

Transitive-Closure Spanners

Arnab Bhattacharyya* Elena Grigorescu* Kyomin Jung* Sofya Raskhodnikova†

David P. Woodruff‡

Abstract

We define the notion of a transitive-closure spanner of a directed graph. Given a directed graph $G = (V, E)$ and an integer $k \geq 1$, a k -transitive-closure-spanner (k -TC-spanner) of G is a directed graph $H = (V, E_H)$ that has (1) the same transitive-closure as G and (2) diameter at most k . These spanners were studied implicitly in access control, property testing, and data structures, and properties of these spanners have been rediscovered over the span of 20 years. We bring these areas under the unifying framework of TC-spanners. We abstract the common task implicitly tackled in these diverse applications as the problem of constructing sparse TC-spanners.

We study the approximability of the size of the sparsest k -TC-spanner for a given digraph. Our technical contributions fall into three categories: algorithms for general digraphs, inapproximability results, and structural bounds for a specific graph family which imply an efficient algorithm with a good approximation ratio for that family.

Algorithms. We present two efficient deterministic algorithms that find k -TC-spanners of size approximating the optimum. The first algorithm gives an $\tilde{O}(n^{1-1/k})$ -approximation for $k > 2$. Our method, based on a combination of convex programming and sampling, yields the first sublinear approximation ratios for (1) DIRECTED k -SPANNER, a well-studied generalization of k -TC-SPANNER, and (2) its variants CLIENT/SERVER DIRECTED k -SPANNER, and the k -DIAMETER SPANNING SUBGRAPH. This resolves the main open question of Elkin and Peleg (IPCO, 2001). The second algorithm, specific to the k -TC-spanner problem, gives an $\tilde{O}(n/k^2)$ -approximation. It shows that for $k = \Omega(\sqrt{n})$, our problem has a provably better approximation ratio than DIRECTED k -SPANNER and its variants. This algorithm also resolves an open question of Hesse (SODA, 2003).

Inapproximability. Our main technical contribution is a pair of strong inapproximability results. We resolve the approximability of 2-TC-spanners, showing that it is $\Theta(\log n)$ unless $P = NP$. For constant $k \geq 3$, we prove that the size of the sparsest k -TC-spanner is hard to approximate within $2^{\log^{1-\epsilon} n}$, for any $\epsilon > 0$, unless $NP \subseteq DTIME(n^{\text{polylog } n})$. Our hardness result helps explain the difficulty in designing general efficient solutions for the applications above, and it cannot be improved without resolving a long-standing open question in complexity theory. It uses an involved application of generalized butterfly and broom graphs, as well as noise-resilient transformations of hard problems, which may be of independent interest.

Structural bounds. Finally, we study the size of the sparsest TC-spanner for H -minor-free digraphs, which include planar, bounded genus, and bounded tree-width graphs, explicitly investigated in applications above. We show that every H -minor-free digraph has an efficiently constructible k -TC-spanner of size $\tilde{O}(n)$, which implies an $\tilde{O}(1)$ -approximation algorithm for this family. Furthermore, using our insight that 2-TC-spanners yield property testers, we obtain a monotonicity tester with $O(\log^2 n/\epsilon)$ queries for any poset whose transitive reduction is an H -minor free digraph. This improves and generalizes the previous $\Theta(\sqrt{n} \log n/\epsilon)$ -query tester of Fischer *et al* (STOC, 2002).

*Massachusetts Institute of Technology, USA. Email: {abhattach, elena_g, kmjung}@mit.edu.

†Pennsylvania State University, USA. Email: sofya@cse.psu.edu. Supported by National Science Foundation (NSF/CCF grant number 0729171).

‡IBM Almaden Research Center, USA. Email: dpwoodru@us.ibm.com.

1 Introduction

A *spanner* can be thought of as a sparse backbone of a graph that approximately preserves distances between every pair of vertices. More precisely, a subgraph $H = (V, E_H)$ is a k -*spanner* of $G = (V, E)$ if for every pair of vertices $u, v \in V$, the shortest path distance $d_H(u, v)$ from u to v in H is at most $k \cdot d_G(u, v)$. Since they were introduced by Peleg and Schäffer [40] in the context of distributed computing, spanners for undirected graphs have been extensively studied. The tradeoff between the parameter k , called the *stretch*, and the number of edges in a spanner is relatively well understood: for every $k \geq 1$, any undirected graph on n vertices has a $(2k - 1)$ -spanner with $O(n^{1+1/k})$ edges [6, 39, 52]. This is known to be tight for $k = 1, 2, 3, 5$ and is conjectured to be tight for all k (see, for example a survey by Zwick [55]). Undirected spanners have numerous applications, such as efficient routing [16, 17, 42, 44, 51], simulating synchronized protocols in unsynchronized networks [41], parallel and distributed algorithms for approximating shortest paths [14, 15, 20], and algorithms for distance oracles [9, 52].

In the directed setting, two notions of spanners have been considered in the literature: the direct generalization of the above definition [40] and *roundtrip spanners* [17, 44]. In this paper, we introduce a new definition of directed spanners that captures the notion that a spanner should have a small diameter but preserve the connectivity of the original graph.

Definition 1.1 (TC-spanner). *Given a directed graph $G = (V, E)$ and an integer $k \geq 1$, a k -**transitive-closure-spanner** (k -**TC-spanner**) is a directed graph $H = (V, E_H)$ with the following properties:*

1. E_H is a subset of the edges in the transitive closure of G .
2. For all vertices $u, v \in V$, if $d_G(u, v) < \infty$, then $d_H(u, v) \leq k$.

Notice that a k -TC-spanner of G is just a directed spanner of the transitive-closure of G with stretch k . Nevertheless, a k -TC-spanner is interesting in its own right due to the numerous TC-spanner-specific applications we present in Section 1.3.

One of the focuses of this paper is the study of the computational problem of finding the size of the sparsest k -TC-spanner for a given digraph, referred to as k -TC-SPANNER. It is a special case of the problem of finding the size of the sparsest directed spanner, called DIRECTED k -SPANNER, that has been previously studied. Both problems are NP-hard (see Appendix F).

1.1 Related Work

Thorup [47] considered a special case of TC-spanners of graphs G that have at most twice as many edges as G , and conjectured that for all directed graphs G with n vertices there are such TC-spanners with stretch polylogarithmic in n . He proved his conjecture for planar graphs [48], but later Hesse [33] gave a counterexample to Thorup’s conjecture for general graphs. TC-spanners were also studied for directed trees: implicitly in [5, 8, 12, 18, 54] and explicitly in [49]. For the directed line, [5] (and later, [8]) showed that the size of the sparsest k -TC-spanner is $\Theta(n \cdot \lambda_k(n))$, where $\lambda_k(n)$ is the k^{th} -row inverse Ackermann function. [5, 12, 49] gave the same bounds for rooted directed trees.

Approximability of directed spanner problems. All algorithms for DIRECTED k -SPANNER immediately yield algorithms for k -TC-SPANNER with the same approximation ratio. Kortsarz and Peleg [37] give an $O(\log n)$ -approximation algorithm for DIRECTED-2-SPANNER, and Kortsarz [35] shows that this approximation ratio cannot be improved unless $P=NP$. For $k = 3$, Elkin and Peleg [21] present an $\tilde{O}(n^{2/3})$ -approximation algorithm. Their algorithm is complicated, and the polylog factor hidden in the \tilde{O} notation is not analyzed. For $k \geq 4$, sublinear factor approximation algorithms are known only in the undirected setting [40]. We note that Dodis and Khanna [19] and Chekuri *et al.* [13] study algorithms that might seem relevant to k -TC-SPANNER. In Appendix A.1 we explain why these algorithms do not work for k -TC-SPANNER.

| Setting of k | Implied by previous work | This paper | Notes |
|-----------------------------------------------------|-------------------------------------|-------------------------------------------------|---------------------------------------|
| $k = 2$ | $O(\log n)$ [37] | $\Omega(\log n)$ | |
| constant $k > 2$ | | $\Omega(2^{\log^{1-\epsilon} n})$ | |
| $k = 3$ | $O(n^{2/3} \text{polylog } n)$ [21] | $O((n \log n)^{2/3})$ | applies to DIRECTED k -SPANNER |
| $k > 3$ | $O(n)$ [trivial] | $O((n \log n)^{1-1/k})$ | |
| $k = \Omega\left(\frac{\log n}{\log \log n}\right)$ | $O(n)$ [trivial] | $O\left(\frac{n \log n}{k^2 + k \log n}\right)$ | separation from DIRECTED k -SPANNER |

Table 1: Summary of Results on Approximability of k -TC-SPANNER

For any constant $k > 2$ and $0 < \epsilon < 1$, it is hard to approximate DIRECTED k -SPANNER within a factor of $2^{\log^{1-\epsilon} n}$, assuming $\text{NP} \not\subseteq \text{DTIME}(n^{\text{poly} \log n})$ [21]. Moreover, [24] extend this result to $3 \leq k = O(n^{1-\delta})$ for any $0 < \delta < 1$. Thus, according to Arora and Lund’s classification [34] of NP-hard problems, DIRECTED k -SPANNER is in class III, for $3 \leq k = O(n^{1-\delta})$. Moreover, [24] show that proving that DIRECTED k -SPANNER is in class IV, that is, inapproximable within n^δ for some $0 < \delta < 1$, would resolve a long standing open question in complexity theory, and cause classes III and IV to collapse into a single class.

1.2 Our Contributions

The contributions of this paper are the following: (1) we bring several diverse applications, including property testing, access control and data structures, under the unifying framework of TC-spanners, (2) we obtain strong bounds on the approximability of k -TC-SPANNER and DIRECTED k -SPANNER as well as some well-studied variants of these problems, and (3) we characterize the exact size of TC-spanners and obtain better bounds for the family of H -minor free graphs, which include planar, bounded-treewidth, and bounded genus graphs. Our results on the approximability of k -TC-SPANNER are summarized in Table 1.

Algorithms for k -TC-SPANNER and Related Problems. We present two deterministic polynomial time approximation algorithms for k -TC-SPANNER. Our first algorithm uses a new combination of convex programming and sampling, and gives an $O((n \log n)^{1-1/k})$ -ratio for k -TC-SPANNER. Moreover, our method yields the same approximation ratio for DIRECTED k -SPANNER and its well-studied variants: CLIENT/SERVER DIRECTED k -SPANNER, and k -DIAMETER SPANNING SUBGRAPH (see [22] for definitions). This resolves the open question of finding a sublinear approximation ratio for these problems for $k > 3$, described as a “challenging direction” for research on directed spanners by Elkin and Peleg [23]. Our algorithm for $k = 3$ is arguably simpler than the $O(n^{2/3} \text{polylog } n)$ -approximation algorithm of [23].

Our second algorithm has an $\tilde{O}(n/k^2)$ ratio for k -TC-SPANNER. This demonstrates a separation between k -TC-SPANNER and DIRECTED k -SPANNER: for $k = \sqrt{n}$, it gives $O(\log n)$ -approximation for k -TC-SPANNER while [24, Theorem 6.6] showed that DIRECTED \sqrt{n} -SPANNER is $2^{\log^{1-\epsilon} n}$ -inapproximable. Moreover, Hesse [33] asks for an algorithm to add $O(|G|)$ “shortcuts” to a digraph and reduce its diameter to \sqrt{n} . Our second algorithm returns a \sqrt{n} -TC-spanner of size $O(|G| + \log n)$, answering his question.

Inapproximability of k -TC-SPANNER. We present two results on the hardness of k -TC-SPANNER. First, we prove for $k = 2$ that the $O(\log n)$ ratio of [37] is optimal unless $\text{P}=\text{NP}$. Next, we show that for constant $k > 2$, k -TC-SPANNER is inapproximable within a factor of $2^{\log^{1-\epsilon} n}$, for any $\epsilon > 0$, unless $\text{NP} \subseteq \text{DTIME}(n^{\text{poly} \log n})$. This result is our main technical contribution. Observe that a stronger inapproximability result for $k > 2$ would imply the same inapproximability for DIRECTED- k -SPANNER, and as shown in [24], collapse classes III and IV in Arora and Lund’s classification.

Our $2^{\log^{1-\epsilon} n}$ -hardness matches the known hardness for DIRECTED k -SPANNER. As is the case for DIRECTED k -SPANNER, we start by building a directed graph from a well-known hard problem called MIN-REP, which has the same inapproximability as SYMMETRIC LABEL COVER. However, as illus-

trated in Section 3, all known hard instances for DIRECTED k -SPANNER cannot imply anything better than $\Omega(1)$ -hardness for k -TC-SPANNER. Intuitively, our lower bound is much harder to prove than the one for DIRECTED k -SPANNER since our instance must be transitively-closed, and thus, many more “shortcut” routes between pairs of vertices exist. Our construction uses a novel application of the generalized butterfly and broom graphs, together with several transformations of the MIN-REP problem, which make it *noise-resilient*. We call a MIN-REP instance noise-resilient to indicate that its structure is preserved under small perturbations. The paths in the generalized butterfly are well-structured, which allows us to analyze the many different routes possible in the transitive closure. To realize these ideas, we perform various transformations on a MIN-REP instance to coordinate it with multiple copies of butterflies and brooms.

Structural Results. Finally, we study the minimum k -TC-spanner size for a specific graph family with sparse k -TC-spanners: H -minor-free graphs. A graph H is a *minor* of G if H is a subgraph of a graph obtained from G by a sequence of edge contractions and deletions. For a fixed graph H (e.g., K_5), the family of H -minor-free graphs is a minor-closed family that excludes H . Examples of such families include planar graphs, bounded treewidth graphs, and bounded genus graphs, explicitly studied in applications in Section 1.3. For H -minor-free graphs, we efficiently construct 2-TC-spanners of size $O(n \log^2 n)$, and k -TC-spanners of size $O(n \cdot \log n \cdot \lambda_k(n))$, where $\lambda_k(\cdot)$ is the k -row inverse Ackermann function. The main idea is to use the path separators for undirected H -minor free graphs due to Abraham and Gavoille [1]. However, although the separators are paths, in our digraph they may be the union of many dipaths, and so we cannot efficiently recurse using the sparse k -TC-spanners for the directed line of Alon and Schieber [5]. We observe that these separators satisfy a stronger property than claimed in [1], effectively allowing us to encode the direction of edges in a cost function associated with the separators.

1.3 Applications of TC-spanners

Monotonicity Testing. Monotonicity of functions [4, 10, 18, 25, 27, 28, 30, 32] is one of the most studied properties in property testing [31, 45]. Fischer *et al.* [28] prove that testing monotonicity is equivalent to several other testing problems. Let V_n be a poset of n elements and $G_n = (V_n, E)$ be the relation graph, i.e., the Hasse diagram, for V_n . A function $f : V_n \rightarrow \mathbb{R}$ is called *monotone* if $f(x) \leq f(y)$ for all $(x, y) \in E$. We say f is ϵ -far from monotone if f has to be changed on $\geq \epsilon$ fraction of the domain to become monotone, that is, $\min_{\text{monotone } g} |\{x : f(x) \neq g(x)\}| \geq \epsilon n$. A monotonicity tester on G_n is an algorithm that, given an oracle for a function $f : V_n \rightarrow \mathbb{R}$, passes if f is monotone but fails with probability $\geq \frac{2}{3}$ if f is ϵ -far from monotone. The optimal monotonicity tester for the directed line L_n , consisting of vertices $\{1, 2, \dots, n\}$ and edges $\{(i, i+1) : 1 \leq i \leq n-1\}$, proposed by Dodis *et al.* [18], is based on the sparsest 2-TC-spanner for that graph. Implicit in the proof of Proposition 9 in [18] is a lemma relating the complexity of a monotonicity tester for L_n to the size of a 2-TC-spanner for L_n . We generalize this by observing that a sparse 2-TC-spanner for any partial order graph G_n implies an efficient monotonicity tester on G_n . In Appendix A.2, we prove the following lemma.

Lemma 1.1. *If a directed acyclic graph G_n has a 2-TC-spanner with $s(n)$ edges, then there exists a monotonicity tester on G_n that runs in time $O\left(\frac{s(n)}{\epsilon n}\right)$.*

Therefore, all the 2-TC-spanner constructions described in this paper yield monotonicity testers for functions defined on the corresponding posets. Moreover, for H -minor free graphs, the resulting tester has much better query complexity than the previously known, due to Fischer *et al.* [28]. Indeed, we achieve testers with $O(\log^2 n/\epsilon)$ queries, whereas previous testers required $\Theta(\sqrt{n}/\epsilon)$ queries.

Key Management in an Access Hierarchy. In the problem of key management in an access hierarchy, i.e., access control, there is a partially ordered set (poset) of access classes and a key associated with each class. This is modeled by a directed graph G whose nodes are classes and whose edges indicate an ordering.

A user is entitled to access a certain class and all classes reachable from it. This problem arises in content distribution, operating systems, and project development (see, e.g., the references in [8]). One approach to the access control problem [7, 8, 46] is to associate public information $P(i, j)$ with each edge $(i, j) \in G$ and a secret key k_i with each node i . There is an efficient algorithm A which takes k_i and $P(i, j)$ and generates k_j . However, for each (i, j) in G , it is computationally hard to generate k_j without knowledge of k_i . To obtain a key k_v from a key k_u , algorithm A is run $d_G(u, v)$ times. To speed this up, [8] suggest adding edges to G to increase connectivity. To preserve the access hierarchy of G , new edges must be from the transitive closure of G . The number of edges added corresponds to the space complexity of the scheme, while the shortest-path distances correspond to the time complexity. Implicit in [8] are TC-spanners for directed trees with $k = 3$ and size $O(n \log \log n)$ and also with $k = O(\log \log n)$ and size $O(n)$. Our results for H -minor free graphs extend the known posets for which access control schemes have $O(n \text{ polylog } n)$ storage and $O(1)$ key derivation time. Our approximation algorithms yield sparser k -TC-spanners for general posets.

