

Bridge Girth: A Unifying Notion in Network Design*

Greg Bodwin¹, Gary Hoppenworth¹, and Ohad Trabelsi²

¹University of Michigan EECS. {bodwin,garytho}@umich.edu

²Toyota Technological Institute at Chicago. ohadt@ttic.edu[†]

Abstract

A classic 1993 paper by Althöfer et al. proved a tight reduction from spanners, emulators, and distance oracles to the extremal function γ of high-girth graphs. This paper initiated a large body of work in network design, in which problems are attacked by *reduction* to γ or the analogous extremal function for other girth concepts. In this paper, we introduce and study a new girth concept that we call the *bridge girth of path systems*, and we show that it can be used to significantly expand and improve this web of connections between girth problems and network design. We prove two kinds of results:

- We write the maximum possible size of an n -node, p -path system with bridge girth $> k$ as $\beta(n, p, k)$, and we write a certain variant for “ordered” path systems as $\beta^*(n, p, k)$. We identify several arguments in the literature that implicitly show upper or lower bounds on β, β^* , and we provide some polynomial improvements to these bounds. In particular, we construct a tight lower bound for $\beta(n, p, 2)$, and we polynomially improve the upper bounds for $\beta(n, p, 4)$ and $\beta^*(n, p, \infty)$.
- We show that many state-of-the-art results in network design can be recovered or improved via black-box reductions to β or β^* . Examples include bounds for distance/reachability preservers, exact hopsets, shortcut sets, the flow-cut gaps for directed multicut and sparsest cut, an integrality gap for directed Steiner forest.

We believe that the concept of bridge girth can lead to a stronger and more organized map of the research area. Towards this, we leave many open problems related to both bridge girth reductions and extremal bounds on the size of path systems with high bridge girth.

*This work was supported by NSF:AF 2153680.

[†]Work partially done at University of Michigan, and partially supported by the NSF grant CCF-1815316 and the NWO VICI grant 639.023.812.

Contents

1	Introduction	1
1.1	Bridge Girth and the Landscape of Girth Reductions	2
1.2	Reductions from Network Design to Bridge Girth	4
1.3	New Extremal Bounds for Bridge Girth	8
1.4	Future Directions and Open Problems	10
2	Preliminaries on Path Systems	13
3	Bounds on β, β^* Functions	14
3.1	Lower Bounds for $k = 2$	14
3.2	Upper Bounds for $k = 4$	16
3.3	Proof Wrapup	24
3.4	Bounds for $k = \infty$	26
4	Extremal Reductions to Bridge Girth	29
4.1	Distance Preservers	29
4.2	Shortest Path Oracles	31
4.3	Reachability Preservers	32
4.4	Path Oracles	34
4.5	Online Reachability Preservers	35
4.6	Shortcut Sets and Exact Hopsets	36
4.7	Approximate Distance Preservers and Evidence for Conjecture 9	37
5	Flow-Cut and Directed Steiner Forest Integrality Gaps	38
5.1	Directed Edge Multicut and Statement of First Result	38
5.2	Directed Vertex Multicut	40
5.3	Multicut Flow-Cut Gap Construction	40
5.4	Integrality Gap Analysis of G	42
5.5	Sparsest Cut Flow-Cut Gap	44
5.6	Integrality Gap of the Flow LP of Directed Steiner Forest	45
A	A Tour through Prior Work on Girth Problems	55
A.1	The Girth Problem	55
A.2	The Weighted Girth Problem	56
A.3	The Bipartite Girth Problem	57
A.4	The Ruzsa-Szemerédi Problem	58
A.5	The Set Girth Problem	58
B	Proof of Cleaning Lemma	59
C	Missing Proofs for Section 5	60
C.1	Lemma 57	60
C.2	Theorem 64	61
C.3	Claim 66	61
C.4	Lemma 67	62
D	Implicit Bounds on β in Prior Work	64
D.1	Upper Bounds for $k = 2$	64
D.2	Bounds for $k = 3$	64

1 Introduction

A common goal in theoretical computer science is to compress a graph into a small-space representation while approximately preserving structural information related to shortest paths, distances, or reachability. Examples include spanners [1, 3, 6, 7, 11, 14, 15, 32, 52, 69, 70, 84, 85, 86, 98, 104, 105], emulators [45, 52, 97], distance oracles [33, 51, 53, 88, 96, 106], distance and reachability preservers [2, 13, 21, 22, 27, 29, 30, 41], hopsets [17, 38, 47, 48, 49, 66, 81], shortcut sets [57, 64, 65, 68, 71, 72, 80, 95, 100], etc.; see survey [5] for more. We shall broadly refer to this research area as *network design*.

A successful strategy has been to reduce network design problems to *girth problems* in extremal combinatorics. Generally speaking, a girth problem asks for the maximum possible size of a combinatorial system that avoids short “cycles” of some kind. The contribution of this paper is to introduce a new girth problem, based on a particular kind of cycle in *path systems* that we call “bridges.” We then use our new girth problem to organize and improve the understanding of several well-studied problems in network design. This paper contains two kinds of results:

1. We polynomially improve upper and lower bounds on the maximum possible size of path systems of high bridge girth (over bounds implicit in the previous literature), and
2. We show reductions from various problems in network design to our new bridge girth problem, and use them to recover or improve state-of-the-art upper and/or lower bounds.




Figure 1: Relationships among some girth problems in the literature, and the problems they capture. Bridge girth, ordered bridge girth, and the associated reductions – all in blue – are new in this paper. See Appendix A for more detail on the prior work reflected in this chart.

1.1 Bridge Girth and the Landscape of Girth Reductions

1.1.1 Previously-Studied Girth Concepts

In Appendix A, we provide a detailed tour through all the objects mentioned in Figure 1. Here we give a much faster overview, to show how our new concept of bridge girth relates to previously-studied girth concepts in the literature.

The strategy of reducing network design problems to girth problems was pioneered in a classic 1993 paper by Althöfer, Das, Dobkin, Joseph, and Soares [11]. They provided an exactly-tight reduction between the extremal functions of *spanners* and *high-girth graphs*. We will rephrase their main result a bit, to give an example of the reduction-based perspective used in this paper.

Definition 1 (Multiplicative Spanners). *A multiplicative k -spanner of a graph G is a subgraph H satisfying $\text{dist}_H(s, t) \leq k \cdot \text{dist}_G(s, t)$ for all nodes s, t . The function $MS(n, k)$ is the least integer such that every undirected weighted n -node graph has a k -spanner on $\leq MS(n, k)$ edges.*

Definition 2 (Graph Girth). *The girth of a graph G is the least number of edges in a cycle in G (or ∞ if G is a forest). The function $\gamma(n, k)$ is the maximum possible number of edges in an n -node graph of girth $> k$.*

Theorem 1 ([11]). $\gamma(n, k) = MS(n, k + 1)$.

It remains a major open question in extremal combinatorics to determine the asymptotic value of γ . Regardless, by tight reduction, Theorem 1 is considered by the community to close the question of the existential size of multiplicative spanners. It has been highly influential in network design, spawning a long line of work similarly reducing spanner or spanner-like problems to γ [23, 25, 34, 44, 50, 56, 77, 89].

More recently, some variants and extensions of the function γ have emerged as similarly fundamental in network design. One example is an elegant paper by Elkin, Neiman, and Solomon, which developed a notion of “weighted girth,” and showed that the corresponding extremal function is equivalent to the tradeoff between stretch and *lightness* for spanners [50] (see Section A.2 for details). Another example is a line of this work [24, 42, 101] that has developed reductions to $\gamma\gamma$, a generalization of γ to bipartite graphs:

Definition 3 (Bipartite Graph Girth). *The function $\gamma\gamma(n, p, k)$ is the maximum possible number of edges in a bipartite graph with n, p nodes on each side of its bipartition and girth $> k$.*

There is a folklore reduction showing that $\gamma(n, k) = \Theta(\gamma\gamma(n, n, k))$ (see Theorem 72), and so any reductions to $\gamma(n, k)$ can be equivalently phrased as a reduction to $\gamma\gamma(n, n, k)$. The importance of $\gamma\gamma$ was further shown in an important paper by de Caen and Székely [42], which proves an equivalence between a special case of $\gamma\gamma$ and the *Ruzsa-Szemerédi function* (see Appendix A.4), another extremal function that captures various problems in network design [22, 76]. The function $\gamma\gamma$ is sometimes expressed in the equivalent language of *set system girth*:

Definition 4 (Set Systems). *A set system is a pair $S = (V, \mathcal{T})$, where V is a ground set of nodes and $\mathcal{T}$ is a multiset of subsets of V . The size of S is given by*

$$\|S\| := \sum_{T \in \mathcal{T}} |T|.$$

Definition 5 (Set System Girth). *A k -cycle in a set system $S = (V, \mathcal{T})$ is a circularly-ordered list of distinct nodes $v_0, v_1, \dots, v_k = v_0 \in V$ and distinct sets $T_0, T_1, \dots, T_k = T_0 \in \mathcal{T}$ for which we have $v_i, v_{i+1} \in T_i$ for all i . The girth of a set system is the smallest integer k for which the system has a k -cycle. The maximum possible size of a set system with n nodes, s sets, and girth $> k$ is written $\Sigma(n, s, k)$.*

For example, a set system in which each set has size 2 can be viewed as an undirected graph.¹ Set systems are in natural bijection with their bipartite *incidence graphs*, and this implies the folklore equivalence $\Sigma(n, p, k) = \gamma\gamma(n, p, 2k)$ (see Theorem 74). Thus, we can consider the set system girth problem as merely a rephrasing of the bipartite girth problem.

¹A set system is equivalent to a (not necessarily uniform) hypergraph. We call these set systems rather than hypergraphs (1) to emphasize the way in which path systems can be viewed as a directed variant, and (2) because there are several competing notions of hypergraph size/girth in the literature, but these terms are unambiguous for set systems.

1.1.2 Bridge Girth and Ordered Bridge Girth

A directed version of a set system is a *path system*, in which we have node sequences instead of sets:

Definition 6 (Path Systems). *A path system is a pair $S = (V, \Pi)$ where V is a ground set of nodes and Π is a multiset of vertex sequences called paths. Each path may contain at most one instance of each node. The size of a path system is written²*

$$\|S\| := \sum_{\pi \in \Pi} |\pi|.$$

For example, a path system in which all paths have length 2 is essentially a directed graph. Our new girth concept is based on the following notion of a “cycle” in a path system:

Definition 7 (*b*-Bridges). *In a path system $S = (V, \Pi)$, a *b*-bridge is a set of *b* distinct nodes $v_1, \dots, v_b \in V$ and *b* distinct paths $\pi_1, \dots, \pi_b$ such that (1) for all $1 \leq i \leq b-1$ we have $v_i, v_{i+1} \in \pi_i$ with v_i preceding v_{i+1} , and (2) we have $v_1, v_b \in \pi_b$ with v_1 preceding v_b . The path π_b is called the *river*, and the other paths $\pi_1, \dots, \pi_{b-1}$ are called *arcs*.*

Note that the nodes v_i, v_{i+1} are not necessarily consecutive on their arc π_i ; there might be many nodes between these, and it still counts as a bridge. Informally, a *b*-bridge resembles a directed *b*-cycle with one of the path directions reversed; the reversed path is called the river, and the non-reversed paths are called the arcs. See Figure 2 for an example.




Figure 2: A 4-bridge

Definition 8 (Bridge Girth). *The bridge girth of a path system S is the smallest integer b for which S contains a *b*-bridge. The function $\beta(n, p, k)$ is the maximum possible size of a path system with n nodes, p paths, and bridge girth $> k$. We allow $k = \infty$, meaning that the system has no bridges of any size.*

We will also consider a related notion of bridge girth, based on path systems with an ordering on their paths.

Definition 9 (Ordered Path Systems and Ordered Bridges). *An ordered path system is a path system $S = (V, \Pi)$ equipped with a total ordering of its paths Π . An ordered bridge in an ordered path system is a bridge in which the river comes after all the arcs in the ordering (and the arcs may occur in any order relative to each other).*

Definition 10 (Ordered Bridge Girth). *The ordered bridge girth of an ordered path system S is the smallest integer b for which S has an ordered *b*-bridge (with the river last in the ordering). The function $\beta^*(n, p, k)$ is the maximum possible size of an ordered path system with n nodes, p paths, and ordered bridge girth $> k$.*

It is immediate from the definitions that $\beta^*(n, p, k) \geq \beta(n, p, k)$, since an ordered path system of ordered bridge girth $> k$ is a strictly less constrained object than an (unordered) path system of (unordered) bridge girth $> k$.

²Note that $|\pi|$ counts the number of *nodes* in π , and so it differs by 1 from the length of π when viewed as a path through a graph.

We discuss our definitions before proceeding to their applications. Notice that a bridge is a directed version of a set cycle, in the sense that a b -bridge becomes a set b -cycle if we forget the order of each path and interpret it as a set. However, there are many other patterns besides bridges that correspond to set cycles in the same way. Perhaps the most natural alternative is a directed cycle, defined like a bridge but with the paths/nodes circularly ordered instead of having a river with reversed direction (see Figure 3). Why focus on bridges rather than directed cycles?




Figure 3: A set system 4-cycle, a 4-bridge, and directed 4-cycle (left to right).

In part, this paper is a lengthy answer to this question. Our *point* is that, out of the many possible directed versions of set cycles, bridges are the ones that realize applications in network design. Table 1 summarizes our findings to this effect, listing objects where state-of-the-art bounds can be improved, or recovered in a black-box manner, via bridge girth reductions.

But a more succinct answer is that directed cycles do not have an interesting extremal function. One can construct a path system of size $\|S\| = np$ without directed cycles, by taking p identical paths that all proceed through a sequence of the n nodes in the same order. This system would have no directed cycles (but it would have many 2-bridges). In contrast, we shall see shortly that the extremal function β for path systems of high bridge girth is highly nontrivial.

1.2 Reductions from Network Design to Bridge Girth

The main conceptual contribution of this paper is a series of reductions from problems in network design to the functions β or β^* . Table 2 lists our results of this type, separated into three main technical threads. For every row of this table, one can recover or improve the current state-of-the-art bounds for the object in question by plugging in bounds for β or β^* (see Table 1 for details). Thus, (1) β, β^* have arguably been under the surface in prior work on all of these problems, and (2) further improved bounds for β, β^* could have widespread, black-box consequences for the area.

While new ideas are often needed to prove the bounds in Table 2, this is overall the less technical part of our paper; our main technical contributions lie in improved bounds for β, β^* , discussed next.

1.2.1 Technical Overview: Preservers

We begin with reachability preservers.

Definition 11 (Reachability Preservers [2]). *Let $G = (V, E)$ be a directed graph and let $P \subseteq V \times V$ be a set of demand pairs. A reachability preserver is a subgraph $H \subseteq G$ in which, for all $(s, t) \in P$ such that there exists an $s \rightsquigarrow t$ path in G , there also exists an $s \rightsquigarrow t$ path in H .*

We define $RP(n, p)$ as the smallest integer such that every n -node graph and set of $|P| = p$ demand pairs has a reachability preserver on $\leq RP(n, p)$ edges.

Extremal bounds for reachability preservers have been studied recently [2, 13, 29, 30], but they had long been studied algorithmically in the context of the Directed Steiner Forest problem, which asks to compute a reachability preserver of minimum total weight of a given input instance G, P . This problem is NP-hard, but the state-of-the-art approximation algorithms use *extremal* bounds for reachability preservers as an ingredient [2, 36, 61]. We prove:

Theorem 2. $RP(n, p) = \Theta(\beta(n, p, \infty))$.

Object	Bound	Implied By	First Proved
Reachability Preservers	$\text{RP} \geq \Omega(n^{\frac{2}{d+1}} p^{\frac{d-1}{d}}), d \text{ any pos int}$	Thm 37, Cor 41	[2, 41]
	$\text{RP} \leq O(n^{3/4} p^{1/2} + n^{5/8} p^{11/16} + n)$	Thms 37, 14	New
Online Reachability Preservers	$\text{RP}^* \geq \Omega(n^{2/3} p^{2/3} + n)$	Thms 44, 29	New
	$\text{RP}^* \leq O(\min\{n^{1/2} p, np^{1/2}\} + n)$	Thms 44, 28	New
	$\text{RP}^* = \tilde{\Theta}(n^{4/3})$ conditionally	Hyp 30, Thms 44, 31	New
Path Oracles	$\text{PO} \geq \Omega(n^{\frac{2}{d+1}} p^{\frac{d-1}{d}}), d \text{ any pos int}$	Thm 43, Cor 41	New
	$\text{PO} \leq \tilde{O}(n^{3/4} p^{1/2} + n^{5/8} p^{11/16} + n)$	Thms 43, 14	New
Shortcut Sets	$\text{SS} \geq \Omega(n^{\frac{2}{d+1}} p^{\frac{d-1}{d}}), d \text{ any pos int}$	Thm 48, Cor 41	[65] ($p = n$)
			New (other p)
Distance Preservers	$\text{DP} \geq \Omega(n^{2/3} p^{2/3} + n)$	Thms 32, 29	[41]
	$\text{DP} \leq O(\min\{n^{2/3} p, np^{1/2}\} + n)$	Thms 32, 78, 79	[22]
Shortest Path Oracles	$\text{SPO} \geq \Omega(n^{2/3} p^{2/3} + n)$	Thms 36, 29	New
	$\text{SPO} \leq \tilde{O}(\min\{n^{2/3} p, np^{1/2}\} + n)$	Thms 36, 78, 79	[22, 41]
Exact Hopsets	$\text{EH} \geq \Omega(n^{2/3} p^{-1/3})$	Thms 49, 29	[72] ($p = n$)
			New (other p)
Directed Flow-Cut Gap	$\text{MCG} \geq \tilde{\Omega}(n^{1/7})$	Thm 52, Cor 41	[37]
Sparsest Cut Gap	$\text{SCG} \geq \tilde{\Omega}(n^{1/7})$	Thm 64, Cor 41	[37]
Directed Steiner Forest Gap	$\text{DSFG} \geq \Omega(n^{1/2-o(1)})$	Thm 65, Cor 41	New

Table 1: Quantitative bounds for the problems in network design considered in this paper, implied by combining our bridge girth reductions with our bounds on β, β^* . For brevity the parameters (n, p) on the functions are omitted. See also Figure 5.

That is, the extremal bounds for reachability preservers are entirely captured by the value of β . This reduction is perhaps our most consequential one: with our improved upper bounds on $\beta(n, p, 4)$ discussed later, this implies a polynomial improvement in the extremal bounds for reachability preservers.

Corollary 3. *Every n -node graph and set of p demand pairs has a reachability preserver on $O(n^{3/4} p^{1/2} + n^{5/8} p^{11/16} + n)$ edges.*

The previous upper bound was $O(n^{2/3} p^{2/3} + n)$ [2]. The lower bound $\text{RP} \geq \Omega(\beta(n, p, \infty))$ is straightforward and perhaps implicit in [2], but the upper bound $\text{RP}(n, p) \leq O(\beta(n, p, \infty))$ takes more work. A natural proof attempt might be to take a hard input instance G, P for reachability preservers requiring $\text{RP}(n, p)$ edges, carefully choose a path for each demand pair, interpret these choices as a path system, and hope that the resulting path system has bridge girth ∞ and therefore size $\leq \beta(n, p, \infty)$. Unfortunately, this attempt fails: for some inputs G, P , it is not possible to choose paths that yield a path system of bridge girth ∞ . Specifically, this may not work on input instances G, P that have several possible paths for each demand pair, or where these paths are not edge-disjoint (the overlapping parts of paths count as 2-bridges).

Our solution is perhaps conceptually unusual: we do not attempt to handle these troublesome input instances G, P at all. Instead, we prove an *independence lemma*, showing that there exist highly structured hard input instances realizing $\text{RP}(n, p)$. This structure allows us to map these particular

Object	Reduction	Theorem
Reachability Preservers	$\text{RP}(n, p) = \Theta(\beta(n, p, \infty))$	Thm 37
Online Reachability Preservers	$\text{RP}^*(n, p) = \Theta(\beta^*(n, p, \infty))$	Thm 44
Path Oracles	$\text{PO}(n, p) = \tilde{\Theta}(\beta(n, p, \infty))$	Thm 43
Shortcut Sets	$\text{SS}(n, p) = \Omega\left(\frac{\beta(n, p, \infty)}{p}\right)$	Thm 48
Distance Preservers	$\Omega(\beta^*(n, p, \infty)) \leq \text{DP}(n, p) \leq \beta^*(n, p, 2)$	Thm 32
Shortest Path Oracles	$\Omega(\beta^*(n, p, \infty)) \leq \text{SPO}(n, p) \leq \tilde{O}(\beta^*(n, p, 2))$	Thm 36
Exact Hopsets	$\text{EH}(n, p) = \Omega\left(\frac{\beta^*(n, p, \infty)}{p}\right)$	Thm 49
Directed Flow-Cut Gap	$\text{MCG}(\beta(n, n, \infty)) = \tilde{\Omega}\left(\frac{\beta(n, n, \infty)}{n}\right)$	Thm 52
Sparsest Cut Gap	$\text{SCG}(\beta(n, n, \infty)) = \tilde{\Omega}\left(\frac{\beta(n, n, \infty)}{n}\right)$	Thm 64
Directed Steiner Forest Gap	$\text{DSFG}(n, p) = \Omega\left(\frac{\beta(n, p, \infty)}{n^{3/2}}\right)$	Thm 65

Table 2: Our results on the relationships between β, β^* , and objects in the literature on succinct network design.

structured worst-case instances to systems of bridge girth ∞ , which is enough for an extremal reduction between $\text{RP}(n, p)$ and $\beta(n, p, \infty)$. We prove an analogous independence lemma for online reachability preservers³. These independence lemmas are also the missing ingredient towards an incompressibility theorem for reachability preservers: we show that no *data structure* (not necessarily a subgraph) can encode paths among demand pairs with better space efficiency than a reachability preservers, which yields our reduction for path oracles.

A *distance preserver* is a subgraph that preserves distance among demand pairs, not just reachability. Distance preservers were introduced by Coppersmith and Elkin [41], and extremal bounds for distance preservers were studied in [22, 26, 27, 31, 36, 41].

Definition 12 (Distance Preservers [41]). *Let $G = (V, E, w)$ be a directed weighted graph and let $P \subseteq V \times V$ be a set of demand pairs. A distance preserver is a subgraph $H \subseteq G$ in which, for all $(s, t) \in P$, we have $\text{dist}_H(s, t) = \text{dist}_G(s, t)$.*

We define $\text{DP}(n, p)$ as the least integer such that every n -node graph and set of $|P| = p$ demand pairs has a distance preserver on $\leq \text{DP}(n, p)$ edges.

We prove:

Theorem 4. $\Omega(\beta^*(n, p, \infty)) \leq \text{DP}(n, p) \leq \beta^*(n, p, 2)$.

This time, our main conceptual contribution is the lower bound $\text{DP}(n, p) \geq \Omega(\beta^*(n, p, \infty))$, based on realizing ordered path systems as unique shortest paths in a graph. The upper bound $\text{DP}(n, p) \leq \beta^*(n, p, 2)$ is arguably implicit in [41], and is based on a well-known connection to *consistent path systems*: that is, it follows from the simple observation that no two unique shortest paths in a graph may intersect, split apart, and then intersect again later.⁴ However, we note that the following section contains a tight lower bound on $\beta^*(n, p, 2) = \beta(n, p, 2)$. This implies a major technical limitation to further progress on distance preservers: the power of consistency has been pushed to its limit, and so if we are to improve the state-of-the-art upper bounds for distance preservers (see Table 1), we *must* rely on more intricate structural properties of shortest paths.

We also prove an analogous independence lemma for distance preservers, which implies an analogous incompressibility theorem: no data structure can record shortest paths among demand pairs with

³Our model of online reachability preservers is a slight variant of the one introduced recently by Grigorescu, Lin, and Quanrud [61].

⁴Since $\beta^*(n, p, 2) = \beta(n, p, 2)$, and so we could have just as well written $\beta(n, p, 2)$ for the upper bound on DP. We chose $\beta^*(n, p, 2)$ because it suggests an open question: since DP is sandwiched between $\beta^*(n, p, \infty)$ and $\beta^*(n, p, 2)$, can it be placed more precisely in the β^* hierarchy?

significantly better space efficiency than a distance preserver. This yields our reduction for shortest path oracles.

1.2.2 Technical Overview: Flow-Cut and Integrality Gaps

Finally, we consider integrality gaps for three problems in network design: Directed Multicut (DMC), Directed Sparsest Cut (DSC), and Directed Steiner Forest (DSF). The standard integrality gaps for DMC and DSC are often interpreted as *flow-cut gaps*, as we explain next.

In the DMC problem, we are given a graph $G = (V, E)$ and a set of demand pairs $P \subseteq V \times V$, and the objective is to find a minimum-size subset of E whose removal separates all pairs of nodes in P . The Maximum Multicommodity Flow (MMF) problem asks for the maximum total flow that can be simultaneously pushed between the demand pairs (under unit edge capacities). MMF is the LP dual of the fractional relaxation $\widehat{\text{DMC}}$ of DMC. Thus, for any input G, P , we have

$$\text{MMF}(G, P) = \widehat{\text{DMC}}(G, P) \leq \text{DMC}(G, P).$$

The famous min-cut max-flow theorem states that we have equality when $|P| = 1$, but we do *not* have equality in general. It is interesting to study the maximum possible ratio between these terms, as an approximate version of the min-cut/max-flow theorem, and which often has applications in approximation algorithms. For undirected graphs, the seminal work of Leighton and Rao showed that the maximum possible ratio is $\Theta(\log p)$ [59, 78]. We will be interested in the corresponding quantity for directed graphs, called the *directed flow-cut gap*:

Definition 13 (MCG). *The function $\text{MCG}(n)$ is the least integer k such that, for every n -node directed graph G and set of demand pairs P (of any size), we have $\text{DMC}(G, P) \leq k \cdot \text{MMF}(G, P)$.*

This function MCG has been studied in [4, 35, 62, 75, 78, 91]. On the lower bounds side, an important paper by Chuzhoy and Khanna was the first to show that the bound is polynomial, with a lower bound of $\text{MCG}(n) = \tilde{\Omega}(n^{1/7})$ [37]. On the upper bounds side, the current bound is $\text{MCG}(n) = \tilde{O}(n^{11/23})$ [4]. We show the following reduction, which recovers the lower bound from [37] (see Table 1):

Theorem 5. $\text{MCG}(\beta(n, n, \infty)) = \tilde{\Omega}\left(\frac{\beta(n, n, \infty)}{n}\right)$.

At a technical level, this theorem closely follows the construction of Chuzhoy and Khanna [37]. Their construction uses a particular path system construction as an internal ingredient, which may be interpreted as a $\beta(n, n, \infty)$ lower bound system that also has many additional convenient properties (e.g., it is layered and highly symmetric). Our contribution is a generalization of their analysis, to show that these convenient properties are not really necessary, and one can plug in any system achieving $\beta(n, n, \infty)$ as a black box. Similarly, we obtain a reduction for the *sparsest cut problem*, again based on [37].

Finally, in the Directed Steiner Forest problem (DSF), we are given a weighted directed graph $G = (V, E, w)$ and a set of demand pairs $P \subseteq V \times V$, and the goal is to find a minimum weight subgraph $H \subseteq G$ that contains a directed $s \rightsquigarrow t$ path for all $(s, t) \in P$. DSF is NP-hard, but it has a natural integer programming formulation (see Section 5.6), and studying the integrality gap of its LP relaxation is a natural step towards designing approximation algorithms.

This integrality gap and related approximation algorithms have been studied in [18, 55, 79]. Letting $\text{DSFG}(n, p)$ be the integrality gap for instances with n nodes and p demand pairs, the previous best bound was that there exists p for which $\text{DSFG}(n, p) = \Omega(n^{0.0418})$ [79], implied by work on Directed Steiner Tree. We show:

Theorem 6. $\text{DSFG}(n, p) = \Omega\left(\frac{\beta(n, p, \infty)}{n^{3/2}}\right)$ for $p \leq n^{2-o(1)}$.

In particular, plugging in $p = n^{2-o(1)}$, we improve the integrality gap to $\Omega(n^{1/2-o(1)})$. This theorem partially addresses an open question in [9], where Alon, Moitra, and Sudakov asked whether bounds on the Ruzsa-Szemerédi function, which is equivalent to bounds on $\beta(n, p, \infty)$ in the setting of large p (see Theorem 81), could be useful towards proving integrality gaps for Directed Steiner Tree.

k	Bound	Justification
2	$\beta = \beta^* = \Theta(\min\{n^{2/3}p, p^{1/2}n\} + n + p)$	Upper Implicit (see App D.1) Lower New (see Thm 11)
3	$\beta = O\left(\min\left\{n^{2/3}p^{2/3}, \frac{n^2}{2^C \log^* n}\right\} + n + p\right)$ $\beta = \Theta(n^{2/3}p^{2/3})$ when $p \in \{n^{4/5}, n^{7/8}, n, n^{8/7}, n^{5/4}\}$	Implicit (see App D.2)
4	$\beta = O(n^{3/4}p^{1/2} + n^{5/8}p^{11/16} + n + p)$	New (see Thm 14)
	$\beta = O\left(\frac{p^2}{2^C \log^* p} + n\right)$	New-ish (see Thm 27)
	$\beta = \Omega\left(n^{\frac{2}{d+1}}p^{\frac{d-1}{d}}\right)$, d any positive integer	Implicit (see Cor 41)
∞	$\beta^* = O(\min\{n^{1/2}p, p^{1/2}n\} + n + p)$	New-ish (see Thm 28)
	$\beta^* = \Omega(n^{2/3}p^{2/3})$	New-ish (see Thm 29)
	$\beta^* = \tilde{\Theta}(n^{4/3})$ when $p = n$, conditional on Hyp 30	New-ish (see Thm 31)

Table 3: Asymptotic Bounds for β and β^* . For brevity, we write β in place of $\beta(n, p, k)$, and similar for β^* . Results are marked as fundamentally new in this paper, implicit in prior work, or “new-ish” meaning that they reuse a major ingredient from prior work but also have a new idea.

1.3 New Extremal Bounds for Bridge Girth

The main technical contributions of this paper are some new upper and lower bounds for β and β^* , polynomially improving over bounds implicit in the previous literature. These improved bounds imply new results for various problems in network design; see Table 1 in the next section. Table 3 gives a quick reference to our new bounds on β, β^* , as well as bounds implicit in prior work, and Figure 4 plots the state-of-the-art bounds on β, β^* following our paper.

