Towards Bypassing Lower Bounds for Graph Shortcuts

Shimon Kogan ✉
Weizmann Institute of Science, Rehovot, Israel

Merav Parter ✉
Weizmann Institute of Science, Rehovot, Israel

Abstract

For a given (possibly directed) graph G, a hopset (a.k.a. shortcut set) is a (small) set of edges whose addition reduces the graph diameter while preserving desired properties from the given graph G, such as, reachability and shortest-path distances. The key objective is in optimizing the tradeoff between the achieved diameter and the size of the shortcut set (possibly also, the distance distortion). Despite the centrality of these objects and their thorough study over the years, there are still significant gaps between the known upper and lower bound results.

A common property shared by almost all known shortcut lower bounds is that they hold for the seemingly simpler task of reducing the diameter of the given graph, D_G, by a constant additive term, in fact, even by just one! We denote such restricted structures by $(D_G - 1)$-diameter hopsets. In this paper we show that this relaxation can be leveraged to narrow the current gaps, and in certain cases to also bypass the known lower bound results, when restricting to sparse graphs (with $O(n)$ edges):

- **Hopsets for Directed Weighted Sparse Graphs.** For every n-vertex directed and weighted sparse graph G with $D_G \geq n^{1/4}$, one can compute an *exact* $(D_G - 1)$-diameter *hopset* of linear size. Combining this with known lower bound results for *dense* graphs, we get a separation between dense and sparse graphs, hence *shortcutting sparse graphs is provably easier*. For reachability hopsets, we can provide $(D_G - 1)$-diameter hopsets of linear size, for sparse DAGs, already for $D_G \geq n^{1/5}$. This should be compared with the diameter bound of $\widetilde{O}(n^{1/3})$ [Kogan and Parter, SODA 2022], and the lower bound of $D_G = n^{1/6}$ by [Huang and Pettie, SIAM J. Discret. Math. 2018].
- **Additive Hopsets for Undirected and Unweighted Graphs.** We show a construction of $+24$ *additive* $(D_G - 1)$-diameter hopsets with linear number of edges for $D_G \geq n^{1/12}$ for sparse graphs. This bypasses the current lower bound of $D_G = n^{1/6}$ obtained for *exact* $(D_G - 1)$-diameter hopset by [HP'18]. For general graphs, the bound becomes $D_G \geq n^{1/6}$ which matches the lower bound of exact $(D_G - 1)$ hopsets implied by [HP'18]. We also provide new additive D-diameter hopsets with linear size, for any given diameter D.

Altogether, we show that the current lower bounds can be bypassed by restricting to sparse graphs (with $O(n)$ edges). Moreover, the gaps are narrowed significantly for any graph by allowing for a constant additive stretch.

2012 ACM Subject Classification Theory of computation → Graph algorithms analysis

Keywords and phrases Directed Shortcuts, Hopsets, Emulators

Digital Object Identifier 10.4230/LIPIcs.ESA.2023.73

Funding This project is funded by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement No. 949083), and by the Israeli Science Foundation (ISF), grant No. 2084/18.

1 Introduction

A shortcut set (a.k.a hopset) is a small collection of edges H that when added to a given (possibly directed and weighted) graph G reduces the diameter substantially, while preserving key properties from G, such as, reachability, shortest-path distances, etc. Since their

31st Annual European Symposium on Algorithms (ESA 2023).
Editors: Inge Li Gørtz, Martin Farach-Colton, Simon J. Puglisi, and Grzegorz Herman; Article No. 73; pp. 73:1–73:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

introduction by Ullman and Yannakakis [25] and Thorup [24], shortcut sets have been studied extensively due to their wide-range of applications for parallel, distributed, dynamic and streaming algorithms [16, 12, 10, 15, 9, 11, 3]. Their applicability is also demonstrated by recent algorithmic results, e.g., [15, 9, 11] that use shortcuts as their core component.

We focus on the fundamental notions of *reachability* D-diameter hopsets[1] and *exact* D-diameter hopsets. These are sets of edges that when added to G reduce the directed diameter to at most D while preserving reachability (resp., exact D-hop distances). For a graph on n vertices and m edges, the literature has mainly focused on how small the diameter D can become after adding $\widetilde{O}(n)$ or $\widetilde{O}(m)$ edges to the graph. Recent years have witnessed a significant progress on the combinatorial and algorithmic aspects of these structures. Despite these efforts, there are still major gaps in our understanding. Thorup conjectured [24] that any n-vertex digraph with m edges, has a D-diameter shortcut of size $\widetilde{O}(m)$ for[2] $D = \widetilde{O}(1)$. This has been shown to hold for a restricted class of graphs, such as planar [24]. Hesse [13] refuted this conjecture for general graphs by presenting a construction of an n-vertex digraph with $m = \Theta(n^{19/17})$ edges that requires $\Omega(mn^{1/17})$ shortcut edges to reduce its diameter to below $n^{1/17}$. Huang and Pettie [14] showed a diameter lower bound of $\Omega(n^{1/6})$ for $O(n)$-size shortcuts, and Lu, Vassilevska-Williams, Wein and Xu [21] have recently improved the lower bound for D-shortcuts with $O(m)$ edges to $D = \Omega(n^{1/8})$.

Up to very recently, the only known upper bound for reachability hopsets was given by a folklore randomized algorithm, attributed to Ullman and Yannakakis [25], that for every integer $D \geq 1$, provides a shortcut of size $\widetilde{O}((n/D)^2)$. Kogan and Parter [18] have improved the tradeoff to $\widetilde{O}(n^2/D^3 + (n/D)^{3/2})$. Hence, for reachability D-diameter hopsets of linear size, we have $D = \widetilde{O}(n^{1/3})$ and $D = \Omega(n^{1/6})$.

The gap becomes even more dramatic when insisting on preserving the *exact* distances. The only known upper bound for exact D-diameter hopsets are still given by the folklore algorithm of [25], even for the presumably simpler setting of *unweighted* and *undirected* graphs. Currently, exact D-diameter hopsets of linear size require $D = \Omega(n^{1/6})$ for unweighted and undirected graphs, while the upper bound is still $O(\sqrt{n})$. The lower bound for weighted graphs has been recently improved to $D = \Omega(n^{1/3})$ by [17].

In this paper we investigate this curious gap between upper and lower bound results for reachability and exact hopsets. We aim at understanding the limitations of the current lower bound constructions, and study whether they can be bypassed, algorithmically. Our viewpoint differs from the classical algorithmic approach in the sense that we aim to provide algorithms mainly for the purpose of understanding the key barriers for narrowing the observed gaps. In the following, for two given shortcutting tasks Π, Π', we say that task Π is *easier* than task Π' if one can provide for task Π an improved diameter vs. size tradeoff over the current tradeoff known for task Π'.

1.1 Our Contribution

Our starting observation is that almost all known lower bounds for hopsets, a partial list includes, [13, 14, 21, 17], in-fact hold for the seemingly easier task of reducing the diameter of the given graph by just 1. For example, the lower bound of Huang and Pettie [14] exhibits

[1] These structures are usually referred to as reachability shortcuts, and the notion of hopset is usually used in the context of preserving also the shortest-path distances. For clarity of presentation, we unify the notation and refer to all structures as variants of hopsets.

