_1$ and (1, 1, y) for $y \geq d_2$ (corresponding to a chain of two edges). The only cases where the average dot product is equal to 1 is (1, 0, 0) (corresponding to a single edge). All other cases have average dot product less than 1. The critical path we defined from s_π to t_π exactly use the edges (or set of edges) whose average dot product per cost is the largest. By construction, if there is an alternative path from s_π to t_π , then it must use a different collection of edge sets. Furthermore, the total dot product between δ and all the

vectors corresponding to the edges on the alternative path must be equal to the total dot product between δ and all the vectors corresponding to the edges on the critical path. As the critical path uses edges that maximize the average dot product per cost, it implies that the alternative path must have a larger cost. Thus, we conclude that the critical path is the unique shortest path between s_π and t_π .

- First, if two critical paths share the same values (d_1, d_2) and start from distinct vertices on the first layer, they will be disjoint. Thus, it suffices to bound the number of distinct (d_1, d_2) , so that there could potentially be a critical path with value (d_1, d_2) that contains σ .

Furthermore, for even i , a critical path with values (d_1, d_2) for $d_1 \geq q_{(i/2) \bmod r}$ and a critical path with values (d'_1, d'_2) for $d'_1 < q_{(i/2) \bmod r}$ cannot share the same edge between layer i and layer $i + 1$. Similarly, for odd i , a critical path with values (d_1, d_2) for $d_2 \geq q_{((i-1)/2) \bmod r}$ and a critical path with values (d'_1, d'_2) for $d'_2 < q_{((i-1)/2) \bmod r}$ cannot share the same edge between layer i and layer $i + 1$.

The rest of the proof follows essentially in the same way as the proof of Lemma 8.2 Item 3.

□

Proof of Theorem 9.1. The overall proof given Lemma 9.2 is similar to the proof of Theorem 8.1, so we omit some details and highlight the differences.

We consider the graph $G = (V, E)$ and the critical paths P constructed above.

Let E_Δ be the set of shortcut edges in G that connects every pair of vertices (u, v) where u can reach v using at most Δ edges, and the weight of each shortcut edge is the distance from u to v .

We define a potential function with respect to any hopset H :

$$\Phi(H) := \sum_{\pi \in P} \text{hopdist}_{G \cup E_\Delta \cup H}(s_\pi, t_\pi).$$

Initially, we clearly have $\Phi(\emptyset) = |P| \cdot cr/\Delta = \Theta(c^3 r^6 / \Delta)$. Similar to the proof of Theorem 8.1,

$$\Phi(H) - \Phi(H \cup \{e\}) = O\left(\frac{cr}{\Delta} + \frac{r^2}{\Delta^2}\right).$$

Therefore, if $|H| = \varepsilon_1 \min\{c^2 r^5, c^3 r^4 \Delta\}$ for some sufficiently small constant $\varepsilon_1 > 0$, $\Phi(\emptyset) - \Phi(H) \leq \varepsilon_2 c^3 r^6 / \Delta$ for some small constant ε_2 , which can be made arbitrarily small depending on ε_1 . As a result, $\Phi(H) = \Theta(c^3 r^6 / \Delta)$, which will imply that there exists a critical path π where the minimum number of hops of any shortest path from s_π to t_π is $\Omega(\Phi(H)/|P|) = \Omega(cr/\Delta)$.

We set $c = r/B$ for some large constant B . By Lemma 9.2 Item 1, $|V|, |E| = \Theta(r^7/B^3)$. The number of edges allowed in the hopset is $\varepsilon_1 \min\{c^2 r^5, c^3 r^4 \Delta\} = \varepsilon_1 \min\{c^7/B^2, c^7 \Delta/B^3\}$, which can be any constant factor larger than $|E|$ by setting B and Δ large enough. The hop bound is $\Omega(cr/\Delta) = \Omega(r^2/B\Delta)$, which is $\Omega(|V|^{2/7})$ for constant values of B and Δ . □

Acknowledgements

We would like to thank Greg Bodwin for helpful discussions and feedback.