1.3.1 New Lower Bounds for $\beta(n, p, 2)$

Our first main new result is a tight lower bound for $\beta(n, p, 2)$:

Theorem 7. $\beta(n, p, 2) = \Theta(\min\{n^{2/3}p, p^{1/2}n\} + n + p)$.

The upper bound in this theorem is implicit in [22, 41], and the lower bound is new in this paper. To explain our contribution, we discuss the previous (implicit) lower bound. The simplest way to ensure that a path system avoids 2-bridges is to simply ensure that any two paths intersect on at most one node. The largest path systems with this property have long been known: they are *finite projective planes*⁵, which imply the lower bound

$$\beta(n, p, 2) = \Omega\left(\min\{n^{1/2}p, p^{1/2}n\} + n + p\right).$$

If one wants to design a path system that polynomially exceeds this bound, it is necessary for a typical pair of paths π_1, π_2 to intersect on polynomially many nodes. To avoid 2-bridges, it would then be required that π_1, π_2 contain the nodes in $\pi_1 \cap \pi_2$ in exactly opposite orders. It seems rather unlikely to obtain this opposite-order property for all pairs of paths simultaneously. Cementing this intuition, the work of Coppersmith and Elkin [41] implies that the finite projective plane lower bound on $\beta(n, p, 2)$ is indeed tight in the parameter regime $p \geq n$. However, our Theorem 7 shows on the contrary that the finite projective plane is *polynomially far* from optimal in the remaining parameter regime $p \ll n$, and a denser construction with exactly this reverse-order property can in fact be achieved. The construction is a (dualized) version of the finite projective plane, based on *quadratics* over finite fields instead of lines.

⁵Finite projective planes are set systems rather than path systems, but one can place an arbitrary ordering on the sets while retaining the property that any pair of resulting paths intersect on at most one node.




Figure 4: Asymptotic bounds for $\beta(n, p, k)$, drawn to a logarithmic scale.

As discussed in Section 1.2.1, the practical consequence of this lower bound is a technical limitation on the tool of *consistent path systems*. In network design, a common strategy to limit the number of edges in a graph is to show that it arises from a path system with the property that no two paths intersect, split apart, and then intersect again later. Theorem 7 settles the worst-case size of such a path system, and thus to obtain better upper bounds than the ones in Theorem 7, more careful technical arguments are needed.

1.3.2 New Upper Bounds for $\beta(n, p, 4)$

Our next main result is a new upper bound for $\beta(n, p, 4)$.

Theorem 8. $\beta(n, p, 4) = O(n^{3/4}p^{1/2} + n^{5/8}p^{11/16} + n + p)$.

It should be noted here that β is inverse-monotonic in k ; that is, $\beta(n, p, k_1) \leq \beta(n, p, k_2)$ if $k_1 \leq k_2$. Hence, Theorem 8 also implies a polynomially improved upper bound for $\beta(n, p, 5), \dots, \beta(n, p, \infty)$. As discussed in Section 1.2.1, one corollary is a new polynomially improved upper bound for reachability preservers. Another consequence of this result is that it provides a clear avenue for further progress towards understanding $\beta(n, p, \infty)$. A point of this paper, reflected more precisely in Table 2, is that new *lower* bounds for $\beta(n, p, \infty)$ would be very consequential in network design. This in turn motivates the study of *upper* bounds for $\beta(n, p, \infty)$, towards determining the extent to which these improved lower bounds might be possible. The previous-best upper bounds on $\beta(n, p, \infty)$ were inherited all the way from the implicit upper bounds on $\beta(n, p, 3)$. Our theorem is proof-of-concept that exploiting larger forbidden bridges is indeed a worthwhile avenue towards improved understanding of the value of $\beta(n, p, \infty)$.

At a technical level, the proof of Theorem 8 is considerably more involved than other upper bounds in the area, including those for $\beta(n, p, 2)$, $\beta(n, p, 3)$, and the extremal functions of high-girth graphs (see Appendix A). All previous bounds are based roughly on a *forward-search* strategy, in which one picks a node, counts the paths intersecting that node, counts the nodes contained in those paths, and so on; bridge-freeness is used to argue that the nodes/paths that are witnessed are all distinct. See

Theorems 70, 78, 80, for examples. Our challenge is that forward-search does not work so well for 4-bridges, in the sense that a lack of 4-bridges does not imply that distinct nodes/paths are discovered at the appropriate level of the forward search. This requires considerable technical work to overcome, and due to space constraints we defer further technical overiewing to Section 3.2.

1.4 Future Directions and Open Problems

Table 1 lists the quantitative bounds obtained by directly mixing the reductions from Table 2 with the bounds on β, β^* from Table 3. One category of open problem is to improve the quantitative upper or lower bounds for any of these objects, whether or not via reductions to β, β^* .

It would also be interesting just to *recover* state-of-the-art quantitative bounds for network design problems via bridge girth reductions, so that we gain black-box improvements if and when the bounds for β, β^* are improved. Some good candidates for this program might include:

- **(Shortcut/Hopset Upper Bounds)** We have proved that one can recover state-of-the-art lower bounds on shortcut sets (SS) by reduction to $\beta(n, p, \infty)$. On the upper bounds side, a recent breakthrough of Kogan and Parter [73] proved that⁶

$$\text{SS}(n, p) = \begin{cases} \tilde{O}(n^{2/3}p^{-1/3}) & \text{when } p \geq n \\ \tilde{O}(np^{-2/3}) & \text{when } p \leq n \end{cases}.$$

Subsequent work by Berenstein and Wein [19] obtained a similar bound for $(1 + \varepsilon)$ hopsets. Obtaining this bound with a bridge girth reduction would be interesting.

- **(Exact Hopset Upper Bounds)** We have proved that one can recover state-of-the-art lower bounds on exact hopsets (EH) by reduction to $\beta^*(n, p, \infty)$. On the upper bounds side, there is a simple folklore algorithm, sometimes attributed to Ullman and Yannakakis [100], that shows $\text{EH}(n, p) = \tilde{O}(np^{-1/2})$. We refer to [72, 73] for discussion of this algorithm. We find the possibility of recovering this upper bound with a bridge girth reduction intriguing.
- **(Flow-Cut Gap Upper Bounds)** We have proved that one can recover state-of-the-art lower bounds on the flow-cut gaps (MCG, SCG) by reduction to $\beta(n, p, \infty)$. On the upper bounds side, the state-of-the-art is $\text{MCG}(n), \text{SCG}(n) = \tilde{O}(n^{11/23})$ obtained by Agarwal, Alon, and Charikar [4]. Can we recover this bound with a bridge girth reduction?

A recent paper by Kogan and Parter [72] perhaps makes some progress on this program by proving reductions among several important objects in network design (although naturally it does not directly consider bridge girth).

Another natural kind of open problem left by this paper is to obtain quantitative improved upper/lower bounds for β, β^* . We would also consider *self-reductions* very interesting, studying how the values of β, β^* evolve as k increases. The following is a concrete open problem in this vein. Notice that the extremal functions of high girth graphs cease to benefit from bridge girth parameters above $\log n$ (for example, $\gamma(n, \log n) = \Theta(\gamma(n, \infty)) = \Theta(n)$). We think it is likely that a similar effect holds for β, β^* :

Conjecture 9. *For all n, p , we have $\beta(n, p, \log n) = \Theta(\beta(n, p, \infty))$, and $\beta^*(n, p, \log n) = \Theta(\beta^*(n, p, \infty))$.*⁷

In Theorem 51 we provide some additional evidence that Conjecture 9 is true, by showing that its first half is implied by a plausible equality between the sizes of approximate distance preservers and reachability preservers, analogous to results already known for undirected preservers [72] and directed hopsets [19].

Finally, we discuss applications of β, β^* with other values of k . This paper directly motivates bridge girth parameters $k = 2$ and $k = \infty$, which are the settings that arise most commonly in our reductions (see Table 2). We also consider the parameter $k = 3$ to be comparably important, because it generalizes the Ruzsa-Szemerédi problem, which in turn captures prior work in network design [22, 76] (see Appendix A for more details). What about finite $k \geq 4$? Currently, we primarily use these

⁶These bounds are stated under a different parametrization than [73]: we use p as the size of the hopset, and $\text{SS}(n, p)$ as its hopbound, and thus SS is decreasing in p .

⁷We are grateful to an anonymous reviewer for suggesting this open problem.

setting as a tool to understand $\beta(n, p, \infty)$. It is an interesting conceptual open problem to find *direct* applications of $\beta(n, p, k), \beta^*(n, p, k)$, with intermediate choices of k , to problems in network design.

A candidate area in which these applications could arise is in the theory of ordered graphs and matrices. We have already applied this theory a bit, in the connections between Hypothesis 30 and Theorem 31. More broadly, there is a line of work in extremal combinatorics on *ordered matrix patterns*, as pioneered by Pach and Tardos [83]. In this problem, we consider binary $n \times n$ matrices, and we receive a collection of one or more forbidden submatrices. The goal is to determine the maximum possible number of 1's that could appear in such a matrix. One can naturally interpret an $n \times n$ binary matrix as the incidence matrix of an ordered, acyclic path system with n nodes and n paths, and high bridge girth in such a system corresponds to a collection of forbidden patterns. These problems have applications in data structures [87]. We refer to survey [94] for more on work in this space.




Figure 5: A diagram representing the relationships between the different notions in this paper. Here, a directed arrow from A to B means that $A \leq B$, and an undirected double line means they are equal.

2 Preliminaries on Path Systems

Here, we quickly review some standard definitions, notations, and technical lemmas for path systems that will be useful.

- For a path system $S = (V, \Pi)$, its *incidence graph* is the bipartite graph where the nodes on one side of the bipartition correspond to V , the nodes on the other side of the bipartition correspond to Π , and there is an edge between $v \in V$ and $\pi \in \Pi$ iff $v \in \pi$.
- A path system $S = (V, \Pi)$ is said to be *acyclic* if it does not have any directed cycles; equivalently, there is a total order of V (called a “topological order”) such that the order of every path $\pi \in \Pi$ is simply the order of V restricted to the nodes in π .
- For a path π , we write $x <_\pi y$ to mean that x, y are both nodes in π , and x strictly precedes y in π . We use the notation $x \leq_\pi y$ similarly.
- For a path π , a *subpath* is a (not necessarily contiguous) subsequence $\pi' \subseteq \pi$.
- A path system $S' = (V', \Pi')$ is a *subsystem* of $S = (V, \Pi)$, written $S' \subseteq S$, if one can obtain S' from S by a sequence of zero or more of the following operations: delete a node from V , delete a path from Π , or delete a single instance of a node from a single path in Π . We say that S' is the *induced subsystem* on V' if it is the system obtained by deleting all nodes in $V \setminus V'$.
- For a path system $S = (V, \Pi)$, the *degree* of a node $v \in V$, written $\deg(v)$, is the number of paths in Π that contain v .
- The *length* of a path $\pi \in \Pi$, written $|\pi|$, is the number of nodes in π (note that this length is bigger by 1 than the length of π when viewed as a path in an unweighted graph).
- For a system S with n nodes and p paths, the *average degree* is the quantity

$$d = \sum_{v \in V} \deg(v) / n$$

and the *average length* is the quantity

$$\ell = \sum_{\pi \in \Pi} |\pi| / p.$$

The *size identity* is that

$$nd = \|S\| = p\ell.$$

The following “cleaning lemma” lets us assume some convenient regularity properties for the path systems realizing $\beta(n, p, k)$ and $\beta^*(n, p, k)$.

Lemma 10 (Cleaning Lemma). *For every triplet n, p, k there exists a path system S with $\leq n$ nodes, $\leq p$ paths, bridge girth $> k$, $\|S\| = \Omega(\beta(n, p, k))$, and the following two additional properties:*

- (Approximately Degree-Regular) All nodes have degree $\Theta(d)$, where d is the average degree in S , and
- (Approximately Length-Regular) All paths have length $\Theta(\ell)$, where ℓ is the average length in S .

An identical claim holds for ordered path systems and ordered bridge girth.

Slight variants of this lemma are standard in the area, so we defer the proof to Appendix B. In the rest of this paper, we will often use the cleaning lemma as a tool to make assumptions about path systems realizing β or β^* . In other words, as our proofs typically start along the lines of “Let S be a path system with n nodes, p paths, bridge girth $> k$, and $\|S\| = \Omega(\beta(n, p, k))$ ”, we may often then use the cleaning lemma without loss of generality to guarantee that, additionally, S is both approximately degree-regular and approximately length-regular.

3 Bounds on β, β^* Functions

In this section, we prove new bounds for β, β^* in the settings $k = 2$, $k = 4$, and $k = \infty$.

3.1 Lower Bounds for $k = 2$

We prove:

Theorem 11. $\beta(n, p, 2) = \beta^*(n, p, 2) = \Theta(\min\{np^{1/2}, n^{2/3}p\} + n + p)$

Since 2-bridges are not sensitive to ordering, we immediately have $\beta(n, p, 2) = \beta^*(n, p, 2)$. The upper bounds for Theorem 11 are implicit in [22, 41]; for completeness, we supply proofs in Appendix D.1. The lower bound is new, and will be the focus of the rest of this section. We recall that

$$\beta(n, p, 2) \geq \Omega(n + p)$$

is an immediate lower bound, by considering either 1 path through all n nodes (giving a lower bound of $\Omega(n)$), or by considering p paths of 1 node each (giving a lower bound of $\Omega(p)$). It thus remains to prove

$$\beta(n, p, 2) \geq \Omega(\min\{np^{1/2}, n^{2/3}p\}).$$

These two minimized bounds meet at $p = n^{2/3}$. We will begin by considering this special case: that is, our goal is to construct a 2-bridge-free path system $S = (V, \Pi)$ with $p = \Theta(n^{2/3})$ paths and $\|S\| = \Omega(n^{4/3})$. We will then generalize to the full bound at the end.

3.1.1 Construction of S

The nodes. Let q be an arbitrary prime and let F_q be the finite field on q elements. Let Q be the set of polynomials over F_q of degree ≤ 2 . The polynomials in Q will ultimately correspond to the *nodes* of the path system: $V = Q$, so $n = |V| = q^3$.

The paths. For each $(x, y) \in F_q^2$, let $Q_{(x,y)} \subseteq Q$ be the set of polynomials that intersect the point (x, y) ; that is,

$$Q_{(x,y)} := \{f \in Q \mid f(x) = y\}.$$

There are q^2 points (x, y) , and so there are q^2 such sets in total. Our plan for defining our paths is to put a circular ordering on the elements of each $Q_{(x,y)}$, and then split the ordering into three parts, giving three paths for each (x, y) . Hence there will be $3q^2$ paths in total.

To define an ordering on $Q_{(x,y)}$: for a polynomial $f(x) = ax^2 + bx + c \in Q$, we define its *derivative* as $f'(x) := 2ax + b$, which we interpret as an element of F_q . Circularly order the polynomials in $Q_{(x,y)}$ by derivative $f'(x)$. We note that some polynomials in $Q_{(x,y)}$ will have tied derivatives; these ties may be broken arbitrarily. We then equitably partition the circular ordering into three contiguous parts, and add all three parts as paths in Π .

Size analysis. We have $q^3 =: n$ nodes (polynomials). We have q^2 points $(x, y) \in F_q^2$; each point is associated to three paths with q^2 nodes between them. Thus we have

$$\|S\| = \sum_{v \in V} \deg(v) = q^4.$$

3.1.2 Proof of 2-bridge-freeness

We will need the following structural lemma:

Lemma 12 (See Figure 6 for intuition). *Fix some $x_1 \neq x_2, y_1, y_2$, and let*

$$Z := Q_{(x_1, y_1)} \cap Q_{(x_2, y_2)}.$$

Then we have:

- For $z \in Z$, the values $z'(x_1)$ are pairwise distinct,
- For $z \in Z$, the values $z'(x_2)$ are pairwise distinct, and
- The circular ordering of $Q_{(x_1, y_1)}$ restricted to Z is exactly the reverse of the circular ordering of $Q_{(x_2, y_2)}$ restricted to Z .

Proof. For any $z(x) = ax^2 + bx + c \in Z$, subtracting the equations $z(x_1) = y_1$ and $z(x_2) = y_2$, we get

$$a(x_2^2 - x_1^2) + b(x_2 - x_1) = y_2 - y_1$$

and so, solving for b , we have

$$b = \frac{y_2 - y_1 - a(x_2^2 - x_1^2)}{x_2 - x_1}.$$

Rearranging $z(x_1) = y_1$, we also have

$$c = y_1 - ax_1^2 - bx_1.$$

Thus the polynomials $z \in Z$ can be written in the following form, parameterized by $a \in F_q$ (only):

$$z(x) = ax^2 + \left(\frac{y_2 - y_1 - a(x_2^2 - x_1^2)}{x_2 - x_1} \right) x + \left(y_1 - ax_1^2 - \left(\frac{y_2 - y_1 - a(x_2^2 - x_1^2)}{x_2 - x_1} \right) x_1 \right)$$

We then have

$$z'(x) = 2ax + \frac{y_2 - y_1 - a(x_2^2 - x_1^2)}{x_2 - x_1} = 2ax + \frac{y_2 - y_1}{x_2 - x_1} - a(x_2 + x_1)$$

and so we compute

$$z'(x_1) = \frac{y_2 - y_1}{x_2 - x_1} + a(x_1 - x_2)$$

and

$$z'(x_2) = \frac{y_2 - y_1}{x_2 - x_1} + a(x_2 - x_1).$$

Both of these functions are affine in a , and hence they take different values for each possible choice of a , proving the first two points. The third point follows from the observation that, letting $z_a, z_{-a} \in Z$ be quadratics with parameters $a, -a \in F_q$ respectively, the previous two equations imply that

$$z'_a(x_1) = z'_{-a}(x_2).$$

In particular: let a_i be the choice of parameter a such that $z'_{a_i}(x_1) = i$ (note that, since $z'(x_1)$ is affine in a , such a choice a_i must exist). Then the circular ordering of quadratics in $Q_{(x_1, y_1)}$ is

$$(z_{a_0}, z_{a_1}, z_{a_2}, \dots, z_{a_{q-2}}, z_{a_{q-1}}, z_{a_q} = z_{a_0}).$$

Meanwhile, using that $z'_{a_i}(x_1) = z'_{-a_i}(x_2)$, the circular ordering of quadratics in $Q_{(x_2, y_2)}$ is

$$(z_{a_0}, z_{a_{-1}}, z_{a_{-2}}, \dots, z_{a_{-(q-2)}}, z_{a_{-(q-1)}}, z_{a_{-q}} = z_{a_0}).$$

Since the parameter a is taken mod q , these are reverse circular orderings, completing the proof. $\square$

Lemma 13. *S is 2-bridge-free.*

Proof. Let $\pi_1, \pi_2 \in \Pi$, and recall that π_1, π_2 are respectively constructed with respect to two points $(x_1, y_1), (x_2, y_2) \in F_q^2$. We consider two cases:

- If $x_1 = x_2$ then by construction π_1, π_2 are node-disjoint, either because $y_1 \neq y_2$ and so $Q_{(x_1=x_2, y_1)}$ and $Q_{(x_1=x_2, y_2)}$ are disjoint, or because $y_1 = y_2$ and π_1, π_2 represent different parts in the node-disjoint partition of $Q_{(x_1=x_2, y_1=y_2)}$. Hence π_1, π_2 do not form a 2-bridge.
- If $x_1 \neq x_2$, then the points common to π_1, π_2 correspond to a subset of the polynomials in

$$Z = Q_{(x_1, y_1)} \cap Q_{(x_2, y_2)}.$$

By Lemma 12, the polynomials in Z have distinct derivatives and opposite circular orderings in $Q_{(x_1, y_1)}, Q_{(x_2, y_2)}$. It follows that, when we partition the points of $Q_{(x_1, y_1)}, Q_{(x_2, y_2)}$ into thirds to form π_1, π_2 , they have opposite orderings of any points $z_1, z_2 \in \pi_1 \cap \pi_2$. (Here it is important that we take $\ll 1/2$ of the circular ordering to form each path, to avoid the possibility of two paths wrapping around either side of the circular ordering to intersect both at the beginning and the end.) Hence π_1, π_2 do not form a 2-bridge. $\square$




Figure 6: In $\mathbb{R}^2$, if one plots quadratics that pass through two fixed points, the orderings by derivative at these two points are exactly opposite. Lemma 12 proves that an analogous fact holds in F_q^2 .

3.1.3 Remaining Lower Bound

We have now completed the lower bound proof in the special case $p = \Theta(n^{2/3})$, and it remains to discuss the extension to general p . To obtain the remaining points on our lower bound curve, we can post-process our construction in one of two ways:

- Suppose we delete nodes from the construction arbitrarily, until only $n' \ll n$ nodes remain. In our original construction, we had n nodes of degree $\Theta(n^{1/3}) = \Theta(p^{1/2})$ each. Thus, after deletions, we have $\|S\| = \Theta(n' \cdot p^{1/2})$, which provides one part of our lower bound curve.
- Alternately, suppose we delete paths from the construction arbitrarily, until only $p' \ll p$ paths remain. In our original construction, we had $p = \Theta(n^{2/3})$ paths of length $\Theta(n^{2/3})$ each. Thus, after deletions, we have $\|S\| = \Theta(p' n^{2/3})$, which provides the other part of our lower bound curve.

3.2 Upper Bounds for $k = 4$

We will prove:

Theorem 14. $\beta(n, p, 4) = O(n^{3/4} p^{1/2} + n^{5/8} p^{11/16} + n + p)$.

3.2.1 Technical Lemma: A Bound on the Sum Square of Path Lengths

A major technical lemma for our proof will be an upper bound on the quantity

$$\|T\|_2^2 := \sum_{\pi \in \Pi} |\pi|^2,$$

i.e., the squared L^2 norm of path lengths in a path system (note that the size notion $\|T\|$ may be viewed as the L^1 norm of path lengths). We name this path system T rather than S here because our plan is *not* to apply this lemma to the entire path system S that we analyze in Theorem 14, but rather to a specific subsystem $T \subseteq S$ that we will construct later. We prove the following bound:

Lemma 15. *Let $T = (V, \Pi)$ be a path system with n nodes, p paths, bridge girth > 3 , maximum path length L , and average path length at least a sufficiently large constant. Then we have*

$$\|T\|_2^2 = O(nL + p^{1/3} n^{4/3}).$$

We will split our proof into a few claims. Let us say that a path $\pi \in \Pi$ is:

- *long* if $|\pi| > Cn^{1/2}$, where C is a sufficiently large absolute constant that we leave implicit,

- *medium* if $\frac{\|T\|_2}{2p^{1/2}} \leq |\pi| \leq Cn^{1/2}$, or
- *short* if $|\pi| \leq \frac{\|T\|_2}{2p^{1/2}}$.

Let us say that the *long paths dominate* if the sum square of long paths is at least as large as the sum square of medium paths and as the sum square of short paths, and the *medium/short paths dominate* if the analogous property holds for the medium/short paths. The two terms added together in Lemma 15 respectively arise from the cases where the long or medium paths dominate. The following lemma dispatches with the remaining case:

Lemma 16. *The short paths do not dominate.*

Proof. There are $\leq p$ short paths, and by definition each one has length $\leq \|T\|_2/(2p^{1/2})$. By unioning, their sum square is at most

$$p \cdot \frac{\|T\|_2^2}{4p} = \frac{\|T\|_2^2}{4}.$$

Thus the short paths contribute at most 1/4 of the total value of $\|T\|_2^2$, so they cannot dominate. $\square$

The following technical lemma will be useful towards bounding $\|T\|_2^2$ in both cases where the long or medium paths dominate:

Lemma 17. *Let j be a parameter that is at least a sufficiently large constant, and let $T_j \subseteq T$ be the subsystem of T that contains exactly the paths π of length $j \leq |\pi| \leq 2j$. Then:*

$$\|T_j\|_2^2 = \begin{cases} O(jn) & \text{if } j \geq n^{1/2} \\ O(j^{-1}n^2) & \text{if } j \leq n^{1/2}. \end{cases}$$

Proof. Let p_j be the number of paths in T_j . Since T_j has bridge girth > 3 , we may apply the bounds on $\beta(n, p, 3)$ implicit in prior work, which give:

$$\|T_j\| = O\left(n^{2/3}p_j^{2/3} + n + p_j\right)$$

(see Theorem 80 in the appendix for a formal proof). Since $p_j \leq n^2$, the term $+p_j$ never dominates, so we may simplify this bound to

$$\|T_j\| = O\left(n^{2/3}p_j^{2/3} + n\right).$$

We then have:

$$j \leq \frac{\|T_j\|}{p_j} = O\left(n^{2/3}p_j^{-1/3} + np_j^{-1}\right).$$

In the case where the first term in the right-hand sum dominates, we continue

$$\begin{aligned} p_j^{1/3} &\leq O\left(n^{2/3}j^{-1}\right) \\ p_j &\leq O\left(n^2j^{-3}\right). \end{aligned}$$

In the case where the second term in the right-hand sum dominates, we continue

$$p_j = O\left(nj^{-1}\right).$$

Combining these, we get

$$p_j = O\left(n^2j^{-3} + nj^{-1}\right).$$

Plugging back into our bound on $\|T_j\|$, we get

$$\begin{aligned} \|T_j\| &= O\left(n^{2/3}\left(n^2j^{-3} + nj^{-1}\right)^{2/3} + n\right) \\ &= O\left(n^{2/3}\left(n^{4/3}j^{-2} + n^{2/3}j^{-2/3}\right) + n\right) \\ &= O\left(n^2j^{-2} + n^{4/3}j^{-2/3} + n\right). \end{aligned}$$

In the case where $j \geq n^{1/2}$, the latter $+n$ term dominates the sum, and so this gives $\|T_j\| = O(n)$. Thus $\|T_j\|_2^2$ is the sum of $O(nj^{-1})$ paths, each of which contribute $\Theta(j^2)$ to the sum, so its total is $O(jn)$. On the other hand, in the case where $j \leq n^{1/2}$, this gives $\|T_j\| = O(n^2j^{-2})$. Thus $\|T_j\|_2^2$ is the sum of $O(n^2j^{-3})$ paths, each of which contributes $\Theta(j^2)$ to the sum, so its total is $O(j^{-1}n^2)$. $\square$

Our next lemma counts the contribution of the long paths:

Lemma 18 (Long Path Gap Bound). $\sum_{\pi \in \Pi \mid \pi \text{ long}} |\pi|^2 = O(Ln)$.

Proof. By Lemma 17, for any parameter $j \geq n^{1/2}$, we have

$$\sum_{\pi \in \Pi \mid \pi \text{ long and } j \leq |\pi| \leq 2j} |\pi|^2 = O(jn).$$

We may therefore control the sum square of long paths by partitioning into subsets of paths of length $j \leq |\pi| \leq 2j$, and summing the contribution of these subsets. This gives:

$$\sum_{\pi \in \Pi \mid \pi \text{ long}} |\pi|^2 = O\left((n^{1/2})n\right) + O\left((2n^{1/2})n\right) + O\left((4n^{1/2})n\right) + \dots$$

This is a geometric sum, which is thus dominated by its largest term. Recall that we have assumed that all paths in T have length $\leq L$, and so the last term has the form $O(Ln)$, proving the lemma. $\square$

Next, we count the contribution of the medium paths.

Lemma 19 (Medium Path Gap Bound). $\sum_{\pi \in \Pi \mid \pi \text{ medium}} |\pi|^2 = O\left(\frac{p}{\|T\|_2} n^2\right)$.

Proof. Let

$$\frac{\|T\|_2}{2p^{1/2}} \leq j \leq n^{1/2}$$

be a parameter. From Lemma 17 and the fact that medium paths have length $\leq Cn^{1/2}$, we have

$$\sum_{\pi \in \Pi \mid \pi \text{ medium and } j \leq |\pi| \leq 2j} |\pi|^2 = O(j^{-1}n^2).$$

As in the long path case, we can bound the sum square of medium path lengths by partitioning the medium paths into parts where all paths in a part have $j \leq |\pi| \leq 2j$. This gives

$$\sum_j O(j^{-1}n^2)$$

where j ranges from $\|T\|_2/(2p^{1/2})$ to $n^{1/2}$ by multiples of 2. This is again a geometric sum which is dominated by its largest term. The largest term occurs when j is smallest, i.e., $j = \|T\|_2/(2p^{1/2})$, and we get

$$O\left(\frac{p^{1/2}}{\|T\|_2} n^2\right),$$

completing the proof. $\square$

Now we put the parts together:

Proof of Lemma 15. We consider two cases:

- If the long paths dominate, then by Lemma 18 we have $\|T\|_2^2 = O(Ln)$.
- If the medium paths dominate, then by Lemma 19 we have

$$\|T\|_2^2 = O\left(\frac{p^{1/2}}{\|T\|_2} n^2\right).$$

Rearranging, we get

$$\begin{aligned} \|T\|_2^3 &= O(p^{1/2} n^2) \\ \|T\|_2^2 &= O(p^{1/3} n^{4/3}). \end{aligned}$$

Finally, we recall by Lemma 16 that the short paths do not dominate, and so this completes the proof. $\square$


3.2.2 Setup and a Stronger Cleaning Lemma

In addition to the usual cleaning lemma, for technical reasons it will be helpful to assume an additional property for the path system S that we analyze. Let us say that a *2-cycle* in a path system is a pair of nodes u, v , and a pair of paths π_1, π_2 with $u <_{\pi_1} v$ and $v <_{\pi_2} u$. We use the following lemma to remove 2-cycles from S :


Lemma 20. *For any n, p and any $k \geq 3$, there exists a path system S with n nodes, p paths, bridge girth $> k$, size $\|S\| = \Theta(\beta(n, p, k))$, and no 2-cycles.*

Proof. Start with a path system $S = (V, \Pi)$ with n nodes, p paths, $\|S\| = \beta(n, p, k)$, and bridge girth $> k$ (which may have 2-cycles). Construct a path system S' as follows.

Initially $S' = (V, \emptyset)$ is empty. For each $\pi \in \Pi$ in an arbitrary order, add a subpath $\pi' \in \pi$ to Π' generated as follows. For each node $v \in \pi$, omit v from π' if there exists a node u and a previously-added path $q \in \Pi'$ for which $v <_{\pi} u$ and $u <_q v$. Otherwise, include $v \in \pi'$. In the following picture, if π is the wavy path on top and q is the straight path at the bottom, the two hollow nodes would be omitted and the four solid nodes would be included in π' (unless another choice of path q causes them to be omitted).