[2] The $\widetilde{O}(.)$ notation hides poly-$\log n$ factors.

an n-vertex graph G with diameter $D_G = \Theta(n^{1/6})$ for which any subset of linear number of shortcut edges cannot reduce the diameter to $D_G - 1$. As this property is shared by many of the lower bound results, we raise the following question:

▶ **Question 1.1** (Additive vs. Multiplicative Diameter Reduction). *Is reducing the graph diameter by a constant* ***additive*** *term easier than by a constant* ***multiplicative*** *factor?*

In this context, we say that task Π is *easier* than task Π', if one can provide a solution for Π' which improves the state-of-the-art bounds known for Π. To answer this question, we focus on the notion of $(D_G - 1)$-diameter hopset: a subset of shortcut edges that reduces the diameter of G from D_G to $D_G - 1$. We provide new algorithmic results for $(D_G - 1)$-diameter hopsets that significantly *narrow* the current gap for these structures.

▶ **Theorem 1.1** (Reachability $(D_G - 1)$-Diameter Hopset). *Every m-edge n-vertex DAG G admits a reachability $(D_G - 1)$-diameter hopset with $\widetilde{O}(m^{4/3}/D_G^{5/3})$ edges.*

By plugging $m = O(n)$, we get a reachability $(D_G - 1)$-hopset with linear number of edges for every $D_G \geq n^{1/5}$. This suggests that at least algorithmically and for the family of sparse graphs, the task of reducing the diameter additively is easier than reducing the diameter by a constant factor. Currently, for the latter task, linear reachability hopsets exist only for $D_G \geq n^{1/3}$, even for sparse DAGs [18, 19].

Another fundamental feature shared by most of the lower bound results [13, 14, 21, 17] is concerned with the density (i.e. number of edges) of the worst-case graph examples. In particular, all these lower bound graphs have a super-linear number of edges (referred to hereafter as *dense* graphs). E.g., the above-mentioned lower bound graph of [14] contains $\Omega(n^{7/6})$ edges. We therefore raise the following question:

▶ **Question 1.2** (Shortcutting Sparse vs. Dense Graphs). *Is reducing graph diameter easier for* ***sparse*** *graphs compared to* ***dense*** *graphs?*

We answer Question 1.2 in the affirmative by providing a *provable* gap between the sparse and dense settings, in the context of *exact* hopsets. We show:

▶ **Theorem 1.2** (Exact $(D_G - 1)$-Diameter Hopset, Directed and Weighted). *Every m-edge n-vertex directed (and possibly weighted) graph G admits an exact $(D_G - 1)$-diameter hopset with $\widetilde{O}((m/D_G)^{4/3})$ edges.*

Taking $m = n$ yields $\widetilde{O}(n)$-size exact hopsets for diameter $D_G \geq n^{1/4}$. This should be compared with the lower bound for exact $(D_G - 1)$-diameter hopsets by [4] (one can adapt Theorem 5, therein to this setting) which requires a super-linear number of shortcut edges for dense graphs G with $\Theta(n^{3/2})$ edges and diameter $D_G = n^{1/4}$.

Finally, we address the setting of exact $(D_G - 1)$-diameter hopsets for unweighted and undirected graphs, which currently admits the largest gap between the known upper and lower bound results. Recall that for exact D-diameter hopsets of linear size, we have $D = O(\sqrt{n})$. As no progress (at least for upper bounds) has been made for exact hopsets over the years, we consider the question of whether it is possible to narrow this gap by introducing an additive stretch:

▶ **Question 1.3** (Exact vs. Additive Stretch). *Is reducing the graph diameter easier when allowing an **additive** stretch compared to **exact** distances?*

To address Question 1.3 we study the notion of additive hopsets, which to the best of our knowledge has not been considered before. For a given graph G, an $+\alpha$-additive D-diameter hopset is a subset of (weighted) shortcut edges H whose addition guarantees the following: for every u, v pair, the graph $G \cup H$ contains a D-hop u-v path (that is the path contains at most D edges) of total length[3] $\ell_{u,v}$ where $\mathrm{dist}_G(u,v) \leq \ell_{u,v} \leq \mathrm{dist}_G(u,v) + \alpha$. We show that the introduction of additive stretch can indeed achieve a lot in terms of improving the state-of-the-art bounds when compared to exact hopsets. We have:

▶ **Theorem 1.3** (Additive Hopsets Upper Bounds). *Every n-vertex m-edge undirected unweighted graph G admits a $+24$-additive $(D_G - 1)$-diameter hopset with $\widetilde{O}(\min\{(n/D_G)^{6/5}, (m/D_G)^{12/11}\})$ edges.*

Theorem 1.3 should be compared with the lower bound of [14]. The latter can be shown to imply that there exists an undirected and unweighted graph G with $\omega(n^{7/6})$ edges for which any *exact* $(D_G - 1)$-diameter hopset requires $\Omega(n)$ edges. For $m = O(n)$, our construction in Theorem 1.3 provides $+24$-additive $(D_G - 1)$-hopset with linear number of edges, for any $D_G \geq n^{1/12}$. That is, by restricting to sparse graphs and introducing a small additive stretch we bypass the lower bound of [14]. Moreover, for any (possibly dense) graph G, we provide a linear-size *additive* $(D_G - 1)$-diameter hopset for $D_G \geq n^{1/6}$, hence matching the lower bound of $n^{1/6}$ implied by [14] for *exact* $(D_G - 1)$-diameter hopsets.

We also show a general construction of α-additive D-hopsets for any given stretch parameter α and diameter bound D (which is possibly much smaller than D_G):

▶ **Theorem 1.4.** *Every n-vertex undirected unweighted graph G with input parameters α, D, admits a $+\alpha$-additive D-diameter hopset with $\widetilde{O}((n/(\alpha \cdot D))^2 + (n/\alpha)^{4/3} + n)$ edges. Hence, of linear size for $D, \alpha = \Omega(n^{1/4})$.*

The construction of Theorem 1.4 is based on an interesting connection between hopsets and additive spanners for weighted graphs [1, 7].

Remark. In our upper bound results, we focus on $(D_G - 1)$-diameter hopsets since this is a recurrent feature of the state-of-the-art lower bounds. Our algorithms can indeed be easily extended to provide a $-c$ reduction for any given term c. In the context of exact hopsets and reachability preservers, the factor c will appear in the final size bound (for constant c, the asymptotic size bound remains as is). For additive hopsets the factor c will appear in the final size as well as in the final additive stretch (i.e., the $+24$ term in Theorem 1.3 might become $+24 \cdot c$).

Recent Breakthrough Lower-Bound Results on Shortcuts and Hopsets. Following the submission of this paper, Bodwin and Hoppenworth [5] provided new lower bound results for shortcuts and hopsets. In particular, they provide a lower bound of $D = \Omega(n^{1/4})$ for

[3] The length of a path is measured by the sum of its edge weights (note that in an unweighted graph the length is the number of edges in the path).

reachability D-diameter hopsets with linear number of edges. Hence, the current gap for these structures is $D = \widetilde{O}(n^{1/3})$ and $D = \Omega(n^{1/6})$. In addition, for exact D-hopsets of linear size, they show a lower bound of $D = \Omega(\sqrt{n})$, implying that the folklore algorithm for these structures is indeed tight.