Partial Products in a Semigroup. Yao [54] and Alon and Schieber [5] study space-efficient data structures for the following problem: Preprocess elements $\{s_1, \dots, s_n\}$ of a semigroup (S, \circ) , such as (\mathbb{R}, \min) , to be able to compute partial products $s_i \circ s_{i+1} \circ \dots \circ s_j$ for all $1 \leq i < j \leq n$ with at most k queries to a small database of pre-computed partial products. This problem reduces to finding a sparsest k -TC-spanner for a directed line L_{n+1} . Chazelle [12] and Alon and Schieber also consider a generalization of the above problem, where the input is an (undirected) tree T with an element s_i of a semigroup associated with each vertex i . The goal is to create a space-efficient data structure that allows one to compute the product of elements associated with all vertices on the path from i to j , for all vertex pairs i, j in T . The generalized problem reduces to finding a sparsest k -TC-spanner for a certain directed tree T' obtained from T . We describe the reduction in Appendix A.3. The same reduction to k -TC-spanners can be used to design space-efficient data structures for any digraph with a unique path between pairs of nodes. Our structural results imply new space-efficient data structures for H -minor free graphs with unique paths, and our approximation algorithms yield more space-efficient data structures for general digraphs with unique paths.

Organization. Section 2 contains an overview of our algorithms. In Section 3, we give an overview of our lower bounds and the techniques involved. Section 4 contains an overview of our bounds for minor-free graphs. We defer the details and proofs of our results to the Appendix. In Appendix A, we discuss the previous work and applications to monotonicity testing and partial products in a semigroup. Appendix B contains our algorithms. In Appendix C, we give our $2^{\log^{1-\epsilon} n}$ -inapproximability for $k > 2$, and in Appendix D we give our $\Omega(\log n)$ -inapproximability for $k = 2$. In Appendix E, we give the proofs of our structural results.

Notation. The *transitive closure* of a graph $G = (V, E)$, denoted $TC(G)$, is defined as the directed graph (V, E') , where $E' = \{(u, v) : u \rightsquigarrow_G v\}$. Vertices u and v are *comparable* if either $(u, v) \in TC(G)$ or $(v, u) \in TC(G)$. The *transitive reduction* of G , denoted $TR(G)$, is a digraph G' with the fewest edges for which $TC(G') = TC(G)$. As shown by Aho *et al* [3], $TR(G)$ can be computed efficiently via a greedy algorithm. For directed acyclic graphs $TR(G)$ is unique, and G is *transitively reduced* if $TR(G) = G$.

For $k \geq 1$, we define: $S_k(G) = \min_H \{|H| : H \text{ is a } k\text{-TC-spanner of } G\}$, and we call H that achieves this minimum a *sparsest k -TC-spanner*. Clearly, $S_k(G) \geq |TR(G)|$. The *Ackermann function* [2] is defined by: $A(1, j) = 2^j$, $A(i+1, 0) = A(i, 1)$, $A(i+1, j+1) = A(i, 2^{A(i+1, j)})$. The inverse Ackermann function is $\alpha(n) = \min\{i : A(i, 1) \geq n\}$ and the i^{th} -row inverse is $\lambda_i(n) = \min\{j : A(i, j) \geq n\}$.

2 Overview of Algorithms for k -TC-SPANNER and Related Problems

Our $O((n \log n)^{1-1/k})$ -approximation for k -TC-SPANNER for arbitrary k is based on a new combination of convex programming and sampling. The technique also achieves an $O((n \log n)^{1-1/k})$ ratio for DIRECTED

k -SPANNER, CLIENT/SERVER DIRECTED k -SPANNER, and k -DIAMETER SPANNING SUBGRAPH. Here we describe the result for DIRECTED k -SPANNER. To achieve the same result for k -TC-SPANNER, it suffices to run the algorithm on the transitive-closure of the input digraph. Missing proofs are in Appendix B.

Theorem 2.1. *For any (not necessarily constant) $k > 2$, there is a deterministic polynomial-time algorithm achieving an $O((n \log n)^{1-1/k})$ -approximation for DIRECTED k -SPANNER.*

We start by formulating the problem as an integer program. We briefly explain the problems with this approach and the ideas required to make it work. One can introduce binary edge variables x_e for each edge e in the transitive closure, and binary path variables y_P for each path P of length $\leq k$ in the transitive closure. One enforces the constraints $y_P \leq x_e$ for each $e \in P$, which allow a path P in the spanner only if all edges along it are present. The final constraint is $\sum_P y_P \geq 1$ for all edges $(u, v) \in G$, where the sum is over paths P of length $\leq k$ from u to v . Finally, one can relax the problem to an LP, and try to round the solution.

The first problem is that the integrality gap is huge, which may be why an LP approach had not been considered before. Indeed, if there are n paths of length at most k between u and v , the LP might assign each of them a value of $\Theta(1/n)$. However, we observe that if there are $r = n^{1-1/k}$ distinct paths from u to v of length $\leq k$, there must be $\geq r^{1/(k-1)}$ distinct vertices w for which $u \rightsquigarrow w \rightsquigarrow v$. Let $BFS(v)$ denote a shortest path tree of edges directed away from v , together with a shortest path tree of edges directed towards v . Then if we sample $\tilde{O}(n/r^{1/(k-1)})$ vertices, and grow $BFS(w)$ of $2(n-1)$ edges around each sample w , we will sample a w for which $u \rightsquigarrow w \rightsquigarrow v$, and the path from u to v along the edges in $BFS(w)$ has length at most k . We take the spanner H to be the union of the outputs of the LP and sampling-based algorithms.

1. $H \leftarrow \emptyset$.
2. For each edge $e \in G$, if $x_e \geq \frac{1/2}{(n \log n)^{1-1/k}}$, $H \leftarrow H \cup \{e\}$.
3. Randomly sample $r = O((n \log n)^{1-1/k})$ vertices $z_1, z_2, \dots, z_r \in G$.
4. $H \leftarrow H \cup (\cup_i BFS(z_i))$. Output H .

With high probability, an edge (u, v) is covered by either the LP relaxation or the sampling.

Lemma 2.2. *With probability at least $1 - 1/n$, H is a k -TC-spanner of G .*

The spanner has at most $r \cdot OPT + \frac{n^2}{r^{1/(k-1)}}$ edges, where OPT is the optimum of the LP. By observing that any spanner must have size $\min(OPT, n-1)$, one can guarantee that this is an $\tilde{O}(n^{1-1/k})$ -approximation. Note that we assume that G is connected, as otherwise we can run the algorithm separately on each component. A more careful analysis gives an $O((n \log n)^{1-1/k})$ -approximation, and a simple greedy algorithm derandomizes the sampling.

Lemma 2.3. $|H| = O((n \log n)^{1-1/k} OPT)$.

The problem with this approach is that the number of variables and the size of each of the constraints grows exponentially with k . We replace the variables y_P with $\min_{e \in P} x_e$, reducing the number of variables to $O(n^2)$. The resulting program is convex, and we use the ellipsoid algorithm with a separation oracle. The oracle, given \vec{x} , just needs to find one pair of vertices (u, v) for which the constraint $\sum_{P: u \rightsquigarrow v} \min_{e \in P} x_e \geq 1$ is violated. It can do this by sorting the coordinates of \vec{x} , and counting the number of u - v paths P for which some particular x_e is the minimum edge variable along P . For this, it iteratively removes edges e from G for which x_e is smallest, and uses matrix multiplication to count the u - v paths that remain in the graph.

Lemma 2.4. *For any k , there exists a separation oracle which runs in time $\text{poly}(n)$.*

Our $\tilde{O}(n/k^2)$ -approximation algorithm, which is specific to k -TC-SPANNER, works by sampling $\tilde{O}(n/k)$ vertices and selectively including $O(n/k)$ edges in the transitive closure adjacent to the samples. We also include the edges of $TR(G)$ in the spanner. A simple greedy algorithm derandomizes the sampling.

Theorem 2.5. *For any k , there exists a deterministic approximation algorithm for the k -TC-SPANNER problem with approximation ratio $O((n \log n)/(k^2 + k \log n))$.*

3 Overview of Hardness Results for k -TC-SPANNER

This section outlines the proof of Theorem 3.1, which is our main technical contribution. Missing details are in Appendix C. At the end we briefly describe the ideas behind the inapproximability result for 2-TC-SPANNER that appears in Appendix D.

Theorem 3.1. *For any fixed $\epsilon \in (0, 1)$, the size of the sparsest k -TC-spanner cannot be approximated to within a factor of $2^{\log^{1-\epsilon} n}$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.*

3.1 The Construction and its Motivation

Since k -TC-SPANNER is a special case of DIRECTED k -SPANNER, which is $\Theta(\log n)$ -inapproximable for $k = 2$ and $2^{\log^{1-\epsilon} n}$ -inapproximable for $k \geq 3$, it is natural to ask whether the hard instances of DIRECTED k -SPANNER from [35, 21, 24] can be used to prove hardness for k -TC-SPANNER. It turns out that all these instances have very small k -TC-spanners. We demonstrate it for the instance used in the proof of $\Omega(\log n)$ -hardness for DIRECTED k -SPANNER, which works via a reduction from SET-COVER.

Let G be a bipartite digraph for SET-COVER with n vertices (“sets”) on the left, n vertices (“elements”) on the right, and edges from left to right. Let I be a set of i new independent vertices, for some value i , and let L be a directed line on $k - 1$ new vertices. Call the first vertex of L the head, and the last vertex the tail. Include directed edges (1) from the tail of L to every set in G , (2) from every vertex of I to the head of L , and (3) from every vertex of I to the sets and the elements of G . Call the constructed digraph G' .

Observe that in G' , all directed edges except those from I to G must be included in the directed k -spanner, as such edges form the unique path between their endpoints. At this point, the only pairs of vertices at distance larger than k are those from a vertex in I to an element of G . Since these vertices are adjacent in G' , there must be a path of length at most k in the spanner. The only possible path is from the vertex in I to a vertex of G . It is easy to see that adding exactly OPT edges from each vertex in I to the sets of G is necessary and sufficient to obtain a spanner, where OPT is the size of the minimum set-cover. By making i sufficiently large, the size of the spanner is easily seen to be $\Theta(i \cdot OPT)$, and thus one can approximate SET-COVER by approximating DIRECTED k -SPANNER, so the problem is $\Omega(\log n)$ -inapproximable.

However, there is a trivial k -TC-spanner for this instance! Indeed, by transitivity we can simply connect the head of L to each of the elements of G . This is a k -TC-spanner of size proportional to the number of vertices in G' . Thus, the best one could hope for with this instance is to show $\Omega(1)$ -hardness for k -TC-SPANNER. For similar reasons, the instance showing $2^{\log^{1-\epsilon} n}$ -inapproximability for DIRECTED k -SPANNER also cannot establish anything beyond $\Omega(1)$ -hardness for k -TC-SPANNER.

In the example above there are many paths to cover (those from I to elements of G), but a few “shortcut” edges cover them all. Ideally, we would have many paths to cover, and each shortcut edge could only cover a single path. Hesse’s digraph requiring a large number of shortcuts to reduce its diameter [33] satisfies the desired condition. His idea was to associate vertices with a subset V of vectors in \mathbb{R}^d such that $(u, v) \in E$ iff $u - v$ is an extreme point of the d -dimensional ball of integer points. By the properties of an extreme point, a shortcut can cover at most one path from a large family of shortest paths.

However, to achieve an inapproximability result, we need better structured graphs. We use *generalized butterflies* defined in [53]. In these digraphs vertices are identified with coordinates $[n^{1/k}]^k \times [k + 1]$, and an edge connects $u = (u_1, \dots, u_k, i)$ to $v = (v_1, \dots, v_k, i + 1)$ iff for all $j \neq i$, $u_j = v_j$. We say a vertex (u_1, \dots, u_k, i) is in *strip* i . It is easy to see that there is a unique shortest path of length k from any u in strip 1 to any v in strip $k + 1$. Moreover, any shortcut is on at most $n^{1-2/k}$ such paths because if it connects a vertex in strip i with a vertex in strip $i + \ell$ (where $\ell \geq 2$) it fixes all but $i - 1$ coordinates of u and all but $k + 1 - (i + \ell)$ coordinates of v . Thus, $\geq n^{1+2/k}$ shortcuts are needed to reduce the diameter to $k - 1$.

Reduction from MIN-REP. To get $2^{\log^{1-\epsilon} n}$ -inapproximability, we reduce from the MIN-REP problem. An (n, r, d, m) -MIN-REP instance is a bipartite graph of maximum degree d in which the left part can be partitioned into sets $\mathcal{A}_1, \dots, \mathcal{A}_r$ and the right part into sets $\mathcal{B}_1, \dots, \mathcal{B}_r$, so that $|\mathcal{A}_i| = |\mathcal{B}_i| = n/r$ for all $i \in [r]$. To describe the last parameter m , call a vertex *isolated* if its degree is 0, and *non-isolated* otherwise. Let $m(\mathcal{A}_i)$ be the inverse of the fraction of non-isolated vertices in \mathcal{A}_i . Then m is the minimum such $m(\mathcal{A}_i)$. Define the *supergraph* to have nodes $\mathcal{A}_1, \dots, \mathcal{A}_r, \mathcal{B}_1, \dots, \mathcal{B}_r$, with a *superedge* $(\mathcal{A}_i, \mathcal{B}_j)$ iff there is a node in \mathcal{A}_i adjacent to a node in \mathcal{B}_j . A *rep-cover* is a vertex set S in the graph such that whenever $(\mathcal{A}_i, \mathcal{B}_j)$ is an edge in the supergraph, there is an edge between some $u, v \in S$ with $u \in \mathcal{A}_i$ and $v \in \mathcal{B}_j$. A solution to MIN-REP is a smallest rep-cover, and its size is denoted by OPT. The problem is $2^{\log^{1-\epsilon} n}$ -inapproximable [21].

As a first attempt, we construct a graph of diameter $k + 2$ as follows. We attach a disjoint copy of a generalized butterfly of diameter $k - 1$ to each \mathcal{A}_i in the MIN-REP instance graph; that is, we identify the vertices in \mathcal{A}_i with the last strip of the butterfly. We call the vertices in the butterfly at distance x from \mathcal{A}_i the x -th *shadow* of \mathcal{A}_i . Next, for each \mathcal{B}_j , we attach what we call a *broom*. This is a 3-layer graph, where the two leftmost layers form a bipartite clique, and the right layer consists of degree-1 nodes, called *broomsticks*, attached to vertices in the middle layer. Each vertex in the middle layer has the same number of broomsticks attached to it. Each \mathcal{B}_j is identified with the left layer of a disjoint broom. All edges are directed from the shadows of the \mathcal{A}_i towards the broomsticks (left to right). Call the resulting digraph G .

We would like to argue that the minimum k -TC-spanner H of G is formed as follows. Let S be a minimum rep-cover of the underlying MIN-REP instance. For each $s \in S$, if s is in an \mathcal{A}_i , include all shortcuts from the 2-shadow of \mathcal{A}_i to s which are in the transitive closure of G . Otherwise (s is in a \mathcal{B}_j), include all shortcuts from s to the broomsticks of \mathcal{B}_j . By balancing the number of broomsticks with the size of 2-shadows, one can show H has size $|S|f(n, k)$, where $f(n, k)$ is an easily computable function. Since S is a rep-cover, H is a k -TC-spanner. If H were optimal, then approximating its size within some factor would approximate MIN-REP within the same factor.

It turns out that H is not optimal, and so our first attempt does not work. However, by modifying G via the transformations below, and by looking at a related k -TC-spanner H of the modified G , we can show that any k -TC-spanner must have size $\Omega(|H|/\log n)$ for constant k . Since MIN-REP is $2^{\log^{1-\epsilon} n}$ -inapproximable, this still gives $2^{\log^{1-\epsilon} n}$ -hardness.

To prove this, we need to argue that most vertices v in the k -shadows do not “benefit” from traversing other shortcuts to reach the broomsticks. This requires a classification of all alternative routes from such v to broomsticks. Given that v is in a generalized butterfly, these routes are well-understood. However, for a generic MIN-REP instance, most of these routes do indeed lead to a much smaller k -TC-spanner!

To rule out the alternative routes, we ensure that OPT and the four parameters of the MIN-REP instance each lie in a narrow range. In Theorem 3.2, we prove that MIN-REP with the required parameter restrictions is inapproximable by giving a reduction from an unrestricted MIN-REP instance. It works by carefully interleaving the following five operations on a “base” MIN-REP instance with unrestricted parameters: (1) disjoint copies, (2) dummy vertices inside clusters, (3) blowup inside clusters with matching supergraph, (4) blowup inside clusters with complete supergraph, and (5) tensoring. Each operation increases one or several parameters by a prespecified factor, and together they give us five degrees of freedom to control the range of OPT and the four parameters of MIN-REP.

Theorem 3.2 (Noise-Resilient MIN-REP is hard). *Fix any $\kappa \in (0, 1)$ and $R, D, M, F \in (0, 1 - \kappa)$ satisfying $F \in (R, 2R)$ and $D + M + F < 1$. Noise-Resilient MIN-REP is a family of (n, r, d, m) -MIN-REP instances with $r \in [n^R, n^{R+\kappa}]$, $d \in [n^D, n^{D+\kappa}]$, $m \in [n^M, n^{M+\kappa}]$, and $OPT \in [n^F, n^{F+\kappa}]$. This problem is $2^{\log^{1-\epsilon} n}$ -inapproximable for all $\epsilon \in (0, 1)$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.*

The variant of MIN-REP in Theorem 3.2 is called “noise-resilient” because even if many vertices in the sets \mathcal{A}_i and \mathcal{B}_j are adversarially deleted in an instance of this problem, the minimum rep-cover does




Figure 1: The TC-spanner instance \mathcal{G} , and an example of a broom.

not shrink significantly. This property helps us rule out many alternative routes in the TC-spanner, though we will need to change our graph G . Our reduction from noise-resilient MIN-REP to k -TC-SPANNER for $k > 2$ consists of two steps: first we produce a *specialized* MIN-REP instance \mathcal{I} from an arbitrary instance \mathcal{I}_0 of noise-resilient MIN-REP, and then we construct a k -TC-SPANNER instance \mathcal{G} by carefully adjoining generalized butterfly graphs on the left and broom graphs on the right of \mathcal{I} .

