Error Amplification for Pairwise Spanner Lower Bounds

Amir Abboud and Greg Bodwin*

Stanford University

{abboud,gbodwin}@cs.stanford.edu

Abstract

A *pairwise spanner* of a graph $G = (V, E)$ and a "pair set" $P \subseteq V \times V$ is a subgraph H that preserves all pairwise distances in P, up to some additive error term $+\beta$. When $\beta = 0$ the object is called a *pairwise distance preserver*.

A large and growing body of work has considered upper bounds for these objects, but lower bounds have been elusive. The only known lower bound results are (1) Coppersmith and Elkin (SODA'05) against preservers, and (2) considerably weaker bounds by Woodruff (FOCS'06) against spanners.

Our main result is an amplification theorem: we prove that lower bounds against pairwise *distance preservers* imply lower bounds against pairwise *spanners*. In other words, to prove lower bounds against *any* constant error spanners, it is enough to consider only subgraphs that are not allowed any error at all!

We apply this theorem to obtain drastically improved lower bounds. Some of these include:

- Linear size pairwise spanners with up to $+(2k-1)$ error cannot span $|P| = \omega(n^{(1+k)/(3+k)})$ pairs. This is a large improvement over Woodruff's $|P| = \omega(n^{2-2/k})$ ($|P|$ is now linear, rather than quadratic, as k gets large).
- $|E(H)| = \Omega(n^{1+1/k})$ edges are required for a $+(2k-1)$ spanner of $|P| = \Omega(n^{1+1/k})$ pairs - this is another large improvement over Woodruff's $|P| = \Omega(n^2)$.
- The first tight bounds for pairwise spanners: for $+2$ error and $P = \Theta(n^{3/2})$ we show that $\Theta(n^{3/2})$ edges are necessary and sufficient (this also reflects a new *upper* bound: we construct $+2$ pairwise spanners on $O(n|P|^{1/3})$ edges, removing a log factor from a prior algorithm).

We also show analogous improved lower bounds against subset spanners (where $P = S \times S$ for some node subset S), and the first lower bounds against D threshold spanners (where P is the set of node pairs at distance at least D).

1 Introduction

An extremely important and highly successful subfield of graph theory and combinatorics is that of *graph sparsification*, in which one attempts to remove many edges from a graph while still approximately preserving certain important properties. One line of sparsification research has considered the fundamental property of *pairwise shortest path distances*. This problem captures natural questions in a large range of fields, including protocol design over unsynchronized networks [PU89a], compact routing schemes [Cow01, CW04, PU89b, RTZ08, TZ01], solving diagonally dominant linear systems [FPZW04], compressed distance oracles [TZ05, BS07, BK06, RTZ08], approximate all pairs shortest paths algorithms [EZ06, Elk05, Elk07, DHZ96], and more. In all of these fields, the quality of the application is directly tied to the tradeoff between the amount the underlying graph has been sparsified and the amount its distances have been perturbed. As such, improving this tradeoff is considered to be a fundamental and important question in theoretical computer science.

In this paper, we consider one of the most extensively studied notions of graph sparsification in recent years, namely *additive spanners*. The first works on additive spanners appear to be by Liestman and Shermer [LS91, LS93] after multiplicative spanners were introduced and studied by Peleg and Shäffer [PS89]. We will be interested in the question: what is the sparsest subgraph of a graph in which the distances between some given important pairs of nodes does not increase by much? Formally, the following definition of *pairwise spanners* captures many of the previously studied objects.

*This work was supported by NSF Grants CCF-1417238, CCF-1528078 and CCF-1514339, and BSF Grant BSF:2012338.

DEFINITION 1. (PAIRWISE SPANNERS) *Given a graph $G = (V, E)$ on n nodes and a set $P \subseteq V \times V$, we say that a subgraph $H = (V, E')$ is a $+\beta$* pairwise spanner *of G, P if $\delta_H(u, v) \leq \delta_G(u, v) + \beta$ for all $(u, v) \in P$.*

These have been extensively studied in the special case $P = V \times V$ (henceforth, *standard spanners*). Some upper bounds in this area are: A $+2$ spanner on $O(n^{3/2})$ edges by Aingworth et al. [ACIM99] (see also the follow-up works [DHZ96, EP04, RTZ05, TZ06]), a $+4$ spanner on $O(n^{7/5})$ edges by Chechik [Che13], and a $+6$ spanner on $O(n^{4/3})$ edges by Baswana et al. [BKMP05, BKMP10] (see also the follow-up works [Woo10, Knu14]). Most recently, Bodwin and Williams [BW16] obtained near-linear size spanners with error up to $+n^{3/7}$.

Lower bounds have been trickier to find. In the early works of [Awe85, PS89, ADD+93] it was noted that the Girth Conjecture of Erdös implies interesting lower bounds for spanners. The conjecture says that there are graphs on $\Omega(n^{1+1/k})$ edges with girth (minimum cycle length) of at least $(2k + 2)$. In such graphs, deleting any edge (u, v) will increase the distance between u and v from 1 to at least $2k+1$, and so no proper subgraph is a $+(2k - 1)$ spanner. This conjecture, however, has been open for many years, and has only been verified for $k = 1, 2, 3, 5$ [Wen91], which led researchers to seek lower bounds for spanners by other means. The only other lower bound is due to Woodruff [Woo06], who constructed graphs on $n^{1+1/k}$ edges in which removing a constant fraction of edges increases the distance between some pair of nodes by $2k$, thus proving the same $\Omega(n^{1+1/k})$ lower bound against $+(2k - 1)$ spanners *unconditionally*.

We will be interested in the case in which we only have to approximately-preserve the distances between a smaller set of important pairs, $P \subset V \times V$, while we can perturb the other distances by any amount.

In this setting it is natural to investigate the $k = 0$ case, i.e. sparsifying without changing the distances between our pairs at all. This gives rise to *pairwise distance preservers*, first formalized by Coppersmith and Elkin in [CE06], based on a similar object from [BCE03]. The authors showed upper bounds of $|E(H)| = O(\min\{n|P|^{1/2}, n + n^{1/2}|P|\})$, and Bodwin and Williams showed upper bounds of $O(n^{2/3}|P|^{2/3} + n|P|^{1/3})$ [BW16] (this is an improvement whenever $|P| = \omega(n^{3/4})$).

When the additive error k can be non-zero, we obtain the more general *pairwise spanners*. This notion was first studied by Cygan, Grandoni, and Telikepalli [CGK13] and has garnered a lot of attention since then, both because it is a fundamental mathematical object and because in many spanner applications we only care about a given set of node pairs. In particular, a natural line of attack against *standard* spanners is via obtaining better pairwise spanners [BW16]. A natural restriction is that of *subset spanners* (or preservers) in which $P = S \times S$ is the set of all pairs of terminals from a node subset S.

Subset spanners had been studied even earlier by Pettie [Pet09] who presented an $O(n|S|^{1/2})$ $+2$ subset spanner that is attributed to an unpublished work of Elkin. In [CGK13], the authors produce $+O(\log n)$ pairwise spanners on $O(n|P|^{1/4})$ edges. Kavitha and Varma [KV13] produce a $+2$ pairwise spanner on $\tilde{O}(n|P|^{1/3})$ edges. Kavitha [Kav15] constructs $+4$ pairwise spanners on $\tilde{O}(n|P|^{2/7})$ edges and $+6$ pairwise spanners on $\tilde{O}(n|P|^{1/4})$ edges. Bodwin and Williams [BW16] give a series of generalizations of the $+2$ subset spanner construction to sparser subset spanners on polynomial error.

This work is mostly concerned with lower bounds. A simple observation is that the Girth Conjecture implies an interesting even stronger lower bound for pairwise spanners: given a graph on $n^{1+1/k}$ edges with girth $2k + 2$, if we set P to be the set of pairs of nodes connected by an edge, we get an $\Omega(n^{1+1/k})$ lower bound against pairwise $+(2k - 1)$ spanners with $|P| = n^{1+1/k}$. Woodruff [Woo06] proved a much weaker bound: $\Omega(\frac{1}{k}|P|^{1/2} \cdot \min\{|P|^{1/2}, \frac{1}{k}n^{1/k}\})$ edges are required, which only matches the $\Omega(n^{1+1/k})$ lower bound when $P = V \times V$ is the set of all pairs of nodes, i.e. much larger than $n^{1+1/k}$. A modification of Woodruff's construction was suggested by Parter [Par14]. In the special case $P = S \times V$, called a *sourcewise spanner*, the author shows a lower bound of $\Omega(n|S|^{1/k})$.