1.2 Preliminaries

Graph Notations. For an a-b dipath P and a b-c dipath P' the concatenation of the paths is denoted by $P \circ P'$. Let $|P|$ denote the number of edges (hop) on P and let $\text{len}(P)$ be the *length* of path P, measured by the sum of its (possibly) weighted edges. In the context of weighted graphs, we refer to number of edges on a path P as the number of *hops*. For a collection of paths $\mathcal{P}$, let $V(\mathcal{P}) = \bigcup_{P \in \mathcal{P}} V(P)$. For an element set X and $p \in [0,1]$, let $X[p]$ be the set obtained by taking each element of X into $X[p]$ independently with probability p. For a subset $V' \subseteq V$, let $G[V']$ be the induced graph on V'. For a (possibly directed and weighted) graph G and $u, v \in V(G)$, let $\text{dist}_G(u,v)$ be the shortest-path distance between u, v in G. Let $N_{in}(u,G) = \{v \in V(G) \mid (v,u) \in E(G)\}$ and $N_{out}(u,G) = \{v \in V(G) \mid (u,v) \in E(G)\}$. Define $\deg_{in}(u,G) = |N_{in}(u,G)|$ and $\deg_{out}(u,G) = |N_{out}(u,G)|$. For an undirected graph G, a vertex $v \in V(G)$ and an integer k, let $N_k(v,G) = \{u \in V(G) \mid \text{dist}_G(u,v) \leq k\}$ and let $\deg_k(v,G) = |N_k(v)|$. When G is clear from the context, we may omit it. Also, when $k = 1$, we may write $N(v,G)$, rather than $N_1(v,G)$. For a given path $P \subseteq G$ and $u \in P$, the k-hop P-neighbor of u is some vertex in $P \cap N_k(u,G)$.

For an n-vertex digraph G, let $TC(G)$ denote its transitive closure. For $(u,v) \notin TC(G)$, $\text{dist}_G(u,v) = \infty$. A *shortcut* edge is an edge in $TC(G)$. These definitions are naturally extended to *undirected* graphs, by treating all edges as bidirectional, consequently, $TC(G)$ is an $n \times n$ clique for connected undirected graphs. For an integer weighted (possibly directed) graph $G = (V, E, \omega)$ where $\omega : E \to [1, W]$, the weighted transitive closure, denoted as $TC_W(G)$, has the same edge set as $TC(G)$ weighted by the G-distances between the edges' endpoints. For any vertices $u, v \in V$ and integer $\beta \leq n$, define $\text{dist}^{(\beta)}_G(u,v)$ to be the minimum length u-v path with at most β edges (hops). If there is no such path, then define $\text{dist}^{(\beta)}_G(u,v) = \infty$. Throughout, we define the *diameter* of the graph, D_G, by the smallest value β satisfying that $\text{dist}^{(\beta)}_G(u,v) = \text{dist}_G(u,v)$ for every $u, v \in V$. In the context of weighted graphs G, D_G as defined above is usually refereed to as the *hopbound*. For ease of readability, we slightly revise the notion of diameter in a way the captures weighted and unweighted hopsets as well as reachability shortcuts.

Reachability, Exact and Additive Hopsets. For a directed graph G, a reachability d-diameter hopset H is a subset of edges from the transitive closure $TC(G)$ such that $D_{G \cup H} \leq d$. For a (possibly directed and weighted) graph G, an α-additive d-diameter hopset H is a subset of weighted edges in $TC_W(G)$ satisfying that for every $u, v \in V(G)$:

$$\text{dist}_G(u,v) \leq \text{dist}^{(d)}_{G \cup H}(u,v) \leq \text{dist}_G(u,v) + \alpha \ . \tag{1.1}$$

An *exact* d-diameter hopset satisfies Eq. (1.1) for $\alpha = 0$.

In this paper, we study the notion of $(D_G - 1)$-diameter hopsets which given a graph G provide the desired hopbound guarantees for $d = D_G - 1$. Our algorithms can easily provide any constant additive reduction, i.e., a $(D_G - c)$-diameter hopsets, for any desired constant c. We use the notion of $(D_G - 1)$ as a reference to the special property of the current lower bounds. That is, $(D_G - 1)$-hopsets are precisely the structures for which we currently have lower bounds for. We note that our bounds also hold in the following alternative formulation:

Given an input parameter D, then the output hopset guarantees a hopbound reduction by 1 (or by a constant term) for any u-v shortest path with hopbound at least D. That is, the reduction can be provided for all paths of length $\geq D$ and not only for $D = D_G$.

Shortcutting Tools. We use the basic shortcutting algorithm for dipaths from [23].

▶ **Lemma 1.5** (Restatement of Lemma 1.1 in [23]). *Any dipath P admits a reachability 2-diameter hopset H with $\widetilde{O}(|P|)$ edges.*

▶ **Lemma 1.6** ([25]). *For every n-vertex (possibly directed and weighted) graph G and integer $D \leq n$, there is an algorithm* ExactHopset *that computes an exact D-diameter hopset H with $\widetilde{O}((n/D)^2)$ edges.*

Our constructions employ a useful variant of reachability hopsets, in which it is required to provide the desired hopbound w.r.t to a given subset of vertices [19].

▶ **Definition 1.1** (Subset Reachability Hopset [19]). Given a graph $G = (V, E)$, a subset $S \subseteq V$ and an integer D, a set of edges $H \subseteq TC(G)$ is an (S, D)-reachability-hopset, if for every $u, v \in V$ such that $(u, v) \in TC(G)$, there is a u-v path $P_{u,v}$ (not necessarily a shortest path) in $G \cup H$ with at most D vertices from S.

For completeness we provide a complete proof for the following which also provides an additional property compared to the construction of [19]. See Appendix A for the proof.

▶ **Lemma 1.7.** *For every n-vertex DAG G, $S \subseteq V(G)$ and input parameter D, one can compute an (S, D)-reachability-hopset $H \subseteq TC(G)$ of size $\widetilde{O}(|S| + (|S|^2/D^3))$. In addition, the hopset H satisfies the following for every $u, v \in V$: If there is a shortest path $P_{u,v} \subseteq G$ which contains k vertices in $V \setminus S$, then in $G \cup H$, there is a u-v path (which is not necessarily the shortest path) that contains at most D vertices from S and at most k vertices from $V \setminus S$.*

Spanners and Emulators. Graph spanners introduced by Peleg and Schäffer [22] are sparse *subgraphs* that preserve shortest path distances up to a small stretch. In contrast to hopsets, these structures are defined only for *undirected* graphs, as no sparsificiation is possible in the worst-case for directed n-vertex graphs with $\Omega(n^2)$ edges.

▶ **Definition 1.2** ((α, β)-Spanner). Given an undirected graph $G = (V, E)$, a subgraph $G^* \subseteq G$ is called an (α, β)-*spanner* if $\text{dist}_{G^*}(u, v) \leq \alpha \cdot \text{dist}_G(u, v) + \beta$ for every $u, v \in V$.

An (α, β)-*emulator* E^* is a weighted set of edges in $V \times V$ (i.e., not necessarily a subgraph of G) that provides the same stretch guarantees as (α, β)-spanners, where $\text{dist}_G(u, v) \leq \text{dist}_{E^*}(u, v) \leq \alpha \cdot \text{dist}_G(u, v) + \beta$ for every $u, v \in V$.

Our constructions also use recent algorithms for computing additive spanners for weighted graphs [1, 8]. For a weighted graph $G = (V, E, \omega)$ where $\omega : E \to \{1, \ldots, W\}$, these spanners provide a $+\beta \cdot W$ stretch guarantees. We use the following theorem for weighted additive spanners by [2] (recently improved by [7]).