From Noise-resilient MIN-REP to Specialized MIN-REP. Set $\delta = \frac{k-1}{k-\frac{1}{4}}$, $\eta = \frac{\delta}{2(4k-4)(4k-2)}$, and $\zeta = \delta \left(\frac{4k-5}{4k-4} + \frac{1}{4k-2} \right)$. Let κ be a sufficiently small positive constant which will be chosen in the course of the proof. We start from an (n_0, r_0, d_0, m_0) -instance \mathcal{I}_0 of noise-resilient MIN-REP with optimum OPT_0 , where $n_0 = n^\delta$, $r_0 \in [n^{\delta/2}, n^{\delta/2+\kappa}]$, $d_0 \in [n^\eta, n^{\eta+\kappa}]$, $m_0 \in [n^{2\eta}, n^{2\eta+\kappa}]$, and $OPT_0 \in [n^\zeta, n^{\zeta+\kappa}]$. By instantiating Theorem 3.2 with $R = \frac{1}{2}$, $D = \frac{\eta}{\delta}$, $M = \frac{2\eta}{\delta}$, $F = \frac{\zeta}{\delta}$ and κ , we obtain that the (n_0, r_0, d_0, m_0) -MIN-REP problem is $2^{\log^{1-\epsilon} n}$ -inapproximable unless $NP \subseteq DTIME(n^{\text{polylog } n})$. The conditions on the parameters in Theorem 3.2 are satisfied because $\zeta \in (\frac{\delta}{2}, \delta)$ and $\eta + 2\eta + \zeta < \delta$.

We transform \mathcal{I}_0 to a specialized (n, r, d, m) -MIN-REP instance \mathcal{I} by applying on \mathcal{I}_0 the transformation T_4 defined in the proof of Theorem 3.2 in Appendix C. More precisely, set $\mathcal{I} = T_4(\mathcal{I}_0, n^{1-\delta})$. By definition of T_4 , graph \mathcal{I} has n vertices, $r = r_0$, $d = d_0 n^{1-\delta}$ and $m = m_0$. The transformation results in a bipartite graph \mathcal{I} with nodes partitioned into clusters $\mathcal{A}_1, \dots, \mathcal{A}_r$ on the left, and $\mathcal{B}_1, \dots, \mathcal{B}_r$ on the right. Each \mathcal{A}_i and \mathcal{B}_j is a union of $n^{1-\delta}$ groups $A_{i,s}$ and $B_{j,s}$, respectively, with $s \in [n^{1-\delta}]$. Each group $A_{i,s}$ and $B_{j,s}$, for $i, j \in [r]$, $s \in [n^{1-\delta}]$, is a copy of \mathcal{A}_i and, respectively, \mathcal{B}_j , from the original instance \mathcal{I}_0 . For each edge (u, v) with $u \in \mathcal{A}_i$ and $v \in \mathcal{B}_j$ of \mathcal{I}_0 , graph \mathcal{I} has edges between the copy of u in A_{i,k_1} and the copy of v in B_{j,k_2} , for all $k_1, k_2 \in [n^{1-\delta}]$. This completes the description of the specialized MIN-REP instance \mathcal{I} .

From Specialized MIN-REP to k -TC-SPANNER. From \mathcal{I} , we construct a graph \mathcal{G} of diameter $k+2$ as follows. We first attach a disjoint generalized butterfly of diameter $k-1$, denoted $BF(A_{i,s})$, to each group $A_{i,s}$ in \mathcal{I} , for all $i \in [r]$, $s \in [n^{1-\delta}]$. That is, we identify vertices in $A_{i,s}$ with the last strip of $BF(A_{i,s})$ in the way discussed below. Denote by $BF(\mathcal{A}_i) = \cup_s BF(A_{i,s})$ the set of all the vertices attached in this manner to the cluster \mathcal{A}_i . Let $BF^j(A_{i,s})$ be the vertices in strip j of the butterfly $BF(A_{i,s})$, where $BF^k(A_{i,s}) = A_{i,s}$, and let $BF^j(\mathcal{A}_i) = \cup_s BF^j(A_{i,s})$. We call the vertices in the butterfly $BF(A_{i,s})$ at distance x from $A_{i,s}$ the x -th shadow of $A_{i,s}$. Call the in-degree as well as out-degree of the vertices in the butterflies $d_* \stackrel{\text{def}}{=} \left(\frac{n^\delta}{r} \right)^{\frac{1}{k-1}}$.

Next, for each $\mathcal{B}_{i,s}$, we attach a broom, denoted $BR(B_{i,s})$. More specifically, each vertex in $B_{i,s}$ is connected to the vertices of a set $BR^{k+2}(B_{i,s})$ of size d_* , and each vertex $v \in BR^{k+2}(B_{i,s})$ is connected to a disjoint set of nodes, called broomsticks, of size d_* . Let $BR^{k+3}(B_{i,s})$ be the set of broomsticks adjacent to $BR^{k+2}(B_{i,s})$. Let $BR^{k+2}(\mathcal{B}_i) = \cup_s BR^{k+2}(B_{i,s})$ and $BR^{k+3}(\mathcal{B}_i) = \cup_s BR^{k+3}(B_{i,s})$. Identify layer V_j with $\cup_{i,s} BF^j(A_{i,s})$ for $j \in [k]$, layer V_{k+1} with $\cup_{i,s} B_{i,s}$, and layer V_j with $\cup_i BR^j(\mathcal{B}_i)$ for $j \in \{k+2, k+3\}$. Direct all the edges from V_i to V_{i+1} . See Figure 1.

Attaching butterflies. Recall that we identify vertices in $A_{i,s}$ with the last strip $BF^k(A_{i,s})$ of a disjoint butterfly, for all $i \in [r]$, $s \in [n^{1-\delta}]$. The mapping from $A_{i,s}$ to $BF^k(A_{i,s})$ is constructed in Appendix C.

Here we explain the requirements we impose on the mapping. Recall that each group $A_{i,s}$ has $\leq \frac{n^\delta}{rm}$ non-isolated vertices. For our analysis, each vertex in $BF^{k-1}(A_{i,s})$ must be adjacent to $\leq \frac{d_*}{m}$ non-isolated vertices in $A_{i,s}$. The isolated vertices help us control the number of routes with shortcut edges from the x -shadows to the $(x-2)$ -shadows, for some $x > 2$, since connecting vertices in the 1-shadow to many isolated vertices decreases the number of comparable pairs in the first and last layers of \mathcal{G} connected by a path containing such a shortcut edge.

A sparse TC-spanner \mathcal{H} for the k -TC SPANNER instance \mathcal{G} . Let S_0 be a smallest rep-cover of \mathcal{I}_0 of size OPT . Recall that each \mathcal{A}_i and \mathcal{B}_j is replicated $n^{1-\delta}$ times in \mathcal{I} . Let S be the set of all replicas in \mathcal{I} of vertices in S_0 . Consider a k -TC-spanner \mathcal{H} of \mathcal{G} that contains shortcuts from the nodes in layer V_{k-2} to their descendants in $S \cap V_k$, and from the nodes in $S \cap V_{k+1}$ to their descendants in V_{k+3} . The Rep-cover Spanner Lemma (Lemma C.2) shows that $|\mathcal{H}| = O(OPT n^{1-\delta} (\frac{n^\delta}{r})^{\frac{2}{k-1}})$.

3.2 Path Analysis and Rerandomization

The next lemma shows that the k -TC-spanner \mathcal{H} defined above and analyzed in Lemma C.2 is nearly optimal.

Lemma 3.3. *Any k -TC-spanner \mathcal{K} of \mathcal{G} has $|\mathcal{K}| = \Omega\left(OPT n^{1-\delta} \left(\frac{n^\delta}{r}\right)^{\frac{2}{k-1}} / \log n\right)$.*

We introduce a bit of notation. A k -TC-spanner for $\mathcal{G} = V_1 \cup V_2 \cup \dots \cup V_{k+3}$ is built by adding shortcut edges (u, v) between comparable u and v , where $u \in V_i, v \in V_{i+\ell}$ and $\ell \geq 2$. For given ℓ, i , we classify such a shortcut edge as *type $\ell\&i$* . Since \mathcal{G} has diameter $k+2$, a k -TC-spanner for \mathcal{G} remains a k -TC-spanner when a type $\ell\&i$ edge (u, v) with $\ell \geq 4$ is replaced by a type $3\&i$ edge (u, v') , where v' is a predecessor of v . Therefore, it is enough to consider k -TC-spanners with shortcut edges only of types $2\&i$ for $1 \leq i \leq k+1$ and $3\&i$ for $1 \leq i \leq k$. Say a path π from V_1 to V_{k+3} is of *type $(\ell\&i)$* if it uses an edge of type $\ell\&i$ (with $\ell \in \{2, 3\}$), and π is of *type $(2\&i, 2\&j)$* if it uses edges of types $2\&i$ and $2\&j$, $i < j$. Notice that the k -TC-spanner constructed in Lemma C.2 contains only edges of type $2\&(k-2)$ and $2\&(k+1)$.

Proof of Lemma 3.3. Given a k -TC-spanner \mathcal{K} of \mathcal{G} with $o\left(\frac{n^{1-\delta}d_*^2}{\log n}\right)OPT$ edges, we show that we can construct a MIN-REP cover for \mathcal{I} of size $o(OPT)$, which is a contradiction (recall that $d_* = (\frac{n^\delta}{r})^{\frac{1}{k-1}}$). We will accomplish this by a series of transformations which modify \mathcal{K} into a k -TC-spanner that uses only shortcut edges of the form $2\&(k-2)$ and $2\&(k+1)$. The process increases the size of the k -TC-spanner only by a logarithmic factor. Finally, we show that from the modified k -TC-spanner, one can extract a MIN-REP cover of size $o(OPT)$ for \mathcal{I} , the desired contradiction.

We call a superedge $(\mathcal{A}_i, \mathcal{B}_j)$, where $i, j \in [r]$, *deletable with respect to \mathcal{K}* if at least $1/4$ of the vertex pairs $(u, v) \in BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j)$ have a path between them in \mathcal{K} of length at most k and of type other than $(2\&(k-2), 2\&(k+1))$. Our first step is to show that such cluster pairs can be essentially ignored.

Lemma 3.4 (Path Analysis Lemma). *The number of deletable superedges with respect to \mathcal{K} is $o(OPT)$.*

Proof Sketch. We call a path *canonical* if it contains shortcut edges of types both $2\&(k-2)$ and $2\&(k+1)$; otherwise, a path is *alternative*. Observe that every alternative path contains a shortcut edge from one of the following three categories: (1) edges that connect vertices in V_i and V_j , where $i \leq k$ and $j \geq k+1$; (2) edges of type $3\&i$ where $i \leq k-3$; (3) edges of type $2\&i$ where $i \leq k-3$. Let S_B be the set of all shortcut edge types included in the three cases. By analyzing the three cases separately, we show that for any $S \in S_B$, the number of superedges $(\mathcal{A}_i, \mathcal{B}_j)$, $(i, j) \in [r]^2$, such that at least a $\frac{1}{4|S_B|}$ fraction of pairs $(u, v) \in BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j)$ have an alternative path containing a shortcut of type S , is $o(OPT)$. Then by a union bound over $S \in S_B$, we prove the lemma. In the analysis for the first type, we use the fact that the degree of each non-isolated vertex of V_k is at least $n^{1-\delta}$ which is bigger than d_* . When S is of the second

type, we need the facts that the out-degree of each vertex in V_k is at most $d_0 n^{1-\delta}$ and that $n^\eta = o(d_*)$. For the third case, we use the facts that for any vertex v in V_{k-1} the number of non-isolated vertices in V_k that v is connected to is at most $\frac{d_*}{m}$, and that $n^\eta = o(n^{2\eta})$. \square

Next, form the graph \mathcal{G}' from \mathcal{G} by deleting all edges of \mathcal{G} connecting \mathcal{A}_i to \mathcal{B}_j , for all the deletable superedges $(\mathcal{A}_i, \mathcal{B}_j)$ with respect to \mathcal{K} . Similarly, obtain a graph \mathcal{K}' from \mathcal{K} as follows: for all deletable superedges $(\mathcal{A}_i, \mathcal{B}_j)$ with respect to \mathcal{K} , delete all edges of \mathcal{K} connecting \mathcal{A}_i to \mathcal{B}_j , and also delete all shortcuts in \mathcal{K} of types other than $2\&(k-2)$ and $2\&(k+1)$. Note that for any cluster pair $(\mathcal{A}_i, \mathcal{B}_j)$ of \mathcal{G}' , either there are no edges between vertices in \mathcal{A}_i and \mathcal{B}_j or at least $\frac{3}{4}$ of the pairs in $BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j)$ are connected by a canonical path. Also define a MIN-REP instance \mathcal{I}' from \mathcal{I} by deleting all edges in \mathcal{I} corresponding to all the deletable superedges with respect to \mathcal{K} .

For $\mu \in [0, 1]$, we say a subgraph of $TC(\mathcal{G})$ is a μ -good k -TC-spanner for \mathcal{G} if for every $(i, j) \in [r]^2$ such that \mathcal{A}_i and \mathcal{B}_j are comparable in \mathcal{G} , at least a μ fraction of pairs $(u, v) \in BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j)$ are connected by canonical paths in the subgraph. E.g., the graph \mathcal{K}' is a $\frac{3}{4}$ -good k -TC-spanner for \mathcal{G} .

Lemma 3.5 (Rerandomization Lemma). *If a $\frac{3}{4}$ -good k -TC-spanner \mathcal{K}' for \mathcal{G}' is given, then there exists \mathcal{K}'' , a 1-good k -TC-spanner for \mathcal{G}' , such that $|\mathcal{K}''| \leq O(|\mathcal{K}'| \cdot \log n)$.*

Proof Sketch. To construct \mathcal{K}'' from \mathcal{K}' , we let \mathcal{K}'' be the union of $O(\log n)$ random transformations of the edges of \mathcal{K}' . Each transformation Π_r will keep the edges of \mathcal{G}' invariant but move the shortcut edges. Thus, when we let $\mathcal{K}'' = \cup_{r=1}^{O(\log n)} \Pi_r(\mathcal{K}')$, the edges of \mathcal{K}'' are still a subset of the edges in $TC(\mathcal{G}')$. The goal of the random transformations is to ensure that in $\Pi_r(\mathcal{K}')$, with a constant probability, each vertex in $BF^1(\mathcal{A}_i)$ can reach a vertex in V_{k-2} incident to a shortcut edge, and each vertex in $BR^{k+3}(\mathcal{B}_j)$ is incident to a shortcut edge from V_{k+1} . We achieve this by randomly permuting the groups inside the clusters \mathcal{A}_i and \mathcal{B}_j and by randomly permuting the edges of the butterfly and broom graphs attached to \mathcal{A}_i and \mathcal{B}_j . After these random transformations, any two vertices u and v in $BF^1(\mathcal{A}_i)$ and $BR^{k+3}(\mathcal{B}_j)$ are connected by a canonical path with probability at least $\frac{1}{16}$. Hence, \mathcal{K}'' has such a path between them with probability $1 - \frac{1}{\text{poly}(n)}$. The union bound over all possible (u, v) and (i, j) shows that the desired \mathcal{K}'' with the claimed size exists. \square

Now that the k -TC-spanner is 1-good, it is easier to reason about rep-covers of the underlying MIN-REP instance. Recall that \mathcal{G}' has $n^{1-\delta}$ copies of MIN-REP instance \mathcal{I}' embedded in it. Moreover, many pairs of vertices in layers V_1 and V_{k+3} rely on each instance to connect. We partition the shortcut edges of \mathcal{K}'' into $n^{1-\delta} d_*^2$ parts, according to which groups of vertex pairs in $V_1 \times V_{k+3}$ they can help to connect. By averaging, one of the parts has $o(OPT)$ shortcut edges, and can be used to extract a rep-cover of \mathcal{I}' of size $o(OPT)$. By including two vertices for each of the $o(OPT)$ deleted superedges, we obtain a rep-cover for \mathcal{I} of size $o(OPT)$. This is a contradiction.

Lemma 3.6 (Rep-cover Extraction Lemma). *Given \mathcal{K}'' , a 1-good k -TC-spanner for \mathcal{G}' , of size $o(OPT \cdot n^{1-\delta} \cdot d_*^2)$, there exists a MIN-REP cover of \mathcal{I} of size $o(OPT)$.* \square

Our $\Omega(\log n)$ -inapproximability for 2-TC-SPANNER, described in Appendix D, is based on a reduction from SET-COVER instead of MIN-REP. Our hard instance is a generalized butterfly of diameter 2 attached to an instance of transformed SET-COVER. We identify strip 3 of the butterfly with the sets in the instance, and using ideas similar to our proof for $k > 2$ for ruling out alternative routes, show that up to a constant factor, the optimal 2-TC-spanner contains only shortcuts from strip 1 to a minimum set-cover in strip 3.

4 Overview of Structural Results

In [28], the authors implicitly give 2-TC-spanners for planar digraphs of size $O(n^{3/2} \log n)$ using Lipton-Tarjan separators. For planar digraphs, our first idea is to instead use Thorup's planar separators [50] in

conjunction with the efficient k -TC-spanners for the directed line of Alon and Schieber [5] to recursively construct k -TC-spanners of size $O(n \log^2 n)$. More generally, for H -minor-free graphs, using an idea in [50], we take an arbitrary rooted spanning tree T of the digraph G and use it to partition G into a union of edge-disjoint digraphs so that in each part G_i , if one undirects the edges of G_i , any undirected root path of T restricted to G_i is the union of at most two dipaths. Next, instead of Thorup’s planar separators, we use a result of Abraham and Gavoille [1] that provides a “path separator” for undirected H -minor-free graphs.

However, the Abraham-Gavoille separators cannot be directly applied, since they do not provide enough flexibility in the structure of the separators. That is, these separators consist of a sequence of unions of minimum cost paths, where the cost function on the edges is arbitrary but specified in advance. We, however, need to adaptively change the cost function during the construction of the separator. Indeed, in the outermost level of recursion we need the path separator to lie on T , as otherwise the path separator may be the union of $\Omega(n)$ dipaths in the underlying digraph, and therefore we cannot use the efficient k -TC-spanner of Alon and Schieber [5] for the directed line in order to efficiently recurse. Thus, we specify the cost of an edge in T to be 1, while outside of T it is ∞ . However, when we partition G into subgraphs in the recursion, it may be that two vertices in the same subgraph no longer have a path contained in T . Since the cost function is fixed and the cost of any path between these two vertices is now ∞ , a path separator in the recursive step need not be contained in T , and so it may not be the union of a small number of dipaths. Thus, we again cannot efficiently recurse. If, however, we could change the cost function in the recursive step, we could define a new rooted tree in each subgraph and base our cost function on that. We observe that the proof of the Abraham-Gavoille separators can be used to show that their path separators satisfy this stronger property.