References

- [AB16] Amir Abboud and Greg Bodwin. Error amplification for pairwise spanner lower bounds. In *Proceedings of the 27th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 841–854, 2016. [1](#)
- [AB17] Amir Abboud and Greg Bodwin. The $4/3$ additive spanner exponent is tight. *J. ACM*, 64(4):28:1–28:14, 2017. [1](#)
- [ABP17] Amir Abboud, Greg Bodwin, and Seth Pettie. A hierarchy of lower bounds for sublinear additive spanners. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 568–576, 2017. [1](#)
- [ADD⁺93] Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993. [6](#)
- [Alo02] Noga Alon. Testing subgraphs in large graphs. *Random Structures & Algorithms*, 21(3-4):359–370, 2002. [1](#)
- [ASZ20] Alexandr Andoni, Clifford Stein, and Peilin Zhong. Parallel approximate undirected shortest paths via low hop emulators. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 322–335, 2020. [1](#)
- [BCE05] Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. *SIAM J. Discrete Math.*, 19(4):1029–1055, 2005. [1](#), [27](#)
- [BGWN20] Aaron Bernstein, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Near-optimal decremental SSSP in dense weighted digraphs. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1112–1122, 2020. [1](#)
- [BH22] Greg Bodwin and Gary Hoppenworth. New additive spanner lower bounds by an unlayered obstacle product. In *Proceedings of the 63rd IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 778–788. IEEE, 2022. [5](#), [6](#), [7](#)
- [BH23] Greg Bodwin and Gary Hoppenworth. Folklore sampling is optimal for exact hopsets: Confirming the $\sqrt{n}$ barrier. In *2023 IEEE 64th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 701–720, 2023. [10](#), [34](#), [36](#), [37](#)
- [BHT23] Greg Bodwin, Gary Hoppenworth, and Ohad Trabelsi. Bridge girth: A unifying notion in network design. In *Proceedings of the 64th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 600–648. IEEE, 2023. [1](#), [2](#), [3](#), [7](#), [13](#), [18](#), [47](#)
- [BL24] Greg Bodwin and Tuong Le. Improved online reachability preservers. *arXiv preprint arXiv:2410.20471*, 2024. [3](#)
- [Bod17] Greg Bodwin. Linear size distance preservers. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 600–615. SIAM, 2017. [2](#)
- [Bod21] Greg Bodwin. New results on linear size distance preservers. *SIAM J. Comput.*, 50(2):662–673, 2021. [2](#), [28](#)

- [BRR10] Piotr Berman, Sofya Raskhodnikova, and Ge Ruan. Finding sparser directed spanners. In *IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, 2010. [5](#), [6](#)
- [BV15] Greg Bodwin and Virginia Vassilevska Williams. Very sparse additive spanners and emulators. In *Proceedings of the 2015 Conference on Innovations in Theoretical Computer Science (ITCS)*, pages 377–382, 2015. [1](#)
- [BV21] Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. *ACM Trans. Algorithms*, 17(4):1–24, 2021. [1](#), [2](#), [3](#), [9](#), [25](#), [26](#), [48](#)
- [BW23] Aaron Bernstein and Nicole Wein. Closing the gap between directed hopsets and shortcut sets. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 163–182, 2023. [2](#), [3](#), [5](#), [6](#)
- [CE06] Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM J. Discrete Math.*, 20(2):463–501, 2006. [1](#), [2](#), [3](#), [10](#), [13](#), [25](#), [28](#), [29](#), [48](#)
- [CFR20] Nairen Cao, Jeremy T. Fineman, and Katina Russell. Efficient construction of directed hopsets and parallel approximate shortest paths. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 336–349, 2020. [1](#)
- [CGMW18] Hsien-Chih Chang, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Near-Optimal Distance Emulator for Planar Graphs. In *Proceedings of the 26th Annual European Symposium on Algorithms (ESA)*, volume 112 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:17, Dagstuhl, Germany, 2018. [1](#)
- [Coh00] Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM*, 47(1):132–166, 2000. [5](#)
- [Die00] Reinhard Diestel. *Graduate texts in mathematics: Graph theory*. Springer, 2000. [9](#)
- [EFN17] Michael Elkin, Arnold Filtser, and Ofer Neiman. Terminal embeddings. *Theor. Comput. Sci.*, 697:1–36, 2017. [1](#)
- [Ele01] György Elekes. Sums versus products in number theory, algebra and erdos geometry. *Paul Erdos and his Mathematics II*, 11(241-290):5, 2001. [33](#)
- [EM21] Michael Elkin and Shaked Matar. Deterministic PRAM approximate shortest paths in polylogarithmic time and slightly super-linear work. In Kunal Agrawal and Yossi Azar, editors, *Proceedings of the 33rd ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 198–207. ACM, 2021. [1](#)
- [EN16] Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. In *Proceedings of the IEEE 57th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 128–137, 2016. [1](#)
- [EN19] Michael Elkin and Ofer Neiman. Linear-size hopsets with small hopbound, and constant-hopbound hopsets in RNC. In *Proceedings of the 31st ACM Symposium on Parallelism in Algorithms and Architectures (SPAA)*, pages 333–341, 2019. [1](#), [6](#)