▶ **Theorem 1.8** (Theorem 3 in [2]). *For any n-vertex weighted graph $G = (V, E, \omega)$ with maximum edge weight W, there is an algorithm* WeightedSpanner *that computes a $+8W$ additive spanner $H \subseteq G$ with $\widetilde{O}(n^{4/3})$ edges.*

2 Directed Shortcuts

In this section we observe that the existing lower bounds for directed hopsets hold for the relaxed task of $(D_G - 1)$-diameter hopsets. We then show that these lower bounds can be bypassed for sparse graphs. Our upper bounds yield a separation between sparse and dense graphs, implying that shortcutting *sparse* graphs might be *simpler* in terms of providing an improved diameter vs. size tradeoffs.

2.1 Exact $(D_G - 1)$-Hopsets

Known Lower Bound. The following lower-bound follows by plugging $\ell = n^{1/3}$ in Theorem 5 of [4]. See also Table 3 in [6]. We observe that this lower bound argument, as all prior lower bounds, holds even when it is required to reduce the diameter of the given graph by 1.

▶ **Theorem 2.1** (Exact Hopsets, Directed, Follows by Theorem 5 of [4]). *There exist an n-vertex (dense) directed graphs with $\Theta(n^{3/2})$ edges for which any exact $(D_G - 1)$-hopset must have $\Omega(n^{5/4})$ shortcut edges provided that $D_G \leq n^{1/4}$.*

New Upper Bound (Proof of Theorem 1.2). We show that the lower bound of Theorem 2.1 can be bypassed for sparse graphs while preserving the exact distances in a weighted digraph. Note that for $m = n$, Theorem 1.2 yields $O(n)$-size exact hopsets for diameter $D \geq n^{1/4}$ while the lower bound of Theorem 2.1 for exact $(D_G - 1)$-diameter requires a super-linear size for $D = n^{1/4}$.

The Construction. Set an integer threshold $k = \lceil (m/D_G)^{1/3} \rceil$ and define a vertex u to be *low-deg* if $\deg_{in}(u) \leq k$ and $\deg_{out}(u) \leq k$. Otherwise, the vertex is *high-deg*. By a simple counting, there are $O(m/k)$ high-deg vertices in the given m-edge graph.

Let $L = V[p]$ for $p = \Theta(\log n/D_G)$ be a random sample of vertices, and let L_h, L_ℓ be the sets of high-deg (resp., low-deg) vertices in L (hence, $L_h \cup L_\ell = L$). The algorithm adds *two* subsets of shortcut edges H_h, H_ℓ which handle the high-deg and low-deg, respectively:

- $H_h = (L_h \times L_h) \cap TC(G)$.
- For every $u \in L_\ell$, $H_\ell(u) = \{(x, y) \mid x \in N_{in}(u), y \in N_{out}(u)\}$.
- $H_\ell = \bigcup_{u \in V} H_\ell(u)$.

All added shortcut edges in H_h, H_ℓ are weighted by their shortest-path distances in G. The output hopset is given by $H = H_h \cup H_\ell$. This completes the description of the construction.

Proof of Theorem 1.2.

Size. Since there are $O(m/k)$ high-deg vertices, w.h.p., $|L_h| = O(m \log n/(kD_G))$. Hence, $|H_h| = \widetilde{O}((m/(kD_G))^2)$. We next bound the size of H_ℓ. Let V_ℓ be the set of all low-deg vertices in G. Then, $\sum_{u \in V_\ell} \deg_{in}(u) \cdot \deg_{out}(u) = O(km)$. Since L_ℓ is obtained by sampling each low-deg vertex with probability of $\Theta(\log n/D_G)$, we get that w.h.p., $|H_\ell| = \widetilde{O}(km/D_G)$. The size argument holds by setting $k = \lceil (m/D_G)^{1/3} \rceil$.

Diameter and Distances. Consider some u-v shortest-path $P \subseteq G$ with D_G edges. First assume that $P \cap L_\ell \neq \emptyset$. I.e., that P contains at least one sampled low-deg vertex, say z. Let z', z'' be the vertex preceding (resp., subsequent to) z on the path P. Since $z' \in N_{in}(z)$ and $z'' \in N_{out}(z)$, the shortcut edge (z', z'') is in H_ℓ. Consequently, the path P can be shortcut into a path P' obtained by replacing the 2-hop segment $P[z', z''] = (z', z) \circ (z, z'')$ by a weighted edge (z', z''). Note that $\text{len}(P') = \text{len}(P)$ and that $|P'| = |P| - 1$.

From now on assume that P has no sampled low-deg vertex. W.h.p., it then holds that P contains at least two high-deg sampled vertices, i.e., $|P \cap L_h| \geq 2$. In addition, we can assume that those two sampled vertices, x, y are at hop-distance at least $D/3$ from each other. Since the weighted edge $(x, y) \in H_h$, the $\Omega(D)$-hop segment $P[x, y]$ can be replaced by the single edge (x, y). The distances are clearly preserved. Note that in this case, the hopbound can be reduced by an even a constant factor, and hence the additive reduction is bottle-necked by the low-deg vertices. ◀

2.2 Reachability $(D_G - 1)$-Hopsets

Known Lower Bound. We start by making the immediate observation that the well-known lower-bound by Huang and Pettie [14] also holds for $(D_G - 1)$-hopsets.

▶ **Theorem 2.2** (Reachability Hopsets, Slight Restatement of [14]). *There exists an n-vertex directed acyclic graph with $\Omega(n^{7/6})$ edges for which any reachability $(D_G - 1)$-hopset must have $\Omega(n)$ shortcut edges provided that $D_G \leq n^{1/6}$.*

Reachability $(D_G - 1)$-Diameter Hopsets (Proof of Theorem 1.1). When settling for reachability, rather than exact distances, one can provide improved bounds. In particular, Theorem 1.1 claims that sparse DAGs admit a reachability $(D_G - 1)$-diameter hopset of linear size for $D_G \geq n^{1/5}$ (compared to $D_G \geq n^{1/4}$ when preserving the exact distances).

Throughout, recall that G is a DAG[4]. The algorithm for reachability $(D_G - 1)$-hopsets is similar to that of Theorem 1.2. The main distinction is that instead of connecting each pair of sampled high-deg vertices in the hopset, we add the subset reachability hopset of Lemma 1.7 w.r.t. to the set of high-degree vertices. In our context, this adds $|L_h|^2/D_G^3$ edges, rather than $|L_h|^2$ edges, where L_h is the set of sampled high-degree vertices in the construction of Theorem 1.2.

Proof of Theorem 1.1. Set $k = \lceil m^{4/3}/D^{5/3} \rceil$. The definition of H_ℓ is the same, hence $|H_\ell| = \widetilde{O}(mk/D) = \widetilde{O}(m^{4/3}/D_G^{5/3})$. The subset H_h is defined by applying the subset reachability hopset of Lemma 1.7 with $S = L_h$ and $D = D_G/2$. Since $|S| = \widetilde{O}(m/(kD))$, the size of H_h can be bounded by $\widetilde{O}(|L_h| + (|L_h|^2/D_G^3)) = \widetilde{O}(m^{4/3}/D_G^{5/3})$.

It remains to consider the diameter argument. By the proof of Theorem 1.2 it is sufficient to consider a u-v shortest path P of hopbound D_G that has at least $D_G/3$ high-deg vertices and at most $D_G/3$ low-deg vertices. By Lemma 1.7, we have that H_h contains a u-v path with at most $D_G/3$ low-deg vertices and at most $D_G/3$ high-deg vertices. Note that in this case, we get a constant reduction in the diameter. The theorem follows. ◀

3 Additive Shortcuts for Undirected Unweighted Graphs

We next consider the gap obtained for exact hopsets in undirected and unweighted graphs. Our goal is to narrow this gap by (i) restricting attention to $(D_G - 1)$-hopsets (for which the known lower bound results hold), and (ii) allow a constant additive stretch in the distances. We will show that for sparse graphs, one can even bypass the current lower bound obtained for exact hopsets, as the latter is based on a dense graph example.