Theorem 4.1. *If G is an H -minor-free graph, then it has a 2-TC-spanner of size $O(n \log^2 n)$ and, more generally, a k -TC-spanner of size $O(n \cdot \log n \cdot \lambda_k(n))$ where $\lambda_k(\cdot)$ is the k -row inverse Ackermann function.*

Acknowledgments. We would like to thank Michael Elkin, Vitaly Feldman, Cyril Gavoille, Piotr Indyk, T.S. Jayram, Elad Hazan, Ronitt Rubinfeld, and Adam Smith for helpful discussions.

References

- [1] I. Abraham and C. Gavoille. Object location using path separators. In *PODC*, pages 188–197, 2006.
- [2] W. Ackermann. Zum Hilbertschen aufbau der reellen zahlen. *Math. Ann.*, 99:118–133, 1928.
- [3] A. V. Aho, M. R. Garey, and J. D. Ullman. The transitive reduction of a directed graph. *SIAM J. Comput.*, 1(2):131–137, 1972.
- [4] N. Ailon and B. Chazelle. Information theory in property testing and monotonicity testing in higher dimension. *Inf. Comput.*, 204(11):1704–1717, 2006.
- [5] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. Technical Report 71/87, Tel-Aviv University, 1987.
- [6] I. Althöfer, G. Das, D. Dobkin, D. Joseph, and J. Soares. On sparse spanners of weighted graphs. *Discrete & Computational Geometry*, 9(1):81–100, 1993.
- [7] M. Atallah, M. Blanton, and K. Frikken. Key management for non-tree access hierarchies. In *ACM Symposium on Access Control Models and Technologies (SACMAT’06)*, pages 11–18, June 2006.
- [8] M. J. Atallah, K. B. Frikken, and M. Blanton. Dynamic and efficient key management for access hierarchies. In *ACM Conference on Computer and Communications Security*, pages 190–202, 2005.

- [9] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in $\tilde{O}(n^2)$ time. In *SODA '04: Proceedings of the fifteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 271–280, Philadelphia, PA, USA, 2004. Society for Industrial and Applied Mathematics.
- [10] T. Batu, R. Rubinfeld, and P. White. Fast approximate PCPs for multidimensional bin-packing problems. In *Proceedings of RANDOM*, pages 245–256, 1999.
- [11] P. Berman and M. Karpinski. On some tighter inapproximability results. Technical Report 85193-CS, DIMACS, 1998.
- [12] B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.
- [13] C. Chekuri, G. Even, A. Gupta, and D. Segev. Set connectivity problems in undirected graphs and the directed steiner network problem. In *SODA*, 2008.
- [14] E. Cohen. Fast algorithms for constructing t-spanners and paths with stretch t. *SIAM J. Comput.*, 28(1):210–236, 1998.
- [15] E. Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *J. ACM*, 47(1):132–166, 2000.
- [16] L. Cowen. Compact routing with minimum stretch. *J. Algorithms*, 38(1):170–183, 2001.
- [17] L. Cowen and C. G. Wagner. Compact roundtrip routing in directed networks. *J. Algorithms*, 50(1):79–95, 2004.
- [18] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In D. S. Hochbaum, K. Jansen, J. D. P. Rolim, and A. Sinclair, editors, *RANDOM-APPROX*, volume 1671 of *Lecture Notes in Computer Science*, pages 97–108. Springer, 1999.
- [19] Y. Dodis and S. Khanna. Designing networks with bounded pairwise distance. In *STOC*, pages 750–759, 1999.
- [20] M. Elkin. Computing almost shortest paths. In *PODC*, pages 53–62, 2001.
- [21] M. Elkin and D. Peleg. Strong inapproximability of the basic k -spanner problem. In *Proceedings of the 27th International Colloquium on Automata, Languages and Programming*, 1999.
- [22] M. Elkin and D. Peleg. The client-server 2-spanner problem with applications to network design. In F. Comellas, J. Fàbrega, and P. Fraigniaud, editors, *SIROCCO*, volume 8 of *Proceedings in Informatics*, pages 117–132. Carleton Scientific, 2001.
- [23] M. Elkin and D. Peleg. Approximating k -spanner problems for $k \geq 2$. *Theor. Comput. Sci.*, 337(1-3):249–277, 2005.
- [24] M. Elkin and D. Peleg. The hardness of approximating spanner problems. *Theory of Computing System*, 41(4):691–729, 2007.
- [25] F. Ergun, S. Kannan, S. R. Kumar, R. Rubinfeld, and M. Viswanathan. Spot-checkers. *JCSS*, 60(3):717–751, 2000.

- [26] U. Feige. A threshold of $\ln n$ for approximating set cover (preliminary version). In *STOC*, pages 314–318, 1996.
- [27] E. Fischer. On the strength of comparisons in property testing. *Inf. Comput.*, 189(1):107–116, 2004.
- [28] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. Monotonicity testing over general poset domains. In *STOC*, pages 474–483, 2002.
- [29] C. Gavoille. Personal Communication, November 2007.
- [30] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.
- [31] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *JACM*, 45(4):653–750, 1998.
- [32] S. Halevy and E. Kushilevitz. Testing monotonicity over graph products. In J. Díaz, J. Karhumäki, A. Lepistö, and D. Sannella, editors, *ICALP*, volume 3142 of *Lecture Notes in Computer Science*, pages 721–732. Springer, 2004.
- [33] W. Hesse. Directed graphs requiring large numbers of shortcuts. In *SODA*, pages 665–669, 2003.
- [34] D. Hochbaum, editor. *Approximation Algorithms for NP-hard Problems*. PWS Publishing Company, Boston, 1997.
- [35] G. Kortsarz. On the hardness of approximating spanners. In *Proceedings of the 1st workshop on Approximation Algorithms for Combinatorial Optimization problems*, pages 135–146. Springer-Verlag, New York, Berlin, 1998.
- [36] G. Kortsarz. On the hardness of approximating spanners. *Algorithmica*, 30(3):432–450, 2001.
- [37] G. Kortsarz and D. Peleg. Generating sparse 2-spanners. *J. Algorithms*, 17:222–236, 1994.
- [38] C. Lund and M. Yannakakis. On the hardness of approximating minimization problems. In *STOC*, pages 286–293, 1993.
- [39] D. Peleg. *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 2000.
- [40] D. Peleg and A. A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.
- [41] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.
- [42] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *J. ACM*, 36(3):510–530, 1989.
- [43] R. Raz and S. Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability pcp characterization of np. In *STOC*, pages 475–484, 1997.
- [44] L. Roditty, M. Thorup, and U. Zwick. Roundtrip spanners and roundtrip routing in directed graphs. In *SODA*, pages 844–851, 2002.
- [45] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

- [46] A. D. Santis, A. L. Ferrara, and B. Masucci. Efficient provably-secure hierarchical key assignment schemes. In *MFCS*, pages 371–382, 2007.
- [47] M. Thorup. On shortcutting digraphs. In *WG '92: Proceedings of the 18th International Workshop on Graph-Theoretic Concepts in Computer Science*, pages 205–211, London, UK, 1993. Springer-Verlag.
- [48] M. Thorup. Shortcutting planar digraphs. *Combinatorics, Probability & Computing*, 4:287–315, 1995.
- [49] M. Thorup. Parallel shortcutting of rooted trees. *J. Algorithms*, 23(1):139–159, 1997.
- [50] M. Thorup. Compact oracles for reachability and approximate distances in planar digraphs. *Journal of the ACM*, 51:993–1024, 2004.
- [51] M. Thorup and U. Zwick. Compact routing schemes. In *ACM Symposium on Parallel Algorithms and Architectures*, pages 1–10, 2001.
- [52] M. Thorup and U. Zwick. Approximate distance oracles. *J. ACM*, 52(1):1–24, 2005.
- [53] D. P. Woodruff. Lower bounds for additive spanners, emulators, and more. In *FOCS*, pages 389–398, 2006.
- [54] A. C.-C. Yao. Space-time tradeoff for answering range queries (extended abstract). In *STOC*, pages 128–136, 1982.
- [55] U. Zwick. Exact and approximate distances in graphs — A survey. *Lecture Notes in Computer Science*, 2161:33+, 2001.

A Missing Details from Section 1

A.1 Previous Work on Other Related Problems

Dodis and Khanna [19] study the problem of finding the minimum-cost subset of missing edges that can be added to a (directed) graph G , with costs and lengths associated to the missing edges, so as to ensure that there is a path of length at most k between every pairs of nodes (not only those connected in G). Observe that k -TC-SPANNER is a special case of that problem: we can let G be the transitive reduction (see definitions below) of the input graph to k -TC-SPANNER, for all edges in the transitive closure of G set the length to 1 and cost to 1, and for the remaining edges set the length to k and cost to 0. Given this instance, the algorithm of Dodis and Khanna will produce a k -TC-spanner. However, the guarantee on the resulting k -TC-spanner size is only $\leq |G| + O(OPT n \log k)$, where OPT is the number of missing edges that need to be added. If $|G| = OPT = \Theta(n)$, their algorithm may return a k -TC-spanner with $\Omega(n^2)$ edges. Thus, in general, the resulting approximation ratio is no better than $O(n)$. Since their problem is more general, their hardness results do not apply to TC-spanners.

Chekuri *et al* [13] give an $O(p^{1/2+\epsilon})$ -approximation algorithm for the directed Steiner network problem where, given a digraph and node pairs $(s_1, t_1), \dots, (s_p, t_p)$, the goal is to connect all pairs with as few edges as possible. We can reduce k -TC-SPANNER to this problem by specifying all comparable pairs of nodes in levels 1 and $k + 2$ in the $k + 1$ -extension of G (see definition 5.5 of [19]). However, their ratio is only $O(n^{1+\epsilon})$ when $p = \Omega(n^2)$, and thus the resulting ratio for k -TC-SPANNER is no better than $O(n)$.

A.2 Sparse 2-TC-spanners Imply Efficient Monotonicity Testers

In this section we restate and prove Lemma 1.1, referred to in the introduction. The proof of the lemma explains how to use 2-TC-spanners to obtain efficient monotonicity testers.

Lemma A.1. *If a directed acyclic graph G_n has a 2-TC-spanner H with $s(n)$ edges, then there exists a monotonicity tester on G_n that runs in time $O\left(\frac{s(n)}{\epsilon n}\right)$.*

Proof. The tester selects $\frac{4s(n)}{\epsilon n}$ edges of the 2-TC-spanner H uniformly at random. It queries function f on the endpoints of all the selected edges and rejects if and only if one of the selected edges is *violated* by f , that is, $f(x) > f(y)$ for an edge (x, y) .

If the function f is monotone on G_n , the algorithm always accepts. The crux of the proof is to show that functions that are ϵ -far from monotone are rejected with probability at least $\frac{2}{3}$. Let $f : V_n \rightarrow \mathbb{R}$ be a function that is ϵ -far from monotone. It is enough to demonstrate that f violates at least $\frac{\epsilon n}{2}$ edges in H . Then each selected edge is violated with probability $\frac{\epsilon n}{2s(n)}$, and the lemma follows by elementary probability theory.

Denote the transitive closure of G by $TC(G)$. We say a vertex $x \in V_n$ is assigned a *bad* label by f if x has an incident violated edge in $TC(G_n)$; otherwise, x has a *good* label. Let V' be a set of vertices with good labels. Observe that f is monotone on the induced subgraph $G' = (V', E')$ of $TC(G)$. This implies ([28], Lemma 1) that f can be changed into a monotone function by modifying it on at most $|V_n - V'|$ vertices. Since f is ϵ -far from monotone, it shows that there are at least ϵn vertices with bad labels.

Every function that is ϵ -far from monotone has a matching M of $\frac{\epsilon n}{2}$ violated edges in $TC(G)$ [18]. We will establish an injection from the set of edges in M to the set of violated edges in H . For each edge (x, y) in the matching, consider the corresponding path from x to y of length at most 2 in the 2-TC-spanner H . If the path is of length 1, (x, y) is the violated edge in H corresponding to the matching edge (x, y) . Otherwise, let (x, z, y) be a path of length 2 in H . At least one of the edges (x, z) and (z, y) is violated, and we map (x, y) to that edge. Since M is a matching, all edges in M have distinct endpoints. Therefore, each edge in M is mapped to a unique violated edge in $TC(G)$. Thus, the 2-TC-spanner H has at least $\frac{\epsilon n}{2}$ violated edges, as required. \square

The fact that H is a 2-TC-spanner is crucial for the proof. If it was a k -TC-spanner for $k > 2$, the path of length k from x to y might not have any violated edges incident to x or y , even if $f(x) > f(y)$. Consider $G_{2n} = (V_{2n}, E)$ where $V_{2n} = \{x_1, \dots, x_{2n}\}$, $E = \{(x_i, x_n) \mid i < n\} \cup (x_n, x_{n+1}) \cup \{(x_{n+1}, x_j) \mid j > n + 1\}$. G_n is a 3-TC-spanner for itself. Now set $f(x_i) = 1$ for $i \leq n$ and $f(x_i) = 0$ otherwise. Clearly, this function is $\frac{1}{2}$ -far from monotone, but only one edge, (x_n, x_{n+1}) is violated in the 3-TC-spanner.

A.3 Partial Products in a Semigroup

Chazelle [12] and Alon and Schieber also consider a generalization of the above problem, where the input is an (undirected) tree T with an element s_i of a semigroup associated with each vertex i . The goal is to create a space-efficient data structure that allows to compute the product of elements associated with all vertices on the path from i to j , for all vertex pairs i, j in T . The generalized problem reduces to finding a sparsest k -TC-spanner for a certain directed tree T' obtained from T by appending a new vertex to each leaf, and then selecting an arbitrary root and directing all edges away from it. A k -TC-spanner for T' with $s(n)$ edges yields a preprocessing scheme with space complexity $s(n)$ for computing products on T with at most $2k$ queries as follows. The database stores a product $s_{v_1} \circ \dots \circ s_{v_t}$ for each k -TC-spanner edge (v_1, v_{t+1}) if the endpoints of that edge are connected by the path v_1, \dots, v_t, v_{t+1} in T' . Let $LCA(u, v)$ denote the lowest common ancestor of u and v in T . Compute the product corresponding to a path from u to v in T as follows: (1) if u is an ancestor of v (or vice versa) in T , query the products corresponding to the k -TC-spanner edges on the shortest path from u to a child of v (from v to a child of u , respectively); (2) otherwise, make queries

corresponding to the k -TC-spanner edges on the shortest path from $LCA(u, v)$ to a child of u and on the shortest path from a child of $LCA(u, v)$ nearest to u to a child of u . This gives a total of at most $2k$ queries.

B Approximation Algorithms for k -TC-SPANNER and Related Problems

B.1 Algorithm for DIRECTED k -SPANNER

We give the algorithm for DIRECTED k -SPANNER, which is a more general problem than k -TC-SPANNER. We then mention the extensions to other problems.

Theorem B.1. *For any (not necessarily constant) $k > 2$, there is a deterministic polynomial-time algorithm achieving an $O((n \log n)^{1-1/k})$ -approximation for DIRECTED k -SPANNER.*

Proof. Consider the following integer programming formulation. Let OPT be the size of an optimal k -spanner of G . For each edge e in the input digraph G , we have a variable x_e indicating whether x_e occurs in the k -spanner. Also, for each (not necessarily simple) path P containing at most k edges, we have a variable y_P indicating whether all of the edges of P occur in the k -spanner.

$$\begin{aligned} & \min \sum_{e \in G} x_e \\ \text{s.t. } & \forall e = (u, v) \in G, \quad \sum_{P \text{ from } u \text{ to } v, |P| \leq k} y_P \geq 1 \end{aligned} \quad (1)$$

$$\forall P = (e_1, e_2, \dots, e_r), y_P \leq x_{e_1}, y_P \leq x_{e_2}, \dots, y_P \leq x_{e_r} \quad (2)$$

$$\forall e \forall P, x_e, y_P \in \{0, 1\} \quad (3)$$

The first constraint ensures that there is at least one path of length at most k spanning each edge (u, v) in the spanner, while the second constraint only allows a path to be included if each of its edges is also in the spanner. Thus, any solution to this program is a k -spanner, and vice versa. Notice, however, that the number of path variables grows exponentially with k . We can instead write this as the following integer program:

$$\begin{aligned} & \min \sum_{e \in G} x_e \\ \text{s.t. } & \forall e = (u, v) \in G, \quad \sum_{P=(e_1, \dots, e_r) \text{ from } u \text{ to } v, |P| \leq k} -\min(x_{e_1}, x_{e_2}, \dots, x_{e_r}) \leq -1 \end{aligned} \quad (4)$$

$$\forall e, x_e \in \{0, 1\} \quad (5)$$

Now the number of variables is m , where m is the number of edges of G . We relax the constraints $x_e \in \{0, 1\}$ to $x_e \in [0, 1]$. The resulting set K we optimize over is convex since $\vec{x} \in [0, 1]^m$ and the functions $-\min(x_{e_1}, \dots, x_{e_r})$ are convex (as is their sum). We reduce the problem to a feasibility one by taking the convex set $K' = K \cap \{\vec{x} : \sum_{e \in G} x_e \leq t\}$, for a parameter t which we do binary search over.

The problem is still that we sum over a number of terms which can be exponential in k . However, we design a separation oracle A which does the following: given a point $\vec{x} \in \mathbb{R}^m$, A decides whether $\vec{x} \in K'$, and if not, provides an $\vec{a} \in \mathbb{R}^m$ and $b \in \mathbb{R}$ for which $\langle \vec{a}, \vec{x} \rangle < b$ but $\langle \vec{a}, \vec{y} \rangle \geq b$ for all $\vec{y} \in K'$. We will design an oracle for this task running in time $\text{poly}(n)$. Later, we explain the details of algorithm A .