Meanwhile, in the zero error case, sophisticated constructions of lower bound graphs have been suggested by Coppersmith and Elkin [CE06]. The authors prove a family of lower bounds: for all $d = 2, 3, 4, \ldots$, there is a lower bound of $\Omega(n^{2d/(d^2+1)}|P|^{d(d-1)/(d^2+1)})$ against pairwise distance preservers. This agrees with our intuition: it is possible to obtain much denser lower bounds against preservers (zero error) than the ones we have against spanners.

1.1 Main Result

We prove an "amplification theorem" that can take *any* lower bound construction against distance preservers and turn it into a lower bound construction against pairwise spanners. That is, we reduce the task of constructing $+k$ pairwise spanner lower bounds to the easier looking task of constructing distance preserver lower bounds. The density of the final graphs decreases with the error parameter k, which is to be expected, and depends on $\mathcal{T}(P)$ - the number of distinct nodes that appear in the pairset (i.e. our set of "terminals").

THEOREM 1.1. (MAIN RESULT) *Suppose that G, P is a lower bound against pairwise distance preservers. Let k be any nonnegative integer constant. Then there exists a graph G' on $\Theta(n|\mathcal{T}(P)|^{k-1})$ nodes and $\Theta(|E(G)||\mathcal{T}(P)|^{k-1})$ edges and a pair set P' of size $\Theta(|P||\mathcal{T}(P)|^{k-1})$ such that G', P' is a lower bound against $+(2k-1)$ pairwise spanners. Additionally, $|\mathcal{T}(P')| = \Theta(|\mathcal{T}(P)|^k)$.*[1]

Our formal definitions for "lower bound graphs" are the natural ones. Another way of stating our Theorem is as follows: you can amplify a pairwise distance preserver lower bound from +0 to +1 for free, and then you can amplify it by an additional +2 any number of times by multiplying the node count, edge count, and pair set size all by $|\mathcal{T}(P)|$. Accordingly, the lower bounds produced by Theorem 1.1 are strongest when the initial preserver lower bound uses very few terminals. It is easy to see that $|\mathcal{T}(P)|$ is minimized when $P = S \times S$ for some node subset S; i.e. the initial lower bound is a lower bound against *subset* distance preservers.

The currently known lower bounds against subset distance preservers are: $|E(H)| = \Omega(n^{9/11}|S|^{6/11} + n^{10/11}|S|^{4/11} + |S|^2)$. The first two terms are due to Coppersmith and Elkin [CE06], and the latter term is a trivial lower bound that arises when the initial graph is a clique. The state of the art suggests that the "true" subset preserver bound might eventually be $\Omega(n^{2/3}|S|^{4/3})$, as stated in the following hypothesis. This hypothesis would hold, for example, if the "true" bound has the simple form $\Omega(n^a|S|^b)$ for some constants a, b, rather than an erratic piecewise behavior. See Figure 1 for a visualization of all these bounds.

HYPOTHESIS 1. *There is a lower bound against subset distance preservers of $\Omega(n^{2/3}|S|^{4/3})$.*

We will now apply our amplification theorem to all of these lower bound constructions, immediately giving new lower bounds against pairwise and subset spanners. These lower bounds shed new light on interesting open questions in the world of additive spanners.

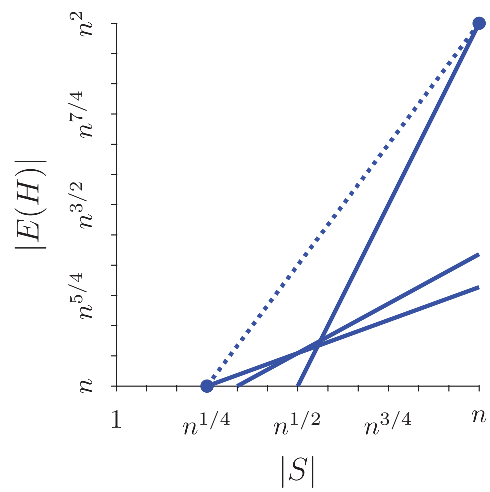

Figure 1: Known lower bounds against subset distance preservers (due to Coppersmith and Elkin [CE06]) are shown by solid lines, while our hypothesized lower bound is shown by the dotted line.

1.2 Consequences: New Lower Bounds for Pairwise Spanners.

We will use the following generic corollary of Theorem 1.1.

COROLLARY 1.1. *Let a, b be constants such that there is a lower bound of $|E(H)| = \Omega(n^a|S|^b)$ against subset distance preservers. Then there is a lower bound of $|E(H)| = \Omega(n^a|P|^{(a+b+k-1-ak)/(k+1)})$ against $+(2k-1)$ pairwise spanners, so long as $|P| = O(n^{1+1/k})$.*

For a visualization of the direct implications of this corollary, see Figure 2.

Replacing the Girth Conjecture. Our first major consequence comes from applying our amplification theorem to a clique on $|S|$ nodes, which is an $\Omega(|S|^2)$ subset preserver lower bound. Fixing $|P| = n^{1+1/k}$ and applying Corollary 1.1, we obtain a construction that unconditionally proves an $\Omega(n^{1+1/k})$ lower bound against $+(2k-1)$ pairwise spanners with only $|P| = \Theta(n^{1+1/k})$ pairs. This matches the lower bound previously only known to hold under the Girth Conjecture. Previously, by Woodruff's result [Woo06], it was only known that this number of edges is required to span *all* distances in the graph.

In particular, this implies that there is a $\Omega(n^{3/2})$ lower bound against +2 pairwise spanners with $|P| = \Theta(n^{3/2})$. This matches an upper bound proven in [KV13] up to log factors. We remove this log factor, proving the first tight bounds for pairwise spanners.

THEOREM 1.2. *For all graphs G and pair sets P, there is a +2 pairwise spanner of G, P on $O(n|P|^{1/3})$ edges.*

The proof mixes ideas from [BW16] and [KV13].

Another consequence of our lower bound for the case $k = 2$ is that subset spanners can be *sparser* than

[1] These Θ's hide multiplicative factors that depend only on k.

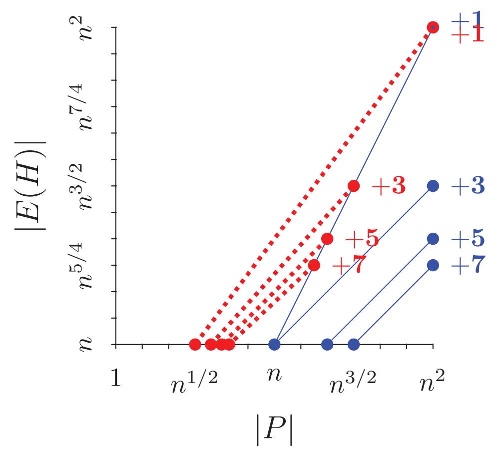

Figure 2: The old lower bounds for pairwise spanners are in solid blue. The new lower bounds for pairwise spanners, conditional on our hypothesized subset preserver lower bound, are in dotted red. Only slightly worse lower bounds are available unconditionally (not pictured). In particular, the red dots at the end of each line hold unconditionally.

pairwise spanners. Known constructions for $+2$ subset spanners beat our pairwise lower bound! For subsets of size $|S| = n^{3/4}$ (and therefore $|S \times S| = n^{3/2}$), it is known [Pet09, CGK13] that we can obtain a $+2$ subset spanner on $O(n^{11/8}) = o(n^{3/2})$ edges. This conclusively proves that subset spanners are fundamentally easier objects to construct than pairwise spanners, at least in certain edge ranges. This is surprising since such behavior disappears entirely when the error allowance is $+0$: our current upper bounds for subset preservers match our current upper bounds for pairwise preservers; they do not exploit the structure $P = S \times S$ at all.