[4] Note that the general DAG reduction introduces a factor of 2 in the diameter, and hence we cannot employ it. All lower-bound graphs for this problem are DAGs as well.

Known Lower Bounds for Exact $(D_G - 1)$-Hopsets. We start by observing that the lower bounds for reachability hopsets by Huang and Pettie [14] also hold for exact hopsets for undirected and unweighted graphs. See Appendix A for a proof.

▶ **Theorem 3.1** (Lower Bound for Exact Hopsets, Undirected Unweighted, Immediate by [14]). *There exist n-vertex undirected and unweighted graphs G with $m = \Omega(n^{7/6})$ edges and $D_G = \Theta(n^{1/6})$, such that any exact $(D_G - 1)$-diameter hopset for G has $\Omega(n)$ edges.*

We next show that the lower bound of Theorem 3.1 can be bypassed for sparse graphs when allowing additive stretch. In addition, for general graphs (possibly dense) we match the bound of $n^{1/6}$ obtained for exact hopsets.

New Additive $(D_G - 1)$-Hopsets (Proof of Theorem 1.3). We start by presenting an algorithm that achieves a size bound of $\widetilde{O}((n/D_G)^{6/5})$ (hence, linear-size for $D_G \geq n^{1/6}$) and then explain how to modify the construction to yield the bound of $\widetilde{O}((m/D_G)^{12/11})$ (which provides the improved results for sparse graphs). For the sake of this extension, we show the following slightly stronger statement:

▶ **Lemma 3.2.** *Every n-vertex m-edge undirected unweighted graph G and any input parameter D, one can compute a hopset H with $\widetilde{O}((n/D)^{6/5})$ edges with the following guarantee: For any u, v pair at distance at least D in G, it holds that $\mathrm{dist}_G(u,v) \leq \mathrm{dist}^{(D-1)}_{G \cup H}(u,v) \leq \mathrm{dist}_G(u,v) + 24$.*

Note that Lemma 3.2 in particular implies a $+24$-additive $(D_G - 1)$-diameter hopsets with $\widetilde{O}((n/D)^{6/5})$ edges. The lemma is stronger in the sense that for any given D (where possibly $D < D_G$), the (-1) reduction in the hopbound holds for any shortest path of length at least D. In contrast, $(D_G - 1)$-hopsets provides the (-1) reduction only for shortest-paths of length *exactly* D_G.

The Construction. Let $\deg_2(u)$ be the number of 2-hop neighbors of u in the given graph G. Set $k = (n/D)^{1/5}$ as a parameter that serves as our 2-degree threshold, as follows. A vertex u is *low-deg* if $\deg_2(u) \leq k$ and it is *high-deg* otherwise. The algorithm has two phases. The first phase handles the low-deg vertices by adding $\widetilde{O}(nk/D)$ shortcut edges. At the end of that phase, the algorithm outputs also a subgraph G' in which each vertex is high-deg. The second phase handles these high-deg vertices by adding an additional subset of $\widetilde{O}(nk/D + (n/(Dk^2))^2)$ shortcut edges. The size bound of $(n/D)^{6/5}$ is obtained by balancing these two size terms, which is achieved for $k = (n/D)^{1/5}$. Throughout, all added shortcut edges are weighted by the corresponding shortest-path distances between the edge endpoints.

Step (1): Handling Low-Degree Vertices. The algorithm iterates over the low-deg vertices in the graph, as long as such exists. Initially, set $G_0 = G$ and $H_0 = \emptyset$. In every iteration $i \geq 1$, it gets as input a subgraph G_{i-1} and H_{i-1} and considers an arbitrary low-deg vertex u. (If no such exists, then $G' = G_{i-1}$, $H' = H_{i-1}$ and the step terminates). First, the algorithm removes u by letting $G_i = G_{i-1} \setminus \{u\}$. Then, with probability of $p = \Theta(\log n/D)$, the algorithm connects u to each of its (current) 2-hop neighbors in G_{i-1} by letting $H_i = H_{i-1} \cup \{(u,v) \mid v \in N_2(u, G_{i-1})\}$. This completes the description of the i^{th} iteration. Denoting the number of iterations by ℓ, then the output of the step is given by $G' = G_\ell$ and $H' = H_\ell$.

Step (2): Handling the Remaining High-Degree Vertices. We now restrict attention only to the graph G'. Letting $V' = V(G')$, the algorithm computes *three* random samples of V'-vertices: $S = V'[p]$, $Q = V'[q]$ and $R = V'[r]$ where $p = \Theta(\log n/D)$, $q = \Theta(\log n/k)$ and $r = \Theta(\log n/(Dk^2))$. Also, initially set $H'' = (R \times R)$, and add to H'' a subset of shortcut edges by applying the following shortcutting procedure for every vertex $u \in S$:

- Build a depth-6 BFS tree $T_6(u, G') \subseteq G'$ rooted at u (i.e., a BFS which spans only $N_6(u, G')$).
- If $|T_6(u, G')| = O(k^2)$, add the edges $E'(u) = \{(u, v) \mid v \in T_6(u, G') \cap Q\}$ to H''.

The output hopset is given by $H' \cup H''$. We next turn to analyze the construction and prove Lemma 3.2.

Size. In the first step, for every low-deg vertex u, with probability of $p = \Theta(\log n/D)$, the algorithm adds $\deg_2(u) \leq k$ shortcut edges. Hence this adds $|H'| = \widetilde{O}(nk/D)$ edges, w.h.p. Consider the second step. W.h.p., $|S| = O(n \log n/D)$ and $|R| = O(n \log n/(Dk^2))$. For every $u \in S$ with $|T_6(u, G')| = O(k^2)$, we have that w.h.p. $|T_6(u, G') \cap Q| \leq k \log n$, and therefore, we have $O(|S| \cdot k \log n)$ edges, due to this step. Overall, we added to H'' a total of $O(|R|^2 + |S| \cdot k \log n)$ edges. The size bound follows by plugging $k = (n/D)^{1/5}$.

Diameter and Stretch Analysis. Throughout, we override notation and redefine a vertex v to be low-deg if $v \notin V(G')$. That is, a vertex is low-deg if there exists an iteration i in the Step (1) in which $N_2(v, G_{i-1}) \leq k$. A vertex v is then high-deg if $v \in V(G')$ (i.e., it is a high-deg in each subgraph G_{i-1} considered in each iteration i of Step (1)). We also need the following classification of the high-deg vertices. A vertex $u \in V(G')$ is *large* if its 6-depth BFS tree has size $|T_6(u, G')| = \omega(k^2)$, and it is *small* otherwise. Let $V_\ell = V \setminus V(G')$ be the low-deg vertices, V_h^{small} (resp., V_h^{large}) denote the small (resp., large) high-deg vertices. It then holds that $V_\ell \cup V_h^{small} \cup V_h^{large} = V(G)$.

Consider now a u-v shortest-path P with D edges (hops). The analysis breaks into three cases depending on the number of vertices in $V(P)$ that belong to each of the three subsets of vertices V_ℓ, V_h^{small} and V_h^{large}.