It is immediate from the construction that we do not complete any 2-cycles in S' , and since $S' \subseteq S$ we still have that S' has bridge girth $> k$. So it only remains to prove that $\|S'\| = \Theta(\|S\|)$. Consider a fixed node $v \in V$. Each time we consider a path π with $v \in \pi$, we either keep $v \in \pi$ or we omit it. If we keep v , then π contributes $+1$ to the degree of v in S' . If we omit v , we do so because of a previously-added path q with $v \in q$. In this case, let us say that q is *marked* by this action. We claim that each path can only be marked once. To see this: suppose for contradiction, that there are two different paths π, π' , which both contain v and which both mark q . This implies that q, π, π' form a 3-bridge, as in the following picture (with v as the first node and π as the river):



Thus, each time we omit a node $v \in \pi$ from π in the construction of S' , we may amortize this against a previously-added path q that kept v in S' . It follows that $\deg_{S'}(v) \geq \deg_S(v)/2$. Since this holds for all nodes v , we have $\|S'\| \geq \|S\|/2$, completing the proof. $\square$

Using this lemma, let $S = (V, \Pi)$ be a path system with bridge girth > 4 , no 2-cycles, n nodes, p paths, and size $\|S\| = \Theta(\beta(n, p, 4))$. By the Cleaning Lemma (Lemma 10), we may further let ℓ, d be the average path length and node degree in S (respectively), and assume without loss of generality that all paths have length $\Theta(\ell)$ and that all nodes have degree $\Theta(d)$.⁸ We assume that ℓ, d are both at least sufficiently large constants (if not, then we immediately have $\|S\| = O(n + p)$). Under all these assumptions, our goal is now to prove that

$$\|S\| = O\left(n^{3/4}p^{1/2} + n^{5/8}p^{11/16}\right).$$

⁸Technically, to assume the properties of the cleaning lemma and 2-cycle-freeness simultaneously, we need to use the fact that the construction in the cleaning lemma cannot create 2-cycles. This is immediate from the proof.

3.2.3 The Random Subsystem S'

Our next step is construct a particular subsystem of S that will be useful in analysis. Consider the following process, parametrized by a positive integer $h \leq \ell$ that we choose later, that generates a random subsystem $S' \subseteq S$:

- Choose a path $\pi_b \in \Pi$ uniformly at random, called the *base path*.
- Let $Q \subseteq \Pi$ be the set of paths that intersect π_b at exactly one node.
- **(Vertices of S')** Flip a coin to choose either “forwards” or “backwards.” Let V' be the set of nodes $v \in V$ with the following property:
 - If we choose “forwards,” then the property is: there exists a node $u \in \pi_b$ and a path $q \in Q$ with $u \leq v$ in q and with $|q[u \rightsquigarrow v]| < h$. (That is, u weakly precedes v along q , and these nodes are at most $h - 1$ positions apart in q .)
 - If we choose “backwards,” the property is similar except that we require $v \leq u$ with $|q[v \rightsquigarrow u]| < h$. (That is, u weakly *follows* v along q , and these nodes are at most $h - 1$ positions apart in q .)
- $S' = (V', \Pi')$ is the induced subsystem of S on the vertex set V' . Recall: this means that Π' contains the subpath $\pi \cap V'$ for each $\pi \in \Pi$.




Figure 7: The random subsystem S' , generated with respect to “height” parameter h , and a “forwards” coin flip.

Let $n' := |V'|$ be the number of surviving nodes in S' , let $p' := |\Pi'|$ be the number of surviving paths in S' , and let $\ell' := \|S'\|/p'$ be the average path length in S' . Note that part of each path in Q survives in S' ; we call the surviving parts $Q' \subseteq \Pi'$. We have the following controls on their values:

Lemma 21 (Properties of S').

- $|Q'| = \Theta(\ell d)$
- All nodes in S' have degree $\Theta(d)$
- $n' \leq O(\ell d h)$
- The maximum path length in S' is $O(\ell)$

Proof. The fact that all nodes in S' have degree $\Theta(d)$ is inherited directly from S , since the degrees of surviving nodes do not change in an induced subsystem. Similarly, the fact that the maximum path length in S' is $O(\ell)$ is inherited from S .

Since $|\pi_b| = \Theta(\ell)$, and each node has degree $\Theta(d)$, and d is a sufficiently large constant, there exist $\Theta(\ell d)$ paths that intersect π_b . We notice that these paths are pairwise-distinct: if a path q hits π_b at two different nodes, then (since S has no 2-bridges) q, π_b must form a 2-cycle, but from Lemma 20 there are no 2-cycles in S . Thus we have $|Q'| = \Theta(\ell d)$. (This is one of two places where the assumption of no 2-cycles will be useful.)


By construction every node in V' is contained in a path in Q' , and every path in Q' has length $O(h)$, so we also have $n' \leq O(\ell d h)$. $\square$

Although the previous lemma acknowledges that the node degrees in S' do not change from S , we unfortunately have no such guarantee for the path lengths. That is, the average path length ℓ' in S' might be very different from the average path length ℓ in S , and moreover the path lengths in S' might fluctuate wildly instead of all being $\Theta(\ell')$. We unfortunately won't be able to enforce approximate length regularity in S' by re-applying the cleaning lemma, either. The problem is that our strategy in analysis will be to bound $\|S'\|_2^2$, and while the cleaning lemma gives approximate length regularity while preserving $\|S'\|$, it can significantly change the value of $\|S'\|_2^2$.

It will, however, be helpful in the following analysis to assume that ℓ' is at least a sufficiently large constant. We enable this assumption using the following lemma, which provides a good bound on the size of S in the case where ℓ' is only a constant:

Lemma 22. *If $\mathbb{E}[\ell'] = O(1)$, then $\|S\| = O(n^{2/3}p^{2/3}h^{-1/3})$.*

Proof. First, we claim that $\mathbb{E}[n'] = \Theta(\ell dh)$. This follows by noticing two facts. First, for any path $q \in Q$, the expected length of the corresponding path in Q' (over the forwards/backwards coin flip) is $\Theta(h)$. Second, the paths in Q' are pairwise node-disjoint, except possibly on nodes in π_b . This holds because, if we have two paths $q_1, q_2 \in Q'$ that intersect on a node $v \notin \pi_b$, then q_1, q_2, π_b form a 3-bridge:



Thus we have $|Q'| = \Theta(\ell d)$ (deterministically), and each path in Q' contributes $\Theta(h)$ nodes to V' in expectation, so we have $\mathbb{E}[n'] = \Theta(\ell dh)$. Additionally, from the previous lemma, we have $n' = O(\ell dh)$ (deterministically). Together, these imply that there is positive constant probability c that $n' = \Theta(\ell dh)$.

Next, assuming that $\mathbb{E}[\ell'] = O(1)$, we may apply Markov's inequality (with a sufficiently large hidden constant in the $O(1)$) to conclude that

$$\Pr[\ell' = O(1)] > 1 - c.$$

Hence, by an intersection bound, there is positive probability that we simultaneously have $n' = \Theta(h\ell d)$ and $\ell' = O(1)$. In this event, we have $\|S'\| = \Theta(h\ell d^2)$, and thus

$$p \geq p' = \Omega(h\ell d^2).$$

Rearranging, we have

$$\frac{p^2 n^2}{h} = \Omega(\ell p \cdot (nd)^2)$$

$$\frac{p^2 n^2}{h} = \Omega(\|S\|^3)$$

$$\|S\| = O(n^{2/3}p^{2/3}h^{-1/3}). \quad \square$$

In the rest of the proof, we will make two simplifying assumptions: (1) that ℓ' is at least a large enough constant (otherwise we apply the previous lemma), and (2) that the forwards/backwards coin flip comes up “forwards.” the only place we need to flip this coin, rather than deterministically choosing “forwards,” is to argue that $\mathbb{E}[n'] = \Theta(\ell dh)$ in the previous lemma. Every lemma in the rest of the proof can be immediately proved by a symmetric argument in the case where the coin comes up “backwards,” so we assume “forwards” for simplicity.


3.2.4 Analysis of $\|S'\|_2^2$ and $\|S\|_2^2$

We next make some structural observations on the intersection patterns exhibited by paths in S or S' . For an ordered pair of paths $(\pi_1, \pi_2) \in \Pi^2$, we define


$$R_S(\pi_1, \pi_2) := \{(x, y) \in \pi_1 \times \pi_2 \mid \text{there exists } \pi \in \Pi \text{ with } \pi \cap \pi_1 = \{x\}, \pi \cap \pi_2 = \{y\}\}.$$

and $R_{S'}$ is defined similarly, with paths taken from Π' rather than Π . We make a few observations in order to motivate this definition.


- Suppose that $\pi_1, \pi_2 \in \Pi'$ and that π_1 intersects π_b before π_2 (i.e. $(\pi_1 \cap \pi_b) <_{\pi_b} (\pi_2 \cap \pi_b)$). Then for a path π witnessing a pair $(x, y) \in R_{S'}(\pi_1, \pi_2)$, we must specifically have that $x <_{\pi} y$. This follows by noticing that, if instead $y <_{\pi} x$, then π_1, π_b, π_2, π together form a 4-bridge, with π_1 as the river.




- Our next observation is that it is *not* possible to have pairs $(x, y), (x', y') \in R_S(\pi_1, \pi_2)$ that strictly cross each other, with $x <_{\pi_1} x'$ and $y' <_{\pi_2} y$ as in the following picture, if these pairs $(x, y), (x', y')$ are witnessed by two different paths. The reason for this is that otherwise, they imply a 4-bridge, with the path intersecting (x, y) as the river.



- At first, one might worry that crossing node pairs $(x, y), (x', y') \in R(\pi_1, \pi_2)$ can arise if both node pairs are caused by a single path π , as in the following picture. However, this can arise only if π intersects π_1 at both x, x' , and π intersects π_2 at both y, y' . This would imply a 2-cycle, which we have removed from S via Lemma 20. So this does not occur. (This is our last use of removing 2-cycles in the argument.)



- The previous observations imply that the node pairs in $R_S(\pi_1, \pi_2)$ are arranged roughly as in the following picture, with their points of intersection with π_1, π_2 increasing along both paths.



Let

$$\mathcal{R}_S := \sum_{(\pi_1, \pi_2) \in \Pi^2} |R_S(\pi_1, \pi_2)|.$$

The size of $\|S\|_2^2$ can be related to $\mathcal{R}_S$ as follows:

Lemma 23. $\|S\|_2^2 = \Theta\left(\frac{\mathcal{R}_S}{d^2}\right)$

Proof. Recall we have used the cleaning lemma to assume that ℓ is at least a large constant, and so by approximate length-regularity, we may assume that all paths in Π have ≥ 2 nodes. A given path π contributes $|\pi|^2$ to the value of $\|S\|_2^2$. We may therefore only count node pairs satisfying $x <_\pi y$, as there are $\binom{|\pi|}{2} = \Theta(|\pi|^2)$ such node pairs.

For each such node pair (x, y) , there are $\Theta(d)$ paths in Π that intersect x and $\Theta(d)$ paths in Π that intersect y . Thus, this node pair (x, y) appears in $\Theta(d^2)$ different sets $R_S(\pi_1, \pi_2)$ with $(\pi_1, \pi_2) \in \Pi^2$. So the pair (x, y) contributes $\Theta(d)^2$ points to the value of $\mathcal{R}_S$. It follows that

$$\|S\|_2^2 \cdot d^2 = \Theta(\mathcal{R}_S),$$

and the lemma follows by rearranging. $\square$

Consider a fixed, ordered pair of paths $(\pi_1, \pi_2) \in \Pi^2$. When we create S' , let us say that an ordered node pair (x, y) is *charged* to the pair (π_1, π_2) if:

- $\pi_1, \pi_2 \in Q$, and hence subpaths $\pi'_1 \subseteq \pi_1, \pi'_2 \subseteq \pi_2$ are in Q' , and
- $(x, y) \in R_{S'}(\pi'_1, \pi'_2)$.

The next lemma gives a lower bound on the expected number of node pairs that get charged to (π_1, π_2) . Note that, in the case where $\pi_1 \notin Q$ or $\pi_2 \notin Q$, then 0 node pairs are charged to (π_1, π_2) .


Lemma 24. *For each $(\pi_1, \pi_2) \in \Pi^2$, the expected number of node pairs charged to (π_1, π_2) is*

$$\Omega\left(|R_S(\pi_1, \pi_2)|^2 \cdot \frac{h}{\ell p}\right).$$

Proof. For pairs $(x, y), (x', y') \in R_S(\pi_1, \pi_2)$, let us say that (x', y') is *close behind* (x, y) if we have

$$1 \leq |\pi_1[x \rightsquigarrow x']| < h \quad \text{and} \quad 1 \leq |\pi_2[y \rightsquigarrow y']| < h.$$

We note that this definition implies that (x, y) is considered to be close behind itself.



The point of this definition is that, if the path containing (x, y) is selected as the base path π_b , then we will charge some node pairs to (π_1, π_2) , and the number of such node pairs is exactly the number of pairs in $R_S(\pi_1, \pi_2)$ that are close behind (x, y) .

Let us say that (x, y) is *typical* if, for some sufficiently large constant C , there are at least

$$|R_S(\pi_1, \pi_2)| \cdot \frac{h}{C\ell}$$

pairs in $R_S(\pi_1, \pi_2)$ that are close behind (x, y) . If we happen to sample a base path π_b that contains a typical pair $(x, y) \in R_S(\pi_1, \pi_2)$, then the number of pairs charged to (π_1, π_2) is

$$\Omega\left(|R_S(\pi_1, \pi_2)| \cdot \frac{h}{\ell}\right).$$

Thus, to prove the lemma, it suffices to prove that we sample a base path π_b that contains a typical pair $(x, y) \in R_S(\pi_1, \pi_2)$ with probability

$$\Omega\left(|R_S(\pi_1, \pi_2)| \cdot \frac{1}{p}\right);$$

that is, a constant fraction of the pairs in $R_S(\pi_1, \pi_2)$ are typical. We show this using the following counting argument:

By our analysis of the structure of path intersections, the pairs $(x, y) \in R_S(\pi_1, \pi_2)$ may be totally ordered by positions of $x \cap \pi_1, y \cap \pi_2$. Consider these pairs (x, y) in *increasing* order, that is, the first pair (x, y) considered is the one that intersects π_1, π_2 closest to their start nodes. When each pair (x, y) is considered, if it is typical, then add 1 to the count of typical pairs. Otherwise, if (x, y) is not typical, then we throw away the next

$$|R_S(\pi_1, \pi_2)| \cdot \frac{h}{C\ell}$$

pairs in the ordering (including (x, y)); these pairs are thrown out without increasing the count whether or not they are typical, and then we continue. Each time we skip a pair (x, y) , notice that the next pair considered (x', y') is *not* close behind (x, y) , which means we have

$$|\pi_1[x \rightsquigarrow x']| \geq h \quad \text{or} \quad |\pi_2[y \rightsquigarrow y']| \geq h.$$

That is, we progress at least h nodes along π_1 , or at least h nodes along π_2 . Since $|\pi_1| = O(\ell)$ and $|\pi_2| = O(\ell)$, we can thus only perform this skip operation $O(\ell/h)$ times. By unioning, it follows that only

$$O\left(|R_S(\pi_1, \pi_2)| \cdot \frac{h}{C\ell}\right) \cdot O(\ell/h) = O\left(|R_S(\pi_1, \pi_2)| \cdot \frac{1}{C\ell}\right)$$

pairs get discarded. By choice of large enough C , this is only a constant fraction of the total pairs in $R_S(\pi_1, \pi_2)$, which means a constant fraction of the pairs in $R_S(\pi_1, \pi_2)$ are counted as typical. By the previous discussion, the lemma follows. $\square$

Lemma 25. $\mathbb{E}[\|S'\|_2^2] = \Omega\left(\frac{h}{\ell p^3} \cdot \mathcal{R}_S^2\right)$

Proof. We can lower bound the expected value of $\|S'\|_2^2$ by counting the expected number of node pairs (x, y) that get charged to some $(q_1, q_2) \in Q^2$. By Lemma 24, this is

$$\begin{aligned} \mathbb{E}[\{|(x, y) \in V^2 \mid (x, y) \text{ charged}|\}] &= \Omega\left(\sum_{(\pi_1, \pi_2) \in \Pi^2} |R_S(\pi_1, \pi_2)|^2 \cdot \frac{h}{\ell p}\right) \\ &= \frac{h}{\ell p^3} \cdot \Omega\left(\sum_{(\pi_1, \pi_2) \in \Pi^2} |R_S(\pi_1, \pi_2)|^2\right) \cdot p^2 \\ &= \Omega\left(\frac{h}{\ell p^3} \cdot \left(\sum_{(\pi_1, \pi_2) \in \Pi^2} |R_S(\pi_1, \pi_2)|\right)^2\right) \quad \text{Cauchy-Schwarz} \\ &= \Omega\left(\frac{h}{\ell p^3} \cdot \mathcal{R}_S^2\right). \end{aligned} \quad \square$$

3.3 Proof Wrapup

The remainder of the proof is essentially just algebra. Using the previous two lemmas, we have:

Lemma 26. $\|S\| = O(n^{3/4}p^{1/2} + n^{8/13}p^{9/13}h^{1/13}).$

Proof. We have

$$\begin{aligned}
\|S\|_2^2 &= \Theta(d^{-2}\mathcal{R}_S) && \text{Lemma 23} \\
&= d^{-2} \cdot O\left(\frac{\ell p^3}{h} \cdot \mathbb{E}[\|S'\|_2^2]\right)^{1/2} && \text{Lemma 25} \\
&= O\left(\frac{\ell p^3}{d^4 h} \cdot \mathbb{E}[\|S'\|_2^2]\right)^{1/2}.
\end{aligned}$$

Recall from Lemma 21 that S' has $O(\ell dh)$ nodes, $\leq p$ paths, maximum path length $O(\ell)$, and bridge girth > 3 . We may thus apply Lemma 15 to conclude that the following bound holds deterministically:

$$\|S'\|_2^2 = O\left((\ell dh) \cdot \ell + p^{1/3}(\ell dh)^{4/3}\right).$$

Using this as an upper bound for expectation, we may continue:

$$\begin{aligned}
\|S\|_2^2 &= O\left(\frac{\ell p^3}{d^4 h} \cdot \left(\ell^2 dh + p^{1/3}(\ell dh)^{4/3}\right)\right)^{1/2} \\
&= O\left(\ell^3 p^3 d^{-3} + \ell^{7/3} p^{10/3} d^{-8/3} h^{1/3}\right)^{1/2}.
\end{aligned}$$

Our next step will be to apply the Cauchy-Schwarz inequality, which gives $\|S\|_2^2 \cdot p \geq \|S\|^2$. Using this, we may continue:

$$\begin{aligned}
\frac{\|S\|^2}{p} &= O\left(\ell^3 p^3 d^{-3} + \ell^{7/3} p^{10/3} d^{-8/3} h^{1/3}\right)^{1/2} \\
\frac{\|S\|^4}{p^2} &= O\left(\ell^3 p^3 d^{-3} + \ell^{7/3} p^{10/3} d^{-8/3} h^{1/3}\right) \\
\|S\|^4 &= O\left(\ell^3 p^5 d^{-3} + \ell^{7/3} p^{16/3} d^{-8/3} h^{1/3}\right) \\
\|S\|^4 &= O\left(n^3 p^2 + \|S\|^{-1/3} n^{8/3} p^3 h^{1/3}\right) && \|S\| = nd = p\ell.
\end{aligned}$$

We next split into two cases, by which of these terms in the right-hand side dominate. If the first term dominates, then we get

$$\|S\| = O\left(n^{3/4} p^{1/2}\right).$$

If the second term dominates, then we get

$$\begin{aligned}
\|S\|^4 &= O\left(\|S\|^{-1/3} n^{8/3} p^3 h^{1/3}\right) \\
\|S\|^{13} &= O\left(n^8 p^9 h\right) \\
\|S\| &= O\left(n^{8/13} p^{9/13} h^{1/13}\right).
\end{aligned}$$

Summing the two cases gives our claimed bound. $\square$

The next step is to choose h to balance terms. We are balancing the term $n^{8/13} p^{9/13} h^{1/13}$ in the previous lemma, which applies in the case where the expected average path length in S' is at least a large constant, with the term $n^{2/3} p^{2/3} h^{-1/3}$ from Lemma 22 which applies in the case where the average path length in S' is bounded by a constant. The proper setting of h is computed as:

$$\begin{aligned}
n^{8/13} p^{9/13} h^{1/13} &= n^{2/3} p^{2/3} h^{-1/3} \\
h^{16/39} &= n^{2/39} p^{-1/39} \\
h &= n^{1/8} p^{-1/16}.
\end{aligned}$$

These terms then balance at

$$\begin{aligned}
&n^{2/3} p^{2/3} \left(n^{1/8} p^{-1/16}\right)^{-1/3} \\
&= n^{5/8} p^{11/16}.
\end{aligned}$$

Thus, the total bound on the size of $\|S\|$ is

$$\|S\| = O\left(n^{3/4}p^{1/2} + n^{5/8}p^{11/16}\right).$$

Finally, we acknowledge that this proof used the assumption that ℓ, d were both at least sufficiently large constants, so the final size of $\|S\|$ also absorbs a $O(n + p)$ term.

3.4 Bounds for $k = \infty$

Theorem 27. $\beta(n, p, \infty) = O\left(\frac{p^2}{2^{O(\log^* p)}} + p\right)$

Proof. In this proof, it will be helpful to write $\vec{\beta}(n, p, k)$ for the extremal function of *acyclic* path systems of high bridge girth. Let S be an acyclic path system with n nodes, p paths, bridge girth > 3 , and $\|S\| = \vec{\beta}(n, p, 3)$. Let G be its incidence graph, which has $|E(G)| = \|S\|$ edges. We claim that G has girth > 6 . To see this, we argue:

- Since G is bipartite it does not have 3- or 5-cycles.
- A 4-cycle in G implies that there are two paths $\pi_1, \pi_2 \in \Pi$ that intersect the same pair of nodes $u, v \in V$. Since S is acyclic, these paths would need to use u, v in the same order, and hence they would form a 2-bridge. But since S has bridge girth > 3 , this means G may not have a 4-cycle.
- Similarly, a 6-cycle in G implies that there are three paths π_1, π_2, π_3 and three nodes t, u, v with $t, u \in \pi_1, u, v \in \pi_2, t, v \in \pi_3$. Since S is acyclic, we may assume without loss of generality that these nodes are ordered t, u, v in a topological sort. Thus π_1, π_2, π_3 form a 3-bridge with π_3 as the river. Since S has bridge girth > 3 , this means G may not have a 6-cycle.

Hence G is a bipartite graph with n, p nodes per side, girth > 6 , and $\vec{\beta}(n, p, 3)$ edges. It follows that⁹

$$\vec{\beta}(n, p, 3) \leq \gamma\gamma(n, p, 6).$$

We also notice that $\gamma\gamma$ is symmetric in its first two parameters, and so

$$\gamma\gamma(n, p, 6) = \gamma\gamma(p, n, 6).$$

With these two facts in mind, we now have:

$$\begin{aligned} \beta(n, p, \infty) &= \Theta\left(\vec{\beta}(n, p, \infty)\right) && \text{Corollary 42} \\ &\leq O\left(\vec{\beta}(n, p, 3)\right) \\ &\leq O(\gamma\gamma(n, p, 3)) \\ &= O(\gamma\gamma(p, n, 3)) \\ &\leq O\left(\frac{p^2}{\text{rs}(p)} + n\right) && [43, 58, 82]. \quad \square \end{aligned}$$

Theorem 28. $\beta^*(n, p, \infty) = O\left(\min\{n^{1/2}p, np^{1/2}\} + n + p\right)$

Proof. This follows as a consequence of Corollary 47. Recall that it suffices to bound the maximum possible size of an *acyclic* ordered path system S with bridge girth ∞ . If we drop the path ordering in S , and treat it as a normal unordered path system, then it still has bridge girth > 2 . Since S is acyclic, this implies that any two paths intersect on at most one node. So if we further drop the order of nodes within each path, and instead treat S as a *set* system, it has girth > 2 . It thus satisfies

$$\|S\| = O\left(\min\{n^{1/2}p, np^{1/2}\} + n + p\right)$$

by well-known bounds on the size of high-girth set systems (or bipartite graphs) [102]. $\square$

⁹In fact this reduction can be reversed, showing that these functions are equal. But we will only need to use the inequality in one direction, so we omit the proof on the other side.

Theorem 29. $\beta^*(n, p, \infty) = \Omega(n^{2/3}p^{2/3} + n + p)$

Proof. This construction is a slight modification of a construction from [93], although the analysis is somewhat new. We will construct an ordered path system $S = (V, \Pi)$ on $n = |V|$ nodes and $p = |\Pi|$ paths realizing this lower bound. Let $1 \leq \ell \leq n^{1/2}$ be an integer parameter; all paths in our system will have length exactly ℓ . Let

$$V := [1, \ell] \times \left[1, \frac{n}{\ell}\right]$$

be a rectangular subset of the integer lattice $\mathbb{Z}^2$.¹⁰ Our paths in Π will correspond to lines in $\mathbb{Z}^2$. The starting points of our lines (paths) are captured by the set

$$X := \{1\} \times \left[1, \frac{n}{2\ell}\right].$$

The slopes of our lines are captured by the set

$$W := \left\{ (1, i) \mid i \in \left[1, \frac{n}{2\ell^2}\right] \right\}.$$

For each $x \in X$ and $\vec{w} \in W$, we add the path

$$(x, x + \vec{w}, x + 2\vec{w}, \dots, x + (\ell - 1)\vec{w})$$

to Π . To order our paths: for paths $\pi, \pi' \in \Pi$, we assign $\pi < \pi'$ if $\vec{w}_y < \vec{w}'_y$, where $\vec{w}_y$ denotes the y -component of vector $\vec{w} \in \mathbb{Z}^2$; if $\vec{w}_y = \vec{w}'_y$ then the tie may be broken arbitrarily. Intuitively, this orders our paths in Π by increasing value of the slope of the corresponding line in $\mathbb{Z}^2$. This completes the construction, and we now check its parameters. We have

$$p = |\Pi| = |X||W| = \Theta\left(\frac{n^2}{\ell^3}\right).$$

Additionally, all paths in Π have length exactly ℓ , and so

$$\|S\| = p\ell = p \cdot \Theta\left(\frac{n^{2/3}}{p^{1/3}}\right) = \Theta\left(n^{2/3}p^{2/3}\right).$$

Now it only remains to verify that S is an ordered bridge-free path system. Let $\pi \in \Pi$ be a path constructed via start point $x \in X$ and slope $\vec{w} \in W$, and suppose for the sake of contradiction that π is the river for an ordered bridge in S . Since S is layered, this implies there is a collection of vectors $\vec{v}_1, \vec{v}_2, \dots, \vec{v}_{\ell-1} \in W$ (not necessarily distinct), which correspond to the differences among adjacent nodes along the path formed by the arcs of the bridge, and such that

$$\sum_{i=1}^{\ell-1} \vec{v}_i = (\ell - 1)\vec{w}.$$

Moreover, the vectors $\{\vec{v}_i\}$ may not all be identical to $\vec{w}$. This implies that there exists at least one vector $\vec{v}_i$ with $(\vec{v}_i)_y > \vec{w}_y$, and there exists at least one vector $\vec{v}_j$ with $(\vec{v}_j)_y < \vec{w}_y$. However, since we order our paths by increasing slope, the arc corresponding to $\vec{v}_i$ would be placed later in the ordering than the path corresponding to $\vec{w}$. Thus this arc cannot participate as an arc of an ordered bridge with $\vec{w}$ corresponding to the river. This completes the contradiction, and we conclude that π has no ordered bridge. $\square$

Finally, we show that this lower bound is conditionally tight in the setting $p = n$. We first give some background on the relevant condition. An *ordered graph* is a simple graph with a total ordering on its vertices. A natural problem is to investigate how classic results from extremal graph theory extend to this setting. Let:

- $\text{ex}(n, H)$ be the Turán function of the graph H ; that is, the maximum possible number of edges in an n -node graph that does not contain H as a subgraph.

¹⁰We will assume for convenience that n/ℓ and similar terms are integral; otherwise, rounding to the nearest integer affects our argument only by lower-order terms which may be ignored.

- $\text{ex}_{<}(n, H)$ be the *ordered* Turán function of the *ordered* graph H , defined analogously.

These extremal functions satisfy the basic inequality $\text{ex}_{<}(n, H) \geq \text{ex}(n, \overline{H})$, where $\overline{H}$ denotes the (unordered) graph underlying the ordered graph H . Tardos [94] asked how high the ratio can be, that is, the value of

$$\max_H \frac{\text{ex}_{<}(n, H)}{\text{ex}(n, \overline{H})}.$$

The current lower bound is $\Omega(n^{1/3-\varepsilon})$, and the current upper bound is $O(n^{1-\varepsilon})$. A reasonable hypothesis could be that the lower bound is closer to the correct answer:

Hypothesis 30. *Let H be an ordered graph with greater than two vertices and at least one edge. Then $\frac{\text{ex}_{<}(n, H)}{\text{ex}(n, \overline{H})} = O(n^{1/3})$.*

Under this hypothesis, our new lower bound for $\beta^*(n, p, \infty)$ is nearly-tight when $p = n$. To be clear, we do not necessarily think there is evidence that Hypothesis 30 is true. However, we do think that it represents a natural limitation on current methods in the theory of ordered graphs, and a significant new idea will be needed to prove or refute it. Thus, our point is simply that it is likely beyond the reach of current techniques to significantly improve our lower bound on $\beta^*(n, n, \infty)$ (if it is improvable at all).

Theorem 31. *Under Hypothesis 30, $\beta^*(n, n, \infty) = \tilde{\Theta}(n^{4/3})$.*

Proof. Note that by Theorem 29, $\beta^*(n, n, \infty) = \Omega(n^{4/3})$. We will prove that under Hypothesis 30, $\beta^*(n, n, \infty) = O(n^{4/3} \log n)$.

Let $S = (V, \Pi)$ be an acyclic ordered path system with n nodes, n paths, no ordered bridges, and size $\|S\| = \Omega(\beta^*(n, n, \infty))$. (S must exist by the cleaning lemma and by Corollary 47.) Since S is acyclic, there is a total order σ_1 on V such that for all $s, t \in V$, $s <_{\sigma_1} t$ if $s <_{\pi} t$ for some $\pi \in \Pi$. Likewise, let σ_2 be the total order on Π in the ordered path system S .