Linear-Size Pairwise Spanners. Next, we apply our amplification theorem to the constructions of Coppersmith and Elkin, obtaining improved lower bounds in many settings. For example, we improve our current answer to the following natural open question:

Open Question 1. *For constant k, what is the largest δ such that one can always produce a $+k$ pairwise spanner of a pair set of size $\Theta(n^\delta)$ on $O(n)$ edges?*

Questions of this type, where one looks for the optimal attainable parameters with near-linear size subgraphs, have been asked in many different contexts: for example, mixed spanners[2] [EP04, TZ06, Pet09], standard additive spanners with polynomial error [BCE03, BKMP05, BKMP10, Pet09, Che13, BW15, BW16], and pairwise preservers [CE06]. In particular, in the $k = 0$ case, Coppersmith and Elkin obtain a complete answer by showing that $|P| = \Theta(n^{1/2})$ pairs can be preserved with $O(n)$ edges, but not more than that [CE06].

Interestingly, the current best upper bounds for $+k$ pairwise spanners, for $k > 0$, do not allow us to support more than $|P| = O(n^{1/2})$ pairs with a linear number of edges. Another way of phrasing this is: for linear size spanners, we don't know how to benefit from any constant error allowance.

Woodruff's lower bound for $+(2k+1)$ pairwise spanners shows that $\omega(n)$ edges are required whenever $|P| = \omega(n^{2-2/k})$, i.e. $\delta \leq 2 - 2/k$ which is almost quadratic as k grows. The Girth Conjecture gives a much better answer, saying that even $|P| = \omega(n)$ pairs cannot be supported, i.e. $\delta \leq 1$.

Applying Corollary 1.1 with the $\Omega(n^{10/11}|S|^{4/11})$ lower bound of Coppersmith and Elkin gives an even better answer than this. It implies that one cannot in general construct $O(n)$ size $+(2k-1)$ pairwise spanners with $|P| = \omega(n^{(1+k)/(3+k)})$. To our knowledge, this is the first lower bounds of any sort against pairwise spanners with sublinear pair sets size.

This also highlights the area in which our current knowledge of pairwise spanners is the most lacking: let δ be the parameter in Open Question 1; this corollary implies $\delta \leq (1+k)/(3+k)$, while it is currently only known that $\delta \geq 1/2$. This gap seems to be largely due to the fact that known upper bounds cannot currently use a non-zero error allowance to support more pairs.

For large enough pair sets (although not for $|P| = \Theta(n^2)$), one can improve the unconditional lower bounds (but not the hypothesized lower bound) by amplifying the lower bounds against *pairwise* preservers rather than these subset preserver lower bounds. We will not discuss this point any further.

Another use of Corollary 1.1 is through its contrapositive: one can obtain better subset preserver *upper* bounds (that refute our hypothesis) simply by obtaining better pairwise spanner upper bounds.

Lower Bounds Against Subset Spanners. Our pairwise spanner bounds can be straightforwardly extended into lower bounds against subset spanners. In particular, we show:

Corollary 1.2. *Let a, b be constants such that there is a lower bound of $|E(H)| = \Omega(n^a|S|^b)$ against subset distance preservers. Then there is a lower bound of $|E(H)| = \Omega(n^a|S|^{(a+b+k-1-ak)/k})$ against $+(2k-1)$ subset spanners.*

[2] Spanners which have both an additive and multiplicative error

Using the subset preserver L.B.	**We get a pairwise spanner L.B. of**
$\Omega(n^{10/11}\|S\|^{4/11})$ ([CE06])	$\Omega(n^{10/11}\|P\|^{(k+3)/(11k+11)})$
$\Omega(n^{9/11}\|S\|^{6/11})$ ([CE06])	$\Omega(n^{9/11}\|P\|^{(2k+4)/(11k+11)})$
$\Omega(\|S\|^2)$ (clique)	$\Omega(\|P\|)$
$\Omega(n^{2/3}\|S\|^{4/3})$ (hypothesized)	$\Omega(n^{2/3}\|P\|^{(k+3)/(3k+3)}$

Table 1: New pairwise spanner lower bounds.

Using the subset preserver L.B.	**We get a subset spanner L.B. of**
$\Omega(n^{10/11}\|S\|^{4/11})$ ([CE06])	$\Omega(n^{10/11}\|S\|^{(k+3)/(11k)})$
$\Omega(n^{9/11}\|S\|^{6/11})$ ([CE06])	$\Omega(n^{9/11}\|S\|^{(2k+4)/(11k)})$
$\Omega(\|S\|^2)$ (clique)	$\Omega(\|S\|^{1+1/k})$
$\Omega(n^{2/3}\|S\|^{4/3})$ (hypothesized)	$\Omega(n^{2/3}\|S\|^{(k+3)/(3k)})$

Table 2: New subset spanner lower bounds.

Again, we can instantiate Corollary 1.2 with any of the lower bounds discussed above; the results can be seen in Table 2. For a visualization, see Figure 3. Again, slightly improved unconditional results are available for large enough S (but not $|S| = \Theta(n)$) if one amplifies the current best pairwise spanner lower bounds instead.

Lower Bounds for D Threshold Spanners. Bollobás, Coppersmith, and Elkin [BCE03] introduced D-preservers, which are sparse subgraphs that preserve the distances between all pairs of nodes that are at distance at least D. The authors show matching upper and lower $\Theta(n^2/D)$ bounds for these objects. Another interesting consequence of our Theorem 1.1 is the first lower bound for D Threshold Spanners – the approximate versions of D-preservers in which we hope to get sparser subgraphs by allowing $+k$ stretch. Upper bounds for these objects were shown in [KV13][3]. When we amplify the lower bound of Bollobás et al. [BCE03], we get:

Proposition 1.1. *For all integer constants $k \geq 1$, for any parameter D, there is a lower bound against D threshold spanners with $|E(G)| = \Theta(n^{1+1/k}/D^{1/k})$.*

When D is constant, the problem reduces to the question of building a standard spanner; here our lower bound matches the usual $\Omega(n^{1+1/k})$ lower bound. This lower bound is quite interesting for large D. The Bollobás et al. lower bound implies that one cannot, in general, make a D-preserver where $D = o(n)$ on a linear number of edges. Corollary **??** shows that this remains true even for D threshold spanners with *any* constant amount of error.

See Figure 4 for a visualization.

Pairwise and Standard Spanner are in Sync. Suppose you design a pairwise spanner upper bound, and you want to see how well it performs for large $|P|$. Naively, you might look at the largest possible P; that is, $|P| = \Theta(n^2)$. However, one of the major insights of Corollary 1.1 is that our lower bounds are piecewise: they improve steadily so long as $|P| \leq |E(H)|$, but then they plateau in the range $|E(H)| \leq |P| \leq n^2$. At least for $k = 2$, *both* pieces of this upper bound are tight. In particular, once $|P| \geq |E(H)|$, the best possible upper bound comes from simply ignoring the given pair set and constructing a standard spanner. In some sense, then, the largest P worth considering is where $|P| = |E(H)|$: for larger P, you start to capture the hardness of the standard spanner question rather than the pairwise spanner question.

How do modern constructions perform under this $|P| = |E(H)|$ benchmark? The results are fairly interesting. For $k = 1$, the upper bound at $|P| = n^{3/2}$ is $|E(H)| = O(n^{3/2})$. For $k = 2$, the upper bound at $|P| = n^{7/5}$ is $|E(H)| = \tilde{O}(n^{7/5})$. For $k = 3$, the upper bound at $|P| = n^{4/3}$ is $|E(H)| = O(n^{4/3})$. These are the

[3] In prior work, these objects were simply called D-spanners. We add the word "threshold" to disambiguate from T-spanners, which is standard notation for a spanner with multiplicative stretch T.

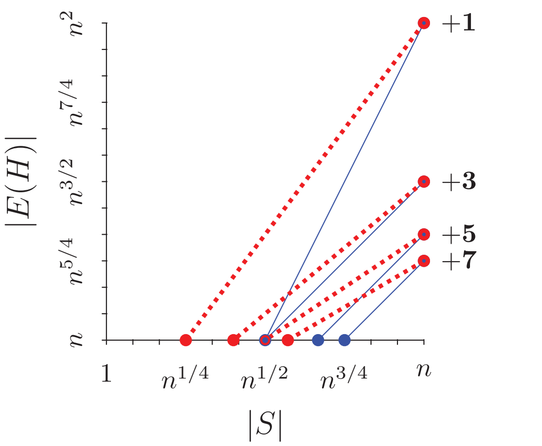

Figure 3: The old lower bounds for subset spanners are in solid blue. The new lower bounds for subset spanners, conditional on our hypothesis about subset preserver lower bounds, are in dotted red. Only slightly worse lower bounds are available unconditionally (not pictured). In particular, the red dots at the end of each line hold unconditionally.