Case 1: $|P \cap V_\ell| = \Omega(D)$. For every low-deg vertex z, let i_z be the iteration in which z is removed in Step (1). I.e., z is the (unique) selected low-deg vertex in G_{i_z-1}. We then say that the low-deg vertex $z \in P \cap V_\ell$ is *bad* if both of its 2-hop P-neighbors z_1, z_2 are not in G_{i_z-1}. This can happen if these two 2-hop neighbors were removed in prior iterations. Otherwise, the vertex is *good*.

▶ **Lemma 3.3.** *The path P can contain at most* two *consecutive bad vertices.*

Proof. The claim follows by noting that a 2-hop P-neighbor of a bad vertex $x \in P$ must be good. To see this, let z be a 2-hop P-neighbor of x. Since x is bad, it implies that $i_x \geq i_z + 1$, i.e., z is removed *before* x in Step (1). This implies that $x \in G_{i_z-1}$ and hence z must be good. ◀

By Lemma 3.3, we get that in this case P contains $\Omega(D)$ good vertices. Since the algorithm add shortcut edges to each good vertex w.p. $p = \Theta(\log n/D)$, we get that w.h.p., the coin flip is successful for at least one good vertex, say x, on P. By the definition of the good vertex, at least one of its 2-hop P-neighbors, say x', is in G_{i_x-1} and therefore the shortcut edge (x', x) is in H'. Let P' be the path obtained by replacing the 2-hop segment $P[x, x']$ with an edge (x, x'). Then, $\text{len}(P) = \text{len}(P')$ but $|P'| \leq |P| - 1$, as required.

Case 2.1: $|P \cap V_h^{large}| = \Omega(D)$. Note that for any for any two vertices u', u'' on P at distance at least 13 from each other, it holds that their 6-hop neighborhoods are disjoint. Also note that since $N_6(u', G') \subseteq N_6(u', G)$, we have that $|N_6(u', G)| = \omega(k^2)$ for every $u' \in V_h^{large}$. As P contains $\Omega(D)$ vertices from V_h^{large}, we get that the size of the 6-hop neighborhood of the path P is $\Omega(Dk^2)$. Since we sample each vertex in V into R independently with probability of $r = \Theta(\log n/(Dk^2))$, we get that w.h.p., the following holds: There are two vertices $u', u'' \in S \cap P \cap V_h^{large}$ at distance $\Omega(D)$ from each other such that there exists $w \in N_6(u', G) \cap R$ and $w' \in N_6(u'', G) \cap R$. Since the algorithm adds to H'' the shortcut edge (w, w'), the hopbound between u and u' is reduced from $\Omega(D)$ to at most 13. By the triangle inequality, this introduces an additive stretch of at most $+24$. See Fig. 1 (Left) for an illustration.

Case 2.2: $|P \cap V_h^{small}| = \Omega(D)$. Since P has $\Omega(D)$ vertices from V_h^{small}, we can also assume the following. The path P contains $\ell = \Omega(D)$ segments $P_1, \ldots, P_\ell$ such that: (i) each $P_i \subseteq G'$, (ii) $|P_i| = 40$ and (iii) the internal 20-length segment of P_i contains a vertex in V_h^{small}. We then get that w.h.p. there exists a vertex $u' \in S \cap P \cap V_h^{small}$ that belongs to the internal 20-length segment of some $P_i \subseteq G'$. Let w be a vertex at distance 4 from u' on P_i. Since each vertex in G' has 2-deg at least k, we get (i) $|N_2(w, G')| \geq k$ and (ii) $N_2(w, G') \subset T_6(u', G')$. Therefore, w.h.p., it holds that there exists some $z \in N_2(w, G') \cap Q$ (where $Q = V'[\Theta(\log n/k)]$) and consequently, $(u', z) \in E'(u)$. The 4-hop path segment $P[u', w]$ can then be replaced by a 3-hop segment $P' = (u', z) \circ (z, z') \circ (z', w)$ for some $z' \in N(w, G') \cap N(z, G')$. It is easy to see that the additive stretch is at most $+4$, as we replace a 4-hop segment by a 3-hop segment of length at most 8. See Fig. 1 (Right) for an illustration. Lemma 3.2 follows.

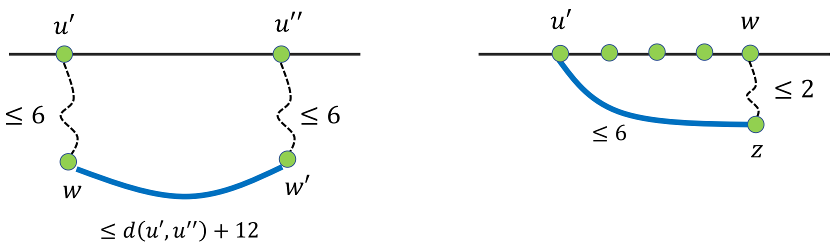

Figure 1 An illustration for stretch and diameter argument of Theorem 1.3. Left: The path P contains $\Omega(D)$ vertices which are high-deg and an in addition with large 6-depth BFS trees. The shortcut edge is shown in blue. Right: The path P has $\Omega(D)$ vertices which are high-deg and with small 6-depth BFS trees.

We are now ready to complete the proof of Theorem 1.3.

Theorem 1.3. It remains to modify the construction to obtain a size bound of $\widetilde{O}((m/D_G)^{12/11})$ edges. We start with a preliminary sparsification step that handles vertices with 1-degree (that is simply the degree) at most $k' = (m/D_G)^{1/11}$. Let $V_{\ell,1}$ be the subset of all low-degree vertices (i.e., with 1-deg at most k') and let $V_{h,1} = V \setminus V_{\ell,1}$. Let H_1 be a subset of shortcut edges that handles the low-degree vertices, as follows. Let $V'_{\ell,1} = V_{\ell,1}[p]$ for $p = \Theta(\log n/D_G)$. Then,

$$H_1 = \bigcup_{v \in V'_{\ell,1}} \{(x, y) \mid x, y \in N(v) \text{ and } x \neq y\} .$$

Next, we apply the construction of Lemma 3.2 on the graph $G_h = G[V_{h,1}]$ with $D = D_G$, which outputs a $+24$-additive $(D_G - 1)$-diameter hopset H_2 on the graph G_h. Observe that the diameter of G_h might be considerably larger than D. This motivates the more dedicated guarantees of Lemma 3.2. The output hopset is given by $H_1 \cup H_2$.

Size. By the Chernoff bound, w.h.p., $|H_1| = \widetilde{O}(mk'/D_G)$. In addition, $|V_{h,1}| = O(m/k')$ by a simple counting argument. Letting $n' = |V_{h,1}| = O(m/k')$, by Lemma 3.2, $|H_2| = \widetilde{O}((n'/D_G)^{6/5})$. This size bound follows by plugging $k' = (m/D_G)^{1/11}$.

Diameter and Stretch. Consider a u-v shortest-path P with D_G edges. If P contains $\Omega(D_G)$ vertices of degree at most k', then w.h.p., $P \cap V'_{\ell,1} \neq \emptyset$, and P is shortcut by one hop, and the distances are preserved. It remains therefore to consider the case where P contains $\ell = \Omega(D_G)$ segments $P_1, \ldots, P_\ell$, each of length 100, that are fully contained in the graph $G_h = G[V_h]$. For every $i \in \{1, \ldots, \ell\}$, letting $P_i = [u_{i1}, \ldots, u_{i100}]$ then define an internal segment $P'_i = [u_{i50}, \ldots, u_{i80}] \subset P_i$.