Let G_S be the *incidence graph* corresponding to S . Recall that this means: $G_S = (L \cup R, E)$ is the bipartite graph such that $L := V$, $R := \Pi$, and for all $v \in V$ and $\pi \in \Pi$, $(v, \pi) \in E$ if vertex v is contained in path π in S . Note that $|E| = \|S\|$. Define the following total order on the vertices $L \cup R$ of G_S :

- If $u, v \in L$, then apply ordering σ_1 .
- If $u, v \in R$, then apply ordering σ_2 .
- If $u \in L$ and $v \in R$, then let $u < v$.

We claim that graph G_S (with vertices ordered as above) does not contain any simple, ordered $2k$ -cycles H_k of form

$$H_k := (v_1, \pi_1, v_2, \pi_2, \dots, v_k, \pi_k),$$

where

- $v_1 < \pi_1$,
- $v_i < v_j$ if $i < j \in [1, k]$,
- $\pi_i < \pi_k$ for $i \in [1, k-1]$,
- and $k \geq 2$.

Suppose for the sake of contradiction that such an ordered graph H_k is contained in G_S . Then since G_S is bipartite and $v_1 < \pi_1$ in H_k , we must have that $v_i \in L$ and $\pi_i \in R$ for $i \in [1, k]$. Additionally, by the choice of edges E in G_S , it follows that $v_i, v_{i+1} \in \pi_i$ for $i \in [1, k-1]$ and $v_1, v_k \in \pi_k$ in S . Moreover, by our choice of ordering σ_1 , since $v_i < v_{i+1}$ in H and $v_i, v_{i+1} \in \pi_i$ it follows that $v_i <_{\pi_i} v_{i+1}$ and $v_1 <_{\pi_k} v_k$ in S . Finally, since $\pi_i < \pi_k$ for all $i \in [1, k-1]$ in H_k , path π_k comes last in the total order σ_2 of paths Π in S , so nodes $v_1, \dots, v_k \in V$ and paths $\pi_1, \dots, \pi_k \in \Pi$ correspond to an ordered k -bridge in S . This contradicts our assumption that S has no ordered bridges, so we conclude that G_S does not contain H_k . Consequently, the size of S is at most

$$\|S\| = |E| \leq \text{ex}_{<}(2n, H_k).$$

By known bounds on the extremal function of cycles, $\text{ex}(n, \overline{H_k}) = O(kn^{1+1/k})$ [28]. Then under Hypothesis 30,

$$\|S\| \leq \text{ex}_{<}(2n, H) \leq \text{ex}(2n, \overline{H}) \cdot O(n^{1/3}) = O(kn^{4/3+1/k}).$$

Taking $k = \log n$ completes the proof. $\square$

4 Extremal Reductions to Bridge Girth

In this section, we discuss various objects in network design where the extremal state-of-the-art upper or lower bounds on size can be reduced to extremal functions of bridge girth.

4.1 Distance Preservers

Recall Definition 12 for the formal definition of distance preservers and their associated extremal function DP . Our goal is to prove:

Theorem 32. $\Omega(\beta^*(n, p, \infty)) \leq \text{DP}(n, p) \leq \beta(n, p, 2)$.

We begin by proving the upper bounds of Theorem 32. As a warmup, let us restrict attention to a very specific kind of distance preserver input $G = (V, E, w), P$ that enjoys a property that we will call *independence* among the demand pairs:

Definition 14 (Independence). *For a graph $G = (V, E, w)$ and set of demand pairs $P \subseteq V \times V$, we say that P is independent in G if for all $(s, t) \in P$ there is a unique shortest path $\pi(s, t)$ in G , and these paths are pairwise edge-disjoint.*

An independent input instance $G = (V, E, w), P$ can be naturally associated to a path system $S = (V, \Pi)$, where the paths in Π are precisely the node sequences corresponding to the unique shortest paths for demand pairs in P . We claim that this system S is 2-bridge-free. To see this, notice that S has a 2-bridge iff there are two distinct nodes $u, v \in V$ and two distinct paths $\pi_1, \pi_2 \in \Pi$ that both contain u and then v (in that order). On one hand, we cannot have π_1, π_2 coincide on their $u \rightsquigarrow v$ subpaths, as this would imply that the associated paths in G share edges, violating independence. On the other hand, we cannot have that π_1, π_2 use distinct $u \rightsquigarrow v$ subpaths, as this would violate the property that π_1, π_2 are each unique shortest paths in the underlying graph. Thus it is not possible for S to have a 2-bridge, and so we have

$$\|S\| \leq \beta(n, p, 2).$$

The natural distance preserver for G, P is obtained by overlaying the unique edge-disjoint shortest paths for the demand pairs in P , and it has exactly $\|S\| - p$ edges.¹¹ Thus any such independent instance G, P has a distance preserver on $\beta(n, p, 2) - p$ edges, which satisfies Theorem 32.

This part of the proof is not exactly surprising, and it is essentially a rephrasing of the well-known fact that unique shortest paths in graphs exhibit *consistency*. The more interesting part of the proof is to show that independence of input instances may be assumed without loss of generality. This is accomplished in the following lemma:

Lemma 33 (Independence Lemma for Distance Preservers). *For any positive integers n, p , there exists an n -node graph G and a set of $|P| \leq p$ **independent** demand pairs such that the minimal distance preserver of G, P has exactly $\text{DP}(n, p)$ edges.*

Proof. Let G, P be a (not necessarily independent) input instance on n nodes and p paths realizing $\text{DP}(n, p)$. We may assume without loss of generality that G itself has exactly $\text{DP}(n, p)$ edges (otherwise, replace G with a distance preserver of G, P and then perform the following analysis). We may also assume without loss of generality that every demand pair in P has a unique shortest path in G . This follows by the standard method of random reweighting: that is, for each edge e , randomly choose a real number in the range $[0, \varepsilon]$ and add this number to $w(e)$. Shortest path ties are broken with

¹¹The $-p$ term arises since $\|S\|$ counts the number of *nodes* in each path, while for the distance preserver we count the number of *edges* in each path.

probability 1, and if we choose $\varepsilon > 0$ small enough, the changes in edge weights will not cause a previously non-shortest path to become a shortest path.

The instance G, P might still not be independent, because the unique shortest paths for the pairs in P might overlap on edges. To remove overlap, we will further modify G, P by executing either of the following two steps until neither one is possible. In the following, for a demand pair $(s, t) \in P$ we will write $\pi(s, t)$ for its unique shortest path, and we will say that $\pi(s, t)$ *uniquely uses* an edge e if $e \in \pi(s, t)$ and there is no other demand pair $(s', t') \in P$ with $e \in \pi(s', t')$.

- If there exists a demand pair $(s, t) \in P$ that does not uniquely use any edges, we delete (s, t) from P . Note that it is still the case that every edge in G is used by at least one unique shortest path for a demand pair.
- Suppose that there exists a demand pair $(s, t) \in P$ and a sequence of three contiguous nodes $(x, y, z) \in \pi(s, t)$, such that $\pi(s, t)$ uniquely uses one of the two edges $\{(x, y), (y, z)\}$, but the other of these two edges is used by another unique shortest path $\pi(s', t')$ as well. For ease of notation we will assume that (x, y) is the edge uniquely used by $\pi(s, t)$; the other case is symmetric. We then add (x, z) to G as a new edge, and we set its weight to $w(x, z) := w(x, y) + w(y, z)$, and we delete the edge (x, y) . Notice that:
 - We add (x, z) to G and we remove (x, y) from G , so the number of edges in G stays the same. The demand pair (s, t) still has a unique shortest path, which now uses the edge (x, z) in place of the 2-path (x, y, z) .
 - No other unique shortest path besides $\pi(s, t)$ is affected by this change to the structure of the edges in G . This holds because (x, y, z) must be the unique $x \rightsquigarrow z$ shortest path, and $\pi(s, t)$ is the only path that uses (x, y) , and therefore $\pi(s, t)$ is the only path that contains the nodes (x, z) in that order.
 - The other edge (y, z) is still used by a unique shortest path, since by hypothesis we have $(y, z) \in \pi(s', t')$ for some other demand pair (s', t') .




Figure 8: The “edge skip” operation used in the second case of the proof of Lemma 33

In either case, the sum of lengths of the paths $\{\pi(s, t)\}_{(s, t) \in P}$ decreases by at least 1, and thus the process eventually terminates. Once it terminates, every unique shortest path $\pi(s, t)$ uniquely uses at least one edge, and it does not contain two consecutive edges where one is uniquely used and the other is not. Therefore $\pi(s, t)$ uniquely uses *all* of its edges, which implies independence. $\square$

This implies:

Lemma 34. $DP(n, p) \leq \beta(n, p, 2)$.

Proof. By the independence lemma (33), there exists an n -node graph and a set of $|P| \leq p$ independent demand pairs that have $DP(n, p)$ edges in their union. By an earlier discussion, any such instance can be associated to a 2-bridge-free path system of size $DP(n, p) + p$. It follows that $DP(n, p) \leq \beta(n, p, 2)$. $\square$

Now we turn to the lower bounds:

Lemma 35. $\Omega(\beta^*(n, p, \infty)) \leq \text{DP}(n, p)$

Proof. Let $S = (V, \Pi)$ be an ordered path system with n nodes, p paths, no ordered bridges, and size $\|S\| = \beta^*(n, p, \infty)$. We will use S to construct an independent instance G, P that requires $\|S\| - p$ edges for any distance preserver.

Starting with an empty graph $G = (V, \emptyset)$, consider the paths in Π in their order in S . When a path $\pi(s, t) \in \Pi$ is considered, add (s, t) as a demand pair to P , and then add all consecutive pairs of nodes on $\pi(s, t)$ as new *directed* edges in G . Since S is 2-bridge-free, the paths $\{\pi(s, t)\}$ are pairwise edge-disjoint, and moreover since S has no ordered bridges, each path $\pi(s, t)$ is the unique $s \rightsquigarrow t$ path in G at the time it is added. Hence it is the unique *shortest* $s \rightsquigarrow t$ path at the time it is added, regardless of the edge weights we assign to its edges. We may therefore choose sufficiently large weights for the new edges on $\pi(s, t)$, so that no previously-added demand pair will gain a new shortest path using any edges in $\pi(s, t)$.

Once all paths in Π have been considered, we have unique edge-disjoint shortest paths in G for all p demand pairs, and the union of these paths contains

$$\beta^*(n, p) - p = \Omega(\beta^*(n, p))$$

edges. Thus, we can interpret this graph G together with the set P holding the endpoints of the paths in Π as an input instance for distance preservers. All edges in paths in Π must remain in a distance preserver, which implies

$$\Omega(\beta^*(n, p)) \leq \text{DP}(n, p). \quad \square$$

4.2 Shortest Path Oracles

We next prove an *incompressibility theorem* for distance preservers. We consider shortest path oracles, which are the natural data structure version of distance preservers:

Definition 15 (Path Oracles). *Given a directed graph $G = (V, E, w)$ and a set of demand pairs P , a shortest path oracle is a data structure that, when queried with $(s, t) \in P$, can report a shortest $s \rightsquigarrow t$ path in G (or “no path” if none exists).*

We define $\text{SPO}(n, p)$ as the smallest integer such that every n -node graph and set of $|P| = p$ demand pairs has a path oracle on $\leq \text{SPO}(n, p)$ bits.

Note that a distance preserver of G, P on m edges implies a shortest path oracle of G, P on $O(m \log n)$ bits. Consequently, we have

$$\text{SPO}(n, p) \leq O(\text{DP}(n, p) \log n).$$

The following theorem states that we cannot expect much smaller shortest path oracles in general.

Theorem 36. $\Omega(\text{DP}(n, p)) \leq \text{SPO}(n, p) \leq O(\text{DP}(n, p) \log n)$.

Proof. Let $G = (V, E, w)$ be an n -node graph and let P be a set of $|P| = p$ independent demand pairs such that the minimal distance preserver of G, P has exactly $\text{DP}(n, p)$ edges. Note that such G and P exist by the independence lemma for distance preservers (Lemma 33). We associate with G, P a path system $S = (V, \Pi)$ where the paths $\pi(s, t)$ in Π are precisely the node sequences corresponding to the unique shortest paths for demand pairs (s, t) in P . Note that

$$\|S\| = \sum_{\pi \in \Pi} |\pi| = \text{DP}(n, p) + p.$$

Let $S' = (V, \Pi')$ be a path system obtained from S by replacing each $s \rightsquigarrow t$ path $\pi(s, t)$ in Π with an arbitrary $s \rightsquigarrow t$ subpath $\pi'(s, t)$. Let $\mathcal{S}$ denote the set of all path systems S' generated in this way. Note that the total number of pairs $(v, \pi) \in V \times \Pi$ such that v is an internal node in π is $\text{DP}(n, p) - p$, so $|\mathcal{S}| = 2^{\text{DP}(n, p) - p}$.

For each path system $S' \in \mathcal{S}$, we construct a graph $G_{S'} = (V, E', w')$ as follows. For each path $\pi'(s, t) \in \Pi'$, we add all consecutive pairs of nodes on $\pi'(s, t)$ to E' . For each edge $e = (u, v) \in E'$, we

assign the weight $\text{dist}_G(u, v)$ to the edge (u, v) in $G_{S'}$. Observe that by our choice of weights, for all $s, t \in V$, $\text{dist}_G(s, t) \leq \text{dist}_{G_{S'}}(s, t)$. Additionally, for every $(s, t) \in P$, the path in $G_{S'}$ corresponding to $\pi'(s, t) \in \Pi'$ has path length exactly $\text{dist}_G(s, t)$. Since G, P is independent, $\pi(s, t)$ is a unique shortest $s \rightsquigarrow t$ path in G , and so $\pi'(s, t)$ is a unique shortest $s \rightsquigarrow t$ path in $G_{S'}$.

We have shown that every demand pair $(s, t) \in P$ has a unique shortest path in $G_{S'}$ that corresponds exactly to path $\pi'(s, t) \in \Pi'$ of S' . Then any shortest path oracle for $G_{S'}, P$ will have to output the path $\pi'(s, t)$ when queried with (s, t) . Now consider the graph family $\mathcal{G} = \{G_{S'}\}_{S' \in \mathcal{S}}$. Any two distinct graphs $G_1, G_2 \in \mathcal{G}$ will require distinct shortest path oracle data structures, since the corresponding path systems $S'_1, S'_2 \in \mathcal{S}$ are distinct. Consequently, at least one of the graphs in $\mathcal{G}$ will need

$$\Omega(\log |\mathcal{G}|) = \Omega(\text{DP}(n, p))$$

bits to represent its shortest path oracle data structure. $\square$

4.3 Reachability Preservers

Here we will prove:

Theorem 37. $\text{RP}(n, p) = \Theta(\beta(n, p, \infty))$.

We note that $\text{RP}(n, p)$ is only well-defined in the range $p \leq O(n^2)$, so naturally we prove Theorem 37 only in this parameter range. We again start with the upper bound, and we will need another independence lemma. We will overload the word “independent” for the analogous definition for reachability preservers:

Definition 16 (Independence). *For a graph $G = (V, E)$ and set of demand pairs $P \subseteq V \times V$, we say that P is independent (in the context of reachability preservers) if for all $(s, t) \in P$ there is a unique path $\pi(s, t)$, and these paths are pairwise edge-disjoint.*

Lemma 38 (Independence Lemma for Reachability Preservers). *For any positive integers n, p , there exists an n -node graph G and a set of $|P| = p$ **independent** demand pairs such that any reachability preserver of G, P has exactly $\text{RP}(n, p)$ edges.*

Proof. The proof is somewhat analogous to Lemma 33, but it requires an additional technical ingredient. The reason for the change in proof is essentially that for distance preservers we can assume that demand pairs have unique shortest paths, but for reachability preservers we cannot immediately make the analogous assumption that each demand pair has a unique path.

Let G, P be a (not necessarily independent) instance realizing $\text{RP}(n, p)$. First, we will use a helpful reduction from [2], allowing us to assume that G is a DAG. If not, we may consider each strongly connected component, add an in- and out-BFS tree from an arbitrary node to preserve reachability among all node pairs in that component, and then contract the component into a single super-node. The resulting contracted graph is a DAG, and it suffices to build a reachability preserver on this graph.¹²

Next, let us introduce some terminology. We will say that a demand pair $(s, t) \in P$ *requires* an edge e if every $s \rightsquigarrow t$ path includes e . We will say that (s, t) *uniquely requires* e if it requires e , and there is no other demand pair that also requires e . We may assume without loss of generality that every edge in G is required by at least one demand pair (or else that edge may be removed from G). Thus G itself is the unique reachability preserver of G, P , so it has $\text{RP}(n, p)$ edges. We then further modify G, P by the following steps:

- For each demand pair (s, t) , considered in arbitrary order, choose any $s \rightsquigarrow t$ path $\pi(s, t)$. Then:
 - If (s, t) does not uniquely require any edge in $\pi(s, t)$, delete (s, t) from P .
 - Otherwise, let $(u, v) \in \pi(s, t)$ be the first edge uniquely required by (s, t) . Replace the demand pair (s, t) with (u, t) , and replace $\pi(s, t)$ with its $u \rightsquigarrow t$ suffix.

We note that every edge in G is still required by at least one demand pair, and in the end every remaining demand pair (s, t) uniquely requires the *first* edge in $\pi(s, t)$.

¹²The contraction step costs $O(n)$ edges, which may be safely ignored since we already have $\text{RP}(n, p) = \Omega(n)$, e.g. by considering a path on input for which a reachability preserver must keep $n - 1$ edges.

- Next, repeat the following until no longer possible. Find a demand pair $(s, t) \in P$ and a contiguous 3-node subpath $(x, y, z) \subseteq \pi(s, t)$, such that the demand pair (s, t) uniquely requires (x, y) but it does not uniquely require (y, z) . If there are several possible choices of $\{(s, t), (x, y, z)\}$, then we will specifically need to consider one in which the node y comes as early as possible in the topological ordering of nodes in G . (There may still be several possible choices using this same minimal node y , in which case we can choose among these arbitrarily.) We then delete (x, y) from G , and add (x, z) to G as a new edge. We modify the path $\pi(s, t)$ by replacing its subpath (x, y, z) with the single edge (x, z) . After this operation, the number of edges in G stays the same.

For correctness, we now need to argue that after this change, it is still the case that every edge in G is required by at least one demand pair. We have:

- The new edge (x, z) is uniquely required by (s, t) . This follows from the fact that (x, y) was required by (s, t) , and so when (x, y) is deleted (but before (x, z) is added), there is no $s \rightsquigarrow t$ path. When (x, z) is added there is an $s \rightsquigarrow t$ path again, which implies that every $s \rightsquigarrow t$ path uses the edge (x, z) .
- We also claim that the edge (y, z) is still required by at least one demand pair. We argue this as follows. Before our modification of G , we know that the demand pairs (s, t) and (s', t') each uniquely requires all edges in their prefixes $s \rightsquigarrow y$ and $s' \rightsquigarrow y$, respectively. This means that s' still cannot reach x , and thus (s', t') cannot use the new edge (x, z) and still requires the edge (y, z) .

Each time we repeat this step, the sum of lengths of the paths $\{\pi(s, t)\}_{(s, t) \in P}$ decreases by 1. Therefore, we halt after finitely many steps.

To summarize, once this process halts, every demand pair (s, t) has the property that $\pi(s, t)$ uniquely requires its first edge, and moreover for any two consecutive edges $(x, y), (y, z) \in \pi(s, t)$, if (s, t) uniquely requires the first edge (x, y) , then it also uniquely requires the second edge (y, z) . Together, these properties imply that each demand pair (s, t) uniquely requires every edge on its path $\pi(s, t)$, which implies independence. $\square$

Lemma 39. $RP(n, p) \leq O(\beta(n, p, \infty))$.

Proof. By our independence lemma, there is an n -node graph $G = (V, E)$ and set of $|P| = p$ independent demand pairs for which any reachability preserver has at least $RP(n, p)$ edges. We can naturally associate G, P to a path system $S = (V, \Pi)$ by associating each demand $(s, t) \in P$ to the unique $s \rightsquigarrow t$ path in G . We thus have $\|S\| = RP(n, p) + p$. Moreover, this path system cannot have bridges, since the paths for demand pairs in P are unique. Thus we have constructed a ∞ -bridge-free path system of size $\|S\| \geq RP(n, p)$, and the lemma follows. $\square$

We now turn to the lower bound:

Lemma 40. $RP(n, p) \geq \Omega(\beta(n, p, \infty))$.

Proof. Let $S = (V, \Pi)$ be a ∞ -bridge-free path system with n nodes, p paths, and size $\|S\| = \beta(n, p, \infty)$. Let $G = (V, E)$ be the directed graph over the same vertex set, where we put an edge $(u, v) \in E$ iff there is a path in Π where the nodes u, v appear consecutively (in that order). Define demand pairs P to be the set of node pairs (s, t) that are endpoints of the paths in Π .

Since S is ∞ -bridge-free, for every demand pair $(s, t) \in P$ there is a unique $s \rightsquigarrow t$ path in G , and these paths are pairwise edge-disjoint. Thus, it is necessary and sufficient for a reachability preserver to keep all edges contained in these paths. The number of edges contained in these paths is exactly

$$\|S\| - p = \beta(n, p, \infty) - p.$$

This proves that

$$\beta(n, p, \infty) - p \leq RP(n, p),$$

and we then notice that always $\beta(n, p, \infty) \geq \Omega(p)$, and so in fact we have

$$\beta(n, p, \infty) \leq O(RP(n, p)). \quad \square$$

The previous two lemmas imply Theorem 37. We now turn to its consequences.

Corollary 41. *For all positive integers d , we have $\beta(n, p, \infty) = \Omega\left(n^{\frac{2}{d+1}} p^{\frac{d-1}{d}}\right)$.*

Proof. Follows from Theorem 37 and from plugging in the state-of-the-art lower bounds on $\text{RP}(n, p)$ from [2] (which, in turn, are directly based on the distance preserver lower bounds from [41]). We remark that one does not *really* need Theorem 37 to prove this lower bound on $\beta(n, p, \infty)$, in the sense that it is straightforward to interpret the reachability preserver lower bound construction from [2] directly as a lower bound against $\beta(n, p, \infty)$. $\square$

The following corollary uses the equivalence between $\beta(n, p, \infty)$ and reachability preservers more directly. It shows that the extremal path systems realizing the lower bound for $\beta(n, p, \infty)$ have some extra structure: they must in fact be *acyclic*.

Corollary 42. *For all n, p , there exists an **acyclic** path system S with n nodes, p paths, bridge girth ∞ , and size $\|S\| = \Theta(\beta(n, p, \infty))$.*

Proof. Let $\vec{\beta}(n, p, \infty)$ be the maximum possible size of an n -node, p -path, *acyclic* path system of bridge girth ∞ . Let $\vec{\text{RP}}(n, p)$ be the maximum number of edges needed for a reachability preserver of an n -node DAG and p demand pairs. We first notice that, by exactly the same reduction as in Theorem 37, we have

$$\vec{\beta}(n, p, \infty) = \Theta\left(\vec{\text{RP}}(n, p)\right).$$

Next, it is proved in [2] that

$$\vec{\text{RP}}(n, p) = \Theta(\text{RP}(n, p)).$$

That is, they show a reduction from finding reachability preservers in general graphs to DAGs. To briefly summarize this reduction, suppose we are given a graph G and demand pairs P , and we wish to construct a reachability preserver. For each strongly-connected component C , choose an arbitrary node $c \in C$ and add two trees in C rooted at c ; one with edges pointing away from c , and one with edges pointing towards c . Thus, reachability is preserved between all pairs of nodes in C , and we can contract C into a single node before proceeding. This reduction costs at most $2n = O(\text{RP}(n, p))$ edges in total.

Finally, from Theorem 37 we have

$$\text{RP}(n, p) = \Theta(\beta(n, p, \infty)).$$

Putting the parts together, we have

$$\vec{\beta}(n, p, \infty) = \Theta(\beta(n, p, \infty)),$$

as required. $\square$

4.4 Path Oracles

We next prove an *incompressibility theorem* for reachability preservers, much like the one proved previously for distance preservers. We consider path oracles, which are the natural data structure version of reachability preservers:

Definition 17 (Path Oracles). *Given a directed graph G and a set of demand pairs P , a path oracle is a data structure that, when queried with $(s, t) \in P$, can report an $s \rightsquigarrow t$ path in G (or “no path” if none exists).*

We define $\text{PO}(n, p)$ as the smallest integer such that every n -node graph and set of $|P| = p$ demand pairs has a path oracle on $\leq \text{PO}(n, p)$ bits.

Note that a reachability preserver of G, P on m edges implies a path oracle of G, P on $O(m \log n)$ bits, by simply writing down a description of the reachability preserver. Consequently, we have

$$\text{PO}(n, p) = O(\text{RP}(n, p) \log n).$$

The following theorem states that we cannot expect much smaller path oracles in general.

Theorem 43. $\Omega(RP(n, p)) \leq PO(n, p) \leq O(RP(n, p) \log n)$.

Proof. The upper bound follows from the previous discussion. The lower bound follows from an argument identical to that of Theorem 36, except we use the independence lemma for reachability preservers (Lemma 38), and we replace DP with RP, and we replace “unique shortest path” with “unique path.” $\square$

4.5 Online Reachability Preservers

Here we discuss the online version of the reachability preserver problem. There are several ways to reasonably define such online versions; this one is a slight variation of the one introduced [61] in the context of online directed Steiner forest algorithms.¹³

Definition 18 (Online Reachability Preservers). *The online reachability preserver game is the following two-player game, between a builder and an adversary:*

- *The adversary starts with an n -node directed graph $A = (V, \emptyset)$, and the builder starts with an n -node directed graph $B = (V, \emptyset)$. Both graphs are initially empty. The builder is trying to minimize the final number of edges in B , and the adversary is trying to maximize the final number of edges in B .*
- *Repeat the following for p rounds:*
 - *(Adversary’s Turn) The adversary adds any number of edges to A , and then names a pair of nodes (s, t) such that an $s \rightsquigarrow t$ path in A exists.*
 - *(Builder’s Turn) The builder must respond by choosing a set of edges that are currently in A , and adding those edges to B . Afterwards, we require that an $s \rightsquigarrow t$ path must exist in B .*
- *The value of the game is the final number of edges in the graph B .*

We define $RP^*(n, p)$ as the min-max value of this game, where the adversary is maximizing and the builder is minimizing (the value of the game), relative to parameters n, p .

In the same way that (offline) reachability preservers are captured by $\beta(n, p, \infty)$, as in Theorem 37, we claim that online reachability preservers are captured by β^* :

Theorem 44. $\beta^*(n, p, \infty) = \Theta(RP^*(n, p))$

First we will prove an upper bound for $\beta^*(n, p, k)$:

Lemma 45. $RP^*(n, p) \geq \Omega(\beta^*(n, p, \infty))$.

Proof. The strategy of the adversary works as follows. First, they think of an ordered path system $S = (V, \Pi)$ with n nodes, p paths, no ordered bridges, and size $\|S\| = \beta^*(n, p, \infty)$. Let $\pi_i \in \Pi$ denote the i^{th} path in the ordering. In each round i of the game, the adversary considers π_i , and adds each consecutive pair of nodes along π_i as a new edge in A . Then, they name the endpoints (s, t) of π as the pair for this round. Since S has no ordered bridges, currently π_i is the unique simple $s \rightsquigarrow t$ path in A . Thus the builder has no choice but to add the $|\pi| - 1$ edges corresponding to π to B . In total, the builder thus adds $\beta^*(n, p, \infty) - p = \Omega(\beta^*(n, p, \infty))$ edges to B . $\square$

Next, we prove a matching lower bound:

Lemma 46. $RP^*(n, p) \leq \beta^*(n, p, \infty)$

Proof. In the online reachability preserver game, we will assume only that the builder adds a *minimal* set of edges in each round. That is, when the builder adds edge set E_i in round i , we assume that there is no proper subset $E'_i \subsetneq E_i$ that could have been added instead, and still satisfy the adversary’s demand. We claim that, so long as the builder’s choices satisfy this property, they will add $\leq \beta^*(n, p, \infty)$ edges to B in total.

¹³More specifically: we allow the adversary to add edges to the graph A throughout the game, whereas [61] essentially require the adversary to commit to a graph in preprocessing.

Indeed, given a sequence of choices made by a builder and an adversary, let us track an auxiliary ordered path system S as follows. The vertex set V of S is the same as the vertex set of the graphs A, B in the game. We next describe the paths of S . By minimality of the builder's choices, their selected edge set E_i may be interpreted by considering a simple path π_i between the adversary's demand pair, and setting $E_i := \pi_i \setminus E(B)$. In round i , we add a path q_i to the auxiliary path system S , where q_i is the sequence of vertices v for which there is an edge in E_i entering v , ordered by appearance in π_i . We have $|q_i| = |E_i|$, and therefore $\|S\| = |E(B)|$. So it only remains to show that S has no ordered bridges, and thus $|E(B)| \leq \beta^*(n, p, \infty)$.

Suppose for contradiction that S has an ordered bridge, with path q_i as its river, and nodes $s, t \in q_i$ participating in the bridge. Since $t \in q_i$, there exists an edge of the form $(u, t) \in E_i$. However, since nodes s, t participate in the bridge, there exists an $s \rightsquigarrow t$ path in B before round i , that is, the one corresponding to the arcs of said bridge. Therefore, the builder could have added $E_i \setminus \{(u, t)\}$ in round i . This contradicts minimality of the builder's choices, and thus S has no ordered bridges. $\square$

Analogous to Corollary 42, the following corollary implies that, without loss of generality, *acyclic* graphs are enough to achieve lower bounds for $\beta^*(n, p, \infty)$:

Corollary 47. *For all n, p , there exists an acyclic ordered path system S with n nodes, p paths, ordered bridge girth ∞ , and size $\|S\| = \Omega(\beta^*(n, p, \infty))$.*

Proof. In the parameter regime where $\beta^*(n, p, \infty) = O(n)$, one can take S to be any ordered path system with one path of length n and the remaining paths of length 1, and so the claim is trivial. In the following, we assume that $\beta^*(n, p, \infty) \geq cn$ for a sufficiently large constant c .