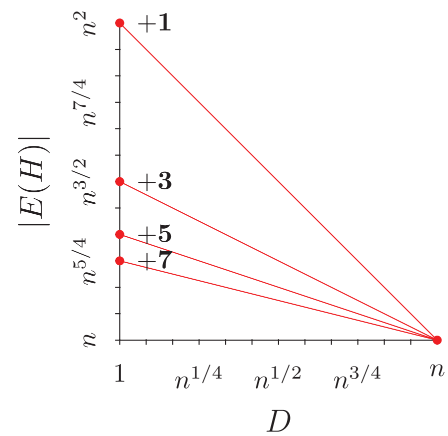

Figure 4: New lower bounds for D threshold spanners. No lower bounds were previously known.

exact same upper bounds that are known for standard spanners. For $k = 1$ our standard spanner upper bound is $O(n^{3/2})$, for $k = 2$ it is $\tilde{O}(n^{7/5})$, and for $k = 3$ it is $O(n^{4/3})$.

From this, it seems likely that any further improvement to our upper/lower bounds on pairwise spanners for *large* $|P|$ will probably require a new set of techniques that is also capable of improving our upper/lower bounds on standard spanners. This is a problem on which we have been stuck for a long time. It seems likely, then, that the easiest way to improve pairwise spanners upper bounds is to improve their behavior for *small* $|P|$, while perhaps maintaining the same behavior at $|P| = |E(H)|$.

2 The Error Amplification Theorem

The rest of this paper is dedicated to the proof of Theorem 1.1.

Overview. We start with a short summary of the high-level ideas in our proof, and a comparison with the previously known lower bounds of Woodruff. Let G, P be a lower bound for pairwise distance preservers, that is, no edge can be deleted without increasing the distance between a pair $(s, t) \in P$. Our new graph G' will be obtained from G by taking a certain kind of *graph product* of G with itself k times while taking the pair set P into account. Roughly speaking, the nodes of G' will have the form $s = (s_1, \ldots, s_k)$ and $t = (t_1, \ldots, t_k)$ where s_i, t_i correspond to nodes of G and the shortest path from s to t will have to travel all of the $s_i \rightsquigarrow t_i$ shortest paths, where $(s_i, t_i) \in P$. This allows us to show that if a subgraph of G' is too sparse, then the new shortest path from s to t must take a +2 detour from the $s_i \rightsquigarrow t_i$ path *for every* $i \in [k]$, which will lead to the distance increasing by at least $2k$. To make these ideas work, the final construction requires further technicalities that will be explained later on.

When we apply our amplification to the clique, which is an $\Omega(S^2)$ lower bound for subset preservers, we get a graph G' that looks similar to Woodruff's previously known lower bound construction. From a distance, both graphs look like a k-layered graph with structured complete bipartite graphs in between subsets of these layers. Surprisingly, despite this similarity, our final lower bounds (even when amplifying the clique) are substantially stronger than before. Our main improvement comes from the second ingredient in the amplification, which is choosing the pair set. Woodruff's argument for standard spanners shows that it is hard to approximately-preserve the distances in G' between the first and last layers, while his improved lower bound for pairwise spanners shows that one can choose a *subgraph* of G' in which a smaller pair set gives the same hardness. Our approach is different: we try to find the smallest possible pair set for which any way of deleting a constant fraction of edges of G' must increase one of the distances beyond the error allowance. To achieve this, we identify strong structural properties of the amplified graphs and the shortest paths within them which allow us to consider only much smaller pair sets. We introduce a process of distributing tokens across the shortest paths of the graph, which, intuitively, gives a small set of pairs that cannot be "fooled" by any adversary who is trying to delete many edges without stretching the distance for any of those pairs.

2.1 Preprocessing

Let G, P be a lower bound against pairwise distance preservers. Our first step is a purely technical preprocessing phase, in which we massage G, P to have some convenient properties.

DEFINITION 2. *A node pair u, v in a graph G* requires *an edge e if every shortest path between u and v uses e.*

We need:

1. The hardness of the lower bound is evenly spread around the pairs in P. This means that one can assign edges to pairs that require them in such a way that (a) a constant fraction of the edges are assigned to some pair, and (b) all pairs own the same number of edges.

2. There is a value $\alpha = \Theta(\frac{|P|}{\mathcal{T}(P)})$ such that every node in $\mathcal{T}(P)$ is used in between α and 2α pairs in P.

3. The graph is bipartite.

We can achieve these three properties simultaneously without asymptotically changing the values of n, $|E(G)|$, $|P|$, or $|\mathcal{T}(P)|$. The algorithm that accomplishes this is given in the next subsection.

2.2 Preprocessing

First Property. First, repeat the following process until you can do so no longer: find an edge that is not required by any pair $p \in P$ and remove that edge from G. The resulting graph is clearly still a distance preserver of G, P; and by our definition of a distance preserver lower bound we could not have removed more than a constant fraction of the edges in G.

We can now assign each edge in G to a pair in P that requires it. First, we will establish a lower bound on the number of edges owned by each pair. Fix $p = |P|$, then delete all pairs from P that own fewer than half the average number of edges (that is, fewer than $\lfloor \frac{|E(G)|}{2p} \rfloor$ edges). Leave the edges owned by deleted pairs unassigned. Note that this leaves at most $p \cdot \lfloor \frac{|E(G)|}{2p} \rfloor \le \frac{|E(G)|}{2}$ edges without an owner; therefore, a constant fraction of the edges still have an owner.

We will now lower the number of edges owned by each pair to this upper bound. We will do this by copying each pair many times and splitting the edges evenly between its copies. Specifically: for each pair $q \in P$, replace q with $\lfloor c / \lfloor \frac{|E(G)|}{2p} \rfloor \rfloor$ copies of q, where c is the number of edges owned by q. Assign each copy exactly $\lfloor \frac{|E(G)|}{2p} \rfloor$ of the edges previously owned by q, and leave the remainder unassigned.

This transformation of P has the following effect on our parameters:

- For each pair q, at least half of the edges previously owned by q are now owned by a copy of q; therefore, at the end of this process, it is still true that a constant fraction of the edges in G have an owner.

- It is now the case that all pairs own exactly $\lfloor \frac{|E(G)|}{2p} \rfloor$ edges.

- The final size of P therefore satisfies $|P| \cdot \lfloor \frac{|E(G)|}{2p} \rfloor \le |E(G)|$, and so $|P| \le 3p$, which means that our pair copying has expanded the size of P by at most a constant factor.

Second Property. The size of p may have changed throughout the above process, so we re-fix $p = |P|$ for the *new* size of P, and we also fix $t = |\mathcal{T}(P)|$. First, we will establish a lower bound on the number of pairs in which each node participates. For each node $u \in \mathcal{T}(P)$ that participates in fewer than half the average number of pairs (that is, $\lfloor \frac{p}{2t} \rfloor$ pairs) in P, delete all pairs from P that include the node u. Note that we delete at most $t \cdot \lfloor \frac{p}{2t} \rfloor \le \frac{p}{2}$ pairs, and each pair owns the same number of edges; therefore, at the end of this process, it is still true that a constant fraction of edges in G have an owner.

Next, we will enforce an upper bound on the number of pairs in which each node participates. For each surviving node $v \in \mathcal{T}(P)$, add $\lfloor c / \lfloor \frac{p}{2t} \rfloor \rfloor$ new nodes $v_1, v_2, \ldots$ to G, where c is the number of pairs in P that v participates in. Connect each of these new nodes with an edge to v (do not assign these new edges to any pair). For each pair $(u, v) \in P$ containing v, replace it with (u, v_i) for some i so that the pairs are divided as evenly as possible between the v_i nodes. Assign each new pair all of the edges owned by the pair it replaced.

This transformation of G, P has the following effect on our parameters:

- We have preserved the property that all pairs own the same number of edges, since each new pair inherits exactly the set of edges previously owned by an old pair.

- The size of P has shrunk by at most a factor of 2.