We next show that this suffices to recover the same argument as obtained in Lemma 3.2. Partition the vertices on $P \cap V_{h,1}$ into three subsets $V_\ell, V_h^{small}, V_h^{large}$ as in the argument of Lemma 3.2. We then consider the same cases as in Lemma 3.2 with minor modifications.

Case 1: V_ℓ intersects with $\Omega(D_G)$ distinct segments of $P'_1, \ldots, P'_\ell$. By Lemma 3.3, we get that w.h.p. the algorithm adds a shortcut edge between some vertex $V_\ell \cap P$ to its 2-hop neighbor on the path. Hence, the diameter (hopbound) is reduced by 1, and the distances are preserved.

Case 2.1: V_h^{large} intersects with $\Omega(D_G)$ distinct segments of $P'_1, \ldots, P'_\ell$. Note that any for any two vertices u, v on P at distance at least 13 from each other, it holds that their 6-hop neighborhoods are disjoint. This implies that the 6-hop neighborhood of the path P is $\Omega(Dk^2)$. By the exact same argument as obtained for Case 2.1 in Lemma 3.2, we get that there is at least one segment P_j whose hopbound is reduced while introducing an additive stretch of at most $+24$. This holds as the argument in Lemma 3.2 is local in the sense that it shortcuts segments of constant length provided that there are sampled $u', u'' \in V_h^{large}$ that are sufficiently apart on P and that each has a 4-hop P-neighbor on the segment.

Case 2.2: V_h^{small} intersects with $\Omega(D_G)$ distinct segments of $P'_1, \ldots, P'_\ell$. The claim follows by noting that the argument for Case 2.2 of Lemma 3.2 shows that any segment P_j gets shortcut (and while introducing an additive stretch of $+4$) with probability of $\Theta(\log n/D_G)$. The probabilities are independent for vertex-disjoint segments. Note that this holds since it is sufficient for the segment P_j to include the 6-hop neighborhood of the vertex. Since V_h^{small} intersects with $\Omega(D_G)$ disjoint segments, at least one of them is successful, w.h.p., and provides the desired (-1) reduction in the hopbound (while introducing an additive stretch of $+4$). ◀

New Additive D-Diameter Hopsets (for any D). We now turn to proving Theorem 1.4. For simplicity we show a construction of $O(D)$-diameter hopsets with additive stretch $O(\alpha)$, but these constant factors can be easily omitted. The algorithm has two main steps. The first step computes a graph G' on $O(n \log n/\alpha)$ vertices at the cost of introducing an additive stretch of $O(\alpha)$. The second step computes an exact D-diameter hopset on G'.

Specifically, the algorithm starts by computing a weighted *net* graph $G' = \mathsf{Net}(G, p)$ for the given (unweighted) graph G where $p = \Theta(\log n/\alpha)$. The algorithm $\mathsf{Net}(G, p)$ outputs a graph $G' = (V', E', \omega')$ defined as follows. Let $V' = V[p]$ be a random sample of V, obtained by sampling each $v \in V$ independently with probability of p. Let $E' = \{(u, v) \in V' \times V' \mid \text{dist}_G(u, v) \leq \Theta(\log n/p) \cdot W\}$ and $\omega'((u, v)) = \text{dist}_G(u, v)$ for every $(u, v) \in E'$. We use the following observation in our constructions:

▶ **Observation 3.1** (Observation 3.4 in [20]). Let $G' = (V', E', \omega')$ be the output net graph of Alg. $\mathsf{Net}(G, p)$ where $G = (V, E, \omega)$ is an n-vertex graph with maximum edge weight W. Then w.h.p., the following holds: (i) $|V'| = O(np \log n)$, (ii) for every $u, v \in V'$, $\text{dist}_{G'}(u, v) = \text{dist}_G(u, v)$, and (iii) the maximum edge weight of G' is bounded by $W' = \Theta(W \log n/p)$.

Denote $S_1 = V(G')$, hence $S_1 = V[p]$. By Obs. 3.1(iii), the maximum edge weight of G' is $W = O(\alpha)$. Let $S_2 = S_1[q]$ for $q = \Theta(\log n/D)$. For a vertex u and a set S_1, let $\mathit{Closest}(u, S_1)$ be the closest vertex to u in S_1, breaking ties arbitrarily. The output hopset H is the union of $H_0 \cup H_1 \cup H_2$ of weighted edges, where:

- $H_0 \leftarrow \{(u, \mathit{Closest}(u, S_1)) \mid u \in V\}$.
- $H_1 \leftarrow \mathsf{WeightedSpanner}(G')$ of Theorem 1.8.
- $H_2 \leftarrow \mathsf{ExactHopset}(H_1, D)$ of Lemma 1.6.

This completes the description of the hopset.

Proof of Theorem 1.4. Clearly, $|H_0| \leq n$. By the Chernoff bound, w.h.p., $|S_1| = O(n \log^2 n/\alpha)$, and therefore by Theorem 1.8, $|H_1| = \widetilde{O}((n/\alpha)^{4/3})$. The size of H_2 can bounded by $|S_2|^2 = \widetilde{O}((n/(\alpha \cdot D))^2)$, w.h.p. We next turn to consider the additive stretch and the hopbound.

Consider a u, v pair and let $u' = \mathit{Closest}(u, S_1)$ and $v' = \mathit{Closest}(v, S_1)$. W.h.p., it holds that $\text{dist}_G(u, u'), \text{dist}_G(v, v') = O(\alpha)$. Since, $u', v' \in S_1$, by Obs. 3.1 we have that $\text{dist}_{G'}(u', v') = \text{dist}_G(u, v)$. Hence, $\text{dist}_{H_1}(u', v') \leq \text{dist}_{G'}(u', v') + O(\alpha)$. Since H_2 is an exact D-diameter hopset for H_1, we get:

$$\text{dist}^{(D)}_{H_1 \cup H_2}(u', v') = \text{dist}_{H_1}(u', v') \leq \text{dist}_G(u', v') + O(\alpha) \ .$$

Letting P be the shortest u'-v' path with at most D hops in $H_1 \cup H_2$, we get that the u-v path $P' = (u, u') \circ P \circ (v', v) \subseteq G \cup H$ has at most $O(D)$ hops and of total length $\text{dist}_G(u, v) + O(\alpha)$. The theorem follows. ◀

References

1 Abu Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen G. Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Comput. Sci. Rev.*, 37:100253, 2020.

2 Abu Reyan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Stephen G. Kobourov, and Richard Spence. Weighted additive spanners. In Isolde Adler and Haiko Müller, editors, *Graph-Theoretic Concepts in Computer Science – 46th International Workshop, WG 2020, Leeds, UK, June 24-26, 2020, Revised Selected Papers*, volume 12301 of *Lecture Notes in Computer Science*, pages 401–413. Springer, 2020.

3 Aaron Bernstein, Maximilian Probst Gutenberg, and Christian Wulff-Nilsen. Near-optimal decremental SSSP in dense weighted digraphs. In *61st IEEE Annual Symposium on Foundations of Computer Science, FOCS 2020, Durham, NC, USA, November 16-19, 2020*, pages 1112–1122. IEEE, 2020.

4 Greg Bodwin. New results on linear size distance preservers. *SIAM J. Comput.*, 50(2):662–673, 2021.

5 Greg Bodwin and Gary Hoppenworth. Folklore sampling is optimal for exact hopsets: Confirming the $\sqrt{n}$ barrier. *CoRR*, abs/2304.02193, 2023. doi:10.48550/arXiv.2304.02193.