By Theorem 44, it suffices to prove that in the online reachability preserver game, the adversary has a (near-)optimal strategy in which the underlying graph G is always acyclic. To show this, imagine the following strategy that the builder could use. Any time the adversary adds an edge to G that completes a directed cycle C , the builder immediately adds all edges in C to their reachability preserver, and then for the rest of the game they treat C as a single contracted supernode. Since each contraction step costs $|C|$ edges and reduces the number of nodes in G by $|C| - 1$, the builder pays only $O(n)$ edges in total for these contraction steps, which is negligible.

We now shift perspective back to the adversary. Any time the adversary *would* add an edge (u, v) to the graph that completes a directed cycle C , they could instead omit (u, v) and contract C into a single supernode in their internal representation of the graph. By the above analysis, this is without loss of generality, and affects the overall min/max value of the game by at most $O(n)$, which is negligible. Thus, the adversary never completes a directed cycle in G , and the theorem follows. $\square$

4.6 Shortcut Sets and Exact Hopsets

Here we show that the lower bounds for reachability and distance preservers can be extended to shortcut sets and exact hopsets, respectively. The proofs are similar to each other in spirit.

Definition 19 (Shortcut Sets). *For a directed graph $G = (V, E)$, a Δ -diameter-reducing shortcut set is a set of additional directed edges H such that every edge $(u, v) \in H$ is in the transitive closure of G , and*

$$\max_{s, t \in V, \text{ exists } s \rightsquigarrow t \text{ path in } G} \text{dist}_{G \cup H}(s, t) \leq \Delta.$$

We write $SS(n, p)$ for the smallest integer Δ such that every n -node graph has an Δ -diameter-reducing shortcut set of $|H| = p$ edges.

Theorem 48. $SS(n, p) = \Omega\left(\frac{\beta(n, p, \infty)}{p}\right)$

Proof. Let $S = (V, \Pi)$ be an ∞ -bridge-free path system on $n = |V|$ nodes and $p + 1 = |\Pi|$ paths, such that¹⁴

$$\|S\| = \Theta(\beta(n, p + 1, \infty)) = \Theta(\beta(n, p, \infty)).$$

The average path length in S is thus

$$\ell = \Theta\left(\frac{\beta(n, p, \infty)}{p + 1}\right) = \Theta\left(\frac{\beta(n, p, \infty)}{p}\right).$$

¹⁴This latter equality is intuitive, but for completeness it is formally proved in Lemma 76 in the appendix.

By the Cleaning Lemma (Lemma 10), we may further assume without loss of generality that all paths in S have length $\Theta(\ell)$. Let us now associate to S a directed graph G as usual, by including a directed edge (u, v) for each pair of nodes u, v that appear consecutively on any path in S (so $|E(G)| = \Theta(\beta(n, p, \infty))$). We also define a set P of $|P| = p + 1$ demand pairs as the endpoints of the paths in S . Recall that shortcut sets require diameter reduction among *all* node pairs, rather than just a specific set of demand pairs, but nonetheless it will be helpful to focus our analysis on these demand pairs.

Since S has bridge girth ∞ , for each demand pair $(s, t) \in P$ there is a unique simple $s \rightsquigarrow t$ path in G ; let us denote this path by $\pi(s, t)$. Since these paths $\{\pi(s, t)\}$ do not have 2-bridges, we additionally have that for any ordered pair of nodes (x, y) , there is at most one demand pair (s, t) with $x <_{\pi(s, t)} y$. In particular, let H be an arbitrary shortcut set of size $|H| = p$. Since $|P| = p + 1$ but $|H| = p$, there exists a demand pair $(s, t) \in P$ such that there is no $(x, y) \in H$ with $x <_{\pi(s, t)} y$. It follows that $\pi(s, t)$ remains the unique simple $s \rightsquigarrow t$ path in the graph $G \cup H$. The number of hops in this path is

$$|\pi(s, t)| - 1 = \Theta(\ell) = \Omega\left(\frac{\beta(n, p, \infty)}{p}\right),$$

and so the hopset H cannot reduce diameter below this threshold. $\square$

We now give an analogous proof for exact hopsets.

Definition 20 (Exact Hopsets). *For a directed weighted graph $G = (V, E)$, a β -hop exact hopset is a set of additional directed weighted edges H such that every edge $(u, v) \in H$ has weight $w(u, v) = \text{dist}_G(u, v)$, and for all node pairs s, t , there exists a shortest $s \rightsquigarrow t$ path in $G \cup H$ that uses at most β edges.*

We write $\text{EH}(n, p)$ for the smallest integer β such that every n -node graph has a β -hop exact hopset of $|H| = p$ edges.

Although we give the following theorem in terms of β^* , we note that the following proof essentially shows that

$$\text{EH}(n, p) \geq \Omega\left(\frac{\text{DP}(n, p)}{p}\right) \geq \Omega\left(\frac{\beta^*(n, p, \infty)}{p}\right)$$

where the latter equality is from Theorem 32.

Theorem 49. $\text{EH}(n, p) = \Omega\left(\frac{\beta^*(n, p, \infty)}{p}\right)$.

Proof. Let $S = (V, \Pi)$ be an ordered path system with n nodes, $p + 1$ paths, ordered bridge girth ∞ , and size $\|S\| = \Omega(\beta^*(n, p + 1, \infty))$. So the average path length in S is

$$\ell = \Omega\left(\frac{\beta^*(n, p + 1, \infty)}{p + 1}\right) = \Omega\left(\frac{\beta^*(n, p, \infty)}{p}\right).$$

By the Cleaning Lemma (Lemma 10), we may assume without loss of generality that all paths in S have length $\Theta(\ell)$. As in Lemma 35, we may associate S to a directed weighted n -node graph G and set of demand pairs P such that there is a unique shortest path for each demand pair, and these shortest paths are pairwise edge-disjoint.

Now let H be an arbitrary exact hopset of G of size $|H| \leq p$. The rest of the proof is identical to Theorem 48. In particular, since $|P| = p + 1$ and $|H| = p$, there exists a demand pair $(s, t) \in P$ such that there is no edge $(x, y) \in H$ where both x, y lie along the unique shortest path $\pi(s, t)$. Thus $\pi(s, t)$ remains the unique shortest $s \rightsquigarrow t$ path in the graph $G \cup H$, and it has

$$\Theta(\ell) = \Omega\left(\frac{\beta^*(n, p, \infty)}{p}\right).$$

edges. So H cannot reduce the number of hops below this value. $\square$

4.7 Approximate Distance Preservers and Evidence for Conjecture 9

Next, we present evidence in favor of the first part of Conjecture 9. We do so by considering approximate distance preservers (also sometimes called pairwise spanners):

Definition 21 (α -Approximate Distance Preservers). Let $G = (V, E, w)$ be a directed weighted graph, $P \subseteq V \times V$ a set of demand pairs, and $\alpha \geq 1$ a parameter. An α -approximate distance preserver is a subgraph $H \subseteq G$ in which, for all $(s, t) \in P$, we have $\text{dist}_H(s, t) \leq \alpha \cdot \text{dist}_G(s, t)$.

We define $\text{ADP}(n, p, \alpha)$ as the least integer such that every n -node graph and set of $|P| = p$ demand pairs has an α -approximate distance preserver on $\leq \text{ADP}(n, p, \alpha)$ edges.

Lemma 50. $\beta(n, p, k) \leq O(\text{ADP}(n, p, 2k/\ell))$, where $\ell = \beta(n, p, k)/p$ is the average path length in a system realizing $\beta(n, p, k)$.

Proof. Let $S = (V, \Pi)$ be a path system realizing $\beta(n, p, k)$. By Lemma 10, we may assume that all path lengths in S are at most

$$\ell' < \frac{\beta(n, p, k)}{2p}.$$

Interpret the endpoints of the paths in Π as demand pairs, and consider the unweighted directed graph G that contains an edge (u, v) iff there is a path that uses the nodes u, v consecutively. The size of G is

$$|E(G)| = \Theta(\|\Pi\|) = \Theta(\beta(n, p, k)).$$

Moreover, we claim that G is the unique approximate distance preserver of G, P with error parameter $\alpha = 2k/\ell$. To see this, let $H \subsetneq G$ be a subgraph that does not contain an edge (u, v) from G , and let $\pi \in \Pi$ be a path with endpoints s, t that uses u, v consecutively. The distance from s to t in H must be at least k . Hence, the distance increases by a factor of

$$\frac{k}{\ell'} > \frac{2pk}{\beta(n, p, k)} = \frac{2k}{\ell}.$$

It follows that H is not a $2k/\ell$ approximate distance preserver of G , and the lemma follows. $\square$

We note that this lemma is *not* interesting for constant k , since we will clearly have $\ell = \Omega(k)$ in this regime and so the approximation factor is meaningless. Rather, the most interesting regime for this lemma is when $k \geq \omega(1)$ but $k \neq \infty$. This lemma *might* imply new bounds in this regime. In particular, if Conjecture 9 is false, then choosing $k = \log n$ implies new lower bounds against ADP (in the regime where $\ell = \Theta(\log n)$ as well). Or, stated in the contrapositive, and using the fact that $\text{RP}(n, p) = \Theta(\beta(n, p, \infty))$ from Theorem 37:

Theorem 51. If $\text{ADP}(n, p, \alpha) = O(\text{RP}(n, p))$ for all fixed $\alpha > 1$ then Conjecture 9 is true.

We think this premise is plausible. In support, we note that results in [72] could be interpreted as the analogous statement for *undirected* weighted distance preservers, and that results in [19] imply that for directed hopsets, the state-of-the-art bounds for error $\alpha = \infty$ (reachability) and for $(1 + \varepsilon)$ essentially match.

5 Flow-Cut and Directed Steiner Forest Integrality Gaps

In this section, we give polynomial integrality gap lower bounds for the standard LP relaxations of the directed multicut, directed sparsest cut, and directed Steiner forest problems as a function of $\beta(n, p, \infty)$. The duals of the LP relaxations of directed multicut and directed sparsest cut correspond to the well-studied maximum multicommodity flow and maximum concurrent flow problems, so the integrality gaps for these LPs correspond to *flow-cut gaps*. We also give new integrality gap lower bounds for the standard LP relaxation of directed Steiner forest. This LP relaxation can be interpreted as a generalization of minimum-cost flow to multiple demand pairs.

5.1 Directed Edge Multicut and Statement of First Result

In the directed multicut problem, we are given a directed graph $G = (V, E)$ and a set of p demand pairs $P = \{(s_i, t_i) \in V \times V \mid i \in [1, p]\}$, and the objective is to find a minimum subset of E whose removal separates all pairs of vertices (s_i, t_i) in P . It will be helpful to phrase this as an integer program, as follows:

- For each edge $e \in E$, let $x_e \in \{0,1\}$ be an indicator variable that takes value 1 if e is in the multicut solution and else 0.
- For each demand pair $(s_i, t_i) \in P$, denote by Π_i the set of directed paths from s_i to t_i in G .
- Then the multicut problem is equivalent to minimizing $\sum_e x_e$, subject to the constraint $\sum_{e \in \pi} x_e \geq 1$ for all $i \in [1, p]$ and $\pi \in \Pi_i$.

We will write $\text{MCut}(G, P)$ for the value of the directed multicut problem on inputs G, P . A natural LP relaxation of directed multicut is to let $x_e \geq 0$, so that cuts on edges can be fractional. We will write the fractional value as $\widehat{\text{MCut}}(G, P)$, and we state the LP formally as:

$(\mathbf{P}_{LP}) \quad \widehat{\text{MCut}}$ $\min \quad \sum_{e \in E} x_e$ $\text{s.t.} \quad \sum_{e \in \pi} x_e \geq 1 \quad \forall i \in [p], \pi \in \Pi_i$ $\quad \quad \quad x_e \geq 0 \quad \forall e \in E$	$(\mathbf{D}_{LP}) \quad \text{MMFlow}$ $\max \quad \sum_{i \in [p]} \sum_{\pi \in \Pi_i} f_\pi$ $\text{s.t.} \quad \sum_{\pi: e \in \pi} f_\pi \leq 1 \quad \forall e \in E$ $\quad \quad \quad f_\pi \geq 0 \quad \forall i \in [p], \pi \in \Pi_i$
--	--

The dual program $\mathbf{D}_{LP}$ of this LP relaxation is equivalent to the maximum multicommodity flow problem. In this problem, we get a directed graph G and set of demand pairs P on input, and for each demand pair $(s_i, t_i) \in P$ we choose a flow f_i that has s_i as its source, t_i as its sink, and which is conserved at all other nodes. Among the p flows $\{f_1, \dots, f_p\}$, their *total* must respect the edge capacity constraints; that is, the sum of flows on each edge must be ≤ 1 . The *value* of flow f_i is the amount of flow created at s_i and destroyed at t_i . Subject to these constraints, the goal is to maximize the sum of flow values.¹⁵ We write $\text{MMFlow}(G, P)$ for this maximized sum of flow values.

For any G, P , we have:

$$\text{MCut}(G, P) \geq \widehat{\text{MCut}}(G, P) = \text{MMFlow}(G, P).$$

The min-cut/max-flow theorem states that we have equality in the special case where $|P| = 1$, but in general the functions can be quite different. The largest possible gap between them is the *flow-cut gap*, captured by the following function:

Definition 22 (MCG). *The function $\text{MCG}(n)$ is the least integer k such that, for every n -node directed graph G and set of demand pairs P (of any size), we have*

$$\text{MCut}(G, P) \leq k \cdot \text{MMFlow}(G, P).$$

Since $\widehat{\text{MCut}}(G, P) = \text{MMFlow}(G, P)$, we may equivalently interpret this function as the integrality gap of MCut. We refer to [37] for further discussion and proofs of all of the claims in the previous discussion. Our goal is to prove:

Theorem 52. *For all n , the flow-cut gap satisfies*

$$\text{MCG}(\beta(n, n, \infty)) = \Omega\left(\frac{\beta(n, n, \infty)}{n \log n}\right).$$

As in [37], we will actually prove something slightly stronger: this gap holds even when the multicut solution only needs to disconnect a $(1 - \epsilon)$ -fraction of the demand pairs, for some $\epsilon > 0$. Our proof of Theorem 52 will closely follow the previous directed multicut integrality gap lower bound construction of [37].

¹⁵The dual LP given here works by choosing a scalar for each $s_i \rightsquigarrow t_i$ path, which implicitly defines flow values on the edges by the sum of the scalars on paths that contain that edge.

5.2 Directed Vertex Multicut

We will begin by lower bounding an integrality gap for the directed *vertex* multicut problem. This problem is defined as MCUT , except instead of deleting edges to disconnect demand pairs, we delete non-terminal *vertices*.¹⁶ The least number of nonterminal vertices required to disconnect demand pairs P in a graph G is written $\text{VMCUT}(G, P)$. We will also consider the LP relaxation $\widehat{\text{VMCUT}}$, defined as follows:

$$\begin{array}{ll} \widehat{\text{VMCUT}} \\ \min & \sum_{v \in V} x_v \\ \text{s.t.} & \sum_{v \in \pi \cap V} x_v \geq 1 \quad \forall i \in [p], \pi \in \Pi_i \\ & x_v \geq 0 \quad \forall v \in V \end{array}$$

As noted in [37], the integrality gap of the vertex multicut problem is at most the integrality gap of the (edge) multicut problem. In particular, we will use the following lemma:

Lemma 53. *[[37]] Suppose we can construct an n -node graph G and set of demand pairs P with*

$$\frac{\text{VMCUT}(G, P)}{\widehat{\text{VMCUT}}(G, P)} \geq k.$$

Then $\text{MCG}(n) \geq k$.

Thus, our lower bound on MCG will work by lower bounding the integrality gap for VMCUT . The authors of [37] used the same strategy. Specifically, they constructed a directed graph G and a set of demand pairs P such that:

1. For every pair $(s, t) \in P$, the shortest $s \rightsquigarrow t$ path in G is of length at least $L = \tilde{\Omega}(n^{1/7})$.
2. A vertex cut of size $\Omega(n)$ is required to separate the set of pairs P .

By property 1, if we assign a $1/L$ -fraction cut to each vertex in G , then we obtain a valid fractional vertex cut of G of size $\tilde{O}(n^{6/7})$. Together with property 2 this implies an integrality gap of $\tilde{\Omega}(n^{1/7})$. The authors of [37] construct the graph G by combining a certain random graph H and a graph $\mathcal{L}$ called a labeling scheme, which carries a set of demand pairs with long, unique paths. Our proof of Theorem 52 largely follows their construction and analysis, but we generalize some piece of the argument to show that the particular properties of the graph $\mathcal{L}$ are not really needed, and instead we can use an *arbitrary* bridge-free path system with n nodes and n paths that realizes the bound $\beta(n, n, \infty)$ bound.

5.3 Multicut Flow-Cut Gap Construction

Let n be a parameter. Let $S = (V, \Pi)$ be a path system with n nodes, $p = n$ paths, bridge girth ∞ , and size $\Omega(\beta(n, n, \infty))$. By the cleaning lemma (Lemma 10) we may assume without loss of generality that S is approximately degree-regular and length-regular. Specifically, we may assume every node has degree $\Theta(d)$, where d is the average degree of S . Since $p\ell = nd$ and $p = n$, d is also the average path length in S ; we thus also have that all paths have length $\Theta(d)$. Note that $d = \omega(1)$, by applying known lower bounds on $\beta(n, n, \infty)$.

Our next step is to use S to build a corresponding graph G_S that will inherit the long, unique paths property of S . We will use graph G_S as a black box version of the labeling scheme $\mathcal{L}$ in [37].

Lemma 54. *The set of paths Π can be partitioned into d nonempty sets $\Pi_1, \Pi_2, \dots, \Pi_d$ such that $|\Pi_i| = n/d$ for $i \in [1, d]$ and at least $\Omega(d)$ sets Π_i satisfy the property that¹⁷*

$$\left| \bigcup_{\pi \in \Pi_i} \pi \right| = \Omega(n).$$

¹⁶A terminal vertex is one that appears as either endpoint of a demand pair. We assume the input is such that no demand pair (s, t) has an edge going directly from s to t , so that the cut exists.

¹⁷We assume for convenience that n is divisible by d . If not, some part sizes may be rounded up or down while only affecting the following argument by lower-order terms.

Proof. Uniformly at random, partition Π into d sets Π_i for $i \in [1, d]$, each containing n/d paths. Fix a node $v \in V$ and an index $i \in [1, d]$. Note that the probability a node v belongs to a randomly chosen path in Π is at least $\Omega(d/n)$, since S is approximately degree-regular. Then v belongs to a path in Π_i with constant probability. It follows that $\mathbb{E}[|\Pi_i|] = \Omega(n)$. Since the maximum possible size is $|\Pi_i| \leq n$, by Markov's inequality we have $|\Pi_i| = \Omega(n)$ with constant probability. Thus, over all choices of $i \in [d]$, the expected number of parts Π_i satisfying $|\Pi_i| = \Omega(n)$ is $\Omega(d)$. So there exists a possible partition in which $\Omega(d)$ parts all satisfy $|\Pi_i| = \Omega(n)$. $\square$

Construction of G_S . We build our graph $G_S = (V_S \cup V'_S, E_S)$ corresponding to path system S as follows. Let V_S denote the set of nonterminal nodes of G_S , and let $V_S := V$. Let V'_S denote the set of terminal nodes $s_{i,j}, t_{i,j}$, where $i \in [1, d]$ and $j \in [1, n/d]$. Add all terminal pairs $(s_{i,j}, t_{i,j})$ to our set of demand pairs P_S . Let $\Pi_1, \Pi_2, \dots, \Pi_d$ be the partition of Π as specified in Lemma 54. For every path $\pi \in \Pi_i$, if edge e is in the transitive closure of path π , then add e to edge set E_i^S . Here, we say that e is in the transitive closure of π if π contains e as a (possibly noncontiguous) subsequence. Additionally, for $i \in [1, d]$, order the n/d paths in Π_i arbitrarily, and let $\pi_{i,j}$ denote the j th path in Π_i , for $i \in [1, d]$ and $j \in [1, n/d]$. Add to edge set E_i^S an edge from $s_{i,j}$ to the first vertex of $\pi_{i,j}$. Likewise, add to edge set E_i^S an edge from the last vertex of $\pi_{i,j}$ to $t_{i,j}$. We define the edge set E_S of G_S to be $E_S = \cup_i E_i^S$. We refer to the edges of G_S belonging to E_i^S as edges of type i . We say that an $s_{i,j} \rightsquigarrow t_{i,j}$ path is canonical if it is composed exclusively of edges of type i . The following properties of G_S follow immediately from our choice of path system S .

Observation 55. *Graph G_S has the following properties:*

1. Every $s_{i,j} \rightsquigarrow t_{i,j}$ path in G_S is a canonical path, for all $(s_{i,j}, t_{i,j}) \in P_S$.
2. There exists an $s_{i,j} \rightsquigarrow t_{i,j}$ path of length at least $\Omega(d)$ in G_S , for all $(s_{i,j}, t_{i,j}) \in P_S$.

Recall that graph G_S is intended to replace the *labeling scheme* graph $\mathcal{L}$ in the argument of [37]. Our final graph G will be the product of graph G_S and an additional graph H that we construct next. Roughly, the properties of G_S summarized above will ensure that our final graph G has long shortest paths between all demand pairs, and therefore has a small fractional vertex multicut. The graph H will be a well-connected random graph, which will roughly ensure that our final graph G has a large minimum vertex multicut of its demand pairs. These two properties together will ensure that G has a large integrality gap between its fractional and integral vertex multicut. The graph H will be essentially identical to the graph H given in [37], but with different construction parameters.

Construction of H (c.f. Section 3.1.2 of [37]). We build $H = (V_H \cup V'_H, E_H)$ as follows. Let V_H denote the set of nonterminal nodes of H , and let $V_H := \{v_1, \dots, v_d\}$. Additionally, graph H will have d distinct pairs of terminal nodes $P_H := \{(s_i, t_i) \mid i \in [1, d]\}$ as demand pairs, with V'_H denoting the set of all terminal nodes. Graph H will be defined as the union of graphs $H_i = (V \cup \{s_i, t_i\}, E_i^H)$ for $i \in [1, d]$.

We construct graph H_i for $i \in [1, d]$ as follows. Graph H_i has terminal nodes s_i, t_i and will contain $d' := d/(2 \log d)$ layers each containing $\log d$ nonterminal nodes. We denote the layers as $L_i^1, L_i^2, \dots, L_i^{d'} \subseteq V_H$ and construct them sequentially as follows. To construct the j th layer L_i^j for $j \in [1, d']$, select uniformly at random $\log d$ distinct nodes from $V_H \setminus (L_i^1 \cup L_i^2 \cup \dots \cup L_i^{j-1})$. Note that by construction, $|\cup_j L_i^j| \leq d/2$. Now define the set of edges E_i^H of H_i as follows. Add an edge from s_i to every vertex in L_i^1 . Likewise, add an edge from every vertex in $L_i^{d'}$ to t_i . Finally, add an edge from every vertex in layer L_i^j to every vertex in layer L_i^{j+1} for $j \in [1, d' - 1]$. This concludes the construction of H_i . We define H to be $\cup_i H_i$. Likewise, we define $E_H = \cup_i E_i^H$ to be the edge set of H . We refer to the edges in H belonging to E_i^H as edges of type i . We say that an $s_i \rightsquigarrow t_i$ path in H is a canonical path if it contains only edges of type i . The following properties of graph H will be used in the analysis.

Observation 56. *Graph H has the following properties:*

1. Every canonical $s_i \rightsquigarrow t_i$ path in H contains at least d' nonterminal nodes.
2. $\Omega(d)$ nonterminal nodes must be removed from H to disconnect a constant fraction of demand pairs in P_H .

Property 1 is immediate from construction, and property 2 holds for H with high probability, as proven in [37] and proven in a slightly stronger form below.

Lemma 57 (c.f. Lemma 3.1 of [37]). *Fix an $\varepsilon > 0$. For sufficiently large d , the following holds with probability $\geq 1 - 2^{-d}$:*

There does not exist a set $\mathcal{S} \subseteq V_H$ of nonterminal nodes of H , of size $|\mathcal{S}| \leq d/16$, such that for more than $(1 - \varepsilon)d$ distinct indices $i \in [1, d]$, removing $\mathcal{S}$ disconnects the demand pair $(s_i, t_i) \in P_H$ in graph H_i .

Proof. We defer the proof of this lemma to Appendix C.1 since it follows from the same argument as Lemma 3.1 of [37]. $\square$

Note that d is sufficiently large, since we assumed that $d = \omega(1)$. Then we may assume that property 2 of H , as formalized in Lemma 57, holds for the specific graph H we will use in our construction of G . Note that while all canonical $s_i \rightsquigarrow t_i$ paths in H are of length at least d' by property 1 of H , this is not true in general for all $s_i \rightsquigarrow t_i$ paths in H . We will see that composing graph H with graph G_S will allow us to ensure that the shortest paths between all demand pairs are of length d' in the final graph G .

Construction of G . We now construct our final graph $G = (V_G \cup V'_G, E)$ by composing G_S and H in a natural way. Let graphs $G_S = (V_S \cup V'_S, E_S)$ and $H = (V_H \cup V'_H, E_H)$ be as defined previously. Let $V_G := V_H \times V_S$ be the set of nonterminal vertices of G . We let V'_G denote the set of terminal vertices of G and let $V'_G := V'_S$. Likewise, we let P denote the set of n demand pairs in G and let $P := P_S$.

The set of edges E of G are defined as follows. Let (x, x') and (y, y') be nonterminal vertices in V_G .

- We add edge $((x, x'), (y, y'))$ to set E_i if $(x, y) \in E_i^H$ and $(x', y') \in E_i^S$.
- For $s_{i,j} \in V'_G$ and $(x, x') \in V_G$ we add edge $(s_{i,j}, (x, x'))$ to E_i if $(s_i, x) \in E_i^H$ and $(s_{i,j}, x') \in E_i^S$.
- Finally, we add edge $((x, x'), t_{i,j})$ to E_i if $(x, t_i) \in E_i^H$ and $(x', t_{i,j}) \in E_i^S$. We let $E = \cup_i E_i$.

As with G_S and H , we refer to the edges in E_i as edges of type i , and we say an $s_{i,j} \rightsquigarrow t_{i,j}$ path in G is canonical if its composed of only type i edges. This completes the construction of G .

The number of nonterminal vertices in graph G is $N := |V_H| \cdot |V_S| = dn = \beta(n, p, \infty)$. Then the value of d' is

$$d' = \frac{d}{2 \log d} \geq \frac{N}{2n \log N} = \Omega\left(\frac{\beta(n, p, \infty)}{n \log n}\right).$$

In the following section, we will show that the multicut integrality gap of G is at least d' , which will complete our lower bound.

5.4 Integrality Gap Analysis of G

We now analyze the gap between the fractional cost of a vertex multicut of G, P and the cost of an integral vertex multicut of G, P . Our analysis largely follows that of [37].

Fractional solution. Assign a fractional cut of $1/d'$ to each nonterminal node in G . We will show that this is a valid fractional vertex multicut of G of size N/d' . It is clear that the size of this prospective cut is N/d' as desired. The validity of this fractional cut will be immediate from the following claim.

Claim 58. *For all $(s_{i,j}, t_{i,j}) \in P$, any $s_{i,j} \rightsquigarrow t_{i,j}$ path in G contains at least d' nonterminal nodes.*

Proof. Observe that every $s_{i,j} \rightsquigarrow t_{i,j}$ path π for $(s_{i,j}, t_{i,j}) \in P$ corresponds to a $s_i \rightsquigarrow t_i$ path in H and a $s_{i,j} \rightsquigarrow t_{i,j}$ path in G_S . Then every $s_{i,j} \rightsquigarrow t_{i,j}$ path in G is canonical, since every $s_{i,j} \rightsquigarrow t_{i,j}$ path in G_S is canonical, by property 1 of Observation 55. If a $s_{i,j} \rightsquigarrow t_{i,j}$ path in G is canonical, then the corresponding $s_i \rightsquigarrow t_i$ path in H is canonical by the definition of G , and so by property 1 of Observation 56, π has at least d' nonterminal nodes. $\square$

Integral solution. We will show that $\Omega(N)$ vertices must be removed from G to separate all demand pairs in P . Before proving this, we must first introduce some notation and prove some intermediate results. For the remainder of this section, fix $\mathcal{S} \subseteq V_G$ to be any subset of non-terminal vertices of G with $|\mathcal{S}| \leq \lambda N$, where $\lambda > 0$ is a sufficiently small constant to be specified later. For any set $U \subseteq V_G$, we let U^H denote the preimage of U in H . Namely,

$$U^H := \{v \in V_H \mid (v, v') \in U\}.$$

For $v \in V_S$, we define the set $\mathcal{S}_v \subseteq V_G$ as

$$\mathcal{S}_v := \mathcal{S} \cap (V_H \times \{v\}).$$

Now for $i \in [1, d]$ and $v \in V_S$, we say that the pair $(i, v) \in [1, d] \times V_S$ is H -good if $\mathcal{S}_v^H$ does not disconnect demand pair $(s_i, t_i) \in P_H$ in graph H_i .

Claim 59. *Fix an $\varepsilon > 0$. There are at least $(1 - \varepsilon)(1 - 16\lambda)N$ pairs $(i, v) \in [1, d] \times V_S$ that are H -good.*

Proof. Note that there are $N = dn$ pairs $(i, v) \in [1, d] \times V_S$. For any $v \in V_S$, we know that if $|\mathcal{S}_v^H| \leq d/16$, then there are at least $(1 - \varepsilon)d$ pairs (i, v) that are H -good by Lemma 57. Moreover, since $|\mathcal{S}| \leq \lambda N$, it follows that there are fewer than $16\lambda n$ vertices $v \in V_S$ such that $|\mathcal{S}_v^H| > d/16$. Then the number of pairs $(i, v) \in [1, d] \times V_S$ that are H -good is at least

$$(1 - \varepsilon)(1 - 16\lambda)N. \quad \square$$

We say that a pair $(i, v) \in [1, d] \times V_S$ is S -good if there is a path π in Π_i such that $v \in \pi$.

Claim 60. *There are $\Omega(N)$ pairs $(i, v) \in [1, d] \times V_S$ that are S -good.*

Proof. By Lemma 54 at least $\Omega(d)$ of the sets Π_i , $i \in [1, d]$, satisfy $|\cup_{\pi \in \Pi_i} \pi| \geq \Omega(n)$. Then the number of pairs $(i, v) \in [1, d] \times V_S$ such that (i, v) is S -good is $\Omega(N)$. $\square$

Let $\lambda_1 > 0$ be a constant such that for sufficiently large N at least $\lambda_1 N$ pairs (i, v) in $[1, d] \times V_S$ are S -good. By Claim 60, such a λ_1 must exist. We choose λ to be

$$\lambda := \lambda_1/32.$$

We say that a pair $(i, v) \in [1, d] \times V_S$ is G -good if (i, v) is H -good and S -good. Let $J \subseteq [1, d]$ be the set of all indices $j \in [1, d]$ such that the number of G -good pairs in $\{j\} \times V_S$ is at least $\lambda_1^2/8 \cdot n$.