- Each node in $\mathcal{T}(P)$ participates in at least $\lfloor \frac{p}{2t} \rfloor$ pairs and at most $2 \cdot \lfloor \frac{p}{2t} \rfloor$ pairs (if a node v_i

participates in more than $2 \cdot \lfloor \frac{p}{2t} \rfloor$ pairs, then we would have chosen to add at least one more node v_j connected to the original node v which would take some of its pairs).

- The size of $\mathcal{T}(P)$ at the end of this transformation has only changed by a constant factor. To see this, let x be the number of nodes (with repeats) that appear in the final pair set: we know that $x = 2|P|$, while at the same time $x \geq |\mathcal{T}(P)| \lfloor \frac{p}{2t} \rfloor$ and $x \leq |\mathcal{T}(P)|(2 \cdot \lfloor \frac{p}{2t} \rfloor)$. This gives

$$|\mathcal{T}(P)| \lfloor \frac{p}{2t} \rfloor \leq 2|P| \leq |\mathcal{T}(P)|(2 \cdot \lfloor \frac{p}{2t} \rfloor)$$

- Every node added to G is now in $\mathcal{T}(P)$. Since $|\mathcal{T}(P)|$ has increased by at most a constant factor, and its original size was at most n, we have increased the number of nodes in G by at most a constant factor.

Converting to Bipartite. For each node v in G, we add nodes v^L and v^R to G' (these will be on the left and right side of the bipartite partition). For each edge (a, b) in G, we add edges (a^L, b^R) and (a^R, b^L) to G'. For each pair $(u, v) \in P$, we add:

1. (u^L, v^L) and (u^R, v^R) to P' if $\delta_G(u, v)$ is even
2. (u^L, v^R) and (u^R, v^L) to P' if $\delta_G(u, v)$ is odd

We have doubled the number of nodes and edges in G, the number of pairs in P, and the number of nodes in $\mathcal{T}(P)$, so we maintain our property that these values have not increased by more than a constant factor since the original lower bound. We will next show that we maintain our new properties: (1) a constant fraction of the edges in G' can be assigned to a pair in P' that requires them in such a way that all pairs own the same number of nodes, and (2) all nodes in $\mathcal{T}(P')$ participate in the same number of pairs, up to a factor of 2.

Let (a, b) be an edge in G required by the pair $(u, v) \in P$. We make the simplifying assumption that $\delta_G(u, v)$ and $\delta_G(u, a)$ are both even; the argument is essentially identical if one or both of these values is odd. If $\rho_G(u, v) = \{u, x_1, x_2, x_3, \ldots, x_k, v\}$, then there must be a shortest path $\rho_{G'}(u^L, v^L) = \{u^L, x_1^R, x_2^L, x_3^R, \ldots, x_k^R, v^L\}$ (the same as $\rho_G(u, v)$, but with alternating L/R superscripts). It follows that every shortest path $\rho_{G'}(u^L, v^L)$ uses the edge (a^L, b^R), and so (a^L, b^R) is required by (u^L, v^L) in G'. Similarly, (a^R, b^L) is required by (u^R, v^R). We can then assign these edges to these pairs. Each pair (u^L, v^L) and (u^R, v^R) will own exactly the same number of edges that (u, v) used to own, and each node u^L will participate in exactly the same number of pairs that u used to participate in. Therefore, our properties have been preserved.

2.3 Main Construction

After we are done preprocessing the preserver lower bound G, P, we start "amplifying" it into a pairwise spanner lower bound. This is the main part of our proof.

Constructing the New Graph. Given G, P (that have been preprocessed), we define a new graph G' as follows.

The nodes of G' are tuples that take one of the following two forms (1) $(t_0, \ldots, t_{k-1}, L)$ or (2) $(t_0, \ldots, t_{\ell-1}, v_\ell, t_{\ell+1}, \ldots, t_{k-1}, \ell \cdot k)$.

The last entry of this tuple L is the *layer value* of the node, and it takes integer values between 0 and $k(k-1)$. A node with $L = \ell \cdot k$ for some integer ℓ is said to be in layer ℓ of the graph; otherwise, the node is *between layers*. If a node is between layers (form 1), then all its entries t_i are nodes in $\mathcal{T}(P)$. If a node is in layer ℓ of the graph (form 2), then all its entries t_i are nodes in $\mathcal{T}(P)$, *except* its ℓ^{th} entry which can be any node in G.

The total number of nodes per layer is $n \cdot |\mathcal{T}(P)|^{k-1}$, and so we have $\Theta(n \cdot |\mathcal{T}(P)|^{k-1})$ nodes in all of G' (since we treat k as a constant, the Θ suppresses multiplicative factors that depend only on k).

We next construct the edges of G'. We add an edge between two nodes in G' if either of the following conditions hold:

1. The nodes are of the form $(t_0, \ldots, t_{k-1}, L)$ and $(t_0, \ldots, t_{k-1}, L + 1)$. That is, the two nodes agree everywhere except on the layer value, all their other entries are nodes in $\mathcal{T}(P)$, and their layer values differ by exactly 1. This type of edge is called a *transition edge*.
2. The nodes are of the form $(t_0, \ldots, t_{\ell-1}, a, t_{\ell+1}, \ldots, t_{k-1}, \ell \cdot k)$ and $(t_0, \ldots, t_{\ell-1}, b, t_{\ell+1}, \ldots, t_{k-1}, \ell \cdot k)$, where (a, b) is an edge in G. That is, the two nodes agree everywhere except the ℓ^{th} entry, there is an edge in G that connects the nodes contained in their ℓ^{th} entries. This type of edge is called a *layer ℓ edge*.

The new graph G' contains as a subgraph many copies of the (preprocessed) lower bound graph G. Specifically:

DEFINITION 3. *Let $u = (t_0, \ldots, t_{k-1}, L_1)$ and $v = (t'_0, \ldots, t'_{k-1}, L_2)$ be a pair of nodes in G' whose entries are all nodes in $\mathcal{T}(P)$. Then $F_\ell(u, v)$ is defined to be the subgraph induced by the node set $\{(t'_0, \ldots, t'_{\ell-1}, x, t_{\ell+1}, \ldots, t_{k-1}, \ell \cdot k) \mid x$ is a node in $G\}$. A subgraph that is equal to $F_\ell(u, v)$ for some choice of ℓ, u, v is called an* F-subgraph.

Intuitively, if u is in the first layer and v is in the last layer, this definition of $F_\ell(u, v)$ corresponds to the layer ℓ part of G' that could appear on a path from u to v. We will make use of this definition in the lemmas below.

Note that every F-subgraph is isomorphic to G: the node sets match by definition, and the set of layer edges in a specific F-subgraph is exactly the set of edges in G. Additionally, each layer node belongs to exactly one F-subgraph, which consists of all layer ℓ nodes that share its non-ℓ entries. We can then define:

Definition 4. *For a node v in some layer of G', we define $\mathcal{F}(v)$ to be the unique F-subgraph that contains v. We define $\phi(v)$ to be the corresponding node in G under the isomorphism between $\mathcal{F}(v)$ and G.*

Each layer of the graph consists of $|\mathcal{T}(P)|^{k-1}$ different F-subgraphs, corresponding to the different settings of the other $k-1$ entries. Since each F-subgraph is isomorphic to G, it has $|E(G)|$ edges. So the total number of layer edges is G' is $\Theta(|E(G)| \cdot |\mathcal{T}(P)|^{k-1})$. The number of transition edges is only a multiplicative factor of k more than the number of *nodes* in G', so this quantity is dominated by the number of layer edges in G'. Therefore, the total number of edges in G' is $\Theta(|E(G)| \cdot |\mathcal{T}(P)|^{k-1})$.

Constructing the New Pair Set. A naive way to choose a pair set P' for our spanner lower bound would be picking every pair of nodes, one from the first layer and one from the last. This strategy would yield a pair set of size $|P'| = \Omega(|\mathcal{T}(P)|^{2k})$ and such an amplification result would not yield any nontrivial lower bounds for pairwise spanners. Instead, we chose our pair set P' in a much more careful way, yielding size that is only $|P'| = O(|P| \cdot |\mathcal{T}(P)|^{k-1})$. The pair set will be chosen via the following process, which is the crux of our proof.

Let α to be the integer such that every node $t \in \mathcal{T}(P)$ participates in between α and 2α pairs in P.