6 Greg Bodwin, Gary Hoppenworth, and Ohad Trabelsi. Bridge girth: A unifying notion in network design. *CoRR*, abs/2212.11944, 2022. doi:10.48550/arXiv.2212.11944.

7 Michael Elkin, Yuval Gitlitz, and Ofer Neiman. Improved weighted additive spanners. In Seth Gilbert, editor, *35th International Symposium on Distributed Computing, DISC 2021, October 4-8, 2021, Freiburg, Germany (Virtual Conference)*, volume 209 of *LIPIcs*, pages 21:1–21:15. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.

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

9 Jeremy T. Fineman. Nearly work-efficient parallel algorithm for digraph reachability. *SIAM J. Comput.*, 49(5), 2020. doi:10.1137/18M1197850.

10 Sebastian Forster and Danupon Nanongkai. A faster distributed single-source shortest paths algorithm. In Mikkel Thorup, editor, *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 686–697. IEEE Computer Society, 2018.

11 Maximilian Probst Gutenberg and Christian Wulff-Nilsen. Decremental SSSP in weighted digraphs: Faster and against an adaptive adversary. In Shuchi Chawla, editor, *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2542–2561. SIAM, 2020.

12 Monika Henzinger, Sebastian Krinninger, and Danupon Nanongkai. Improved algorithms for decremental single-source reachability on directed graphs. In Magnús M. Halldórsson, Kazuo Iwama, Naoki Kobayashi, and Bettina Speckmann, editors, *Automata, Languages, and Programming – 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, volume 9134 of *Lecture Notes in Computer Science*, pages 725–736. Springer, 2015.

13 William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms, January 12-14, 2003, Baltimore, Maryland, USA*, pages 665–669. ACM/SIAM, 2003.

14 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. doi:10.1137/19M1306154.

15 Arun Jambulapati, Yang P. Liu, Yang P. Liu, and Aaron Sidford. Parallel reachability in almost linear work and square root depth. In David Zuckerman, editor, *60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, Baltimore, Maryland, USA, November 9-12, 2019*, pages 1664–1686. IEEE Computer Society, 2019. doi:10.1109/FOCS.2019.00098.

16 Philip N. Klein and Sairam Subramanian. A randomized parallel algorithm for single-source shortest paths. *J. Algorithms*, 25(2):205–220, 1997.

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

18 Shimon Kogan and Merav Parter. New diameter-reducing shortcuts and directed hopsets: Breaking the barrier. In Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022*, pages 1326–1341. SIAM, 2022.

19 Shimon Kogan and Merav Parter. Faster and unified algorithms for diameter reducing shortcuts and minimum chain covers. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 212–239. SIAM, 2023. doi:10.1137/1.9781611977554.ch9.

20 Shimon Kogan and Merav Parter. New additive emulators. In *50th International Colloquium on Automata, Languages, and Programming (ICALP 2023)*. Schloss Dagstuhl-Leibniz-Zentrum für Informatik, 2023.

21 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 Joseph (Seffi) Naor and Niv Buchbinder, editors, *Proceedings of the 2022 ACM-SIAM Symposium on Discrete Algorithms, SODA 2022, Virtual Conference / Alexandria, VA, USA, January 9 – 12, 2022*, pages 3311–3331. SIAM, 2022.

22 David Peleg and Alejandro A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.

23 Sofya Raskhodnikova. Transitive-closure spanners: A survey. In Oded Goldreich, editor, *Property Testing – Current Research and Surveys*, volume 6390 of *Lecture Notes in Computer Science*, pages 167–196. Springer, 2010. doi:10.1007/978-3-642-16367-8_10.

24 Mikkel Thorup. On shortcutting digraphs. In Ernst W. Mayr, editor, *Graph-Theoretic Concepts in Computer Science, 18th International Workshop, WG '92, Wiesbaden-Naurod, Germany, June 19-20, 1992, Proceedings*, volume 657 of *Lecture Notes in Computer Science*, pages 205–211. Springer, 1992.

25 Jeffrey D. Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM J. Comput.*, 20(1):100–125, 1991. doi:10.1137/0220006.

A Missing Proofs

Proof of Lemma 1.7. Let $TC' = TC(G)[S]$ be the induced transitive closure on the subset S. Compute a collection of $O(|S|/D)$ vertex-disjoint paths $\mathcal{P}$ in TC' such that $TC' \setminus V(\mathcal{P})$ has no directed path of length $D/10$. Let $H_1 = \bigcup_{P \in \mathcal{P}} H(P)$ where $H(P)$ is a 2-diameter hopset for P that consists of $O(|P| \log |P|)$ edges, using Lemma 1.5. Let $S' = S[p]$ and $\mathcal{P}' = \mathcal{P}[p]$ for $p = \Theta(\log n/D)$. Let $H_2 = \{e(v,P) \mid v \in S', P \in \mathcal{P}'\}$ where $e(v,P)$ is an edge connecting v to its first outgoing neighbor in TC' on P. The size bound is immediate.

Consider the diameter argument. Let Q be a u-v shortest path in G for $u, v \in S$ such that $|Q \cap S| \geq D$. By the properties of the paths $\mathcal{P}$, we can assume that $|Q \setminus V(\mathcal{P})| \leq D/10$. We next show that Q can be transformed into a path $Q' \subseteq G \cup H_1$ such that the following holds: (i) for each $P \in \mathcal{P}$, $|P \cap Q'| \leq 3$ and $V(Q') \setminus V(\mathcal{P}) \subseteq V(Q)$. This can be obtained by traversing Q and at each point of observing a vertex $z \in P \cap Q$, we add the shortcut edges $H(P)$ to connect z with the far most vertex $z' \in P \cap Q$. It is easy to see that $V(Q') \subseteq V(Q)$ as by shortcutting Q we can only omit vertices.

Finally, in the case where $|Q' \cap S| \geq D/2$, by property (i), we have that Q' intersects with $\Omega(D)$ distinct paths from $\mathcal{P}$. Let $w \in S$ be the first sampled vertex in $Q' \cap S'$ (w.h.p., such exists among the first $D/10$ many S' vertices on P). Let $w' \in S$ be the last vertex on Q' that belongs to a sampled path P' in $\mathcal{P}'$ (w.h.p., such exists among that last $D/10$ many S' vertices on P). The diameter argument holds by noting that the shortcut edge $e(w, P')$ is in H_2. ◀

Proof of Theorem 3.1. The lower bound graph of [14] is a DAG with D_G layers which contain $\Omega(n)$ critical pairs. For each critical pair $\langle u, v\rangle$, u belongs to layer 1 and v belongs to layer D_G. Furthermore, there is a unique directed path in G between u to v and this path contains exactly one vertex from each layer. We now remove the directions of the edges in G to get an undirected and unweighted graph G'. Notice that now for each critical pair $\langle u, v\rangle$ in G', we have a unique shortest path of length $D_G - 1$ between u and v, that is there might be other paths between u and v but the length of such path will be greater than $D_G - 1$ because such path will necessarily contain at least one vertex from each level, and for a certain level i there will be two vertices from level i. In particular, the path will contain vertices u, w, v consecutively in the path, where u is in level i, w is in level $i-1$ and v is in level i. We therefore have that the unique directed u-v paths in G are translated into unique undirected unweighted shortest paths in G'. Therefore, the claim holds in the exact same manner as in [14]. ◀