Claim 61. $|J| \geq \lambda_1^2/8 \cdot d$.

Proof. Let $\varepsilon = \lambda_1/2$. Since there are at least $(1 - \varepsilon)(1 - 16\lambda)N$ H -good pairs and at least $\lambda_1 N$ S -good pairs, by an overlap argument it follows that there are at least

$$((1 - \varepsilon)(1 - 16\lambda) + \lambda_1 - 1)N = ((1 - \lambda_1/2)^2 + \lambda_1 - 1)N = \lambda_1^2/4 \cdot N$$

G -good pairs. Now note that for any $i \in J$, there are at most n G -good pairs in $\{i\} \times V_S$. Likewise, for any $i \notin J$, there are at most $\lambda_1^2/8 \cdot n$ G -good pairs in $\{i\} \times V_S$. Then there are at most

$$n|J| + \lambda_1^2/8 \cdot n(d - |J|)$$

G -good pairs. We obtain the following inequality:

$$n|J| + \lambda_1^2/8 \cdot n(d - |J|) \geq \lambda_1^2/4 \cdot N$$

Solving for $|J|$, we conclude that $|J| \geq \lambda_1^2/8 \cdot d = \Omega(d)$. $\square$

We need one more claim before we can prove that $\mathcal{S}$ is not a valid multicut.

Claim 62. *For every $i \in J$, there are $\Omega(\frac{n}{d})$ paths $\pi \in \Pi_i$ such that at least d' pairs in $\{i\} \times \pi$ are G -good.*

Proof. Fix an $i \in J$. By the definition of J , there are at least $\lambda_1^2/8 \cdot n$ G -good pairs in $\{i\} \times V_S$. Now let $J_i \subseteq \Pi_i$ denote the set of all paths π in Π_i such that at least d' pairs in $\{i\} \times \pi$ are G -good. Since S is approximately length-regular by the cleaning lemma, each path π in J_i is of length at most $|\pi| \leq cd$ for some constant $c \geq 1$, and therefore $\{i\} \times \pi$ has at most cd G -good pairs. Likewise, each path $\pi \in \Pi_i \setminus J_i$ has at most d' G -good pairs. Recall that $(i, v) \in \{i\} \times V_S$ is G -good only if $v \in \pi$ for some $\pi \in \Pi_i$. Then by the above discussion there are at most

$$cd|J_i| + d' \left(\frac{n}{d} - |J_i| \right)$$

G -good pairs in $\{i\} \times V_S$. We obtain the following inequality:

$$cd|J_i| + d' \left(\frac{n}{d} - |J_i| \right) \geq \lambda_1^2/8 \cdot n$$

Using the fact that $d' = d/(2 \log d) \leq \lambda_1^2/16 \cdot d$ for sufficiently large d , we conclude that

$$|J_i| \geq \frac{\lambda_1^2}{16c} \cdot \frac{n}{d} \quad \square$$

We will now show that a constant fraction of the demand pairs P remain connected in $G \setminus S$.

Lemma 63. $\Omega(n)$ demand pairs in P are connected in $G \setminus S$.

Proof. Fix an $i \in J$ and a $j \in [1, n/d]$ such that path $\pi_{i,j}$ satisfies the property of Claim 62. (Recall that $\pi_{i,j}$ is the canonical $s_{i,j} \rightsquigarrow t_{i,j}$ path in G_S .) Now let $u_1, u_2, \dots, u_{d'}$ be a set of d' vertices in $\pi_{i,j}$ (listed in the order in which they appear in $\pi_{i,j}$) such that all pairs in $\{i\} \times \{u_1, \dots, u_{d'}\}$ are G -good. Recall that by the construction of G_S , all edges in the transitive closure of $\pi_{i,j}$ are added to E_i^S in graph G_S , so in particular, $(s_{i,j}, u_1), (u_k, t_{i,j}) \in E_i^S$ and $(u_k, u_{k+1}) \in E_i^S$ for all $k \in [1, d' - 1]$.

Now for all $k \in [1, d']$, we claim that there exists a vertex in $L_i^k \times \{u_k\} \subseteq V_G$ that survives in $G \setminus S$, i.e. $(L_i^k \times \{u_k\}) \setminus S \neq \emptyset$. Recall that L_i^k denotes the vertices in the k th layer of H_i . We know that (i, u_k) is a G -good pair, so it is also an H -good pair, which implies that $L_i^k \not\subseteq S_v^H$. (If $L_i^k \subseteq S_v^H$, then S_v^H would disconnect (s_i, t_i) in H_i , a contradiction.) Let w_k denote a vertex in $L_i^k \setminus S_v^H \in V_H$ for $k \in [1, d']$. Note that since $w_k \in L_i^k$ for $k \in [1, d']$, it follows from the construction of H that $(w_k, w_{k+1}) \in E_i^H$ for $k \in [1, d' - 1]$.

Let $x_k = (w_k, u_k) \in V_G$ for $k \in [1, d']$. By the discussion in the last paragraph, $x_k \in V_G$ survives after S is removed from G , i.e. $x_k \in G \setminus S$ for $k \in [1, d']$. Furthermore, since $(u_k, u_{k+1}) \in E_i^S$ and $(w_k, w_{k+1}) \in E_i^H$ for $k \in [1, d' - 1]$, by the construction of G it follows that edge $(x_k, x_{k+1}) \in E$ survives in G after S is removed, for $k \in [1, d' - 1]$. Finally, note that by the construction of G , edges $(s_{i,j}, x_1)$ and $(x_{d'}, t_{i,j})$ are in E . Consequently, $(s_{i,j}, x_1, x_2, \dots, x_{d'}, t_{i,j})$ is a valid $s_{i,j} \rightsquigarrow t_{i,j}$ path in $G \setminus S$, so demand pair $(s_{i,j}, t_{i,j}) \in P$ is connected in $G \setminus S$. Since $|J| = \Omega(d)$ by Claim 61 and for all $i \in J$ there are $\Omega(\frac{n}{d})$ paths π in Π_i satisfying the property of Claim 62, we conclude that $\Omega(n)$ of the $|P| = n$ demand pairs in P are connected in $G \setminus S$. $\square$

We assumed S was an arbitrary vertex set of size $|S| \leq N/1600$, so we conclude by Lemma 63 that any vertex multicut of G, P is of size $\Omega(N)$. Then since G, P has a fractional vertex multicut of cost N/d' , we obtain an integrality gap of $\Omega(d')$ for the minimum vertex multicut problem. Additionally, even if the multicut solution only needs to disconnect a constant fraction $(1 - \varepsilon)n$ of the n demand pairs in P for some sufficiently small $\varepsilon > 0$, the size of the vertex multicut of G, P remains $\Omega(N)$ by Lemma 63. (We will make use of this fact in the sparsest cut flow-cut gap argument.) Theorem 52 is immediate from the above discussion and Lemma 53.

5.5 Sparsest Cut Flow-Cut Gap

In the directed sparsest cut problem, we are given a directed graph $G = (V, E)$ and a set of p demand pairs $P = \{(s_i, t_i) \in V \times V \mid i \in [1, p]\}$, and the objective is to find a subset E' of E that minimizes the ratio $|E'|/|P_{E'}|$, where $P_{E'}$ is the subset of P that is disconnected in graph $G \setminus E'$. It will be helpful to phrase this as an integer program, as follows:

- For each edge $e \in E$, let $x_e \in \{0, 1\}$ be an indicator variable that takes value 1 if e is in the solution E' .

- For each demand pair $(s_i, t_i) \in P$, denote by Π_i the set of directed paths from s_i to t_i in G .
- For each $i \in [1, p]$, let $h_i \in \{0, 1\}$ be an indicator variable that takes value 1 if source-sink pair (s_i, t_i) is disconnected in $G \setminus E'$, i.e. $(s_i, t_i) \in P_{E'}$.
- Let $D := \sum_{i=1}^k h_i$ be the total number of disconnected pairs $|P_{E'}|$. For $e \in E$, let $x'_e := x_e/D$, and for $i \in [1, p]$, let $h'_i := h_i/D$.
- Then the sparsest cut problem is equivalent to minimizing $\sum_e x'_e$ subject to $\sum_i h'_i \geq 1$ and $\sum_{e \in \pi} x'_e \geq h'_i$ for all $i \in [1, p]$ and $\pi \in \Pi_i$.

We will write $\text{SCUT}(G, P)$ for the value of the directed sparsest cut problem on inputs G, P . A natural LP relaxation of directed sparsest cut is to let $x_e \geq 0$, so that cuts on edges can be fractional. We will write the fractional value as $\widehat{\text{SCUT}}(G, P)$, and we will state the LP formally as:

(P_{LP})	$\widehat{\text{SCUT}}$	(D_{LP})	MCFLOW
min	$\sum_{e \in E} x'_e$	max	λ
s.t.	$\sum_{e \in \pi} x'_e \geq h'_i \quad \forall i \in [p], \pi \in \Pi_i$	s.t.	$\sum_{\pi \in \Pi_i} f_\pi \geq \lambda \quad \forall i \in [p]$
	$\sum_{i=1}^p h'_i \geq 1$		$\sum_{\pi: e \in \pi} f_\pi \leq 1 \quad \forall e \in E$
	$x'_e, h'_i \geq 0 \quad \forall e \in E, \forall i \in [p]$		$f_\pi \geq 0 \quad \forall i \in [p], \pi \in \Pi_i$

The dual program D_{LP} of this LP relaxation is equivalent to the maximum concurrent flow problem. In this problem, we are given a directed graph G and a set of demand pairs P on input, and for each demand pair $(s_i, t_i) \in P$ we choose a flow f_i that has s_i as its source, t_i as its sink, and which is conserved at all other nodes. As with minimum multicut, the sum of the flows on each edge must be ≤ 1 , and the value $|f_i|$ of each flow f_i is the amount of flow created at s_i and destroyed at t_i . Subject to these constraints, the goal is to maximize $\lambda = \min_i |f_i|$, the least amount of flow routed from s_i to t_i for any $(s_i, t_i) \in P$. We write $\text{MCFLOW}(G, P)$ to denote this maximized λ . We define the following function to capture the flow-cut gap between maximum concurrent flow and sparsest cut:

Definition 23. *The function $\text{SCG}(n)$ is the least integer k such that, for every n -node directed graph G and set of demand pairs P (of any size), we have*

$$\text{SCUT}(G, P) \leq k \cdot \text{MCFLOW}(G, P).$$

Since $\widehat{\text{SCUT}}(G, P) = \text{MCFLOW}(G, P)$ by LP duality, we may equivalently interpret this function as the integrality gap of $\widehat{\text{SCUT}}$. We refer to [37] for further discussion. Using our graph G from the proof of Theorem 52 and a standard reduction argument from [37], we can prove:

Theorem 64. *For all n , the flow-cut gap between maximum concurrent flow and sparsest cut satisfies*

$$\text{SCG}(\beta(n, n, \infty)) = \Omega\left(\frac{\beta(n, n, \infty)}{n \log n}\right).$$

We defer the proof of this theorem to Appendix C.2 since it is essentially identical to the argument in Section 3.2 of [37].

5.6 Integrality Gap of the Flow LP of Directed Steiner Forest

In the Directed Steiner Forest problem, we are given a weighted, directed graph $G = (V, E, w)$ with weight $w : E \mapsto \mathbb{R}_{\geq 0}$, and a set of p demand pairs $P = \{(s_i, t_i) \in V \times V \mid i \in [1, p]\}$. We are asked to return a subgraph $H \subseteq G$ minimizing $\sum_{e \in E(H)} w_e$, subject to there being a directed path from s_i to t_i in H for all $(s_i, t_i) \in P$. For each edge $e \in E$, let $x_e \in \{0, 1\}$ be an indicator variable that takes value 1 if e is in the solution subgraph and else 0. We can rephrase the Directed Steiner Forest problem as the following (informal) integer program:

DSF IP

$$\begin{array}{ll}
\min & \sum_{e \in E} w_e x_e \\
\text{s.t.} & \text{edge capacities } \{x_e\}_{e \in E} \text{ support a one unit } s\text{-}t \text{ flow } f_{s,t} \text{ in } G \quad \forall (s, t) \in P \\
& x_e \in \{0, 1\}
\end{array}$$

We will write $\text{DSF}(G, P)$ for the value of the Directed Steiner Forest problem on inputs G, P . A natural LP relaxation of Directed Steiner Forest is to let $x_e \geq 0$, so that edge capacities can be fractional. We refer to this LP relaxation as the “flow LP” of Directed Steiner Forest, since we can interpret it as a generalization of the s - t minimum-cost flow problem to multiple demand pairs. We will write $\widehat{\text{DSF}}(G, P)$ for the value of the flow LP on inputs G, P . Now let Π_i denote the set of all $s_i \rightsquigarrow t_i$ -paths in G for all $i \in [1, p]$. We formally state the flow LP as:

Flow LP for DSF

$$\begin{array}{llll}
\min & \sum_{e \in E} w_e x_e \\
\text{s.t.} & \sum_{\pi \in \Pi_i} f_{\pi}^i & \geq & 1 \quad \forall i \in [1, p] \\
& \sum_{\pi \in \Pi_i} f_{\pi}^i & \leq & x_e \quad \forall e \in E, i \in [1, p] \\
& x_e & \geq & 0 \quad \forall e \in E \\
& f_{\pi}^i & \geq & 0 \quad \forall i \in [1, p], \pi \in \Pi_i
\end{array}$$

We define the following function to capture the integrality gap of the flow LP of Directed Steiner Forest as a function of the sizes of the inputs G and P .

Definition 24 (DSFG). *The function $\text{DSFG}(n, p)$ is the least integer k such that for every n -node weighted, directed graph G and set of demand pairs P of size $|P| = p$, we have*

$$\text{DSF}(G, P) \leq k \cdot \widehat{\text{DSF}}(G, P).$$

Our goal is to prove:

Theorem 65. *For all n and $p \in [1, n^{2-o(1)}]$,*

$$\text{DSFG}(n, p) = \Omega\left(\frac{\beta(n, p, \infty)}{n^{3/2}}\right).$$

In particular, $\text{DSFG}(n, n^{2-o(1)}) = \Omega(n^{1/2-o(1)})$.

As with our flow-cut gap lower bounds, we will actually achieve this integrality gap by lower bounding the integrality gap of *Vertex* Directed Steiner Forest, which we define below.

Definition 25 (Vertex Directed Steiner Forest). *In the Vertex Directed Steiner Forest problem, we are given an n -node graph $G = (V \cup V', E)$, where V denotes the nonterminal vertices of G and V' denotes the terminal vertices of G and $V \cap V' = \emptyset$. We are also given a set of demand pairs $P \subseteq V' \times V'$. We are asked to return a subgraph $H \subseteq G$ minimizing $|V(H) \cap V|$, subject to there being a directed path from s to t in H for all $(s, t) \in P$.*

We write $\text{VDSF}(G, P)$ for the value of the Vertex Directed Steiner Forest problem on inputs G, P . We will also consider a natural LP relaxation $\widehat{\text{VDSF}}$, defined as follows.

$\widehat{\text{VDSF}}$

$$\begin{array}{ll}
\min & \sum_{v \in V} x_v \\
\text{s.t.} & \text{node capacities } x_v \text{ support a one unit } s\text{-}t \text{ flow } f_{s,t} \text{ in } G \quad \forall (s, t) \in P \\
& x_v \geq 0
\end{array}$$

Let $\widehat{\text{VDSF}}(G, P)$ denote the value of the $\widehat{\text{VDSF}}$ LP on inputs G, P . To lower bound DSFG , it will suffice to lower bound the integrality gap of $\widehat{\text{VDSF}}$. Specifically, we will need the following lemma:

Claim 66. *Suppose we can construct an n -node graph G and a set of demand pairs P of size p with*

$$\frac{\text{VDSF}(G, P)}{\widehat{\text{VDSF}}(G, P)} \geq k.$$

Then $\text{DSFG}(2n, p) \geq k$.

Proof. This claim follows from a standard reduction of maximum flow in node-capacitated graphs to maximum flow in edge-capacitated graphs. We defer the proof to Appendix C.3. $\square$

Our Vertex Directed Steiner Forest instance G, P will have the following two properties:

1. For all $(s, t) \in P$, there are at least $\Omega(\beta(n, p, \infty)/n)$ pairwise internally vertex-disjoint $s \rightsquigarrow t$ paths in G .
2. Any feasible subgraph $H \subseteq G$ must satisfy $|V(H) \cap V| = \Omega(\sqrt{n})$.

Note that by property 1, if we assign a fractional node capacity of $x_v = cn/\beta(n, p, \infty)$ for some sufficiently large constant $c > 0$ to each vertex $v \in V$, then we obtain a feasible solution to $\widehat{\text{VDSF}}$ LP of size $O(n^2/\beta(n, p, \infty))$. This, together with property 2, implies an integrality gap for Vertex Directed Steiner Forest on the G, P of size $\Omega(\sqrt{n}/(n^2/\beta(n, p, \infty))) = \Omega\left(\frac{\beta(n, p, \infty)}{n^{3/2}}\right)$, as desired.

To construct G, P , we will start with a path system S on n nodes and p paths, and with bridge girth ∞ and size $\|S\| = \beta(n, p, \infty)$. To obtain our desired construction, we will need to modify $S = (V, \Pi)$ so that it is *source-restricted* with respect to a set of source nodes $X \subseteq V$. We say:

Definition 26 (Source-Restricted Path Systems). *A path system $S = (V, \Pi)$ is source-restricted with respect to some $X \subseteq V$ if every path $\pi \in \Pi$ has its first node in X and every following node in $V \setminus X$.*

To obtain our desired source-restricted path system, we will use the following modified cleaning lemma.

Lemma 67 (Source-Restricted Cleaning Lemma). *For all n, p, k , there exists a path system S with $\leq n$ nodes, $\leq p$ paths, bridge girth $> k$, $\|S\| = \Omega(\beta(n, p, k))$, and the following two additional properties:*

- *S satisfies the properties of the original Cleaning Lemma (Lemma 10); that is:*
 - (Approximately Degree-Regular) *All nodes have degree $\Theta(d)$, where d is the average degree in S , and*
 - (Approximately Length-Regular) *All paths have length $\Theta(\ell)$, where ℓ is the average length in S .*
- *S is source-restricted with respect to a set X of size $|X| = \Theta(p/d)$.*

Proof. We defer the proof of this lemma to Appendix C.4 since it's similar to the original cleaning lemma (Lemma 10). $\square$

Now, using the source-restricted cleaning lemma, we may assume $S = (V, \Pi)$ is a path system on n nodes and $p \in [1, n^{2-o(1)}]$ paths that is source-restricted with respect to a set $X \subseteq V$ of size $|X| = \Theta(p/d)$; S has bridge girth ∞ and size $\|S\| = \Theta(\beta(n, p, \infty))$; and S is approximately degree-regular and approximately length-regular. Let d be the average node degree of S and ℓ be the average path length of S ; we will assume that $\ell = \omega(1)$. We can easily guarantee this assumption by requiring that $p = \frac{n^2}{e^{\omega(\sqrt{\log n})}} \leq n^{2-o(1)}$ (this follows from existing reachability preserver lower bounds implied by [16]).

For each $x \in X$, let Π_x denote the set of paths in Π that start with node x . Note that by construction, $\{\Pi_x\}_{x \in X}$ is a partition of the set of paths in Π . We will use path system S , along with set X and collection $\{\Pi_x\}_{x \in X}$, to construct our directed graph G and set of demand pairs P . Roughly, each node $x \in X$ will be a terminal source node in G , and for each $x \in X$ we will add a new terminal sink node y_x to G . The paths in Π_x will become the internally vertex-disjoint $x \rightsquigarrow y_x$ paths in G for $(x, y_x) \in P$. We will explicitly construct G using the following procedure.

- Let $G \leftarrow (V, \emptyset)$, let $P \leftarrow \emptyset$, and let $V' \leftarrow X$.
- Fix a path $\pi \in \Pi_x$ for some $x \in X$. For each consecutive pair of vertices in π , add a directed edge between the corresponding pair of vertices in G . Repeat this procedure for each $\pi \in \Pi_x$, where $x \in X$.
- For each $x \in X$, add a new terminal vertex y_x to G , so that $G \leftarrow G \cup \{y_x\}$. Let $V' \leftarrow V' \cup \{y_x\}$. For each path $\pi[x \rightsquigarrow y] \in \Pi_x$, add a directed edge (y, y_x) to G . Add the demand pair (x, y_x) to P , so that $P \leftarrow P \cup \{(x, y_x)\}$.
- Return the directed graph G and set of demand pairs P . The terminal vertices of G will be V' , and the nonterminal vertices of G will be $V \setminus V'$.

We will now prove that the resulting graph G and set of demand pairs P has our desired properties.

Lemma 68. *The above procedure outputs a directed graph $G = (V \cup V', E)$ and a set of demand pairs $P \subseteq V' \times V'$ of size $|P| = \Theta(p/d)$ satisfying the following properties:*

- for all $(x, y_x) \in P$, there are $\Theta(d)$ pairwise internally vertex-disjoint $x \rightsquigarrow y_x$ paths in G .
- for all $(x, y_x) \in P$, any solution $H \subseteq G$ to vertex directed Steiner forest on G, P must contain a subpath of length $\Omega(\ell)$ of some path in Π_x .

Proof. Fix a node $x \in X$. Note that x has degree $\Theta(d)$ in S by Lemma 67, so $|\Pi_x| = \Theta(d)$. Note that each path $\pi \in \Pi_x$ implies an $x \rightsquigarrow y_x$ path in G . These paths are pairwise internally node-disjoint, since this would otherwise imply a 2-bridge in S . This proves the first property.

To see why the second property is true, observe that every $x \rightsquigarrow y_x$ -path in G contains as a subpath a path $\pi \in \Pi_x$ (otherwise, this would imply a bridge in S). Then since every path $\pi \in \Pi_x$ is of length $|\pi| = \Omega(\ell)$, the second property immediately follows. $\square$

Now we are ready to lower bound the integrality gap of $\widehat{\text{VDSF}}$ on G, P using Lemma 68.

Fractional solution. For each vertex $v \in V$, let the fractional node capacity be $x_v = c/d$ for a sufficiently large constant $c > 0$. Then by property 1 of Lemma 68, since there are $\Theta(d)$ pairwise internally vertex-disjoint $x \rightsquigarrow y_x$ paths in G for all $(x, y_x) \in P$, our node capacities $\{x_v\}_{v \in V}$ support one unit of flow for all demand pairs $(x, y_x) \in P$. Then our fractional solution is feasible and has size $cn/d = O(n/d)$.

Integral solution. By property 2 of Lemma 68, for all $(x, y_x) \in P$, any solution subgraph $H \subseteq G$ to Vertex Directed Steiner Forest on G, P must contain a subpath of length $\Omega(\ell)$ of some path in Π_x . Note that $\Pi_{x_1} \cap \Pi_{x_2} = \emptyset$ for distinct $x_1, x_2 \in X$. Additionally, note that for distinct $\pi_1, \pi_2 \in \Pi$, the corresponding paths in G are edge-disjoint, since S has bridge girth ∞ . Then we conclude that any feasible solution subgraph $H \subseteq G$ to Vertex Directed Steiner Forest must have at least

$$|E(H)| = |P| \cdot \Omega(\ell) = \Omega\left(\frac{\ell p}{d}\right) = \Omega(n)$$

edges. Moreover, $\Omega(n)$ of these edges must be in the induced subgraph $H[V]$, since path system S is source-restricted with respect to X . If the number of edges in $H[V]$ is $|E(H[V])| = \Omega(n)$, then

$$|V(H) \cap V| = |V(H[V])| = \Omega(|E(H[V])|^{1/2}) = \Omega(\sqrt{n}).$$

We conclude that the integral solution of Vertex Directed Steiner Forest must have size at least $\Omega(\sqrt{n})$.

By the above analysis, the integrality gap of Vertex Directed Steiner Forest on G, P is at least $\Omega(\sqrt{n}/(n/d)) = \Omega(d/\sqrt{n}) = \Omega(\beta(n, p, \infty)/n^{3/2})$, as desired. We note that in particular, the generalized Ruzsa-Szemerédi lower bound constructions implied by [16] prove that for $p = n^{2-o(1)}$, we have that $\beta(n, p, \infty) = \Omega(n^{2-o(1)})$, and in particular, the average path length is $\ell = \beta(n, p, \infty)/p = \omega(1)$. Consequently, $\text{DSFG}(n, n^{2-o(1)}) = \Omega(n^{1/2-o(1)})$. Theorem 65 follows from the above discussion and Claim 66.

Acknowledgments

We are grateful to Omer Reingold, Vivek Madan, Matthew Fahrbach, and Idan Shabat for helpful technical discussions.

References

- [1] Amir Abboud and Greg Bodwin. The $4/3$ additive spanner exponent is tight. *Journal of the ACM (JACM)*, 64(4):28:1–28:14, 2017. [1](#)
- [2] Amir Abboud and Greg Bodwin. Reachability preservers: New extremal bounds and approximation algorithms. In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1865–1883. Society for Industrial and Applied Mathematics, 2018. [1](#), [4](#), [5](#), [32](#), [34](#)
- [3] 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. Society for Industrial and Applied Mathematics, 2017. [1](#)
- [4] Amit Agarwal, Noga Alon, and Moses S Charikar. Improved approximation for directed cut problems. In *Proceedings of the thirty-ninth annual ACM symposium on Theory of computing*, pages 671–680, 2007. [7](#), [10](#)
- [5] Reyhan Ahmed, Greg Bodwin, Faryad Darabi Sahneh, Keaton Hamm, Mohammad Javad Latifi Jebelli, Stephen Kobourov, and Richard Spence. Graph spanners: A tutorial review. *Computer Science Review*, 37:100253, 2020. [1](#)
- [6] Donald Aingworth, Chandra Chekuri, Piotr Indyk, and Rajeev Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM Journal on Computing*, 28(4):1167–1181, 1999. [1](#)
- [7] Bandar Al-Dhalaan. Fast construction of 4-additive spanners. *arXiv preprint arXiv:2106.07152*, 2021. [1](#)
- [8] Noga Alon. Testing subgraphs in large graphs. *Random Structures & Algorithms*, 21(3-4):359–370, 2002. [58](#)
- [9] Noga Alon, Ankur Moitra, and Benny Sudakov. Nearly complete graphs decomposable into large induced matchings and their applications. *Journal of the European Mathematical Society*, 15(5):1575–1596, 2013. [7](#)
- [10] Noga Alon and Asaf Shapira. Testing subgraphs in directed graphs. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 700–709, 2003. [58](#)
- [11] 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. [1](#), [2](#), [55](#), [56](#), [57](#)
- [12] Nikhil Bansal and Ryan Williams. Regularity lemmas and combinatorial algorithms. In *2009 50th Annual IEEE Symposium on Foundations of Computer Science*, pages 745–754. IEEE, 2009. [58](#)
- [13] Surender Baswana, Keerti Choudhary, and Liam Roditty. Fault tolerant subgraph for single source reachability: generic and optimal. In *Proceedings of the forty-eighth annual ACM symposium on Theory of Computing*, pages 509–518. ACM, 2016. [1](#), [4](#)
- [14] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (α, β) -spanners. *ACM Transactions on Algorithms (TALG)*, 7(1):5, 2010. [1](#)
- [15] Surender Baswana and Sandeep Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Random Structures & Algorithms*, 30(4):532–563, 2007. [1](#)

- [16] Felix A Behrend. On sets of integers which contain no three terms in arithmetical progression. *Proceedings of the National Academy of Sciences*, 32(12):331–332, 1946. [47](#), [48](#), [58](#)
- [17] Uri Ben-Levy and Merav Parter. New (α, β) spanners and hopsets. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 1695–1714. SIAM, 2020. [1](#)
- [18] Piotr Berman, Arnab Bhattacharyya, Konstantin Makarychev, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Approximation algorithms for spanner problems and directed steiner forest. *Inf. Comput.*, 222:93–107, 2013. [7](#)
- [19] 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. SIAM, 2023. [10](#), [38](#)
- [20] Yitzhak Birk, Nathan Linial, and Roy Meshulam. On the uniform-traffic capacity of single-hop interconnections employing shared directional multichannels. *IEEE Transactions on Information Theory*, 39(1):186–191, 1993. [58](#)
- [21] Greg Bodwin. On the structure of unique shortest paths in graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2071–2089. SIAM, 2019. [1](#)
- [22] Greg Bodwin. New results on linear size distance preservers. *SIAM Journal on Computing*, 50(2):662–673, 2021. [1](#), [2](#), [5](#), [6](#), [8](#), [10](#), [14](#), [58](#), [64](#)
- [23] Greg Bodwin, Michael Dinitz, Merav Parter, and Virginia Vassilevska Williams. Optimal vertex fault tolerant spanners (for fixed stretch). In *Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1884–1900. Society for Industrial and Applied Mathematics, 2018. [1](#), [2](#), [55](#), [57](#)
- [24] Greg Bodwin, Michael Dinitz, and Caleb Robelle. Partially optimal edge fault-tolerant spanners. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3272–3286. SIAM, 2022. [1](#), [2](#), [57](#)
- [25] Greg Bodwin and Shyamal Patel. A trivial yet optimal solution to vertex fault tolerant spanners. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 541–543, 2019. [1](#), [2](#), [55](#)
- [26] Greg Bodwin and Virginia Vassilevska Williams. Better distance preservers and additive spanners. *ACM Transactions on Algorithms (TALG)*, 17(4):1–24, 2021. [6](#)
- [27] Béla Bollobás, Don Coppersmith, and Michael Elkin. Sparse distance preservers and additive spanners. *SIAM Journal on Discrete Mathematics*, 19(4):1029–1055, 2005. [1](#), [6](#)
- [28] John A Bondy and Miklós Simonovits. Cycles of even length in graphs. *Journal of Combinatorial Theory, Series B*, 16(2):97–105, 1974. [29](#)
- [29] Diptarka Chakraborty, Kushagra Chatterjee, and Keerti Choudhary. Pairwise Reachability Oracles and Preservers Under Failures. In Mikołaj Bojańczyk, Emanuela Merelli, and David P. Woodruff, editors, *49th International Colloquium on Automata, Languages, and Programming (ICALP 2022)*, volume 229 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:16, Dagstuhl, Germany, 2022. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. [1](#), [4](#)
- [30] Diptarka Chakraborty and Keerti Choudhary. New Extremal Bounds for Reachability and Strong-Connectivity Preservers Under Failures. In Artur Czumaj, Anuj Dawar, and Emanuela Merelli, editors, *47th International Colloquium on Automata, Languages, and Programming (ICALP 2020)*, volume 168 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:20, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum für Informatik. [1](#), [4](#)