- [EP16] Michael Elkin and Seth Pettie. A linear-size logarithmic stretch path-reporting distance oracle for general graphs. *ACM Trans. Algorithms*, 12(4):50, 2016. [1](#)
- [ES23] Michael Elkin and Idan Shabat. Path-reporting distance oracles with logarithmic stretch and size $O(n \log \log n)$. In *Proceedings of the 64th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 2278–2311, 2023. [1](#), [2](#), [3](#)
- [Fin19] Jeremy T. Fineman. Nearly work-efficient parallel algorithm for digraph reachability. *SIAM J. Comput.*, 49(5), 2019. [1](#)
- [FN18] Sebastian Forster and Danupon Nanongkai. A faster distributed single-source shortest paths algorithm. In *Proceedings of the 2018 IEEE 59th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 686–697, 2018. [1](#)
- [GK17] Manoj Gupta and Shahbaz Khan. Multiple Source Dual Fault Tolerant BFS Trees. In *Proceedings of the 44th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 80, pages 127:1–127:15, Dagstuhl, Germany, 2017. [1](#)
- [Hes03] William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the 14th Annual ACM-SIAM Symposium on Discrete algorithms (SODA)*, pages 665–669, 2003. [5](#), [6](#)
- [HKN14] Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Decremental single-source shortest paths on undirected graphs in near-linear total update time. In *Proceedings of the 2014 IEEE 55th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 146–155, 2014. [1](#)
- [HP19] Shang-En Huang and Seth Pettie. Thorup-zwick emulators are universally optimal hopsets. *Information Processing Letters*, 142:9–13, 2019. [6](#)
- [HP21] Shang-En Huang and Seth Pettie. Lower bounds on sparse spanners, emulators, and diameter-reducing shortcuts. *SIAM J. Discret. Math.*, 35(3):2129–2144, 2021. [6](#)
- [JLS19] Arun Jambulapati, Yang P. Liu, and Aaron Sidford. Parallel reachability in almost linear work and square root depth. In *Proceedings of the 2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1664–1686, 2019. [1](#)
- [KP22a] Shimon Kogan and Merav Parter. Having hope in hops: New spanners, preservers and lower bounds for hopsets. In *Proceedings of the 2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 766–777, 2022. [1](#), [2](#), [3](#), [6](#)
- [KP22b] Shimon Kogan and Merav Parter. New diameter-reducing shortcuts and directed hopsets: Breaking the barrier. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1326–1341, 2022. [5](#), [6](#)
- [KS97] Philip N. Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *J. Algorithms*, 25(2):205–220, 1997. [1](#)
- [KS21] Adam Karczmarz and Piotr Sankowski. A deterministic parallel APSP algorithm and its applications. In *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 255–272, 2021. [1](#)

- [KSSY16] Saarik Kalia, Micha Sharir, Noam Solomon, and Ben Yang. Generalizations of the szemerédi–trotter theorem. *Discrete & Computational Geometry*, 55:571–593, 2016. [1](#)
- [LVWX22] Kevin Lu, Virginia Vassilevska Williams, Nicole Wein, and Zixuan Xu. Better lower bounds for shortcut sets and additive spanners via an improved alternation product. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3311–3331, 2022. [6](#)
- [Tho92] Mikkel Thorup. On shortcutting digraphs. In *Proceedings of the 18th International Workshop on Graph-Theoretic Concepts in Computer Science (WG)*, pages 205–211, 1992. [1](#)
- [UY91] Jeffrey D. Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM J. Comput.*, 20(1):100–125, 1991. [5](#), [6](#)
- [VW18] Virginia Vassilevska Williams and R. Ryan Williams. Subcubic equivalences between path, matrix, and triangle problems. *J. ACM*, 65(5):27:1–27:38, 2018. [4](#)
- [VXX24] Virginia Vassilevska Williams, Yinzhan Xu, and Zixuan Xu. Simpler and higher lower bounds for shortcut sets. In David P. Woodruff, editor, *Proceedings of the 2024 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2643–2656. SIAM, 2024. [5](#), [6](#), [7](#), [10](#), [11](#), [35](#), [38](#)

A Proof of Assumptions 3 and 4 from Section 4

Our proof of Assumptions 3 and 4 will follow from an argument similar to the Cleaning Lemma of [BHT23]. Note that the graph obtained after applying Lemma A.1 may be weighted. Consequently, this lemma is not useful for our unweighted directed distance preserver upper bound in Section 5.