There are several folklore polynomial-time algorithms for solving a convex program given a separation oracle A . We use the ellipsoid algorithm, which, given $\epsilon > 0$, runs in time $\text{poly}(n) \log \frac{1}{\epsilon}$ and, if the program is feasible, returns an \vec{x}^* for which the ℓ_2 -norm $|\vec{x}^* - \vec{x}|_2$ is at most ϵ , where \vec{x} is a feasible solution. Setting $\epsilon = n^{-\Theta(k)}$ guarantees that $\vec{x}^* \in [0, 1]^m$, that $\sum_{v \in G} x_v^* \leq n^{-\Theta(k)} + \sum_{v \in G} x_v$, and for all $e = (u, v) \in G$,

$$\begin{aligned} \sum_{P=(e_1, \dots, e_r) \text{ from } u \text{ to } v, |P| \leq k} -\min(x_{e_1}^*, x_{e_2}^*, \dots, x_{e_r}^*) & \leq n^k \epsilon - \sum_{P=(e_1, \dots, e_r) \text{ from } u \text{ to } v, |P| \leq k} \min(x_{e_1}, \dots, x_{e_r}) \\ & \leq n^{-\Theta(k)} - 1. \end{aligned}$$

Assuming we have an oracle A described above, the following is our algorithm k -SPANNER GENERATION to construct a directed k -spanner H of G . For a vertex $v \in G$, we use $BFS(v)$ to denote the set of edges along a shortest path tree¹ rooted at v . Clearly $|BFS(v)| = O(n)$.

k -Spanner Generation(G):

1. $H \leftarrow \emptyset$.
2. For each edge $e \in G$, if $x_e^* \geq \frac{1/2}{(n \log n)^{1-1/k}}$, $H \leftarrow H \cup \{e\}$.
3. Randomly sample $r = O((n \log n)^{1-1/k})$ vertices $z_1, z_2, \dots, z_r \in G$.
4. $H \leftarrow H \cup (\cup_i BFS(z_i))$. Output H .

Lemma B.2. *With probability at least $1 - 1/n$, H is a k -TC-spanner of G .*

Proof. Consider an edge $(u, v) \in G$. Suppose there are at most $(n \log n)^{1-1/k}$ different $u - v$ paths P of length at most k . By constraint (4) of the convex program and the relationship between \vec{x} and \vec{x}^* , there exists such a $P = (e_1, \dots, e_r)$ for which $\min(x_{e_1}^*, \dots, x_{e_r}^*) \geq 1/(n \log n)^{1-1/k} - n^{-\Theta(k)} \geq 1/(2(n \log n)^{1-1/k})$, for some $r \leq k$. Thus, this path P is included in H in step 2 of k -SPANNER GENERATION.

Now suppose there are more than $(n \log n)^{1-1/k}$ different $u - v$ paths P of length at most k . Let $W_{u,v} = \{w_1, \dots, w_s\}$ be the set of vertices lying on at least one such path. The number of $u - v$ paths of length at most k that can be formed from s vertices is at most s^{k-1} . So, $s^{k-1} = \Omega((n \log n)^{1-1/k})$, or $s = \Omega((n \log n)^{1/k})$. The probability that $\{z_1, z_2, \dots, z_r\} \cap W_{u,v} = \emptyset$ is at most $(1 - s/n)^r \leq e^{-rs/n} \leq e^{-\Omega(\log n)} \leq 1/n^3$, for an appropriate choice of constants.

By a union bound, with probability at least $1 - 1/n$, all edges $(u, v) \in G$ for which there are more than $(n \log n)^{1-1/k}$ different $u - v$ paths P of length at most k satisfy $\{z_1, z_2, \dots, z_r\} \cap W_{u,v} \neq \emptyset$. Conditioned on this event, for each such $(u, v) \in G$ let $z(u, v)$ be an arbitrary element in $\{z_1, z_2, \dots, z_r\} \cap W_{u,v}$. Then the path $u \rightsquigarrow z(u, v) \rightsquigarrow v$ along the edges of $BFS(z(u, v))$ is of length at most k . Indeed, there is a path P of length at most k from u to v which contains $z(u, v)$, and the path from u to v along the edges of $BFS(z(u, v))$ cannot be any longer than the length of P . \square

Lemma B.3. $|H| = O((n \log n)^{1-1/k} OPT)$.

Proof. Let OPT' be the optimum of the convex program. Clearly $OPT' \leq OPT$. In step (2) of k -SPANNER GENERATION, at most $2(n \log n)^{1-1/k} OPT' \leq 2(n \log n)^{1-1/k} OPT$ edges are added to H . In step (4), $O(rn) = O((n \log n)^{2-1/k})$ edges are added to H . So, $|H| = O((n \log n)^{1-1/k})(OPT + n)$. We may assume that $OPT \geq n - 1$, as otherwise G is not connected and we can run k -SPANNER GENERATION on each of its connected components. Therefore, $|H| = O((n \log n)^{1-1/k} OPT)$. \square

As $|H| \geq OPT$, these lemmas show that k -SPANNER GENERATION is a randomized $O((n \log n)^{1-1/k})$ -approximation algorithm for DIRECTED k -SPANNER. The algorithm can be derandomized by greedily choosing the z_i in step 3.

Lemma B.4. *For any constant $k > 2$, DIRECTED k -SPANNER has a deterministic $O((n \log n)^{1-1/k})$ -approximation algorithm.*

¹For a directed graph, this means we take a shortest path tree of edges directed away from v , together with a shortest path tree of edges directed towards v .

Proof. In step (3) of k -SPANNER GENERATION, instead of sampling r random vertices, we do the following. For each edge $(u, v) \in G$ with more than $(n \log n)^{1-1/k}$ simple $u - v$ paths P of length at most k , find the set $W_{u,v}$ of all vertices lying on such a path between u and v . This can be done by computing $BFS(w)$ for each vertex $w \in G$, and checking if $u \rightsquigarrow w \rightsquigarrow v$ along the edges of $BFS(w)$ is a path of length at most k . By averaging, there is a vertex z_1 which occurs in an $\Omega((\log n)^{1/k}/n^{1-1/k})$ fraction of the sets $W_{u,v}$. Choose z_1 , delete the sets $W_{u,v}$ containing z_1 , and repeat. This greedy algorithm finds z_1, \dots, z_r with $r = O((n \log n)^{1-1/k})$. \square

The technique can also be extended to other variants of spanners.

Lemma B.5. *For all constant $k > 2$, there are deterministic $O((n \log n)^{1-1/k})$ -approximation algorithms for CLIENT/SERVER DIRECTED k -SPANNER, k -DIAMETER SPANNING SUBGRAPH, and k -TC-SPANNER.*

Proof. In the client/server problem, we only wish to span a subset of edges of G , called *client edges*, and we may only use a subset of edges of G for spanning, called *server edges*. To modify our algorithm, we have a constraint in the linear program for each client edge rather than for all edges, and we only consider paths along server edges. In k -DIAMETER SPANNING SUBGRAPH, all pairs of vertices (u, v) for which v is reachable from u need to be connected by a path of length at most k . For this we impose constraint (1) for all pairs rather than just all edges. Finally, k -TC-SPANNER is a special case of DIRECTED k -SPANNER when the input is transitively closed. \square

Moreover, k -SPANNER GENERATION is polynomial time provided that we can find \vec{x}^* in polynomial time. For this, it suffices to show that the running time of the separation oracle is polynomial.

The separation oracle A first checks whether $\vec{x} \in [0, 1]^m$ and $\sum_{e \in G} x_e \leq t$ in $\text{poly}(n)$ time, and provides an appropriate hyperplane if any of these constraints are violated. Assume, then, that all of these constraints are satisfied. Let $\text{Count}(G, u, v, k)$ be an algorithm which outputs the number of $u - v$ paths of length at most k . The number of $u - v$ paths of length exactly i is just the (u, v) -th entry of M^i , where M is the adjacency matrix of G . Thus, we can implement $\text{Count}(G, u, v, k)$ in $\text{poly}(n)$ time. For a subset S of edges of G , $\text{Count}(G \setminus S, u, v, k)$ counts the number of $u - v$ paths of length at most k in G which do not use the edges in S .

The oracle sorts the coordinates of x , obtaining $x_{e_1} \leq x_{e_2} \leq \dots \leq x_{e_m}$. Let $S_0 = \emptyset$, and for $i \geq 1$, $S_i = S_{i-1} \cup \{e_i\}$. The oracle computes $\text{Count}(G \setminus S_i, u, v, k)$ for all $i \geq 0$. From this information, for each $j \geq 1$ the oracle can extract c_j , the number of $u - v$ paths of length at most k whose minimum is achieved by x_{e_j} . Indeed, observe that c_j is just the number of $u - v$ paths of length at most k in $G \setminus S_{j-1}$ minus the number of $u - v$ paths of length at most k in $G \setminus S_j$. Algorithm A can now check if the constraint corresponding to (u, v) is satisfied, and it does this for each $(u, v) \in G$. If the constraint for some (u, v) is not satisfied, for all i we set the i -th coordinate of the hyperplane \vec{a} to be c_i , and the scalar b to be 1. This \vec{a}, b pair satisfy the desired constraints, and are output by A . Note that A runs in $\text{poly}(n)$ time for any k . \square

B.2 k -TC-SPANNER Algorithm for Large k

For large k , we have the following better approximation, which is specific to the k -TC Spanner problem.

Theorem B.6. *For any $k \geq 6$, there exists a deterministic approximation algorithm for the k -TC Spanner problem with approximation ratio $O((n \log n)/(k^2 + k \log n))$.*

Proof. Let G be the input digraph. Assume, w.l.o.g., that G is connected. We construct S , a k -TC-spanner for the graph G such that $|S|/S_k(G) \leq O((n \log n)/(k^2 + k \log n))$. Set $k' = ck$ for c to be determined. Let G' be G with each directed cycle contracted to a vertex, and let $H = TR(G')$. For each vertex $v \in V(H)$,

define a set S_v of vertices such that: (i) $|S_v| \leq 4n/k'$, (ii) for each $u \in S_v$, $v \rightsquigarrow_H u$, and (iii) for any vertex w such that $v \rightsquigarrow_H w$, there exists $w' \in S_v$ with $d_H(w, w') \leq k'/4$. One can easily see such a set exists by averaging and it can be efficiently constructed. Next, we define another set of vertices $W \subseteq V(H)$ such that for every pair of vertices u and v such that $u \rightsquigarrow_H v$, either there is a path of length at most $3k'/8$ from u to v in H or there is a path from u to v that contains a vertex in W . The natural greedy algorithm for this problem constructs W to be of size at most $O(\frac{n \log n}{k + \log n})$. To construct the k -TC-spanner S , add to S the edges in H and for each vertex $w \in W$, add edges from w to all vertices in S_w . Also, for each contracted cycle in G , add an undirected star T_v centered at one arbitrary vertex of the cycle. The size of S is at most $|H| + O((n^2 \log n)/(k^2 + k \log n)) + O(n)$. Since $S_k(G) = \Omega(n)$ if G is connected, $|S|/S_k(G) = O((n \log n)/(k^2 + k \log n))$. To see that S is a k -TC-spanner, observe that for any pair of (u, v) with $u \rightsquigarrow_G v$, there will be a vertex $w \in W$ within distance $3k'/8$ of u and a vertex $w' \in S_w$ within distance $k'/4$ of v ; so, if none of the involved vertices are cycles, the distance between u and v in S is at most $3k'/8 + k'/4 + 1 = 5k'/8 + 1$. If the vertices u, v correspond to contracted cycles, it is easy to see that the path length will be at most $5k'/8 + 5$. We choose $k' = ck$ to ensure $5k'/8 + 5$ is at most k ; this is always possible because $k \geq 6$. \square

C $2^{\log^{1-\epsilon} n}$ -Hardness of k -TC-SPANNER for constant $k > 2$

Proof of Theorem 3.2. We give a reduction from MIN-REP with unrestricted parameters, considered in [21]:

Fact C.1 ([21]). *For all $\epsilon \in (0, 1)$, there is no polynomial time algorithm for the MIN-REP problem with approximation ratio $2^{\log^{1-\epsilon} n}$ unless $NP \subseteq DTIME(n^{\text{polylog } n})$.*

We reduce an arbitrary MIN-REP instance on $n^{\kappa'}$ vertices to a MIN-REP instance on n vertices with parameters in the desired range (where κ' is a suitably small constant). Since MIN-REP with unrestricted parameters is $2^{\log^{1-\epsilon} n}$ -inapproximable and the reduction is polynomial time, the theorem follows. The reduction consists of a sequence of five transformations on the original instance. We describe each of the transformations and specify how the parameters of the input and output MIN-REP instances are related.

1. (Disjoint copies)

Given an (n_0, r_0, d_0, m_0) -MIN-REP instance G_0 with OPT_0 as the solution value, $T_1(G_0, n^{\delta_1})$ is defined to be the MIN-REP instance G_1 with n^{δ_1} disjoint copies of G_0 . G_1 is a (n_1, r_1, d_1, m_1) -MIN-REP instance with $n_1 = n^{\delta_1} n_0$, $r_1 = n^{\delta_1} r_0$, $d_1 = d_0$, and $m_1 = m_0$. The solution value of G_1 is $OPT_1 = n^{\delta_1} OPT_0$ because if $OPT_1 < n^{\delta_1} OPT_0$, one could, by averaging over the n^{δ_1} copies of G_0 , extract a MIN-REP cover for G_0 of size smaller than OPT .

2. (Dummy vertices inside clusters)

Given an (n_1, r_1, d_1, m_1) -MIN-REP instance G_1 with OPT_1 as the solution value, $T_2(G_1, n^{\delta_2})$ is defined to be the MIN-REP instance G_2 obtained by increasing the size of each cluster by a factor of n^{δ_2} and not attaching any edges to the new vertices. G_2 is a (n_2, r_2, d_2, m_2) -MIN-REP instance with $n_2 = n^{\delta_2} n_1$, $r_2 = r_1$, $d_2 = d_1$, and $m_2 = n^{\delta_2} m_1$. The solution value of G_2 remains $OPT_2 = OPT_1$ because the minimum cover of G_2 does not include any isolated vertices.

3. (Blowup inside clusters with matching supergraph)

Given an (n_2, r_2, d_2, m_2) -MIN-REP instance G_2 with OPT_2 as the solution value, $T_3(G_2, n^{\delta_3})$ is defined to be the MIN-REP instance G_3 obtained as follows. For each cluster \mathcal{A}_i in G_2 , construct a cluster \mathcal{A}'_i in G_3 consisting of n^{δ_3} copies of \mathcal{A}_i . Let $(\mathcal{A}'_i)_k$ denote the k th copy of \mathcal{A}_i inside \mathcal{A}'_i . Whenever there is an edge in G_2 between $u \in \mathcal{A}_i$ and $v \in \mathcal{B}_j$, for each $1 \leq k \leq n^{\delta_3}$, add an edge

between the copy of u in $(\mathcal{A}'_i)_k$ and the copy of v in $(\mathcal{B}'_j)_k$. This procedure yields a (n_3, r_3, d_3, m_3) -MIN-REP instance G_3 where $n_3 = n^{\delta_3} n_2$, $r_3 = r_2$, $d_3 = d_2$, and $m_3 = m_2$. The solution value of G_3 remains $OPT_3 = OPT_2$ because the supergraph corresponding to G_3 and G_2 are identical.

4. (*Blowup inside clusters with complete supergraph*)

Given an (n_3, r_3, d_3, m_3) -MIN-REP instance G_3 with OPT_3 as the solution value, $T_4(G_3, n^{\delta_4})$ is defined to be the MIN-REP instance G_4 obtained as follows. For each cluster \mathcal{A}_i in G_3 , construct a cluster \mathcal{A}'_i in G_4 consisting of n^{δ_4} copies of \mathcal{A}_i . Let $(\mathcal{A}'_i)_k$ denote the k th copy of \mathcal{A}_i inside \mathcal{A}'_i . Whenever there is an edge in G_3 between $u \in \mathcal{A}_i$ and $v \in \mathcal{B}_j$, for each $1 \leq k_1, k_2 \leq n^{\delta_4}$, add an edge between the copy of u in $(\mathcal{A}'_i)_{k_1}$ and the copy of v in $(\mathcal{B}'_j)_{k_2}$. This procedure yields a (n_4, r_4, d_4, m_4) -MIN-REP instance G_4 where $n_4 = n^{\delta_4} n_3$, $r_4 = r_3$, $d_4 = n^{\delta_4} d_3$, and $m_4 = m_3$. The solution value of G_4 remains $OPT_4 = OPT_3$ because the supergraph corresponding to G_3 and G_4 are identical.

5. (*Tensoring*)

Given an (n_4, r_4, d_4, m_4) -MIN-REP instance G_4 with OPT_4 as the solution value, $T_5(G_4, n^{\delta_5})$ is defined to be the MIN-REP instance G_5 obtained by repeating the following construction $\log_2 n^{\delta_5}$ times². For each cluster \mathcal{A}_i in G_4 , construct two clusters \mathcal{A}'_i and \mathcal{A}''_i in G_5 . Furthermore, \mathcal{A}'_i contains two copies of \mathcal{A}_i and \mathcal{A}''_i contains two copies of \mathcal{A}_i . Denote the two copies inside \mathcal{A}'_i as $(\mathcal{A}'_i)_1$ and $(\mathcal{A}'_i)_2$ and similarly the two copies inside \mathcal{A}''_i as $(\mathcal{A}''_i)_1$ and $(\mathcal{A}''_i)_2$. For each edge (u, v) in G_4 with $u \in \mathcal{A}_i$ and $v \in \mathcal{B}_j$, add the following four edges in G_5 : between the copy of u in $(\mathcal{A}'_i)_1$ and copy of v in $(\mathcal{B}'_j)_1$, between the copy of u in $(\mathcal{A}'_i)_2$ and copy of v in $(\mathcal{B}'_j)_2$, between the copy of u in $(\mathcal{A}''_i)_1$ and copy of v in $(\mathcal{B}''_j)_1$, and between the copy of u in $(\mathcal{A}''_i)_2$ and copy of v in $(\mathcal{B}''_j)_2$.

The procedure yields a (n_5, r_5, d_5, m_5) -MIN-REP instance G_5 where $n_5 = n^{2\delta_5} n_4$, $r_5 = n^{\delta_5} r_4$, $d_5 = d_4$, and $m_5 = m_4$. Also, we argue that $OPT_5 = n^{2\delta_5} OPT_4$. Clearly, $OPT_5 \leq n^{2\delta_5} OPT_4$ because one could choose copies of the vertices in the cover for G_4 in each of the $n^{2\delta_5}$ copies of the clusters of G_4 . For the other direction, notice that G_5 contains $n^{2\delta_5}$ vertex disjoint copies of G_4 , and so, if $OPT_5 < n^{2\delta_5} OPT_4$, then by averaging, there would be a copy of G_4 covered using less than OPT vertices, a contradiction.