- [31] Hsien-Chih Chang, Pawel Gawrychowski, Shay Mozes, and Oren Weimann. Near-Optimal Distance Emulator for Planar Graphs. In Yossi Azar, Hannah Bast, and Grzegorz Herman, editors, *26th Annual European Symposium on Algorithms (ESA 2018)*, volume 112 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 16:1–16:17, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. 6
- [32] Shiri Chechik. New additive spanners. In *Proceedings of the 24th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 498–512. SIAM, 2013. 1
- [33] Shiri Chechik. Approximate distance oracles with constant query time. In *Proceedings of the forty-sixth annual ACM symposium on Theory of computing*, pages 654–663. ACM, 2014. 1
- [34] Shiri Chechik and Christian Wulff-Nilsen. Near-optimal light spanners. In *Proceedings of the twenty-seventh annual ACM-SIAM symposium on Discrete algorithms*, pages 883–892. Society for Industrial and Applied Mathematics, 2016. 2
- [35] Joseph Cheriyan, Howard Karloff, and Yuval Rabani. Approximating directed multicuts. In *Proceedings 42nd IEEE Symposium on Foundations of Computer Science*, pages 320–328. IEEE, 2001. 7
- [36] Eden Chlamtáč, Michael Dinitz, Guy Kortsarz, and Bundit Laekhanukit. Approximating spanners and directed steiner forest: Upper and lower bounds. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 534–553. SIAM, 2017. 4, 6
- [37] Julia Chuzhoy and Sanjeev Khanna. Polynomial flow-cut gaps and hardness of directed cut problems. *Journal of the ACM (JACM)*, 56(2):1–28, 2009. 5, 7, 39, 40, 41, 42, 45, 61
- [38] Edith Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *Journal of the ACM (JACM)*, 47(1):132–166, 2000. 1
- [39] David Conlon and Jacob Fox. Graph removal lemmas. *Surveys in combinatorics*, 1(2):3, 2013. 58
- [40] David Conlon, Jacob Fox, Benny Sudakov, and Yufei Zhao. The regularity method for graphs with few 4-cycles. *Journal of the London Mathematical Society*, 104(5):2376–2401, 2021. 57
- [41] Don Coppersmith and Michael Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, 20(2):463–501, 2006. 1, 5, 6, 8, 14, 34, 64
- [42] D De Caen and László A Székely. *The maximum size of 4-and 6-cycle free bipartite graphs on m, n vertices.* 1, 2, 57, 58
- [43] D De Caen and László A Székely. On dense bipartite graphs of girth eight and upper bounds for certain configurations in planar point–line systems. *journal of combinatorial theory, Series A*, 77(2):268–278, 1997. 26, 65
- [44] Michael Dinitz and Caleb Robelle. Efficient and simple algorithms for fault-tolerant spanners. In *Proceedings of the 39th Symposium on Principles of Distributed Computing*, pages 493–500, 2020. 2
- [45] Dorit Dor, Shay Halperin, and Uri Zwick. All-pairs almost shortest paths. *Siam Journal on Computing (SICOMP)*, 29(5):1740–1759, 2000. 1
- [46] Michael Elkin. An improved construction of progression-free sets. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete Algorithms*, pages 886–905. Society for Industrial and Applied Mathematics, 2010. 58
- [47] Michael Elkin and Ofer Neiman. Hopsets with constant hopbound, and applications to approximate shortest paths. In *Foundations of Computer Science (FOCS), 2016 IEEE 57th Annual Symposium on*, pages 128–137. IEEE, 2016. 1
- [48] Michael Elkin and Ofer Neiman. Linear-size opsets with small hopbound, and distributed routing with low memory. *arXiv preprint arXiv:1704.08468*, 2017. 1

- [49] Michael Elkin and Ofer Neiman. Near-additive spanners and near-exact hopsets, a unified view. *arXiv preprint arXiv:2001.07477*, 2020. [1](#)
- [50] Michael Elkin, Ofer Neiman, and Shay Solomon. Light spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 442–452. Springer, 2014. [1](#), [2](#), [56](#)
- [51] Michael Elkin, Ofer Neiman, and Christian Wulff-Nilsen. Space-efficient path-reporting approximate distance oracles. *Theoretical Computer Science*, 651:1–10, 2016. [1](#)
- [52] Michael Elkin and David Peleg. $(1+\epsilon, \beta)$ -spanner constructions for general graphs. *SIAM Journal on Computing*, 33(3):608–631, 2004. [1](#)
- [53] Michael Elkin and Seth Pettie. A linear-size logarithmic stretch path-reporting distance oracle for general graphs. *ACM Transactions on Algorithms (TALG)*, 12(4):50, 2016. [1](#)
- [54] Paul Erdős. On some problems of elementary and combinatorial geometry. *Annali di Matematica pura ed applicata*, 103(1):99–108, 1975. [56](#)
- [55] Moran Feldman, Guy Kortsarz, and Zeev Nutov. Improved approximation algorithms for directed steiner forest. *Journal of Computer and System Sciences*, 78(1):279–292, 2012. [7](#)
- [56] Arnold Filtser and Shay Solomon. The greedy spanner is existentially optimal. In *Proceedings of the 2016 ACM Symposium on Principles of Distributed Computing*, pages 9–17, 2016. [2](#)
- [57] Jeremy T Fineman. Nearly work-efficient parallel algorithm for digraph reachability. *SIAM Journal on Computing*, 49(5):STOC18–500, 2019. [1](#)
- [58] Jacob Fox. A new proof of the graph removal lemma. *Annals of Mathematics*, pages 561–579, 2011. [26](#), [58](#), [66](#)
- [59] Naveen Garg, Vijay V Vazirani, and Mihalis Yannakakis. Approximate max-flow min-(multi) cut theorems and their applications. *SIAM Journal on Computing*, 25(2):235–251, 1996. [7](#)
- [60] Ashish Goel, Michael Kapralov, and Sanjeev Khanna. On the communication and streaming complexity of maximum bipartite matching. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 468–485. SIAM, 2012. [58](#)
- [61] Elena Grigorescu, Young-San Lin, and Kent Quanrud. Online Directed Spanners and Steiner Forests. In Mary Wootters and Laura Sanità, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2021)*, volume 207 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:25, Dagstuhl, Germany, 2021. Schloss Dagstuhl – Leibniz-Zentrum für Informatik. [4](#), [6](#), [35](#)
- [62] Anupam Gupta. Improved results for directed multicut. In *SODA*, volume 3, pages 454–455, 2003. [7](#)
- [63] Johan Håstad and Avi Wigderson. Simple analysis of graph tests for linearity and pcp. *Random Structures & Algorithms*, 22(2):139–160, 2003. [58](#)
- [64] William Hesse. Directed graphs requiring large numbers of shortcuts. In *Proceedings of the fourteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 665–669. Society for Industrial and Applied Mathematics, 2003. [1](#)
- [65] Shang-En Huang and Seth Pettie. Lower Bounds on Sparse Spanners, Emulators, and Diameter-reducing shortcuts. In David Eppstein, editor, *16th Scandinavian Symposium and Workshops on Algorithm Theory (SWAT 2018)*, volume 101 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 26:1–26:12, Dagstuhl, Germany, 2018. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1](#), [5](#)
- [66] Shang-En Huang and Seth Pettie. Thorup–zwick emulators are universally optimal hopsets. *Information Processing Letters*, 142:9–13, 2019. [1](#)

- [67] Michael Kapralov. Better bounds for matchings in the streaming model. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 1679–1697. SIAM, 2013. [58](#)
- [68] 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. SIAM, 2021. [1](#)
- [69] Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. In *Scandinavian Workshop on Algorithm Theory*, pages 277–281. Springer, 2014. [1](#)
- [70] Mathias Bæk Tejs Knudsen. Additive Spanners and Distance Oracles in Quadratic Time. In Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, and Anca Muscholl, editors, *44th International Colloquium on Automata, Languages, and Programming (ICALP 2017)*, volume 80 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 64:1–64:12, Dagstuhl, Germany, 2017. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1](#)
- [71] Shimon Kogan and Merav Parter. Low-congestion shortcuts in constant diameter graphs. In *Proceedings of the 2021 ACM Symposium on Principles of Distributed Computing*, pages 203–211, 2021. [1](#)
- [72] Shimon Kogan and Merav Parter. Having hope in hops: New spanners, preservers and lower bounds for hopsets. In *2022 IEEE 63rd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 766–777. IEEE, 2022. [1](#), [5](#), [10](#), [38](#)
- [73] 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. SIAM, 2022. [10](#)
- [74] Christian Konrad. Maximum matching in turnstile streams. In *Algorithms-ESA 2015*, pages 840–852. Springer, 2015. [58](#)
- [75] Yana Kortsarts, Guy Kortsarz, and Zeev Nutov. Greedy approximation algorithms for directed multicuts. *Networks: An International Journal*, 45(4):214–217, 2005. [7](#)
- [76] Adrian Kosowski, Przemyslaw Uznanski, and Laurent Viennot. Hardness of exact distance queries in sparse graphs through hub labeling. In *Proceedings of the 2019 ACM Symposium on Principles of Distributed Computing*, pages 272–279, 2019. [1](#), [2](#), [10](#), [58](#)
- [77] Hung Le and Shay Solomon. Near-optimal spanners for general graphs in (nearly) linear time. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3332–3361. SIAM, 2022. [1](#), [2](#), [56](#)
- [78] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999. [7](#)
- [79] Shi Li and Bundit Laekhanukit. Polynomial integrality gap of flow lp for directed steiner tree. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 3230–3236. SIAM, 2022. [7](#)
- [80] Yang P Liu, Arun Jambulapati, and Aaron Sidford. Parallel reachability in almost linear work and square root depth. In *2019 IEEE 60th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1664–1686. IEEE, 2019. [1](#)
- [81] Gary L Miller, Richard Peng, Adrian Vladu, and Shen Chen Xu. Improved parallel algorithms for spanners and hopsets. In *Proceedings of the 27th ACM symposium on Parallelism in Algorithms and Architectures*, pages 192–201. ACM, 2015. [1](#)
- [82] Guy Moshkovitz and Asaf Shapira. A sparse regular approximation lemma. *Transactions of the American Mathematical Society*, 371(10):6779–6814, 2019. [26](#), [58](#), [66](#)
- [83] János Pach and Gábor Tardos. Forbidden paths and cycles in ordered graphs and matrices. *Israel Journal of Mathematics*, 155(1):359–380, 2006. [11](#)

- [84] David Peleg and Jeffrey Ullman. An optimal synchronizer for the hypercube. *SIAM Journal on Computing (SICOMP)*, 18(4):740—747, 1989. [1](#)
- [85] David Peleg and Eli Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM (JACM)*, 36(3):510–530, 1989. [1](#)
- [86] Seth Pettie. Low distortion spanners. *ACM Transactions on Algorithms (TALG)*, 6(1):7, 2009. [1](#)
- [87] Seth Pettie. Applications of forbidden 0–1 matrices to search tree and path compression-based data structures. In *Proceedings of the twenty-first annual ACM-SIAM symposium on Discrete algorithms*, pages 1457–1467. SIAM, 2010. [11](#)
- [88] Liam Roditty, Mikkel Thorup, and Uri Zwick. Deterministic constructions of approximate distance oracles and spanners. In *International Colloquium on Automata, Languages, and Programming*, pages 261–272. Springer, 2005. [1](#)
- [89] Liam Roditty and Uri Zwick. On dynamic shortest paths problems. *Algorithmica*, 61(2):389–401, 2011. [2](#)
- [90] Imre Z Ruzsa and Endre Szemerédi. Triple systems with no six points carrying three triangles. *Combinatorics (Keszthely, 1976)*, *Coll. Math. Soc. J. Bolyai*, 18:939–945, 1978. [58](#)
- [91] Michael Saks*, Alex Samorodnitsky, and Leonid Zosin. A lower bound on the integrality gap for minimum multicut in directed networks. *Combinatorica*, 24(3):525–530, 2004. [7](#)
- [92] Endre Szemerédi. Regular partitions of graphs. Technical report, Stanford Univ Calif Dept of Computer Science, 1975. [58](#)
- [93] Endre Szemerédi and William T. Trotter. Extremal problems in discrete geometry. *Combinatorica*, 3(3-4):381–392, 1983. [27](#)
- [94] Gábor Tardos. Extremal theory of ordered graphs. In *Proceedings of the International Congress of Mathematicians: Rio de Janeiro 2018*, pages 3235–3243. World Scientific, 2018. [11](#), [28](#)
- [95] Mikkel Thorup. On shortcutting digraphs. In *International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 205–211. Springer, 1992. [1](#)
- [96] Mikkel Thorup and Uri Zwick. Compact routing schemes. In *Proceedings of the 13th Annual ACM Symposium on Parallel Algorithms and Architectures (SPAA)*, pages 1–10. Association of Computing Machinery, 2001. [1](#)
- [97] Mikkel Thorup and Uri Zwick. Approximate distance oracles. *Journal of the ACM (JACM)*, 52(1):1–24, 2005. [1](#), [55](#)
- [98] Mikkel Thorup and Uri Zwick. Spanners and emulators with sublinear distance errors. In *Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 802–809. Society for Industrial and Applied Mathematics, 2006. [1](#)
- [99] Jacques Tits. Sur la trialité et certains groupes qui s’ en déduisent. *Publications Mathématiques de l’Institut des Hautes Études Scientifiques*, 2(1):14–60, 1959. [56](#)
- [100] Jeffrey D Ullman and Mihalis Yannakakis. High-probability parallel transitive-closure algorithms. *SIAM Journal on Computing*, 20(1):100–125, 1991. [1](#), [10](#)
- [101] Manuel Fernández V, David P. Woodruff, and Taisuke Yasuda. Graph Spanners in the Message-Passing Model. In Thomas Vidick, editor, *11th Innovations in Theoretical Computer Science Conference (ITCS 2020)*, volume 151 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 77:1–77:18, Dagstuhl, Germany, 2020. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik. [1](#), [2](#), [57](#)
- [102] Hendrik Van Maldeghem. *Generalized polygons*. Springer Science & Business Media, 2012. [26](#), [57](#), [66](#)

- [103] R Wenger. Extremal graphs with no c_4 's, c_6 's, or c_{10} 's. *J. Combin. Theory Ser. B*, 52(1):113–116, 1991. [56](#)
- [104] David P Woodruff. Lower bounds for additive spanners, emulators, and more. In *Foundations of Computer Science, 2006. FOCS'06. 47th Annual IEEE Symposium on*, pages 389–398. IEEE, 2006. [1](#)
- [105] David P Woodruff. Additive spanners in nearly quadratic time. In *International Colloquium on Automata, Languages, and Programming*, pages 463–474. Springer, 2010. [1](#)
- [106] Christian Wulff-Nilsen. Approximate distance oracles with improved preprocessing time. In *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages 202–208. Society for Industrial and Applied Mathematics, 2012. [1](#)

A A Tour through Prior Work on Girth Problems

A.1 The Girth Problem

We first recall the pioneering work on girth reductions by Althöfer, Das, Dobkin, Joseph, and Soares [\[11\]](#):

Definition 27 (Multiplicative Spanners). *A (multiplicative) k -spanner of a graph G is a subgraph H satisfying $\text{dist}_H(s, t) \leq k \cdot \text{dist}_G(s, t)$ for all nodes s, t . The function $\text{MS}(n, k)$ is the least integer such that every undirected weighted n -node graph has a k -spanner on $\leq \text{MS}(n, k)$ edges.*

Definition 28 (Graph Girth). *The girth of a graph G is the least number of edges in a cycle in G (or ∞ if G is a forest). The function $\gamma(n, k)$ is the maximum possible number of edges in an n -node graph of girth $> k$.*

Theorem 69 ([\[11\]](#)). $\text{MS}(n, k) = \gamma(n, k + 1)$.

Proof Sketch. First we show that $\gamma(n, k + 1) \leq \text{MS}(n, k)$. Let G be an unweighted graph with n nodes, girth $> k + 1$, and $\gamma(n, k + 1)$ edges. If one removes any edge (u, v) from G , then $\text{dist}_G(u, v)$ changes from 1 to $> k$. Thus G is the only k -spanner of itself. So if G is taken as an input to the multiplicative spanner problem, one must keep $\gamma(n, k + 1)$ edges in the spanner, so $\text{MS}(n, k) \geq \gamma(n, k + 1)$.

Next we show that $\text{MS}(n, k) \leq \gamma(n, k + 1)$. Let $G = (V, E, w)$ be an n -node graph for which we want to build a k -spanner. Consider the following greedy algorithm to build a k -spanner. Initially $H = (V, \emptyset)$. Consider the edges of G in nondecreasing order of weight. When each edge (u, v) is considered, we add it to H iff $w(u, v) \leq k \cdot \text{dist}_H(u, v)$, i.e., the edge is currently needed in the spanner. One can show that (1) in the end H is indeed a k -spanner of G , and (2) for any cycle C in G that contains $\leq k + 1$ edges, not all edges in cycle will be added to the spanner H . This is roughly because, when we consider the last edge $(u, v) \in C$, then if all previous edges from C were added to H then there is already a $u \rightsquigarrow v$ path of length $\leq k \cdot w(u, v)$ using these edges. Thus H has girth $> k + 1$, so it has $\leq \gamma(n, k + 1)$ edges. So $\text{MS}(n, k) \leq \gamma(n, k + 1)$. $\square$

The reduction of Althöfer et al. [\[11\]](#) generalizes also to *emulators* and more generally to *distance oracles*, which are arbitrary data structures that can approximate the distances of the input graph on query (see also [\[97\]](#)). Recently, tight reductions to γ have been achieved for vertex fault tolerant spanners as well, which ask for the size bounds for k -spanners that retain their distance approximation even after a bounded number of vertices fail in both the spanner and the original graph [\[23, 25\]](#).

The “girth problem” asks for the asymptotic value of γ , which hence would also determine the asymptotic value of MS . This is a major open question in extremal combinatorics and theoretical computer science. The following upper bound is known:

Theorem 70 (Moore Bounds, Folklore). *For any integers n, k , we have $\gamma(n, 2k) = O(n^{1+1/k})$.*

Proof Sketch. Let G be an n -node graph of average degree d , and assume that d is at least a sufficiently large constant. A *non-backtracking k -path* is a path in G , containing exactly $k + 1$ nodes and k edges, which may repeat nodes or edges but which never uses an edge (u, v) followed consecutively by its reverse (v, u) . The following are facts from graph theory:

- G has $n \cdot \Omega(d)^k$ non-backtracking k -paths, and
- If G has two different non-backtracking k -paths with the same pair of endpoints (s, t) , then G has a cycle on $\leq 2k$.

Together, these imply that if G has girth $> 2k$, then it can only have $O(n^2)$ non-backtracking k -paths, and hence $n \cdot \Omega(d)^k = O(n^2)$. Rearranging gives $d = O(n^{1/k})$, proving the theorem. $\square$

Unfortunately, lower bounds are not as well understood. The Moore bounds are known to be asymptotically tight when $k \in \{1, 2, 3, 5\}$ [99, 103]. The *girth conjecture*, attributed to Erdős [54], posits that the Moore bounds are tight for all other values of k as well. The girth conjecture is controversial, with no clear consensus from experts on whether it is likely to be true.

A.2 The Weighted Girth Problem

Besides number of edges, in some applications one wants to minimize the *total weight* of a spanner. This is often measured as the *lightness* of the spanner, relative to the input graph:

Definition 29 (Spanner Lightness). *The lightness of a subgraph H of a graph G is*

$$\ell(H \mid G) := \frac{w(H)}{w(\text{MST}(G))}$$

where $\text{MST}(G)$ is any minimum spanning tree of G (or spanning forest if disconnected). We write $\text{LMS}(n, k)$ for the least¹⁸ L such that every n -node graph G has a k -spanner H of lightness $\ell(H \mid G) \leq L$.

In their study of light spanners, Elkin, Neiman, and Solomon [50] made the interesting point that an extension of the Althöfer et al [11] reduction between γ and MS also gives equivalence between the extremal function of graph lightness and *weighted girth*, defined as follows:

Definition 30 (Weighted Girth). *The weighted girth of a graph G is defined as*

$$\min_C \frac{w(C)}{\max_{e \in C} w(e)}$$

where the $\min$ is over the set of cycles C in G . We define $\lambda(n, k)$ as the maximum¹⁹ lightness over n -node graphs of weighted girth $> k$.

Note that weighted girth generalizes girth, in the sense that the concepts coincide for an unweighted graph. Elkin et al. [50] proved:

Theorem 71 ([50]). $\text{LMS}(n, k + 1) = \lambda(n, k)$.

A natural next question is to ask for the relative values of λ and γ . It follows by considering the unweighted graph realizing γ that

$$\lambda(n, k) = \Omega\left(\frac{\gamma(n, k)}{n}\right)$$

(note: we divide by n on the right, since an unweighted graph has an MST of weight $n - 1$). A fascinating conjecture by Elkin et al. [50], known as the *weighted girth conjecture*, implies that these bounds are asymptotically equal. This remains open, but recent work of Le and Solomon [77] implies that they are *approximately* equal.

¹⁸Formally, one takes the inf of the values L satisfying this condition

¹⁹Formally, λ is determined by the sup of the lightness of graphs satisfying this property.

A.3 The Bipartite Girth Problem

The function γ has a natural generalization to the setting of bipartite graphs:

Definition 31 (The Extremal Function of Bipartite High-Girth Graphs). *The function $\gamma\gamma(n, p, k)$ is the maximum possible number of edges in a bipartite graph with n nodes on one side of the bipartition, p nodes on the other side, and girth $> k$.*

We say that $\gamma\gamma$ *generalizes* γ , rather than merely being different, due to the following fact:

Theorem 72 (Folklore). $\gamma\gamma(n, n, k) = \Theta(\gamma(n, k))$.

Proof Sketch. In one direction, we have

$$\gamma\gamma(n, n, k) \leq \gamma(2n, k) = \Theta(\gamma(n, k))$$

where the first inequality is immediate from the definitions, and the second inequality is by observing that γ depends at most polynomially on its first parameter. In the other direction, we show

$$\gamma(n, k) \leq 2 \cdot \gamma\gamma\left(\frac{n}{2}, \frac{n}{2}, k\right) \leq O(\gamma\gamma(n, n, k)).$$

The second inequality is immediate from the definitions. For the first inequality, we start with a graph G realizing $\gamma(n, k)$, and randomly bipartition its nodes into two parts of size $n/2$ each. Let G' be the bipartite subgraph that keeps only edges crossing the random bipartition. Each edge survives in G' with probability $\geq 1/2$. Thus we have constructed a bipartite graph with $n/2, n/2$ nodes per side, girth $> k$, and $\geq \gamma(n, k)/2$ edges in expectation, which implies the first inequality. $\square$

Thus every extremal reduction to γ can also be expressed as a reduction to a special case of $\gamma\gamma$. However, there are some further problems in distance sketching and extremal combinatorics that can *only* be reduced to $\gamma\gamma$, rather than the non-bipartite version. First, the obvious bipartite generalizations of multiplicative spanners and related objects can be reduced to $\gamma\gamma$, again by the reduction of Althöfer et al. [11]. More interestingly:

- Fernández, Yasuda, and Woodruff [101] constructed lower bounds against the communication complexity of spanner construction, converting $\gamma\gamma$ lower bounds to lower bound instances, and
- Bodwin, Dinitz, and Robelle [24] used $\gamma\gamma$ to provide lower bounds against edge fault tolerant spanners and edge distance sensitivity oracles (based on a construction from [23]).

The following natural extension of the Moore bounds holds for bipartite graphs:

Theorem 73 (Bipartite Moore Bounds). *For all n, p, k , we have*

$$\gamma\gamma(n, p, 2k) = \begin{cases} O((np)^{1/2+1/(2k)} + n + p) & \text{if } k \text{ is odd} \\ O(n^{1/2+1/k}p^{1/2} + n + p) & \text{if } k \text{ is even} \end{cases}$$

The proof is in the same spirit as Theorem 70, but with sensitivity to the average degree on either side of the bipartite graph. The bipartite Moore bounds are known to be fully tight for girth parameters $2k$ when $k \in \{1, 2\}$, and they are also tight for various relative values of n, p when $k \in \{3, 5, 7\}$ [102]. Analogizing the girth conjecture, one might conjecture that the Moore bounds are tight for all n, p, k . However, this was *refuted* in an important paper by de Caen and Székely [42], which showed that the Ruzsa-Szemerédi theorem (discussed next) is equivalent to a (subpolynomial) improvement to the upper bounds on $\gamma\gamma(n, p, 6)$, and thus it implies an improvement on the bipartite Moore bounds in a particular parameter setting. At a technical level, this proof is very similar to Theorem 81, so we shall not repeat it here.

Recent work of Conlon, Fox, Sudakov, and Zhao [40] implies an analogous improvement to $\gamma\gamma(n, p, 10)$; it is an interesting open problem to obtain an analogous improvement to $\gamma\gamma(n, p, 2k)$ for any other odd k .

A.4 The Ruzsa-Szemerédi Problem

The Ruzsa-Szemerédi problem was introduced by Ruzsa and Szemerédi [90], in the context of a combinatorial problem about hypergraphs. Their result was one of the first major uses of the famous Szemerédi regularity lemma [92]. Although it has been interpreted and reinterpreted over the years, the standard phrasing is as follows:

Definition 32 (Induced Matchings and $\mathbf{rs}(n)$). *In a graph $G = (V, E)$, an induced matching is an edge subset $M \subseteq E$ that is a matching, and also the edge subset of an induced subgraph. In other words, for any two edges $(u_1, v_1), (u_2, v_2) \in M$, we have $(u_1, u_2), (u_1, v_2), (v_1, u_2), (v_1, v_2) \notin E$.*

We define $\mathbf{rs}(n)$ as the largest integer such that, for every n -node graph G whose edge set can be partitioned into n induced matchings, we have

$$|E(G)| \leq \frac{n^2}{\mathbf{rs}(n)}.$$

Besides induced matchings, there are many other natural ways to interpret $\mathbf{rs}(n)$ [39]. The following transformation can be used to connect $\mathbf{rs}(n)$. One is: let G be a graph that can be decomposed into n induced matchings, which has $n^2/\mathbf{rs}(n)$ edges. Direct the edges of G arbitrarily, and then add a new node $m_1, \dots, m_n$ for each of the n induced matchings. Then, for each i and for each directed edge (u, v) in the i^{th} induced matching, interpret the triple (m_i, u, v) as a 3-path. One can verify that this yields a bridge-free path system. This transformation is well known, even though the description as a “bridge-free path system” is new.

It is not at all obvious from the definition of $\mathbf{rs}(n)$ that the function is nontrivial, i.e., super-constant. But indeed, Ruzsa and Szemerédi proved that $\mathbf{rs}(n) = \Omega(\log^* n)$. The state-of-the-art upper bound is due to Fox [58]; a notable alternate proof was discovered by Moshkovitz and Shapira [82]. The state-of-the-art lower bound is due to Behrend [16] (see also [46]). These bounds are:

$$2^{\Omega(\log^* n)} \leq \mathbf{rs}(n) \leq 2^{O(\log^c n)}.$$

While it is not clear from the definition that the Ruzsa-Szemerédi problem should be regarded as a girth concept, an important paper by de Caen and Székely [42] explains its inclusion, by tightly reducing between $\mathbf{rs}(n)$ and $\gamma\gamma$. Specifically, their reduction may be interpreted as follows. Given a value of n , let p^* be the largest integer such that $\gamma\gamma(n, p^*, 6) \geq 3p^*$. Then:

$$\gamma\gamma(n, p^*, 6) = \Theta\left(\frac{n^2}{\mathbf{rs}(n)}\right).$$

Thus, the Ruzsa-Szemerédi problem is a special case of the bipartite girth problem. In network design, we mention three applications of the Ruzsa-Szemerédi problem:

- Given an n -node undirected unweighted graph and a set of p demand pairs, one can construct a distance preserver on $O(n^2/\mathbf{rs}(n) + p)$ edges [22].
- For undirected unweighted input graphs with n nodes and $O(n)$ edges, one can construct a distance labeling scheme with average label size $O(n/\mathbf{rs}(n)^c)$ [76].
- Bansal and Williams [12] developed a combinatorial algorithm for All-Pairs Shortest Paths in unweighted graphs, by reducing to a certain *algorithmic* version of the Ruzsa-Szemerédi problem.

Some other miscellaneous uses of $\mathbf{rs}(n)$ or Ruzsa-Szemerédi graphs in theoretical computer science include connections to the PCP theorem by Håstad and Wigderson [63], applications in Channel Scheduling by Birk, Linial, and Meshulam [20], a line of work on maximum matching in streams [60, 67, 74], and a line of work on subgraph testing algorithms [8, 10].

A.5 The Set Girth Problem

The function $\gamma\gamma(n, p, k)$ has an equivalent interpretation in the language of *set systems* rather than graphs. We consider:

Definition 33 (Set Systems).

- A set system is a pair $S = (V, \mathcal{T})$, where V is a ground set of “nodes” and $\mathcal{T}$ is a multiset of node subsets.
- A k -cycle in a set system S is a circularly-ordered list of distinct nodes $v_0, v_1, \dots, v_k = v_0 \in V$ and sets $T_0, T_1, \dots, T_k = T_0 \in \mathcal{T}$ with each $v_i, v_{i+1} \in T_i$.
- The girth of a set system is the smallest integer k for which the system has a k -cycle.
- The size of a set system is written $\|S\| := \sum_{T \in \mathcal{T}} |T|$.