1. Place $2^k\alpha$ tokens on each layer 0 node v such that $\phi(v) \in \mathcal{T}(P)$.

2. For each node v with tokens on it, let $U = \{u \in \mathcal{F}(v) \,|\, (\phi(v), \phi(u)) \in P\}$. Assign each token on v to a different node $u \in U$ in such a way that the number of tokens owned by any $u, u' \in U$ differs by at most 1. Then, slide each token from v to the node u that owns it. Perform this process simultaneously for all the layer 0 nodes.

3. Slide each token along k transition edges to layer 1 of the graph.

4. Repeat until you can no longer slide the tokens to the next layer of the graph (because you have reached layer $k-1$). For each token, add its start node and its end node as a pair to P'.

The number of pairs is equal to the number of tokens, which is $2^k\alpha$ times the number of layer 0 nodes with $\phi(v) \in \mathcal{T}(P)$. The latter quantity is equal to the number of tuples of k nodes from $\mathcal{T}(P)$, which is $|\mathcal{T}(P)|^k$. From our preprocessing step we have that $\alpha = \Theta(\frac{|P|}{|\mathcal{T}(P)|})$. Therefore, the total number of tokens is $2^k \cdot \Theta(\frac{|P|}{|\mathcal{T}(P)|}) \cdot |\mathcal{T}(P)|^k = \Theta(|P||\mathcal{T}(P)|^{k-1})$, so we have this many pairs in P'.

Intuitively, these 2^k factors can be traced back to the fact that our preprocessing algorithm enforced the property that every node in $\mathcal{T}(P)$ is used in between α and 2α pairs in P. We could have just as easily enforced the property that each node in $\mathcal{T}(P)$ is used in α to $(1+\varepsilon)\alpha$ pairs, and then our 2^k factors would become $(1+\varepsilon)^k$ factors.

This completes the construction. We will now argue that $\{G', P'\}$ is now a lower bound against $+(2k-1)$ pairwise spanners.

2.4 Proof of Correctness

Our first lemma states that every node that is eligible to receive a token necessarily receives $\Theta(\alpha)$ tokens.

Lemma 2.1. *Let v be a node in layer ℓ of the graph such that $\phi(v) \in \mathcal{T}(P)$. The total number of tokens v will receive from nodes in the previous layer (i.e. Step 3 of the algorithm) is between $2^{k-\ell}\alpha$ and $2^{k+\ell}\alpha$. The number of tokens v will receive from nodes in its own layer (i.e. Step 2 of the algorithm) is between $2^{k-\ell-1}\alpha$ and $2^{k+\ell+1}\alpha$.*

Proof. We will only prove the lower bound; the upper bound follows from an essentially identical argument. The proof is by induction on the layer number. The base case follows from the fact that we place $2^k\alpha$ tokens on each v in layer 0 with $\phi(v) \in \mathcal{T}(P)$. We will next prove the inductive step. Assume that the layer ℓ node that receives the fewest tokens from the previous layer receives at least $2^{k-\ell}\alpha$ tokens. Note that each node participates in at most 2α pairs, so it pushes at least $\frac{2^{k-\ell}\alpha}{2\alpha} = 2^{k-\ell-1}$ tokens to each of its pairs. Each node participates in at least α pairs, so it receives a total of at least $2^{k-\ell-1}\alpha$ from other nodes. Finally, each node in the next layer $\ell+1$ will receive tokens from exactly one node in the previous layer, so it receives at least $2^{k-(\ell+1)}$ tokens. □

All pairs in P' consist of a node in the first layer and a node in the last layer. Our next lemma shows how the shortest paths in G translate into the shortest path for each of these pairs.

Lemma 2.2. *Let u be a node in layer 0 of G', and let v be a node in layer $k-1$ of G'. For all ℓ, every shortest path between u and v contains as a subpath some shortest path between u_ℓ and v_ℓ in the subgraph $F_\ell(u,v)$.*

Proof. A path between $u = (t_0, \ldots, t_{k-1}, 0)$ and $v = (t'_0, \ldots, t'_{k-1}, k(k-1))$ can be generated by the following process. Travel a shortest path between t_0 and t'_0 in $F_1(u,v)$, arriving at node $(t'_0, t_1, \ldots, t_{k-1}, 0)$. Next, travel k transition edges to the next layer of teh graph, arriving at node $(t'_0, t_1, \ldots, t_{k-1}, k)$. Next travel a shortest path between t_1 and t'_1 in $F_2(u,v)$, and then move to the next layer of the graph, and so on until you have reached v.

Note that this path starts in $F_1(u,v)$, then travels to $F_2(u,v)$, and so on up to $F_{k-1}(u,v)$. In each of these layer ℓ F-subgraphs, its ℓ^{th} coordinate moves from u_ℓ to v_ℓ. Therefore, any path generated by this process contains as a subpath $\rho_{F_\ell(u,v)}(u_\ell, v_\ell)$ for all ℓ. We will prove that *every* shortest path between u and v can be generated by the above process.

First, note that every path between u and v must use at least $\delta_G(t_\ell, t'_\ell)$ layer ℓ edges for all ℓ. This is true because the first such edge must have ℓ^{th} entry equal to t_ℓ, and the last such edge must have ℓ^{th} entry equal to t'_ℓ, so these edges correspond to a path in G between t_ℓ and t'_ℓ. Second, note that every path must use at least $k(k-1)$ transition edges, because otherwise its final node will have layer value less than $k(k-1)$. These bounds are both tight in the case of a path generated by the above process; therefore, no path can be shorter.

Suppose that a path uses exactly $k(k-1)$ transition edges. A straightforward consequence of this is that for all ℓ, all its layer ℓ edges must occur consecutively; that is, no transition edges is used between the first and last layer ℓ edge on the path. It follows that all layer ℓ edges occur in the subgraph $F_\ell(u,v)$. If the path uses exactly $\delta_G(t_\ell, t'_\ell)$ of these edges, then they must correspond to a shortest path between t_ℓ and t'_ℓ in $F_\ell(u,v)$. If instead the path uses more than $k(k-1)$ transition edges, or it uses more than $\delta_G(t_\ell, t'_\ell)$ layer ℓ edges for any ℓ, then the path is necessarily longer than the path generated by the above process, and so it is not a shortest path between u and v. Therefore, for all ℓ, any path $\rho_{G'}(u,v)$ contains as a subpath $\rho_{F_\ell(u,v)}(u_\ell, v_\ell)$. □

From the next lemma, we can infer that the properties we obtained in the preprocessing subroutine have been preserved in G', in some strong sense, due to our process of distributing the tokens. In particular, the second and third properties guarantee that despite having a relatively small pair set, most edges in G' are necessary, since the pairs are "well-spread" across the edges of the layers.

Lemma 2.3. *There is a partial assignment of edges in G' to pairs in P' such that, for each layer ℓ, all of the following properties hold: (1) If an edge e is assigned to a pair p, then p requires e. (2) A constant fraction of the layer ℓ edges are assigned. (3) There is a value β such that all paths own $\Theta(\beta)$ layer ℓ edges.*

Proof. To prove the lemma, we describe such assignment of edges to pairs in P' as follows. Consider any individual F-subgraph in layer ℓ. This subgraph is isomorphic to G and therefore, by our preprocessing step, we can assign a constant fraction of the edges of this subgraph to the pairs in P such that each pair owns the exact same number of (required) edges; let β be the number of edges owned by each pair under this assignment. Next, by Lemma 2.1, for each pair $(u,v) \in P$, we will slide $\Theta(1)$ tokens (because k is a constant) from u to v. We can therefore split the edges owned by (u,v) as evenly as possible between these tokens, and each token will own $\Theta(\beta)$ layer ℓ edges: these edges will be *assigned* to the pair p defined by this token.

We now argue that the three properties hold for this assignment. By Lemma 2.2, the pair defined by any such token must pass through u and then v, so any edge required by (u,v) is necessarily required by the pair defined by the token and the first property holds. The second property holds because we have assigned a constant fraction of edges in *every* F-subgraph in layer ℓ, while the third property holds with the β defined above. □

We are now ready to prove our main structural lemma. All of the work we have done so far has built up to the following result:

Lemma 2.4. *There is a constant $\varepsilon > 0$ such that if $H \subset G'$ with $|E(H)| \leq \varepsilon |E(G')|$, then there is a pair $(u,v) \in P'$ such that for all ℓ, there is a layer ℓ edge required by (u,v) that is missing in H.*

The intuition behind the following proof is as follows. First, we note that any way of deleting edges from G' that leaves only a small $\varepsilon \ll 1/k$ fraction of the edges cannot leave more than an εk fraction of the edges at *any* layer ℓ, because there are only k layers and each layer has an equal number of edges. Since we can choose ε to be a small enough constant, and ε depends on k, we can assume that this is the case. Then, if a large constant fraction of the layer ℓ edges are deleted, then a large fraction of the pairs in P' will be missing one of their assigned layer ℓ edges - this will follow from the properties of the assignment in the previous lemma. Finally, since there are only k layers, we argue that there must be a pair that is missing an edge in *every* layer.