Lemma A.1. *Let G, P be an n -node directed, weighted graph and a set of $|P| = p$ demand pairs, such that the minimal distance preserver H of G, P has*

$$|E(H)| = \Theta(\text{DDP}(n, p))$$

edges. Let $\pi(\cdot, \cdot)$ be a tiebreaking scheme associated with G, P . We can make the following assumptions on G, P, H , and $\pi(\cdot, \cdot)$ without loss of generality:

3. *Every node $v \in V(G)$ has degree at least $\deg_H(v) \geq \frac{|E(H)|}{4n}$ in H .*
4. *Every path $\pi \in \pi(P)$ has at least $|\pi| \geq \frac{|E(H)|}{4p}$ edges.*

Proof. Firstly, we can assume that for all $(s, t) \in P$, $\pi(s, t)$ is a unique shortest path in G . Then, necessarily,

$$H = \bigcup_{\pi \in \pi(P)} \pi.$$

Secondly, we can assume without loss of generality that paths in $\pi(P)$ are pairwise edge-disjoint. Thirdly, we can assume that $G = H$, since deleting edges in $E(G) \setminus E(H)$ from H does not affect the size of H . We will ensure Assumptions 3 and 4 hold without loss of generality by a simple procedure to modify G, P . Let $\Pi = \pi(P)$ denote the collection of shortest paths in G corresponding to demand pairs P . We modify G, P , and Π as follows.

1. Repeat the following two operations until no longer possible:

- (a) While there exists a path $\pi = \pi(s, t) \in \Pi$ such that $|\pi| < |E(G)|/(4p)$, delete path π from Π .
- (b) While there exists a node $v \in V(G)$ such that $\deg_G(v) < |E(G)|/(4n)$, remove node v from G . For each path $\pi \in \Pi$ that contains node v , we repair path π as follows. If v is the first node or the last node in path π , then we repair π by deleting node v from π . Otherwise, let u and w be nodes immediately preceding and succeeding v in π , respectively. We delete node v from π and add an edge from u to w with weight $w((u, w)) = \text{dist}_G(u, w)$.

2. We define our modified version of graph G as

$$G = \bigcup_{\pi \in \Pi} \pi.$$

3. We define our modified version of demand pairs P as

$$P = \{(s, t) \in V(G) \times V(G) \mid \text{there exists an } s \rightsquigarrow t \text{ path in } \Pi\}.$$

It is straightforward to verify that after our modifications, paths in Π are unique shortest paths in G . Moreover, Step 1(a) of our modification deletes at most $p \cdot |E(G)|/(4p) = |E(G)|/4$ edges from H , and Step 1(b) of our modification deletes at most $n \cdot |E(G)|/(4n)$ edges from H , so we conclude that after Step 1 terminates, there are at least $|E(G)|/2$ edges in G , so G is nonempty. Then by our termination conditions, we must have that

- Every path $\pi \in \Pi$ must satisfy $|\pi| \geq |E(G)|/(4p)$.
- Every node $v \in V(G)$ must satisfy $\deg_G(v) \geq |E(G)|/(4n)$.

Finally, we observe that $|E(G)| = \Theta(\text{DDP}(n, p))$, since we deleted at most half of the edges originally in G . This establishes the lemma. $\square$

B Proof of Lemma 5.7 from Section 5

The goal of this section is to prove Lemma 5.7, which we restate below for convenience.

Lemma B.1 (cf. Lemma 8 of [BV21]). *Let G be an n -node directed graph, and let $S \subseteq V(G)$ be a set of $|S| = s$ nodes of weak diameter h (i.e., $\text{dist}_G(s, t) \leq h$ for all $s, t \in S$). Let $P \subseteq S \times V$ be a set of $|P| = p$ sourcewise demand pairs in G . Then there exists a distance preserver of G, P of size*

$$O\left((nsp h)^{1/2} + n\right).$$

The proof of this lemma follows from an argument similar to that of Lemma 8 of [BV21]. We will require the following lemma due to [CE06].

Lemma B.2 ([CE06]; cf. Lemma 7 of [BV21]). *If there are b branching events associated with a tiebreaking scheme $\pi(\cdot, \cdot)$ of G, P , then the associated distance preserver H has size $|E(H)| = O((nb)^{1/2} + n)$.*

We will now prove Lemma 5.7.