For example, a set system in which each set has size 2 can be considered as an undirected graph. Set systems are merely a rephrasing of bipartite graphs, and one can switch between them via *incidence graphs*. In particular:

Theorem 74 (Folklore). *Over set systems S with n nodes, p sets, and girth $> k$, the maximum possible value of $\|S\|$ is exactly $\gamma\gamma(n, p, 2k)$.*

Proof Sketch. A set system $S = (V, \mathcal{T})$ can be naturally bijected with its *incidence graph* G_S . This is a bipartite graph whose nodes on the left correspond to V , whose nodes on the right correspond to $\mathcal{T}$, and whose edges correspond to set membership; that is, we put an edge between $v \in V$ and $T \in \mathcal{T}$ iff $v \in T$. Set systems carry the same information as their incidence graph. One can verify that (1) the size $\|S\|$ of the set system is the same as the number of edges $|E(G_S)|$ of its incidence graph, and (2) if the set system has girth k , then its incidence graph has girth $2k$. The theorem follows from these properties. $\square$

B Proof of Cleaning Lemma

We now prove the Cleaning Lemma (Lemma 10). We will state the proof only for unordered bridge girth; the proof for ordered bridge girth is completely identical. We split the proof into the following two lemmas:

Lemma 75. *Suppose S is a path system with n nodes, p paths, bridge girth b , average node degree d , and average path length ℓ . Then there exists a path system S' that has:*

- $n' = \Theta(n)$ nodes,
- $p' = \Theta(p)$ paths,
- size $\|S'\| = \Theta(\|S\|)$,
- bridge girth $\geq b$,
- average degree $d' = \Theta(d)$, and all nodes have degree $\Theta(d')$,
- average length $\ell' = \Theta(\ell)$, and all paths have length $\Theta(\ell')$.

Proof. We construct S' by the following process. Start with $S = (V, \Pi)$ as a path system with n nodes, p paths, bridge girth $> k$, and size $\|S\| = \beta(n, p, k)$. Fix ℓ, d as the initial average length and degree of S . Then, perform the following sequence of operations on S :

1. While there exists a path $\pi \in \Pi$ of length $|\pi| \geq \ell/2$, split π into two node-disjoint paths π_1 and π_2 in any way such that $\pi = \pi_1 \circ \pi_2$, $|\pi_1| \geq \ell/4$, and $|\pi_2| \geq \ell/4$.
2. While there exists a node $v \in V$ of degree $\geq d/2$, split v into two new nodes v_1, v_2 . Replace each occurrence of v in a path with either v_1 or v_2 , in any way such that $\deg(v_1) \geq d/4$ and $\deg(v_2) \geq d/4$.
3. While there exists a node of degree $< d/4$, or a path of length $< \ell/4$, delete that node or path from S .

Let S' be the resulting path system on n' nodes and p' paths. First note that the construction must terminate, since no step of the construction increases the size of $\|S'\|$. Our only operations are to delete nodes/paths and to split nodes/paths, which do not create bridges; thus, since S does not have a bridge of size $< b$, S' also has no bridge of size $< b$. Operations that split nodes and paths do not change the size of S . It is immediate from the construction that all surviving paths π have length $\Theta(\ell)$ and that all surviving nodes have degree $\Theta(d)$. Meanwhile, we only delete nodes of degree $< d/4$ and paths of length $< \ell/4$, so we have

$$\|S'\| > \|S\| - (nd/4 + p\ell/4) = \|S\|/2.$$

Thus we have $\|S'\| = \Theta(\|S\|)$. Moreover, we notice that

$$\ell'p' = n'd' = \|S'\| > \|S\|/2 = nd/2 = p\ell/2.$$

Since $\ell' = \Theta(\ell)$ and $d' = \Theta(d)$, this implies that $p' = \Theta(p)$ and $n' = \Theta(n)$, completing the proof. $\square$

Lemma 76. *For any absolute constant $0 < c < 1$, we have $\beta(cn, cp, k) = \Omega(\beta(n, p, k))$.*

Proof. We will prove for β ; the proof for β^* is identical. Let $S = (V, \Pi)$ be path system with n nodes, p paths, bridge girth $> k$, and $\|S\| = \beta(n, p, k)$. Let $S' \subseteq S$ be a subsystem obtained by choosing exactly cn nodes in V uniformly at random and cp paths in Π uniformly at random, and keeping these nodes and the paths induced on these nodes in S' , while deleting the rest of S . Then S' has cn nodes, cp paths, bridge girth $> k$ and expected size

$$\mathbb{E}[\|S'\|] = c^2\|S\| = \Omega(\beta(n, p, k)),$$

which completes the proof. $\square$

We can now state the proof of the cleaning lemma. Using the latter lemma, we can choose $c > 0$ as a sufficiently small constant, and then start with S as a path system with $\lceil cn \rceil$ nodes, $\lceil cp \rceil$ paths, bridge girth $> k$, and $\|S\| = \Omega(\beta(n, p, k))$. Then, applying the former lemma, we can find a path system S' that has $n' = \Theta(cn)$ nodes, $p' = \Theta(cp)$ paths, bridge girth $> k$, size $\|S'\| = \Omega(\beta(n, p, k))$, all nodes have degree $\Theta(d')$, and all paths have length $\Theta(\ell')$. By choice of sufficiently small c , we have $n' \leq n$ and $p' \leq p$, and thus S' satisfies the cleaning lemma.

C Missing Proofs for Section 5

C.1 Lemma 57

Fix an $\varepsilon > 0$. Let $\mathcal{S} \subseteq V_H$ be a set of vertices of size $|\mathcal{S}| \leq d/16$ in H . Fix an $i \in [1, d]$, and observe that when we are choosing the vertices in layer L_i^j , the size of the set $V_H \setminus (L_i^1 \cup L_i^2 \cup \dots \cup L_i^{j-1})$ is at least $d/2$. Therefore, the probability that $L_i^j \subseteq \mathcal{S}$ is at most

$$\left(\frac{|\mathcal{S}|}{d/2}\right)^{\log d} \leq \left(\frac{1}{8}\right)^{\log d} = d^{-3}$$

Note that $\mathcal{S}$ separates (s_i, t_i) if and only if $L_i^j \subseteq \mathcal{S}$ for some $j \in [1, d']$. Then the probability that $\mathcal{S}$ separates (s_i, t_i) is at most $d' \cdot d^{-3} \leq d^{-2}$, by the union bound and the fact that $d' \leq d$. Now since our constructions of each graph H_i are independent, the probability that at least εd distinct demand pairs (s_i, t_i) , $i \in [1, d]$ are disconnected by $\mathcal{S}$ is at most

$$\binom{d}{\varepsilon d} (d^{-2})^{\varepsilon d} = \left(\frac{1}{d}\right)^{\varepsilon d} = 2^{-\varepsilon d \log d} < 2^{-d}$$

for sufficiently large d . We have established our desired claim.

C.2 Theorem 64

Recall that our goal is to lower bound the flow-cut gap between concurrent multicommodity flow and (non-bipartite) sparsest cut. We will accomplish this by lower bounding an integrality gap for the directed sparsest *vertex* multicut problem, $\widehat{\text{SVCUT}}$. This problem is defined identically to sparsest cut except we choose a set of non-terminal *vertices* $\mathcal{S}$ that minimizes the ratio $|\mathcal{S}|/|P_{\mathcal{S}}|$, where $P_{\mathcal{S}}$ is the set of demand pairs disconnected in $G \setminus \mathcal{S}$. Below we describe a natural LP relaxation $\widehat{\text{SVCUT}}$ of the directed sparsest *vertex* cut problem.

$$\begin{array}{c} \widehat{\text{SVCUT}} \\ \min \quad \sum_{v \in V} x'_v \\ \text{s.t.} \quad \sum_{v \in \pi \cap V} x'_v \geq h'_i \quad \forall i \in [p], \pi \in \Pi_i \\ \sum_{i=1}^p h'_i \geq 1 \\ x'_v, h'_i \geq 0 \quad \forall v \in V, \forall i \in [p] \end{array}$$

By the discussion in Section 2.2 of [37], the integrality gap between $\widehat{\text{SVCUT}}$ and SVCUT is at most the integrality gap between $\widehat{\text{SCUT}}$ and SCUT . By lower bounding the integrality gap of $\widehat{\text{SVCUT}}$, we will immediately obtain lower bounds for the directed sparsest cut flow-cut gap.

We will lower bound the integrality gap for $\widehat{\text{SVCUT}}$ using our construction from the proof of Theorem 52 and a standard argument from [37]. Let $G = (V \cup V', E)$ be the graph on $N := \beta(n, n, \infty)$ non-terminal nodes defined in Section 5.3, and let $P \subseteq V' \times V'$ be the corresponding set of demand pairs of size $|P| = n$. Observe the following solution to the sparsest vertex cut LP. For every $(s_i, t_i) \in P$, let $h'_i := 1/n$. For every non-terminal vertex $v \in V$, let $x'_v := 1/(nd')$. This is a feasible solution to $\widehat{\text{SVCUT}}$ of size $N/(nd')$ by Claim 58.

Now assume that $\widehat{\text{SVCUT}}$ has integrality gap less than $g(N)$ for some function g , and fix a sufficiently small $\varepsilon > 0$. We will show that there is an (integral) vertex cut $\mathcal{S}$ of G of size $|\mathcal{S}| = O(N/d')g(N)$ that disconnects more than a $(1 - \varepsilon)$ -fraction of the demand pairs P in G . Recall that by Theorem 52, there is an $\varepsilon > 0$ such that $|\mathcal{S}| = \Omega(N)$ if $\mathcal{S}$ disconnects more than a $(1 - \varepsilon)$ -fraction of demand pairs P in G . Then we will conclude that $g(N) = \Omega(d')$.

Fix a sufficiently small $\varepsilon > 0$. Our construction of $\mathcal{S}$ will proceed in rounds, where in each round we will add nodes in G to $\mathcal{S}$ and disconnect some demand pairs in $G \setminus \mathcal{S}$. We will repeat our procedure until $\mathcal{S}$ disconnects more than a $(1 - \varepsilon)$ -fraction of the demand pairs P . Let $G_1 := G$ and $P_1 := P$. In round 1, G_1, P_1 has a feasible solution to $\widehat{\text{SVCUT}}$ of size $\varphi := N/(nd')$, so there is an integral solution to SVCUT of size $\varphi \cdot g(N)$. This means there is a set $\mathcal{S}_1 \subseteq V$ in G of size $|\mathcal{S}_1| = \varphi \cdot g(N) \cdot p_1$ that disconnects $p_1 \geq 1$ pairs $P_{\mathcal{S}_1} \subseteq P_1$. Add $\mathcal{S}_1$ to $\mathcal{S}$, and let $G_2 := G \setminus \mathcal{S}_1$ and $P_2 := P_1 \setminus P_{\mathcal{S}_1}$.

In round i , we are given a graph $G_i \subseteq G$ and a set of demand pairs $P_i \subseteq P$ that are connected in G_i . We halt when $|P_i| < \varepsilon|P|$, so we may assume that $|P_i| \geq \varepsilon|P|$. Then if we let $h'_j := \varepsilon^{-1}n^{-1}$ for every $(s_j, t_j) \in P_i$ and let $x'_v := (\varepsilon nd')^{-1}$, then this is a feasible solution to $\widehat{\text{SVCUT}}$ for G_i, P_i of size at most $N/(\varepsilon nd')$ by Claim 58. Then there is an integral vertex cut $\mathcal{S}_i$ of size at most $\varepsilon^{-1}\varphi \cdot g(N) \cdot p_i$ that disconnects $p_i \geq 1$ pairs $P_{\mathcal{S}_i} \subseteq P_i$. Add $\mathcal{S}_i$ to $\mathcal{S}$, and let $G_{i+1} := G_i \setminus \mathcal{S}_i$ and $P_{i+1} := P_i \setminus P_{\mathcal{S}_i}$. When our procedure ends in round k , we will have a set $\mathcal{S}$ of size at most $\varepsilon^{-1}\varphi \cdot g(N) \cdot n = O(N/d')g(N)$ that disconnects at least $|P| - |P_k| > (1 - \varepsilon)|P|$ pairs in P . Then $|\mathcal{S}| = \Omega(N)$ by Lemma 63 and the subsequent discussion, so we conclude that $g(N) = \Omega(d')$, as desired. The theorem is immediate from the above discussion.

C.3 Claim 66

Fix an n -node unweighted directed graph $G = (V \cup V', E)$ with non-terminal nodes V and terminal nodes V' , and a set of demand pairs $P \subseteq V' \times V'$ of size p . Inputs G, P correspond to an instance of the Vertex Directed Steiner Forest problem. Given G, P , we will construct an instance G_1, P_1 of the Directed Steiner Forest problem. We will then show that the integrality gap of $\widehat{\text{VDSF}}$ on G, P is at most the integrality gap of $\widehat{\text{DSF}}$ on G_1, P_1 , which will imply Claim 66.

We construct the weighted directed graph $G_1 = (V_1, E_1, w)$ as follows. Initialize G_1 as $G_1 := (V \cup V', E)$. For each vertex $v \in V$, we replace v in G_1 with a special directed edge (v^+, v^-) , which

we assign weight $w_{(v^+, v^-)} = 1$. For each original edge $(u, v) \in E \cap (V \times V)$, we replace it with edge (u^-, v^+) , which we assign weight $w_{(u^-, v^+)} = 0$. For each original edge $(u, v) \in E \cap (V \times V')$, we replace it with (u^-, v) and assign it weight $w_{(u^-, v)} = 0$; and for each original edge $(u, v) \in E \cap (V' \times V)$, we replace it with (u, v^+) and assign it weight $w_{(u, v^+)} = 0$. This concludes our construction of G_1 . We let $P_1 := P$. Then G_1, P_1 will be our corresponding instance of Directed Steiner Forest.

Let $\{x_v\}_{v \in V}$ be a feasible solution to $\widehat{\text{VDSF}}$ on G, P . Now for every $v \in V$, assign capacity x_v to edge (v^+, v^-) in G_1 . For all other edges in G_1 , assign capacity 1. Observe that the resulting solution to $\widehat{\text{DSF}}$ on inputs G_1, P_1 is feasible and has capacity

$$\sum_{e \in E(G_1)} w_e x_e = \sum_{v \in V} x_v \leq \widehat{\text{VDSF}}(G, P).$$

Then the size of the optimal solution to $\widehat{\text{DSF}}(G_1, P_1)$ is at most

$$\widehat{\text{DSF}}(G_1, P_1) \leq \widehat{\text{VDSF}}(G, P).$$

Now consider a feasible solution to Directed Steiner Forest on inputs G_1, P_1 . This optimal solution corresponds to a subgraph H_1 of G_1 . Now we define a corresponding feasible solution to Vertex Directed Steiner Forest as follows. Let H be the induced subgraph

$$H = G[\{v \in V \mid (v^+, v^-) \in E(H_1)\} \cup V'].$$

Observe that H is a feasible solution for Vertex Directed Steiner Forest on inputs G, P and has size at most

$$|V(H) \cap V| = |\{v \in V \mid (v^+, v^-) \in E(H_1)\}| = \sum_{e \in E(H_1)} w_e \leq \text{DSF}(G_1, P_1).$$

Then the size of the optimal solution to Vertex Directed Steiner forest on G, P is at most

$$\text{VDSF}(G, P) \leq \text{DSF}(G_1, P_1).$$

Now suppose that for some integer k , inputs G, P satisfy

$$k \leq \frac{\text{VDSF}(G, P)}{\widehat{\text{VDSF}}(G, P)}.$$

Then

$$k \leq \frac{\text{VDSF}(G, P)}{\widehat{\text{VDSF}}(G, P)} \leq \frac{\text{DSF}(G_1, P_1)}{\widehat{\text{DSF}}(G_1, P_1)} = \text{DSFG}(|V(G_1)|, p) \leq \text{DSFG}(2n, p).$$

C.4 Lemma 67

We now prove the source-restricted cleaning lemma (Lemma 67). The proof will require the following lemma:

Lemma 77. *Suppose S is a path system with n nodes, p paths, bridge girth b , average node degree d , and average path length ℓ . Then there exists a path system S' that has:*

- $n' = \Theta(n)$ nodes,
- $p' = \Theta(p)$ paths,
- $\|S'\| = \Theta(\|S\|)$,
- $\text{bridge girth} \geq b$,
- $\text{average degree } d' = \Theta(d)$, and all nodes have degree $\Theta(d')$,
- $\text{average length } \ell' = \Theta(\ell)$, and all paths have length $\Theta(\ell')$,
- S' is source-restricted with respect to a set X of size $|X| = \Theta(p/d)$.

Proof. By the cleaning lemma (Lemma 10), we may assume that $S = (V, \Pi)$ is approximately degree-regular and length-regular.²⁰ In particular, we may assume that $c_1\ell \leq |\pi| \leq c_2\ell$ for $\pi \in \Pi$, and $\deg(v) \leq c_3d$ for $v \in V$, where c_1, c_2, c_3 are positive universal constants. Additionally, we will assume that ℓ is greater than a sufficiently large constant; when $\ell \leq c$ for a constant $c > 0$, we can simply take $X := V$ and make our path system source-restricted with respect to X by shortening all paths in Π until they contain only one node. Now perform the following sequence of operations on S :

1. Uniformly at random, sample a subset $X \subseteq V$ of size $|X| = \frac{p}{d}$. Delete all paths from S that do not contain a node in X .
2. For each path $\pi \in \Pi$, let $x_\pi \in \pi \cap X$ be the first node in π that is also in X . Delete all nodes from π preceding x_π ; that is, all nodes y such that $y <_\pi x_\pi$.
3. For each path $\pi \in \Pi$, delete all occurrences of nodes in $X \setminus \{x_\pi\}$ from π .
4. While there exists a node v of degree $\deg(v) < \lambda d$ or a path π of length $|\pi| < \lambda\ell$, where $\lambda > 0$ is a sufficiently small constant, delete that node or path from S . If a node $v \in X$ is deleted, then delete all paths in Π that contain v .

Let $S' := (V', \Pi')$ be the resulting path system. We will now prove that with nonzero probability, $\|S'\| = \Theta(\|S\|)$. All other properties of S' claimed in the lemma are immediate or follow from arguments identical to those of Lemma 75.

Fix a path π in Π , and let π_1 be the prefix of π corresponding to the first $c_1/2 \cdot \ell$ nodes. Observe that π_1 contains a node in X with probability

$$\Pr[\pi_1 \cap X \neq \emptyset] = 1 - \left(1 - \frac{p/d}{n}\right)^{|\pi_1|} \geq 1 - e^{-\frac{c_1/2 \cdot \ell p}{dn}} = 1 - e^{-c_1/2}.$$

Now suppose that $\pi_1 \cap X \neq \emptyset$, so that $x_\pi \in \pi_1$. Then π survives step 1 of our procedure. Let π_2 be a subpath of π such that $\pi = \pi_1 \circ \pi_2$. Since $x_\pi \in \pi_1$, we are guaranteed that π_2 is a subpath of a path surviving after the second step of our procedure; moreover, $|\pi_2| \geq c_1/2 \cdot \ell$, since $|\pi| \geq c_1\ell$. Let S_2 be the path system obtained after performing the first two steps of our procedure on S . Then the expected size of $\|S_2\|$ is

$$\mathbb{E}[\|S_2\|] \geq \sum_{\pi \in \Pi} \Pr[\pi_1 \cap X \neq \emptyset] \cdot |\pi_2| \geq (1 - e^{-c_1/2})p \cdot c_1/2 \cdot \ell = c_1/2 \cdot (1 - e^{-c_1/2})\|S\|.$$

Consequently, we may assume that S_2 satisfies $\|S_2\| \geq c_1/2 \cdot (1 - e^{-c_1/2})\|S\|$. Now we just need to bound the amount that $\|S_2\|$ decreases in steps 3 and 4 of our procedure. Note that the total decrease of $\|S_2\|$ in step 3 is at most

$$\sum_{x \in X} \deg(x) \leq |X| \cdot c_3d = c_3p.$$

As stated earlier, we may assume ℓ is greater than a sufficiently large constant. If we assume $\ell > \frac{4c_3}{c_1/2 \cdot (1 - e^{-c_1/2})}$, then the total decrease of $\|S_2\|$ in step 3 is at most

$$c_3p \leq \left(\frac{c_1/2 \cdot (1 - e^{-c_1/2})}{4} \cdot \ell\right) \cdot p \leq \frac{c_1/2 \cdot (1 - e^{-c_1/2})}{4} \cdot \|S\| \leq \|S_2\|/4.$$

In step 4, if we delete a node $v \in V \setminus X$ of degree $\deg(v) < \lambda d$, then the decrease in $\|S_2\|$ is at most λd . Else if we delete a path π of length $|\pi| < \lambda\ell$, then the decrease in $\|S_2\|$ is at most $\lambda\ell$. Finally, if we delete a node $x \in X$, then the decrease in $\|S_2\|$ is at most $\lambda d \cdot c_2\ell$, since we delete at most λd paths each of length at most $c_2\ell$. Then the total decrease in $\|S_2\|$ in step 4 is at most

$$|V| \cdot \lambda d + |\Pi| \cdot \lambda\ell + |X| \cdot \lambda d \cdot c_2\ell \leq \lambda(2 + c_2)\|S\|.$$

If we choose our constant λ to be $\lambda := \frac{c_1/2 \cdot (1 - e^{-c_1/2})}{4(2 + c_2)}$, then this decrease is at most $\|S_2\|/4$. We conclude that $\|S'\| \geq \|S_2\|/2 = \Theta(\|S\|)$, as desired. $\square$

Now the proof of Lemma 67 follows by using Lemma 77 and Lemma 76 in an argument identical to the proof of the cleaning lemma (Lemma 10) in Appendix B.

²⁰This initial application of the cleaning lemma is not technically needed to make the following analysis work, but it simplifies the analysis a bit.

D Implicit Bounds on β in Prior Work

D.1 Upper Bounds for $k = 2$

Here, we repeat some arguments from prior work that implicitly show upper bounds on the value of β , translated into language directly about β . We note that $\beta(n, p, 2) = \beta^*(n, p, 2)$, since 2-bridges are not sensitive to ordering, and so we typically prove only the upper bounds on $\beta(n, p, 2)$.

Theorem 78 ([41]). $\beta(n, p, 2) = \beta^*(n, p, 2) = O(np^{1/2} + p)$.

Proof. Let $S = (V, \Pi)$ be a path system with n nodes, p paths, and bridge girth > 2 . By the Cleaning Lemma (Lemma 10), we may assume without loss of generality that all paths have length $\Theta(\ell)$, where ℓ is the average length in S . We may also assume that ℓ is at least a sufficiently large constant, as otherwise the bound $O(p)$ is immediate.

There are $O(n^2)$ ordered pairs of distinct nodes in S . Since S does not have 2-bridges, for each such ordered pair (x, y) , there is at most one path $\pi \in \Pi$ with $x <_\pi y$. On the other hand, each path π contains $\binom{|\pi|}{2} = \Theta(\ell^2)$ such node pairs (note: this equality uses that ℓ is a large constant, and so $|\pi| \geq 2$). We therefore have:

$$\begin{aligned} p\ell^2 &= O(n^2) \\ \|S\|^2 = p^2\ell^2 &= O(n^2p) \\ \|S\| &= O(np^{1/2}). \end{aligned} \quad \square$$

Theorem 79 ([22]). $\beta(n, p, 2) = \beta^*(n, p, 2) = O(n^{2/3}p + n)$.

Proof. Let $S = (V, \Pi)$ be a path system with n nodes, p paths, and bridge girth > 2 . By the Cleaning Lemma (Lemma 10), we may assume without loss of generality that all nodes have degree $\Theta(d)$, where d is the average degree in S . We may also assume that d is at least a sufficiently large constant, as otherwise the bound $O(n)$ is immediate.

We first claim that, for any triple of distinct paths $\pi_1, \pi_2, \pi_3 \in \Pi$, there exists at most one node v in $\pi_1 \cap \pi_2 \cap \pi_3$. To see this, suppose for contradiction that there are distinct nodes $u, v \in (\pi_1 \cap \pi_2 \cap \pi_3)$. Notice that there must be two paths that use u, v in the same order; e.g., without loss of generality, we have $u <_{\pi_1} v$ and also $u <_{\pi_2} v$. But this implies that π_1, π_2 form a 2-bridge, giving contradiction.

Meanwhile, consider an arbitrary node v . There are $\binom{\deg(v)}{3} = \Theta(d^3)$ triples of paths that intersect at v (note: this equality uses that d is a large enough constant, and so $\deg(v) \geq 3$). We therefore have

$$\begin{aligned} nd^3 &= O(p^3) \\ \|S\|^3 = n^3d^3 &= O(n^2p^3) \\ \|S\| &= O(n^{2/3}p). \end{aligned} \quad \square$$

D.2 Bounds for $k = 3$

Theorem 80 (Tweaked Folklore Argument). $\beta(n, p, 3) = O((np)^{2/3} + n + p)$

Proof. This argument can be viewed as a slightly more careful version of the standard $O((np)^{2/3} + n + p)$ upper bound on $\gamma\gamma(n, p, 6)$. While technically slightly different, it follows the same rhythms and overall does not contain a significant new idea.

Let $S = (V, \Pi)$ be a path system with n nodes, p paths, and bridge girth > 3 . By the Cleaning Lemma (Lemma 10), we may assume without loss of generality that all nodes have degree $\Theta(d)$ and all paths have length $\Theta(\ell)$, where d, ℓ are respectively the average degree and length in S . We may also assume that both d, ℓ are sufficiently large constants, as otherwise the bound $O(n + p)$ is immediate.

Choose an arbitrary path $\pi \in \Pi$, which we will call the *main path*. We have that π intersects $\Omega(\ell)$ nodes, and each of these nodes have $\Omega(d)$ paths of length ℓ each. For a path q that intersects π , let us say that the *downstream part* of q is the suffix following the first point at which q, π intersect, and the *upstream part* of q is the prefix preceding the last point at which q, π intersect (so if q, π intersect at several nodes, which is conceivable so long as they use those nodes in opposite order, the downstream/upstream parts of q overlap).

Let U, D be the set of upstream, downstream parts of paths (respectively) that intersect π . We claim that the subpaths in U are pairwise node-disjoint from each other, and also the paths in D are pairwise node-disjoint from each other. To see that the paths in D are pairwise node-disjoint, suppose for contradiction that there are paths $q_1, q_2 \in D$ that intersect at a node v . Also suppose that the full paths containing q_1, q_2 intersect π at nodes x, y , respectively. Then we notice that these paths form a 3-bridge with π , on the nodes $\{v, x, y\}$, reaching contradiction. A similar argument works to show node-disjointness of paths in U .

This node-disjointness implies that $\|U\| + \|D\| = \Omega(d\ell^2)$, and so without loss of generality we may assume $\|D\| = \Omega(d\ell^2)$. Since the paths in D are node-disjoint we have $d\ell^2 = O(n)$, and so

$$(nd)(p^2\ell^2) = O(n^2p^2).$$

Since $nd = p\ell = \|S\|$, this implies

$$\|S\| = O((np)^{2/3})$$

as claimed. □




Figure 9: The upper bound for $\beta(n, p, 3)$. In the top is an illustration of the upstream and the downstream parts of a path π , and in the bottom the 3-bridge on the paths q_1, q_2 , and π .

Theorem 81 (Based on [43]). $\beta(n, p, 3) = O\left(\frac{n^2}{rs(n)} + p\right)$.

Proof. Let $S = (V, \Pi)$ be a path system with n nodes, p paths, and bridge girth > 3 . We may assume that the average path length ℓ is at least a large enough constant, as otherwise the bound of $O(p)$ is immediate. We associate S to an auxiliary graph G as follows:

- Split each path $\pi \in \Pi$ into as many node-disjoint subpaths as possible of length exactly 3 each. (We may discard one or two nodes at the end of the path.) Note that, since ℓ is a large enough constant, we change the size of $\|S\|$ by at most a constant factor over this splitting process.
- Take a uniform-random equitable tripartition $V = V_1 \cup V_2 \cup V_3$. For each path $\pi = (x, y, z) \in \Pi$, keep π iff $x \in V_1, y \in V_2, z \in V_3$; otherwise delete π from Π . Each path survives with constant probability, and so in expectation we again change the size of $\|S\|$ by at most a constant factor.
- Let G be the bipartite graph between vertex sets V_2, V_3 , where we include an edge (v_2, v_3) iff there exists a path $\pi \in \Pi$ with $v_2 <_\pi v_3$. Note that there is one edge in G per path in Π , and thus $|E(G)| = \Theta(\|S\|)$, so it suffices to bound $|E(G)|$.

- For each node $v_1 \in V_1$, define an edge subset $M[v_1] \subseteq E(G)$ as all edges (v_2, v_3) where there exists a path $(v_1, v_2, v_3) \in \Pi$.

In order to bound $|E(G)|$, we will show that each edge subset $M[v_1]$ is an induced matching. To see this, suppose for contradiction that there are distinct edges $(u_2, u_3), (v_2, v_3) \in M[v_1]$, and also an edge $(v_2, u_3) \in E(G)$. Suppose this other edge is caused by a path (v'_1, v_2, u_3) . Then we notice that the three paths $(v_1, u_2, u_3), (v_1, v_2, v_3), (v'_1, v_2, u_3)$ form a 3-bridge, with the first path as the river, giving contradiction.

Thus each set $M[v_1]$ is an induced matching, and so $E(G)$ may be partitioned into $|V_1| < n$ induced matchings. Thus, by definition of $\mathbf{rs}(n)$ we have $|E(G)| = O(n^2/\mathbf{rs}(n))$, completing the proof. $\square$

Corollary 82. $\beta(n, p, 3) = O\left(\frac{n^2}{2^{O(\log^* n)}} + n\right)$.

Proof. Follows from the previous theorem, and by plugging in the state-of-the-art bounds on $\mathbf{rs}(n)$ from [58, 82]. $\square$

Theorem 83. $\beta(n, p, 3) = \Theta((np)^{2/3})$ when $p \in \{n^{4/5}, n^{7/8}, n, n^{8/7}, n^{5/4}\}$.

Proof. The upper bound follows from Theorem 80. For the lower bound, for each of the given values of p , it is known [102] that

$$\gamma\gamma(n, p, 6) = \Omega\left((np)^{2/3}\right);$$

that is, there are constructions of bipartite graphs with n nodes on one side, p nodes on the other side, and girth > 6 . We may convert any such graph to a path system $S = (V, \Pi)$ by taking V as the n nodes on one side, taking Π as the p nodes on the other side, and including a node v in a path π iff (v, π) is an edge in the graph. The order of the nodes in each path can be chosen arbitrarily. Notice that a 2-bridge in S corresponds to a 4-cycle in the graph, and a 3-bridge in S corresponds to a 6-cycle in the graph. Since neither such cycle exists, S has bridge girth > 3 , and its size is $\|S\| = \Omega((np)^{2/3})$, completing the proof. $\square$