Proof. As an intermediate step, we will prove the following statement: for all $d > 0$, there is an $e > 0$ such that if H has only an e fraction as many layer ℓ edges as G' for some fixed ℓ, then at least a $(1-d)$ fraction of the paths in P' are missing a required layer ℓ edge in H.

First, assign the edges of G' to pairs in P' as described in Lemma 2.3. Lemma 2.3 then states that there is a constant c and a value β such that all pairs own between β and $c\beta$ layer ℓ edges. The upper bound of $c\beta$ edges per pair implies that $c\beta|P'| \geq |E(G')|$, so we have $\beta \geq \frac{|E(G')|}{c|P'|}$. Given this constant c, we choose $e = \frac{d}{c}$. Let x be the number of pairs such that H still contains every layer ℓ edge owned by that pair. We then have $x\beta \leq e \cdot |E(G')|$. Combining our two inequalities, we can write:

$$x\frac{|E(G')|}{c|P'} \leq x\beta \leq e \cdot |E(G')| = \frac{d}{c}|E(G')|$$

and so $x \leq d \cdot |P'|$, which is our intermediate step.

To complete the proof: let e be the value that corresponds to $d = \frac{1}{2k}$. We then set $\varepsilon = \frac{e}{k}$. It then must be true that every layer ℓ is missing at least an e fraction of its edges in H, and so a $1-\frac{1}{2k}$ fraction of the pairs are missing a required level ℓ edge. Since there are only k levels, an intersection bound now implies that at least one pair is missing a required edge in every level. ☐

We can then translate this lemma into our error bound.

LEMMA 2.5. *Let $H \subset G'$. Suppose that there is a pair $(u, v) \in P'$ such that for all ℓ, there is a level ℓ edge required by (u, v) that is missing in H. Then $\delta_H(u, v) \geq \delta_G(u, v) + 2k$.*

Proof. From Lemma 2.2, we know that $\delta_G(u, v) = k(k-1) + \sum_{\ell=0}^{k-1} \delta_G(u_\ell, v_\ell)$.

First, suppose that there are layers ℓ, ℓ' such that $\rho_H(u, v)$ travels a level ℓ edge, then a level ℓ' edge, then another level ℓ edge. The path must then travel at least k transition edges that *decrease* the layer value of the node, so it must travel at least $k(k-1) + 2k$ transition edges in total. Then by Lemma 2.2, we have $\delta_H(u, v) \geq k(k-1) + 2k + \sum_{\ell=0}^{k-1} \delta_G(u_\ell, v_\ell)$ which implies the lemma.

It now remains to handle the case in which, for all ℓ, all layer ℓ edges traveled by $\rho_H(u, v)$ are traveled consecutively. Let $H_\ell \subset G$ be the subgraph $F_\ell(u, v)$ with only the edges that remain in H. We now have $\delta_H(u, v) = k(k-1) + \sum_{\ell=0}^{k-1} \delta_{H_\ell}(u_\ell, v_\ell)$.

From Lemma 2.2, we know that the set of level ℓ edges required by (u, v) is precisely the set of edges required by (u_ℓ, v_ℓ) in $F_\ell(u, v)$. We have assumed that at least one such edge is missing from H; therefore, we have $\delta_{H_\ell}(u_\ell, v_\ell) \geq \delta_G(u_\ell, v_\ell) + 1$. However, G is bipartite, and so a path in G that is exactly one longer than the optimal path ends on the wrong side of the bipartite partition; so this implies that $\delta_{H_\ell}(u_\ell, v_\ell) \geq \delta_G(u_\ell, v_\ell) + 2$. We now have $\delta_H(u, v) \geq k(k-1) + \sum_{\ell=0}^{k-1}(\delta_G(u_\ell, v_\ell) + 2) = k(k-1) + 2k + \sum_{\ell=0}^{k-1} \delta_G(u_\ell, v_\ell)$, and the lemma follows. ☐

The proof of Theorem 1.1 is now immediate from Lemmas 2.5 and 2.4.

3 Corollaries of the Error Amplification Theorem

Proof. [of Corollary 1.1] Let $P = S \times S$, and so our subset preserver lower bound is also a pairwise preserver lower bound for G, P. By Theorem 1.1, there is a graph G' that is a lower bound against $+(2k-1)$ pairwise spanners on $\Theta(n^a|S|^b|\mathcal{T}(P)|^{k-1})$ edges and $\Theta(n|\mathcal{T}(P)|^{k-1})$ nodes, along with a new pair set P' with $|P'| = \Theta(|P||\mathcal{T}(P)|^{k-1})$. Note that $\mathcal{T}(P) = S$, and so the density is $|E(G')| = \Theta(n^a|S|^{b+k-1})$, the node count is $n' = \Theta(n|S|^{k-1})$, and the pair set size is $|P'| = \Theta(|S|^{k+1})$.

Let s be a value such that $|S| = n^s$. Let $n' = |V(G')| = \Theta(n^a|S|^{b+k-1})$. We then have: $|E(G')| = \Theta(n'^{(a+sb+sk-s)/(1+sk-s)})$, and $|P'| = \Theta(n'^{(sk+s)/(1+sk-s)})$. It is perhaps easiest to view this as a parametric curve parameterized by $s \in [0, 1]$. The exponent of $|E(G')|$ and the exponent of $|P'|$ are then nonlinear in s, but they have an affine relationship with respect to each other (i.e. one exponent is equal to a constant plus a scalar multiple of the other). We can then compute their relationship by interpolation: set $s = 0$ and for the point $|P'| = \Theta(1), |E(G')| = \Theta(n^a)$, or set $s = 1$ for the point $|P'| = \Theta(n'^{1+1/k}), |E(G')| = \Theta(n'^{(a+b+k-1)/k})$. Interpolating between these two points implies the desired tradeoff. ☐

The proof of Corollary 1.2 follows from Theorem 1.1 in an essentially identical way, and so we omit an explicit proof. The proof of Proposition 1.1 is also quite similar, but contains a few minor details:

Proof. [of Proposition 1.1] Start with a D threshold *preserver* lower bound graph G, which has $\Theta(n^2/D)$ edges, as described in [BCE03]. Let P be the set of pairs in G at distance D or more. The number of terminals in the construction of [BCE03] is $\mathcal{T}(P) = \Theta(n/D)$. Additionally, one can verify from the proof of Theorem 1.1 that if all pairs in P have distance D, then all pairs in P' will have distance more than D. The amplification in the proof of Theorem 1.1, applied to this particular graph, then states: for all positive integers k, there is a graph G' on $\Theta(n \cdot (n/D)^{k-1})$ nodes and $\Theta(n^2/D \cdot (n/D)^{k-1})$ edges and a pair set P' with all pairs at distance D such that G', P' is a lower bound against $+(2k-1)$ pairwise spanners. To prove Proposition 1.1 from this, the calculations now proceed exactly as in the proof of Corollary 1.1 above. □

4 A New Upper Bound for +2 Pairwise Spanners

To obtain a completely tight bound for +2 pairwise spanners we remove a log factor from the upper bound of Kavitha and Varma [KV13].

THEOREM 4.1. *For all graphs G and pair sets P, there is a +2 pairwise spanner of G, P on $O(n|P|^{1/3})$ edges.*

Proof. Let e be a parameter that we will tune later. Let H be a copy of G with all of the nodes and none of the edges; we will build H iteratively as a +2 pairwise spanner of G, P. Consider the pairs $(u, v) \in P$ in some order, and add $\rho_G(u, v)$ to H one by one. Pause if, after adding a path, the following condition ever becomes true: for some node v, there are n^e nodes within distance 1 of v such that these n^e nodes have a total of n^{2e} or more incident edges in H. Once this becomes true, remove all paths from H that were incident on one of these n^e nodes, and then add a complete BFS tree to H centered at v.