For some positive κ' sufficiently smaller than κ , consider an arbitrary $(n^{\kappa'}, r_0, d_0, m_0)$ -MIN-REP instance G_0 with optimum OPT_0 , where the only constraints on the parameters are nontriviality conditions: $r_0 \in [1, n^{\kappa'}]$, $d_0 \in [1, n^{\kappa'}]$, $m_0 \in [1, n^{\kappa'}]$, and $OPT_0 \in [1, 2n^{\kappa'}]$. Let $G = T_5(T_4(T_3(T_2(T_1(G_0, n^{\delta_1}), n^{\delta_2}), n^{\delta_3}), n^{\delta_4}), n^{\delta_5}))$. We choose $\delta_1, \delta_2, \delta_3, \delta_4, \delta_5$ such that G is a (n, r, d, m) -MIN-REP instance with $r \in [n^R, n^{R+\kappa'}]$, $d \in [n^D, n^{D+\kappa'}]$, $m \in [n^M, n^{M+\kappa'}]$ and $OPT \in [n^F, n^{F+\kappa'}]$. By definitions of transformations, $n = n^{\kappa'+\delta_1+\delta_2+\delta_3+\delta_4+2\delta_5}$, $r \in [n^{\delta_1+\delta_5}, n^{\kappa'+\delta_1+\delta_5}]$, $d \in [n^{\delta_4}, n^{\delta_4+\kappa'}]$, $m \in [n^{\delta_2}, n^{\kappa'+\delta_2}]$, and $OPT \in [n^{\delta_1+2\delta_5}, n^{\kappa'+\delta_1+2\delta_5}]$. Therefore, choose $\delta_4 = D$, $\delta_2 = M$, $\delta_5 = F - R$, and $\delta_1 = 2R - F$. All of these values are in $(0, 1)$ by restriction of the parameters in the theorem statement. Now, since $\kappa' + \delta_1 + \delta_2 + \delta_3 + \delta_4 + 2\delta_5 = D + M + F + \delta_3 + \kappa'$ and since $D + M + F < 1$ and κ' can be made as small as we want, we can choose $\delta_3 \in (0, 1)$ such that $n = n^{\kappa'+\delta_1+\delta_2+\delta_3+\delta_4+2\delta_5}$. Therefore, G is a MIN-REP instance with parameters in the desired range. \square

Attaching butterflies in the construction of a hard k -TC-SPANNER instance \mathcal{G} . We further discuss the way the butterflies are attached to the groups. Recall that \mathcal{I}_0 is a (n_0, r, d_0, m) -MIN-REP instance with $n_0 = n^\delta$, $r \in [n^{\delta/2}, n^{\delta/2+\kappa}]$, $d_0 \in [n^\eta, n^{\eta+\kappa}]$ and $m \in [n^{2\eta}, n^{2\eta+\kappa}]$. Thus, for each group $A_{i,j}$ there are at most $\frac{n^\delta}{rm}$ non-isolated vertices. We will attach the butterfly $BF(A_{i,j})$ in such a way that each vertex in $BF^{k-1}(A_{i,j})$ is adjacent to at most $\frac{d_*}{m}$ non-isolated vertices in $A_{i,j}$, out of a total out-degree of size d_* . This

²For simplicity, we assume n^{δ_5} is a power of 2.

is the crucial property exploited later in the proof. We can achieve this property in the following way. Recall that each vertex of $BF^s(A_{i,j})$ is labeled (a_1, \dots, a_{k-1}, s) , where $a_l \in [d_*]$ for all $l \in [k-1]$, $s \in [k]$, and each vertex $v = (a_1, \dots, a_{k-1}, k)$ connects to $v' = (a_1, \dots, a_{k-2}, a'_{k-1}, k-1)$. Thus, for a fixed prefix $b = (b_1, b_2, \dots, b_{k-2})$ all vertices $(b_1, \dots, b_{k-2}, b_{k-1}, k-1)$ connect to the same set A_b of vertices in $A_{i,j}$, and $|A_b| = d_*$. Choose the set A_b to contain at most d_*/m non-isolated vertices, which is possible since the total fraction of non-isolated vertices in $A_{i,j}$ is $\leq \frac{1}{m}$.

Lemma C.2 (Rep-cover Spanner Lemma). *There is a k -TC-spanner \mathcal{H} s.t. $|\mathcal{H}| = O(OPT n^{1-\delta} (\frac{n^\delta}{r})^{\frac{2}{k-1}})$, where OPT is the minimum rep-cover of the underlying \mathcal{I}_0 (and of \mathcal{I} as well). Moreover, \mathcal{H} contains only paths of type $(2\&(k-2), 2\&(k+1))$.*

Proof. We construct the graph \mathcal{H} by adding some shortcut edges to \mathcal{G} . Let S_0 be a minimum rep-cover of \mathcal{I}_0 of size OPT . Recall that each \mathcal{A}_i and \mathcal{B}_j is replicated $n^{1-\delta}$ times in \mathcal{I} . Let S be the set of all replicas in \mathcal{I} of vertices in S_0 . Let $A_{i,j}$ and $B_{k,l}$ be two comparable groups of vertices. Recall that $d_* = (\frac{n^\delta}{r})^{\frac{1}{k-1}}$. To get a k -TC-spanner on $BF(A_{i,j}) \cup BR(B_{k,l})$ connect each vertex v from the restriction of S to $A_{i,j}$ with all its d_*^2 comparable vertices in $BF^{k-2}(A_{i,j})$. Similarly, connect each vertex in the restriction of S to $B_{k,l}$ to its d_*^2 comparable vertices in $BR^{k+3}(B_{k,l})$. Since every vertex $u \in BF^1(A_{i,j})$ is comparable to every vertex $v \in A_{i,j}$, it follows that there is a vertex $w \in BF^{k-2}(A_{i,j})$ comparable to both u and v . Thus, between any such u and v there is a path using an edge of type $2\&(k-2)$. Similarly, every vertex in $BR^{k+3}(B_{k,l})$ is comparable to every vertex of $B_{k,l}$. By our construction, any pair of vertices $(u_1, u_{k+3}) \in BF^1(A_{i,j}) \times BR^{k+3}(B_{k,l})$ is connected by a path of type $(2\&(k-2), 2\&(k+1))$. In addition, any pair of vertices $(u_1, u_{k+2}) \in BF^1(A_{i,j}) \times BR^{k+2}(B_{k,l})$, as well as $(u_2, u_{k+3}) \in BF^2(A_{i,j}) \times BR^{k+3}(B_{k,l})$ are connected by a path of length at most k using shortcut edges of types $2\&(k-2)$ and $2\&(k+1)$, respectively. By connecting all the comparable groups $A_{i,j}$ and $B_{k,l}$ in this manner, we obtain a k -TC-spanner on \mathcal{G} .

Since there are $n^{1-\delta}$ copies of each $A_{i,j}$ and $B_{k,l}$ the total number of shortcut edges added is $OPT n^{1-\delta} d_*^2 = OPT n^{1-\delta} (\frac{n^\delta}{r})^{\frac{2}{k-1}}$ and we only used shortcut edges of types $2\&(k-2)$ and $2\&(k+1)$. In addition, since \mathcal{G} is transitively reduced, \mathcal{H} must include all the edges of \mathcal{G} . We bound the size of \mathcal{G} by inspecting the total number of edges in the butterflies (knd_*), the MIN-REP instance ($\leq nd$), and the brooms ($nd_* + n^{1-\delta} rd_*^2$). Thus, $|\mathcal{G}| \leq k r n^{1-\delta} (\frac{n^\delta}{r})^{1+\frac{1}{k-1}} + n^{2+\eta+\kappa-\delta} + n (\frac{n^\delta}{r})^{\frac{1}{k-1}} + n^{1-\delta} r (\frac{n^\delta}{r})^{\frac{2}{k-1}}$. The following conditions, satisfied by the parameters of our construction, suffice to show that each term of the preceding sum is respectively $o(|\mathcal{H}|)$. (The parameter κ is omitted from the conditions, since if the inequalities are satisfied without κ then κ can be made sufficiently small to ensure that they are satisfied with κ .)

$$\zeta + (1-\delta) + \frac{2}{k-1} \left(\delta - \frac{\delta}{2} \right) > \frac{\delta}{2} + (1-\delta) + \frac{k}{k-1} \left(\delta - \frac{\delta}{2} \right) \quad , \text{ or } \zeta > \delta \frac{2k-3}{2(k-1)} \quad (6)$$

$$\zeta + (1-\delta) + \frac{2}{k-1} \left(\delta - \frac{\delta}{2} \right) > 2 + \eta - \delta \quad , \text{ or } \zeta > 1 + \eta - \frac{\delta}{k-1} \quad (7)$$

$$\zeta + (1-\delta) + \frac{2}{k-1} \left(\delta - \frac{\delta}{2} \right) > 1 + \frac{1}{k-1} \left(\delta - \frac{\delta}{2} \right) \quad , \text{ or } \zeta > \delta \frac{2k-3}{2(k-1)} \quad (8)$$

$$\zeta + (1-\delta) + \frac{2}{k-1} \left(\delta - \frac{\delta}{2} \right) > (1-\delta) + \frac{\delta}{2} + \left(\delta - \frac{\delta}{2} \right) \frac{2}{k-1} \quad , \text{ or } \zeta > \frac{\delta}{2} \quad (9)$$

□

Lemma C.3 (Path Analysis Lemma). *There are $o(OPT)$ deletable superedges $(\mathcal{A}_i, \mathcal{B}_j)$, where $i, j \in [r]^2$.*

Proof. We call a path *canonical* if it contains shortcut edges of types both $2\&(k-2)$ and $2\&(k+1)$; otherwise, a path is *alternative*. Observe that any alternative path contains at least one shortcut edge

from among the following three cases: (1) shortcut edges crossing both V_k and V_{k+1} , i.e. one of the shortcut edge types: $3\&(k-2)$, $3\&(k-1)$, $2\&(k-1)$, $2\&k$, and $3\&k$; (2) shortcut edges of type $3\&\ell$ where $\ell \leq k-3$; (3) shortcut edges of type $2\&\ell$ where $\ell \leq k-3$. Let S_B be the set of all the shortcut edge types contained in the above three cases. Then $|S_B| = \Theta(k)$. Now, for each shortcut edge type $S \in S_B$, let $Del(S) = \{(i, j) \in [r]^2 \mid \text{at least } \frac{1}{4|S_B|} \text{ fraction of pairs } (u, v) \in BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j) \text{ have an alternative path containing a shortcut edge of type } S\}$. By a union bound, the total number of deletable superedges is at most $\sum_{S \in S_B} Del(S)$. Hence it suffices to show that for all $S \in S_B$, $Del(S) = o(OPT)$.

Let $C(S) = \{(u, v) \in BF^1(\mathcal{A}_i) \times BR^{k+3}(\mathcal{B}_j) \mid \exists \text{ an alternative path between } u \text{ and } v \text{ containing a shortcut edge of type } S\}$. By the definition of $Del(S)$, since for all $i \in [r]$, $|BF^1(\mathcal{A}_i)| = \frac{n}{r}$ and $|BR^{k+3}(\mathcal{B}_i)| = n^{1-\delta} d_*^2$, we have

$$|C(S)| \geq |Del(S)| \frac{1}{4|S_B|} \frac{n}{r} n^{1-\delta} d_*^2. \quad (10)$$

Now we will obtain upper bounds of $|C(S)|$ in terms of OPT for each of three cases of shortcut edges, thus obtaining upper bounds on $Del(S)$. Recall that $\delta = \frac{k-1}{k-\frac{1}{4}}$, $\eta = \frac{\delta}{2(4k-4)(4k-2)}$, and $\zeta = \delta \left(\frac{4k-5}{4k-4} + \frac{1}{4k-2} \right)$. Also, recall $r \in [n^{\frac{\delta}{2}}, n^{\frac{\delta}{2}+\kappa}]$, $d \in [n^{(1-\delta)+\eta}, n^{(1-\delta)+\eta+\kappa}]$, and $m \in [n^{2\eta}, n^{2\eta+\kappa}]$ for some small enough constant κ and $d_* = \left(\frac{n^\delta}{r} \right)^{1/(k-1)}$. We mostly ignore κ below since we can make it as small a constant as we like.

Suppose that S is a shortcut edge from the first case. Then S is a shortcut of type $\ell_1\&(k-\ell_2)$, where $2 \leq \ell_1 \leq 3$, $0 \leq \ell_2$, and $\ell_1 - \ell_2 \geq 1$. Now we obtain that for any shortcut of type S , the shortcut can be used for at most $d_*^{k-1-\ell_2} d_*^{3+\ell_2-\ell_1} = d_*^{k+2-\ell_1}$ many pairs $(u, v) \in C(S)$. Hence, $|C(S)| \leq d_*^{k+2-\ell_1} \cdot OPT \frac{n^{1-\delta} d_*^2}{\log n} \leq d_*^k \cdot OPT \frac{n^{1-\delta} d_*^2}{\log n}$. From (10), we obtain that

$$|Del(S)| \leq 4|S_B| \frac{d_*^k OPT n^{1-\delta} d_*^2}{\frac{n}{r} n^{1-\delta} d_*^2 \log n} = O\left(\frac{d_*}{n^{1-\delta} \log n}\right) OPT.$$

Then because $\delta < \frac{k-1}{k-\frac{1}{2}}$, $1-\delta > \frac{\delta}{2(k-1)}$, and so we obtain that $n^{1-\delta}$ is a polynomial factor larger than $d_* = \left(\frac{n^\delta}{r} \right)^{1/(k-1)}$, which proves that $|Del(S)| = o(OPT)$.

Now suppose that S is a shortcut type of the second case. Let S be type $3\&\ell$, where $1 \leq \ell \leq k-3$. Now, from the fact that out-degree of each vertex in V_k is at most $n^{1-\delta+\eta+\kappa}$, we obtain that for any shortcut of type S , the shortcut can be used for at most $d_*^{\ell-1} d_*^{k-3-\ell} n^{(1-\delta)+\eta} d_*^2 = d_*^{k-4} n^{(1-\delta)+\eta} d_*^2$ many pairs $(u, v) \in C(S)$ (ignoring κ as mentioned above). Hence, upto small polynomial factors,

$$|C(S)| \leq d_*^{k-4} n^{(1-\delta)+\eta} d_*^2 \cdot OPT \frac{n^{1-\delta} d_*^2}{\log n} = \frac{d_*^k n^{2-2\delta+\eta}}{\log n} OPT \quad (11)$$

From (10) and (11), we obtain $|Del(S)| \leq 4|S_B| \frac{n^\eta}{d_* \log n} OPT$. Now, $\eta < \frac{\delta}{2(k-1)}$, and so, $|Del(S)| = o(OPT)$.

Now suppose that S is a shortcut type of the third case. Let S be type $2\&\ell$, where $\ell \leq k-3$. Note that for any vertex v in V_{k-1} the number of non-isolated vertices in V_k that v is connected to is $\frac{d_*}{m}$. Hence, together with the fact that out-degree of each vertex in V_k is at most $n^{1-\delta+\eta+\kappa}$, we obtain that for any shortcut of type S , the shortcut can be used for at most $\frac{d_*^{k-3}}{n^{2\eta}} n^{(1-\delta)+\eta} d_*^2$ many pairs (u, v) in $C(S)$ (upto small polynomial factors). Then

$$|C(S)| \leq \frac{d_*^{k-3}}{n^{2\eta}} n^{(1-\delta)+\eta} d_*^2 \cdot OPT \frac{n^{1-\delta} d_*^2}{\log n} = \frac{d_*^{k+1} n^{2-2\delta+\eta}}{n^{2\eta} \log n} OPT. \quad (12)$$

From (10), (12), and the fact that $n^\eta = o(n^{2\eta} \log n)$, we get $|Del(S)| \leq 4|S_B| \frac{n^\eta}{n^{2\eta} \log n} OPT = o(OPT)$. \square

Lemma C.4 (Rerandomization Lemma). *If a $\frac{3}{4}$ -good k -TC-spanner \mathcal{K}' for \mathcal{G}' is given, then there exists \mathcal{K}'' , a 1-good k -TC-spanner for \mathcal{G}' , such that $|\mathcal{K}''| \leq O(|\mathcal{K}'| \cdot \log n)$.*

Proof. First, we fix some notation. Consider some $(i, j) \in [r]^2$ such that there is an edge between a vertex in \mathcal{A}_i and a vertex in \mathcal{B}_j in \mathcal{G}' . Let $S_{i,j}$ be the set of vertices in \mathcal{A}_i that are adjacent to \mathcal{B}_j , and let $T_{i,j}$ be the set of vertices in \mathcal{B}_j that are adjacent to \mathcal{A}_i . We know that at least $\frac{3}{4}$ of the vertices in $BF^1(\mathcal{A}_i)$ have a path of type $(2\&(k-2))$ to $S_{i,j}$ and at least $\frac{3}{4}$ of the vertices in $BR^{k+3}(\mathcal{B}_j)$ have a path of length 1 from $T_{i,j}$. By a Markov argument, for at least $\frac{1}{2}$ of the groups $A_{i,s}$ in \mathcal{A}_i , at least $\frac{1}{2}$ of the vertices in $BF^1(A_{i,s})$ must have a path of type $(2\&(k-2))$ to $S_{i,j}$. Call the butterfly attached to such a group $A_{i,s}$ an (i, j) -good butterfly, and call the set of vertices in $BF^{k-2}(A_{i,s})$ that have shortcut edges to $S_{i,j}$ (i, j) -helpful vertices. Similarly, for at least $\frac{1}{2}$ of the groups $B_{j,t}$, at least $\frac{1}{2}$ of the vertices in $BR^{k+2}(B_{j,t})$ have shortcut edges to $T_{i,j}$. We call the brooms attached to such groups $B_{j,t}$ (i, j) -good brooms and we again call the vertices in $BR^{k+2}(B_{j,t})$ that have shortcut edges to $T_{i,j}$ (i, j) -helpful vertices. It will be clear from context whether a helpful vertex is to the left or right of the MIN-REP instance.