Proof of Lemma 5.7. Let $\pi(\cdot, \cdot)$ be a tiebreaking scheme whose associated distance preserver

$$H = \bigcup_{(s,t) \in P} \pi(s, t)$$

is minimal. For every edge $e \in E(H)$, we assign edge e a label $\ell(e) \in P$. Specifically, we let $\ell(e) = (s, t) \in P$ for an arbitrary demand pair $(s, t) \in P$ such that every $s \rightsquigarrow t$ shortest path in H uses edge e . We note that every edge $e \in E(H)$ can be labeled in this way – otherwise, we could delete edge e from H , contradicting the minimality of H . Likewise, for each branching event $b = (e_1, e_2) \in E(H) \times E(H)$, we assign the branching event label $\ell(b) = (\ell(e_1), \ell(e_2))$.

We will prove Lemma 5.7 by bounding the number of branching events and applying Lemma B.2. Suppose for the sake of contradiction that H has greater than $sp(2h + 1)$ branching events. Then by applying the pigeonhole principle, there is a node $t \in S$ and a demand pair $(u, v) \in P$ such that there are at least $2h + 2$ distinct branching events $b_1, \dots, b_{2h+2}$ such that $\ell(b_i) = ((u, v), (t, w_i))$ for some $w_i \in V(H)$, for all $i \in [2h + 2]$. For each branching event $b_i = (e_i^1, e_i^2)$, let $v_i \in V(H)$ be the unique node incident to e_i^1 and e_i^2 . We may assume wlog that nodes $\{v_1, \dots, v_{2h+2}\}$ are ordered by nondecreasing distance $\text{dist}_H(u, v_i)$. By the triangle inequality, we know that

$$\text{dist}_H(u, v_i) \leq \text{dist}_H(u, t) + \text{dist}_H(t, v_i).$$

Subtracting $\text{dist}_H(t, v_i)$ from both sides and using the fact that S has diameter at most h , we obtain

$$\text{dist}_H(u, v_i) - \text{dist}_H(t, v_i) \leq \text{dist}_H(u, t) \leq h.$$

Likewise, by the triangle inequality,

$$\text{dist}_H(t, v_i) \leq \text{dist}_H(t, u) + \text{dist}_H(u, v_i),$$

and using the fact that S has diameter at most h , we obtain

$$\text{dist}_H(t, v_i) - \text{dist}_H(u, v_i) \leq \text{dist}_H(t, u) \leq h.$$

Combining our two bounds on $\text{dist}_H(u, v_i) - \text{dist}_H(t, v_i)$, we arrive at

$$-h \leq \text{dist}_H(u, v_i) - \text{dist}_H(t, v_i) \leq h.$$

Then by the pigeonhole principle, there exists branching events b_i, b_j where $i \neq j \in [2h + 2]$ such that

$$\text{dist}_H(t, v_i) - \text{dist}_H(u, v_i) = \text{dist}_H(t, v_j) - \text{dist}_H(u, v_j).$$

Rearranging, we obtain

$$\text{dist}_H(t, v_j) = \text{dist}_H(t, v_i) + \text{dist}_H(u, v_j) - \text{dist}_H(u, v_i).$$

Because we assumed that $\text{dist}_H(u, v_i) \leq \text{dist}_H(u, v_j)$ and $v_i, v_j \in \pi(u, v)$, it follows that

$$\text{dist}_H(u, v_j) - \text{dist}_H(u, v_i) = \text{dist}_H(v_i, v_j).$$

And therefore,

$$\text{dist}_H(t, v_j) = \text{dist}_H(t, v_i) + (\text{dist}_H(u, v_j) - \text{dist}_H(u, v_i)) = \text{dist}_H(t, v_i) + \text{dist}_H(v_i, v_j).$$

This implies that the $t \rightsquigarrow v_j$ path

$$\pi(t, v_i) \circ \pi(u, v)[v_i, v_j]$$

is a shortest $t \rightsquigarrow v_j$ path in H , where $\pi(u, v)[v_i, v_j]$ denotes the $v_i \rightsquigarrow v_j$ subpath of path $\pi(u, v)$. We conclude that there is a shortest $t \rightsquigarrow v_j$ path that uses an edge in $\pi(u, v)$ to reach node v_j . Consequently, not every shortest $t \rightsquigarrow v_j$ path includes an edge e_j in the branching event b_j . This contradicts our labeling scheme of branching events.

We conclude that H has $O(sph)$ branching events. Plugging in Lemma B.2, we get that

$$|E(H)| = O((n sph)^{1/2} + n),$$

completing the proof of Lemma 5.7. □