Set $n^e = |P|^{1/3}$. Each time we add a BFS tree we remove $\Omega(P^{2/3})$ paths from H; therefore, we can repeat the process only $|P|^{1/3}$ times. Each BFS tree costs n edges, so the total cost from this process is at most $n|P|^{1/3}$. Additionally, we claim that the cost of the paths *not* removed from H is at most $n|P|^{1/3}$. To see this, unmark all nodes and all edges, then repeat the following process. Choose an unmarked node. If it has at least n^e unmarked neighbors, then take this node and any n^e of its unmarked neighbors and mark all these nodes and all their incident edges. Otherwise, mark this node and all edges that connect it to unmarked nodes. Repeat until there are no more unmarked nodes. In the former case, we mark n^e nodes and $< n^{2e}$ edges. In the latter case, we mark 1 node and $< n^e$ edges. It is clear that once all nodes are marked, all edges are marked too. However, we mark at most n^e times as many edges as nodes, so the total number of edges remaining in the graph is at most $n^{1+e} = n|P|^{1/3}$.

Finally, we must prove that every pair (u, v) is spanned within error +2. If $\rho_G(u, v)$ was never removed from the spanner, then it is clearly spanned with error +0. If the pair corresponding to the path was removed from the graph, then there is a node x adjacent to some node $w \in \rho_G(u, v)$ such that we have added a BFS tree from x. We can then apply the triangle inequality:

$$\delta_G(u, x) + \delta_G(x, v) \leq (\delta_G(u, w) + 1) + (\delta_G(w, v) + 1)$$

$$\delta_H(u, x) + \delta_H(x, v) \leq \delta_G(u, v) + 2$$

$$\delta_H(u, v) \leq \delta_G(u, v) + 2$$

□

Acknowledgement. We are grateful to Virginia Vassilevska Williams for encouraging us to think about these problems and providing guidance along the way. We also thank an anonymous reviewer whose comments greatly improved the writing of this paper.

References

[ACIM99] D. Aingworth, C. Chekuri, P. Indyk, and R. Motwani. Fast estimation of diameter and shortest paths (without matrix multiplication). *SIAM J. Comput.*, 28:1167–1181, 1999.

[ADD+93] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9:81–100, 1993.

[Awe85] B. Awerbuch. Complexity of network synchronization. *Journal of the ACM*, pages 32, 804–823, 1985.

[BCE03] B. Bollobás, D. Coppersmith, and M. Elkin. Sparse distance preservers and additive spanners. *Proc. 14th SODA*, pages 414–423, 2003.

[BK06] S. Baswana and T. Kavitha. Faster algorithms for approximate distance oracles and all-pairs small stretch paths. *Proc. 47th FOCS*, pages 591–602, 2006.

[BKMP05] S. Baswana, T. Kavitha, K. Mehlhorn, and S. Pettie. New constructions of (α, β)-spanners and purely additive spanners. *Proc. 16th SODA*, pages 672–681, 2005.

[BKMP10] Surender Baswana, Telikepalli Kavitha, Kurt Mehlhorn, and Seth Pettie. Additive spanners and (α, β)-spanners. *ACM Transactions on Algorithms*, 7(1):5, 2010.

[BS07] S. Baswana and S. Sen. A simple and linear time randomized algorithm for computing sparse spanners in weighted graphs. *Journal of Random Structures and Algorithms 30*, pages 4, 532–563, 2007.

[BW15] G. Bodwin and V. Vassilevska Williams. Very sparse additive spanners and emulators. *Innovations in Theoretical Computer Science (ITCS)*, 2015.

[BW16] G. Bodwin and V. Vassilevska Williams. Better distance preservers and additive spanners. *Also in SODA*, 2016.

[CE06] D. Coppersmith and M. Elkin. Sparse sourcewise and pairwise distance preservers. *SIAM Journal on Discrete Mathematics*, pages 463–501, 2006.

[CGK13] M. Cygan, F. Grandoni, and T. Kavitha. On pairwise spanners. *Symposium on Theoretical Aspects of Computer Science (STACS)*, 2013.

[Che13] Shiri Chechik. New additive spanners. *SODA*, pages 498–512, 2013.

[Cow01] L. J. Cowen. Compact routing with minimum stretch. *Journal of Algorithms*, pages 28, 170–183, 2001.

[CW04] L. J. Cowen and C. G. Wagner. Compact roundtrip routing in directed networks. *Journal of Algorithms*, pages 50, 1, 79–95, 2004.

[DHZ96] Dorit Dor, Shay Halperin, and Uri Zwick. All pairs almost shortest paths. *Proc. 37th FOCS*, pages 452–461, 1996.

[Elk05] M. Elkin. Computing almost shortest paths. *ACM Trans. Algorithms*, pages 1(2):283–323, 2005.

[Elk07] M. Elkin. A near-optimal distributed fully dynamic algorithm for maintaining sparse spanners. *Proc. 26th PODC*, pages 185–194, 2007.

[EP04] M. Elkin and D. Peleg. $(1+\epsilon, \beta)$-spanner constructions for general graphs. *SIAM J. Comput., 33(3):608–631*, 2004.

[EZ06] M. Elkin and J. Zhang. Efficient algorithms for constructing $(1+\epsilon, \beta)$-spanners in the distributed and streaming models. *Distributed Compting 18*, pages 5, 375–385, 2006.

[FPZW04] A. M. Farley, A. Proskurowski, D. Zappala, and K. Windisch. Spanners and message distribution in networks. *Discrete Applied Mathematics*, pages 137(2):159–171, 2004.

[Kav15] T. Kavitha. New pairwise spanners. *Proc. 42nd ICALP*, 2015.

[Knu14] Mathias Bæk Tejs Knudsen. Additive spanners: A simple construction. *Symposium and Workshop on Algorithm Theory (SWAT)*, pages 277–281, 2014.

[KV13] T. Kavitha and N. Varma. Small stretch pairwise spanners. *Proc. 40th ICALP*, 2013.

[LS91] Arthur L Liestman and Thomas C Shermer. Additive spanners for hypercubes. *Parallel Processing Letters*, 1(01):35–42, 1991.

[LS93] Arthur L Liestman and Thomas C Shermer. Additive graph spanners. *Networks*, 23(4):343–363, 1993.

[Par14] M. Parter. Bypassing erdős' girth conjecture: Hybrid stretch and sourcewise spanners. In *Automata, Languages, and Programming - 41st International Colloquium, ICALP 2014, Copenhagen, Denmark, July 8-11, 2014, Proceedings, Part II*, pages 608–619, 2014.

[Pet09] Seth Pettie. Low distortion spanners. *ACM Transactions on Algorithms, 6(1)*, 2009.

[PS89] D. Peleg and A. Schaffer. Graph spanners. *Journal of Graph Theory*, pages 13:99–116, 1989.

[PU89a] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM Journal of Computing*, pages 18, 740–747, 1989.

[PU89b] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *Journal of the ACM*, pages 36(3):510–530, 1989.

[RTZ05] L. Roditty, M. Thorup, and U. Zwick. Deterministic constructions of approximate distance oracles and spanners. *Proc. 32nd ICALP*, pages 261–272, 2005.

[RTZ08] L. Roditty, M. Thorup, and U. Zwick. Roundtrip spanners and roundtrip routing in directed graphs. *ACM Trans. Algorithms*, page 3(4): Article 29, 2008.

[TZ01] M. Thorup and U. Zwick. Compact routing schemes. *Proc. 13th SPAA*, pages 1–10, 2001.

[TZ05] M. Thorup and U. Zwick. Approximate distance oracles. *Journal of the ACM 52*, pages 1, 1–24, 2005.

[TZ06] M. Thorup and U. Zwick. Spanners and emulators with sublinear distance errors. *Proc. 17th SODA*, pages 802–809, 2006.

[Wen91] R. Wenger. Extremal graphs with no c^4's, c^6's, or c^{10}'s. *Journal of Combinatorial Theory*, pages 113–116, 1991.

[Woo06] D. P. Woodruff. Lower bounds for additive spanners, emulators, and more. *Proc. 47th FOCS*, pages 389–398, 2006.

[Woo10] D. P. Woodruff. Additive spanners in nearly quadratic time. *Proc. 37th ICALP*, pages 463–474, 2010.