Our construction of \mathcal{K}'' ensures that in \mathcal{K}'' , for any two comparable clusters $(\mathcal{A}_i, \mathcal{B}_j)$, each vertex in $BF^1(\mathcal{A}_i)$ is comparable to a helpful vertex in $BF^{k-2}(\mathcal{A}_i)$ and each vertex in $BR^{k+3}(\mathcal{B}_j)$ is a helpful vertex. This is enough to ensure that \mathcal{K}'' is 1-good k -TC-spanner for \mathcal{G}' . We will construct \mathcal{K}'' to be equal to $\bigcup_{r=1}^{O(\log n)} \Pi_r(\mathcal{K}')$ where each Π_r is a random transformation of \mathcal{K}' that moves the shortcut edges.

Each Π_r will be the composition of several transformations on the edges of \mathcal{K}' . The transformations move only shortcut edges, but not transitive reduction edges, in \mathcal{K}' . Informally, the first transformation randomly permutes the groups in each cluster on the left side of the MIN-REP instance, the second randomly permutes the groups in each cluster of the right side of the MIN-REP instance, the third randomly permutes the edges of the butterfly graph, and the fourth randomly permutes the broomsticks. Formally:

- *Left Group permutations:* Π^{lg}

For each $i \in [r]$, independently choose a random permutation $\pi_i : [n^{1-\delta}] \rightarrow [n^{1-\delta}]$. For each cluster \mathcal{A}_i , if (u, v) is an edge in \mathcal{K}' with $u, v \in BF(A_{i,s})$, then there is an edge (u', v') in $\Pi^{lg}(\mathcal{K}')$, where u' and v' are the copies of u and v respectively in $BF(A_{i,\pi_i(s)})$.

- *Right Group permutations:* Π^{rg}

For each $j \in [r]$, independently choose a random permutation $\pi_j : [n^{1-\delta}] \rightarrow [n^{1-\delta}]$. For each cluster \mathcal{B}_j , if (u, v) is an edge in \mathcal{K}' with $u, v \in BR(B_{j,s'})$, then there is an edge (u', v') in $\Pi^{rg}(\mathcal{K}')$, where u' and v' are the copies of u and v respectively in $BR(B_{j,\pi_j(s')})$.

- *Butterfly permutations:* Π^{bf}

For each $i \in [r]$ and $s \in [n^{1-\delta}]$, label a vertex u in $BF(A_{i,s})$ as $(a_1, a_2, \dots, a_{k-1}, m) \in [d_*]^{k-1} \times [k]$, where $u \in V_m$ and $(a_1, a_2, \dots, a_{k-1})$ is the usual vertex labelling that defines a generalized butterfly graph. Now, for every (i, s) and every $(a_1, \dots, a_{k-3}) \in [d_*]^{k-3}$, independently choose two random permutations $\pi_{i,s}^{(a_1, \dots, a_{k-3})} : [d_*] \rightarrow [d_*]$ and $\sigma_{i,s}^{(a_1, \dots, a_{k-3})} : [d_*] \rightarrow [d_*]$. For any edge $(u, w) \in BF^{k-2}(A_{i,s}) \times BF^k(A_{i,s})$ where $u = (a_1, \dots, a_{k-3}, a_{k-2}, a_{k-1}, k-2)$ and $w = (a_1, \dots, a_{k-3}, a'_{k-2}, a'_{k-1}, k)$, there exists the edge (u', w) in $\Pi^{bf}(\mathcal{K}')$ where $u' = (a_1, \dots, a_{k-3}, \pi_{i,s}^{(a_1, \dots, a_{k-3})}(a_{k-2}), \sigma_{i,s}^{(a_1, \dots, a_{k-3})}(a_{k-1}), k-2)$. All other edges in the butterfly stay fixed.

- *Broom permutations:* Π^{br}

For each $j \in [r]$ and $s' \in [n^{1-\delta}]$, independently choose a random permutation $\pi_{j,s'} : [p] \rightarrow [p]$ and $\sigma_{j,s'} : [t] \rightarrow [t]$. Label a vertex $v \in BR^{k+2}(B_{j,s'})$ as an element of $[p]$ and label a vertex $w \in BR^{k+3}(B_{j,s'})$ as an element of $[p] \times [t]$ in the natural way. If $(u, w) \in BR^{k+1}(B_{j,s'}) \times BR^{k+3}(B_{j,s'})$ is an edge in \mathcal{K}' , then $(u, w') \in BR^{k+1}(B_{j,s'}) \times BR^{k+3}(B_{j,s'})$ is an edge in $\Pi^{br}(\mathcal{K}')$, where $w' = (\pi_{j,s'}(w_1), \sigma_{j,s'}(w_2))$ if the label of w is (w_1, w_2) . All other edges in the broom stay fixed.

Now, for each $r = 1, \dots, O(\log n)$, define Π_r to be the composition of Π^{lg} , Π^{rg} , Π^{bf} , and Π^{br} . For each r , choose all the permutations independently. As we said before, we set $\mathcal{K}'' = \cup_r \Pi_r(\mathcal{K}')$.

Claim C.5. For each $(i, j) \in [r]^2$ such that \mathcal{A}_i and \mathcal{B}_j are comparable, for any $u \in BF^1(\mathcal{A}_i)$ and $v \in BR^{k+3}(\mathcal{B}_j)$,

$$\Pr_{\Pi_r}[u \text{ is in a } (i, j)\text{-good butterfly in } \Pi_r(\mathcal{K}')] \geq \frac{1}{2}, \quad \Pr_{\Pi_r}[v \text{ is in a } (i, j)\text{-good broom in } \Pi_r(\mathcal{K}')] \geq \frac{1}{2}$$

Proof. At least half the butterflies attached to \mathcal{A}_i are good from above, and hence, for every vertex $u \in BF^1(\mathcal{A}_i)$, the left group permutations ensure that with probability at least $\frac{1}{2}$, the edges of a good butterfly are mapped to the butterfly that u belongs to. The right group permutations provide the same function for a $v \in BR^{k+3}(\mathcal{B}_j)$. \square

Claim C.6. For each $(i, j) \in [r]^2$ such that \mathcal{A}_i and \mathcal{B}_j are comparable, then for any $v \in BR^{k+3}(\mathcal{B}_j)$:

$$\Pr_{\Pi_r}[v \text{ is a } (i, j)\text{-helpful vertex} \mid v \text{ is in a } (i, j)\text{-good broom}] \geq \frac{1}{2}$$

Proof. At least half of the broomsticks of a good broom are helpful (i.e., incident to a shortcut edge), and hence for every vertex $v \in BR^{k+3}(\mathcal{B}_j)$, the broom permutations ensure that with probability at least $\frac{1}{2}$, v is incident to a shortcut edge. \square

Claim C.7. For each $(i, j) \in [r]^2$ such that \mathcal{A}_i and \mathcal{B}_j are comparable, then for any $u \in BR^1(\mathcal{A}_i)$:

$$\Pr_{\Pi_r}[u \text{ is comparable to a } (i, j)\text{-helpful vertex} \mid u \text{ is in a } (i, j)\text{-good butterfly}] \geq \frac{1}{2}$$

Proof. Suppose $BF(A_{i,s})$ is a (i, j) -good butterfly. For any $(a_{k-2}, a_{k-1}) \in [d_*]^2$, let $S_{(a_{k-2}, a_{k-1})}$ be the set of vertices u in $BF^1(A_{i,s})$ such that u is labelled as $(a_1, \dots, a_{k-3}, a_{k-2}, a_{k-1}, 1)$ where (a_1, \dots, a_{k-3}) are arbitrary elements of $[d_*]^{k-3}$. Note that all the vertices in a given $S_{(a_{k-2}, a_{k-1})}$ are comparable to the same set of vertices in $BF^{k-2}(A_{i,s})$ and hence, either they are all comparable to an (i, j) -helpful vertex or none of them are. Hence, at least $\frac{1}{2}$ of the $S_{(a_{k-2}, a_{k-1})}$'s must have every vertex comparable to a helpful vertex in $BF^{k-2}(A_{i,s})$. Now, because of the random butterfly permutations, a given vertex $v \in BF^1(A_{i,s})$ falls in such a $S_{(a_{k-2}, a_{k-1})}$ with probability at least $\frac{1}{2}$. \square

Thus, for two vertices u and v in comparable \mathcal{A}_i and \mathcal{B}_j respectively, the probability that u and v are connected by a canonical path in $\Pi_r(\mathcal{K}')$ is at least $\frac{1}{16}$. Since we take $O(\log n)$ independent random transformations Π_r , the probability that u and v will be connected by a canonical path in at least one $\Pi_r(\mathcal{K}')$ is at least $1 - \frac{1}{\text{poly}(n)}$. Taking a union bound over all vertex pairs in \mathcal{A}_i and \mathcal{B}_j as well as all possible i and j , we find that with probability at least $\frac{1}{2}$, \mathcal{K}'' has a canonical path between any comparable $u \in V_1$ and $v \in V_{k+3}$. Therefore, the desired \mathcal{K}'' exists and is of size at most $O(|\mathcal{K}'| \cdot \log n)$. \square

Lemma C.8 (Rep-cover Extraction Lemma). Given \mathcal{K}'' , a 1-good k -TC-spanner for \mathcal{G}' , of size $o(OPT \cdot n^{1-\delta} \cdot d_*^2)$, there exists a MIN-REP cover of \mathcal{I} of size $o(OPT)$.

Proof. For $s \in [n^{1-\delta}]$, define \mathcal{K}_s'' to be the subgraph of \mathcal{K}'' induced by $\cup_{i=1}^r (BF(A_{i,s}) \cup BR(B_{i,s}))$. The \mathcal{K}_s'' are clearly disjoint. By averaging, there exists an \bar{s} such that $|\mathcal{K}_{\bar{s}}''| \leq o(OPT \cdot d_*^2)$.

We further partition the shortcut edges in $\mathcal{K}_{\bar{s}}''$ into d_*^2 parts. For each $x, y \in [d_*]$, let $U_{x,y}$ denote the set of all the nodes in $\cup_{i=1}^r BF^1(A_{i,\bar{s}})$ with butterfly coordinates $(u_1, \dots, u_{k-2}, x, y, 1)$, where $u_1, \dots, u_{k-2} \in [d_*]$. To partition the corresponding broomsticks, identify the nodes in $BR^{k+2}(B_{i,\bar{s}})$ with $[d_*]$, and for each such node $x \in [d_*]$, identify its descendants in $BR^{k+3}(B_{i,\bar{s}})$ with $(x, 1), \dots, (x, d_*)$. For each $x, y \in [d_*]$, let $U'_{x,y}$ denote the set of all the broomsticks $\cup_{i=1}^r BR^{k+3}(B_{i,\bar{s}})$ with coordinates (x, y) . Define $\mathcal{K}_{\bar{s},x,y}''$ to be the subgraph of $\mathcal{K}_{\bar{s}}''$ induced by the nodes comparable to the nodes in $U_{x,y} \cup U'_{x,y}$.

Observe that the shortcut edges in different $\mathcal{K}_{\bar{s},x,y}''$ are disjoint because (a) different $U'_{x,y}$ are disjoint and (b) the descendants in V_{k-2} of different $U_{x,y}$ are also disjoint. Thus, by averaging, there exist \bar{x}, \bar{y} such that $\mathcal{K}_{\bar{s},\bar{x},\bar{y}}''$ contains $o(OPT)$ shortcut edges.

Let S be the set of vertices in V_k and V_{k+1} that are incident to shortcut edges in $\mathcal{K}_{\bar{s},\bar{x},\bar{y}}''$. Then $|S| \leq o(OPT)$. Observe that S is a rep-cover for the MIN-REP instance $\mathcal{I}'_{\bar{s}}$ obtained by restricting \mathcal{I}' to the edges between $A_{i,\bar{s}}$ and $B_{j,\bar{s}}$. This holds because in $\mathcal{K}_{\bar{s}}''$, each comparable pair of nodes in $U_{\bar{x},\bar{y}} \times U'_{\bar{x},\bar{y}}$ is connected by a canonical path. But a MIN-REP cover for $\mathcal{I}'_{\bar{s}}$ is also a MIN-REP cover for \mathcal{I}' by definition of \mathcal{I} . Finally, given a rep-cover S of \mathcal{I}' , we can get a rep-cover of \mathcal{I} by adding at most 2 vertices per super-edge deleted from \mathcal{I} to obtain \mathcal{I}' . Since $o(OPT)$ super-edges were deleted and since $|S| \leq o(OPT)$, we obtain a MIN-REP cover for \mathcal{I} of size $o(OPT)$. \square

D $\Omega(\log n)$ -Hardness of 2-TC-Spanner

Theorem D.1. *For any $k \geq 2$, it is NP-hard to approximate the size of the sparsest k -TC-spanner within a ratio of $O(\frac{1}{k} \log n)$. In particular, 2-TC-SPANNER is $\Omega(\log n)$ -inapproximable.*

Our proof uses a reduction from a variant of Set Cover, called (a, b, c) -Nice Set Cover. Before defining this problem we define other variants of Set Cover. An instance of (a, b) -Set Cover, consists of a bipartite graph $G = A \cup B$, with $|A| = a$ and $|B| = b$. An instance of a -Balanced Set Cover consists of a bipartite graph $G = A \cup B$, with $|A| = |B| = a$. An instance of (a, c) -Balanced Bounded Set Cover consists of a bipartite graph $G = A \cup B$, with $|A| = |B| = a$ and such that the degrees of the vertices in A are at most c . Finally, an instance of (a, b, c) -Nice Set Cover consists of a bipartite graph $G = A \cup B$, with $|A| = a$, $|B| = b$. B can be partitioned into disjoint sets B_i such that $B = \cup_{i=1}^a B_i$, $|B_i| = \frac{b}{a}$, assuming $\frac{b}{a}$ is an integer. G must satisfy the property that if $v \in A$ is adjacent to $w \in B_i$, for some $1 \leq i \leq a$, then v is adjacent to every element of B_i . Moreover, v is adjacent to at most c sets B_i . A solution to all these Set Cover variants is a minimum number of vertices in A that cover all the vertices in B .

Lemma D.2. *It is NP-hard to approximate a solution to (n^a, n^b, n^c) -Nice Set Cover to within a ratio of $\gamma a c \log n$ for some constant γ , where $0 < c \leq a \leq b$.*

Proof. We will need the following fact, proved in [43]. Earlier, this result was shown under the weaker assumption that $NP \not\subseteq DTIME(n^{O(\log \log n)})$ [26, 38].

Fact D.3. *There is a $d > 0$ for which it is NP-hard to approximate a solution to (n^d, n) -Set Cover to within a ratio of $\gamma \log n$, for some $\gamma > 0$.*

Claim D.4. *It is NP-hard to approximate a solution to n -Balanced Set Cover to within a ratio of $\gamma \log n$, for the same γ as above.*

Proof. By Fact D.3, (n^d, n) -Set Cover is not approximable within a factor of $\gamma \log n$, unless $P = NP$. Using a reduction from (n^d, n) -Set Cover, if $|A| = n^d < n$, transform this instance into an instance where $|A| = |B|$ by padding A with dummy vertices. If $|A| > n$, transform this instance into an instance where $|A| = |B|$ by padding the set B with dummy vertices and connecting them to all vertices in A . \square

Applying Lemma 2.3 of [36] to an instance of n^c -Balanced Set Cover, and using Claim D.4, we obtain the following.

Claim D.5. *It is NP-hard to approximate a solution to (n^a, n^c) -Balanced Bounded Set Cover to within a ratio of $\gamma a c \log n$, where γ is from Claim D.3 above.*

To complete the proof of the lemma, notice that a set M is a solution to an instance of (n^a, n^b, n^c) -Nice Set Cover iff M is a solution to the instance of (n^a, n^c) -Balanced Bounded Set Cover, resulted from compressing each set B_i into a single vertex b_i , $1 \leq i \leq n^a$. By Claim D.5 above, it follows that (n^a, n^b, n^c) -Nice Set Cover is not approximable within a ratio of $\gamma a c \log n$, unless P=NP. \square

We now prove the main theorem of this section.

Proof of Theorem D.1. Let $\alpha = 1 + \frac{3}{2k}$ and $\beta = \frac{1}{5k}$. Given $G_1 = V_{k+1} \cup V_{k+2}$, an instance of (n, n^α, n^β) -Nice Set Cover, transform it into the following $k + 2$ -partite graph $G = V_1 \cup V_2 \cup \dots \cup V_{k+1} \cup V_{k+2}$, with edges directed from V_i to V_{i+1} . Let $|V_i| = n$, $1 \leq i \leq k + 1$, and $V_{k+2} = n^\alpha$. The induced subgraph on $V_1 \cup V_2 \cup \dots \cup V_{k+1}$ is $BF(k, n)$, the butterfly graph of diameter k and width n . Then $|G| \leq kn^{1+\frac{1}{k}} + n^{\alpha+\beta} = \Theta(n^{1+\frac{17}{10k}})$ edges. Notice that there are indeed at most $n^{\alpha+\beta}$ edges from V_{k+1} to V_{k+2} since there are n vertices in V_{k+1} , each of degree at most $n^{\beta+\alpha-1}$.

Lemma D.6. $OPT_S = \Theta(OPT_{NSC} n^{\frac{2}{k}})$.

Proof. First, we show that there is a k -TC-spanner H of G s.t. $|H| = \Theta(OPT_{NSC} n^{\frac{2}{k}})$ edges. Then we show that any k -TC-spanner of G must have $\Omega(OPT_{NSC} n^{\frac{2}{k}})$ edges.

Notice that the only pairs of vertices of G that are not already at distance at most k are the comparable vertices u, v , with $u \in V_1$ and $v \in V_{k+2}$. In order to connect such pairs by a directed path of length at most k , we need ‘‘shortcut’’ edges between different levels V_i and V_j , $i + 2 \leq j$. W.l.o.g., we may assume that the only shortcut edges used are those connecting vertices in V_i to V_{i+2} , for some i 's. Indeed, a shortcut edge connecting a vertex $u \in V_i$ to a vertex $v \in V_j$, where $j > i + 2$ can be replaced with one edge connecting $u \in V_i$ to a vertex $w \in V_{i+2}$ that is an ancestor of v . In this way, all paths from V_1 to V_{k+2} that previously had a path of length at most k still have a path of length at most k . Define an edge $e = (u, v)$ to be a *type i edge* if $u \in V_i$ and $v \in V_{i+2}$. We next build a k -TC-spanner of G with $\Theta(OPT_{NSC} n^{\frac{2}{k}})$ edges. Let H be the smallest k -TC-spanner of G which only uses shortcut edges of type $k - 1$.

Claim D.7. $|H| = \Theta(n^{\frac{2}{k}} OPT_{NSC})$

Proof. Let O be a set of vertices in V_{k+1} that is an optimal solution to the (n, n^α, n^β) -Nice Set Cover instance. Connect each vertex $v \in O$ to the set A_v of all the $n^{2/k}$ ancestors of v from level V_{k-1} . Direct these edges from A_v to v . Notice that we added $OPT_{NSC} n^{\frac{2}{k}}$ edges and the new graph H' is a k -TC-spanner. Indeed, each vertex $u \in V_1$ is comparable to each vertex $v \in O$, and thus, there is a vertex $w \in A_v$ that is comparable to u . This implies that for every $u \in V_1$ there is a path of length $k - 1$ to each of the vertices of O , resulting in a path of length k to each vertex in V_{k+2} . To show that H (the minimum size k -TC-spanner with shortcuts only of type $k - 1$) needs at least $OPT_{NSC} n^{\frac{2}{k}}$ edges on top of those in G , assume otherwise. For $v \in V_{k-1}$, let $n(v)$ be the number of type $k - 1$ edges leaving from v . By assumption, $\sum_{v \in V_{k-1}} n(v) < OPT_{NSC} n^{\frac{2}{k}}$. Each vertex in $v \in V_{k-1}$ has exactly $a(v) = n^{1-\frac{2}{k}}$ ancestors in V_1 . For $u \in V_1$, let $e(u)$ be the total number of type $k - 1$ shortcuts leaving from its descendants in V_{k-1} . Since there exists a path of length k from u to each vertex in V_{k+2} , it follows that $e(u) \geq OPT_{NSC}$. Notice that $\sum_{v \in V_{k-1}} n(v)a(v) = \sum_{u \in V_1} e(u) \geq OPT_{NSC} n$. This implies that $\sum_{v \in V_{k-1}} n(v) \geq OPT_{NSC} n^{\frac{2}{k}}$, a contradiction to our assumption, concluding that $|H| = |G| + OPT_{NSC} n^{\frac{2}{k}}$. Next we show that $|H| =$

$\Theta(n^{\frac{2}{k}} OPT_{NSC})$. Indeed, OPT_{NSC} is, by construction, the same as the size of the optimal solution to an (n, n^β) -Balanced Bounded Set Cover instance, where we must cover n vertices on the right with n vertices of degree at most n^β on the left. This implies that $OPT_{NSC} \geq n^{1-\beta} = n^{1-\frac{1}{5k}}$. Now, $|G| = \Theta(n^{1+\frac{17}{10k}})$ and $|H| = |G| + OPT_{NSC} n^{\frac{2}{k}}$. Since $OPT_{NSC} n^{\frac{2}{k}} \geq n^{\frac{2}{k}+1-\frac{1}{5k}} = n^{1+\frac{18}{10k}}$, this implies that $|H| = \Theta(OPT_{NSC} n^{\frac{2}{k}})$. \square

Let M be a sparsest spanner on G which possibly uses shortcut edges of types other than $k-1$. Assume for the sake of contradiction that $|M| < \frac{1}{4} n^{\frac{2}{k}} OPT_{NSC}$. A vertex $u \in V_1$ can reach $v \in V_{k+2}$ in at most k steps by using shortcut edges either of type $i \leq k-1$ or of type k . We will show that, under our assumption, there are many vertices in V_1 that can reach at most $\frac{1}{2} OPT_{NSC}$ vertices in V_{k+1} by using only edges of some types $i < k$. Moreover, there are many vertices in V_1 that reach only $n^{\frac{1}{k}} OPT_{NSC}$ edges of type k . That will be enough to argue that a contradiction must occur, allowing us to conclude that $|M| = \Theta(n^{\frac{2}{k}} OPT_{NSC})$.

Claim D.8. *Let R be the set of vertices in V_1 that can reach less than $\frac{1}{2} OPT_{NSC}$ vertices $v \in V_{k+1}$ in at most $k-1$ steps in M . Then $|R| > \frac{n}{2}$.*

Proof. For each vertex $u \in V_1$ and $v \in V_{k+1}$, define an indicator variable $X_{u,v}$ which is 1 iff there is a shortcut edge along the unique path from u to v in G . Consider a type i shortcut edge $e = (v_i, v_{i+2})$, with $v_i \in V_i$ and $v_{i+2} \in V_{i+2}$. Then there are $n^{\frac{i-1}{k}}$ vertices u in V_1 with $u \leq v_i$. Moreover, there are $n^{\frac{k-i-1}{k}}$ vertices $v \in V_{k+1}$ with $v_{i+2} \leq v$. Thus, this shortcut edge e can set at most $n^{\frac{i-1}{k} + \frac{k-i-1}{k}} = n^{1-\frac{2}{k}}$ different $X_{u,v}$ to 1. By assumption, there are less than $\frac{1}{4} n^{\frac{2}{k}} OPT_{NSC}$ shortcut edges of types i , where $i \leq k-1$. It follows that less than $\frac{1}{4} n OPT_{NSC}$ different $X_{u,v}$'s can be set to 1. For $u \in V_1$, let $n(u)$ be the number of vertices $v \in V_{k+1}$ that u can reach in less than k steps. Thus, $\mathbb{E}_{u \in V_1} [n(u)] < \frac{1}{4} OPT_{NSC}$. By Markov's inequality, $Pr_{u \in V_1} [n(u) \geq \frac{1}{2} OPT_{NSC}] < \frac{1}{2}$. This implies that more than $\frac{1}{2}$ of the vertices $u \in V_1$ can reach less than $\frac{n}{2} OPT_{NSC}$ vertices $v \in V_{k+1}$ in less than k steps. Therefore, $|R| > \frac{n}{2}$. \square

We say that a vertex u reaches an edge $e = (v, w)$, if there is a path from u to v . For $u \in V_1$ let $t(u)$ be the number of type k edges that u reaches in M .

Claim D.9. *Let S be the set of vertices $u \in V_1$ s.t. $t(u) < \frac{1}{2} n^{\frac{1}{k}} OPT_{NSC}$. Then $|S| > \frac{n}{2}$.*

Proof. Assuming $|M| < \frac{1}{4} n^{\frac{2}{k}} OPT_{NSC}$, there are at most $\frac{1}{4} n^{\frac{2}{k}} OPT_{NSC}$ edges of type k . Each $v \in V_k$ has exactly $n^{1-\frac{1}{k}}$ ancestors in V_1 , and therefore $\sum_{u \in V_1} t(u) < n^{1-\frac{1}{k}} \frac{1}{4} n^{\frac{2}{k}} OPT_{NSC} = \frac{1}{4} n^{1+\frac{1}{k}} OPT_{NSC}$.

Thus, $\mathbb{E}_{u \in V_1} [t(u)] < \frac{1}{4} n^{\frac{1}{k}} OPT_{NSC}$ and by Markov's inequality, $Pr_{u \in V_1} [t(u) < \frac{1}{2} n^{\frac{1}{k}} OPT_{NSC}] > \frac{1}{2}$. \square

Let $T = R \cap S$. The two claims above imply $|T| \geq 1$. Now we argue that a vertex $v \in T$ cannot reach some vertices in V_{k+2} . Recall that an instance of (n, n^α, n^β) -Nice Set Cover was obtained from an instance of (n, n^β) -Balanced Bounded Set Cover, by copying each vertex on the right $n^{\alpha-1}$ times, which means that the optimal solution to one of them is also an optimal solution to the other. Suppose we remove $\frac{1}{2} n^{\frac{1}{k}} OPT_{NSC}$ vertices from V_{k+2} . This corresponds to removing at most $\frac{1}{2} n^{\frac{1}{k}+1-\alpha} OPT_{NSC} = \frac{1}{2} n^{-\frac{1}{2k}} OPT_{NSC} = o(1) OPT_{NSC}$ vertices from the universe of the related (n, n^β) -Balanced Bounded Set Cover instance. Let OPT_{BSC} be the size of a solution to this new Set Cover problem. Then $OPT_{BSC} \geq (1 - o(1)) OPT_{NSC}$.

Suppose then that $v \in T$ could cover all of the elements in V_{k+2} . Each such vertex $v \in T$ can cover vertices in V_{k+2} in exactly two ways: (1) from the $\frac{1}{2} OPT_{NSC}$ vertices it reaches in V_{k+1} via paths of length $< k$ using type $i < k$ edges, and (2) by at most $\frac{1}{2} n^{\frac{1}{k}} OPT_{NSC}$ type k edges it can reach. Thus we must have $OPT_{BSC} \leq \frac{1}{2} OPT_{NSC}$, which is a contradiction since $OPT_{BSC} \geq (1 - o(1)) OPT_{NSC}$. Thus, $v \in T$

cannot reach all of V_{k+2} , and so the optimal k -TC-spanner on G must have size at least $n^{\frac{2}{k}} OPT_{NSC}/4$. We can then conclude that $|M| = \Theta(n^{\frac{2}{k}} OPT_{NSC})$. \square

Suppose now that we could approximate the size of the sparsest k -TC-spanner within $\gamma_1 \log n$ for some $\gamma_1 > 0$. Then, since $|M| = \Theta(n^{\frac{2}{k}} OPT_{NSC})$, we could approximate a solution to (n, n^α, n^β) -Nice Set Cover within $\gamma_2 \log n$, for some $\gamma_2 > 0$. By Lemma D.2 above, (n, n^α, n^β) -Nice Set Cover cannot be approximated within $\gamma\beta \log n = O(\frac{1}{k}) \log n$, unless $P=NP$. Therefore, the size of the sparsest k -TC-spanner cannot be approximated within a factor $\gamma_3 \frac{1}{k} \log n$, for some $\gamma_3 > 0$, unless $P=NP$. \square

E Constructing Sparse k -TC-Spanners for Path Separable Graphs

Definition E.1 ([1]). *Let G be a connected undirected graph with n vertices. G is (s, m) -path separable (for $m \geq n/2$) if for any rooted spanning tree T of G either (1) there exists a set P of at most s monotone paths³ in T so that each connected component of $G \setminus P$ is of size at most m , or (2) for some $s' < s$, there exists a set P of s' monotone paths in T so that the largest connected component of $G \setminus P$ is $(s - s', m)$ -path separable. G is said to be s -path separable if G is $(s, n/2)$ -path separable. If G is s -path separable, let S be the union of at most s paths in G such that each connected component of $G \setminus S$ is of size at most $n/2$. S is called the s -path separator of G . A digraph G' is called an s -path separable digraph if the undirected graph underlying $TR(G')$ is s -path separable.*

In the above definition, the number of vertices in the path separator is left unspecified. Trees are 1-path separable, since S can be taken to be the centroid. Similarly, graphs of treewidth w are $(w + 1)$ -path separable. Thorup [50] showed that every planar graph is 3-path separable. Indeed, in the case of any planar graph G and any rooted spanning tree for it, Thorup proved that there exists a set of 3 root paths of the tree whose removal disconnects the graph into components of size at most $n/2$. Abraham and Gavoille [1] studied the more general case of H -minor-free graphs and proved the following.

Theorem E.1 (Theorem 1 of [1]). *Every H -minor-free⁴ graph is s -path separable, for $s = s(H)$, and an s -path separator can be computed in polynomial time.*

The definition of path separability used by Abraham and Gavoille is slightly different from Definition E.1. However, the separators produced in their proof of the theorem satisfy our notion of path separability [29]. Our main theorem in this section is the following:

Theorem E.2. *If G is a graph drawn from a minor-closed graph family that is s -path separable, for $s = \Theta(1)$, then G has a 2-TC-spanner of size $O(n \log^2 n)$ and, more generally, a k -TC-spanner of size $O(n \cdot \log n \cdot \lambda_k(n))$ where $\lambda_k(\cdot)$ is the k -row inverse Ackermann function.*

Since the families of bounded treewidth, planar graphs, and H -minor-free graphs (where H is a fixed minor) satisfy the hypotheses of the theorem, these families have 2-TC-spanners of size $O(n \log^2 n)$.

Proof of Theorem E.2. First we describe a preprocessing step resembling [50] in which the digraph is divided into subgraphs so that constructing TC-spanners for each subgraph individually results in a TC-spanner for the entire graph. Then, we show how to efficiently construct sparse 2-TC-spanners for each of these path separable subgraphs. Lastly, we give the construction for general k .

Preprocessing Step. Let G be a transitively reduced digraph. Choose an arbitrary vertex $r \in V(G)$. Let L_0 be the set containing r and all vertices reachable from r by a directed path. For $i \geq 1$, let $L_{2i} \stackrel{\text{def}}{=} \{v \in$

³A monotone path in a rooted tree is a subpath of a path with one endpoint at the root.

⁴A graph is called H -minor-free if it belongs to a minor-closed graph family that excludes H .

$G \setminus \cup_{j=0}^{2i-1} L_j : \exists u \in \cup_{j=0}^{2i-1} L_j$ s.t. $u \rightsquigarrow v$ and $L_{2i-1} \stackrel{\text{def}}{=} \{v \in G \setminus \cup_{j=0}^{2i-2} L_j : \exists u \in \cup_{j=0}^{2i-2} L_j$ s.t. $v \rightsquigarrow u\}$. Then L_0, L_1, \dots, L_t partition the vertices in G , for some integer $t \leq n$. Evidently,

Claim E.3. *For any vertices $u, v \in G$, if $u \rightsquigarrow_G v$ and if $u \in L_i$ and $v \in L_j$, then $|i - j| \leq 1$.*

For $1 \leq i \leq t$, let $G_i \stackrel{\text{def}}{=} L_{i-1} \cup L_i$. By claim E.3, any two vertices with a dipath between them must both be contained in some G_i . Moreover, any dipath between them must lie entirely in G_i . Therefore, a k -TC-spanner for G is the union of k -TC-spanners for each G_i . Notice that $\sum_i |V(G_i)| \leq 2|V(G)|$.

We next construct a spanning tree T_G for the undirected graph underlying G that is rooted at r and has the following property: for any undirected path in T_G from the root, the restriction of the path to a single level L_i consists of a single directed path.

T_G can be constructed inductively. First, since by definition r reaches all the vertices in L_0 , a spanning tree of L_0 rooted at r can be constructed with all edges oriented away from r . Now suppose we have a tree T_{i-1} that is rooted at r , spans all vertices in $\cup_{j=0}^{i-1} L_j$, and whose restriction to each level $0, \dots, i-1$ consists of a single directed path. If i is odd, T_{i-1} can be extended to a tree T_i where all the new edges are oriented towards $\cup_{j=0}^{i-1} L_j$. The case when i is even is symmetric. Our desired spanning tree T_G is T_t . The following lemma is immediate by the construction.

Lemma E.4. *A monotone path in T_G restricted to G_i , for any $i \in [t]$, is a concatenation of ≤ 2 dipaths.*

We assume G is transitively reduced and connected. If G is not connected, we can apply our algorithm on each component. We describe how to construct H , a 2-TC-spanner for G of size $O(n \log^2 n)$.

The recursive graph fragmentation. First, we apply the preprocessing step described above to $G^0 \stackrel{\text{def}}{=} G$; that is, we obtain a spanning tree T_{G^0} and a collection of subgraphs, G_1^0, G_2^0, \dots . By definition of path separability, there exists a set P^0 of monotone paths on T_{G^0} such that one of two situations happens: (1) all the connected components in $G^0 \setminus P^0$ are of size at most $n/2$, (2) the largest component of $G^0 \setminus P^0$ is of size greater than $n/2$ and is path separable. Let G^1 denote the induced subgraph of G^0 on the largest component of $G^0 \setminus P^0$. We can apply the preprocessing to G^1 to obtain a collection of subgraphs G_1^1, G_2^1, \dots and a spanning tree T_{G^1} rooted at some arbitrary vertex in G^1 . Again, we find an appropriate set of paths P^1 in T_{G^1} and we recurse if necessary on the largest component of $G^1 \setminus P^1$. The recursion ends when the graph has been disconnected into components of size at most $n/2$. Notice that the total number of paths in $P^0 \cup P^1 \cup \dots$ is at most $s = \Theta(1)$, and we then recurse only a constant number of times. Let $S \stackrel{\text{def}}{=} P^0 \cup P^1 \cup \dots$.

Connecting the cut pairs in G . Call a pair of vertices (u, v) a *cut pair* if $u \rightsquigarrow_G v$ and every directed path from u to v intersects a path in S . We show how to connect every cut pair by a path of length at most 2.

Repeat the following for every vertex $v \in V(G)$. Let $I = \{i : v \in V(G^i)\}$ and, additionally, for each $i \in I$, let $J_i = \{j : v \in V(G_j^i)\}$. Do the following for each $i \in I$ and each $j \in J_i$. Let P_j^i denote the restriction of the paths in P^i to G_j^i . Each undirected path in P_j^i is a concatenation of at most 2 directed paths by Lemma E.4. Break up the paths in P_j^i into dipaths. Consider some dipath $P \in P_j^i$ which visits the vertices p_1, p_2, \dots, p_m in that order, where $m \leq |V(G_j^i)|$. For simplicity of presentation, assume m is a power of 2. For each $1 \leq z \leq \log_2 m$, add the following two edges in H : (i) an edge from v to $p_{y_1 \cdot m/2^z}$ where $y_1 = \min_y \{1 \leq y < 2^z : v \rightsquigarrow p_{y \cdot m/2^z} \text{ in } G\}$ and (ii) an edge to v from $p_{y_2 \cdot m/2^z}$ where $y_2 = \max_y \{1 \leq y < 2^z : p_{y \cdot m/2^z} \rightsquigarrow v \text{ in } G\}$. If any of the sets inside the min or max is empty, do not add the respective edge. Finally, add an edge (v, p_m) if $v \rightsquigarrow p_m$ in G and (p_m, v) if $p_m \rightsquigarrow v$ in G . Repeat this process for every separator dipath that is a subpath of an undirected path in P_j^i .

Outer Recursion. For each connected component C of $G \setminus S$, recurse on the subgraph induced by C . C is also path separable since the graph family is minor-closed.

Lemma E.5. *The above construction efficiently produces a 2-TC-spanner on G of size at most $O(n \log^2 n)$.*

Proof Sketch. In each G_j^i , the separators are nicely structured as only a constant number of directed paths. Hence, we can add only $O(|V(G_j^i)| \log |V(G_j^i)|)$ edges in order to connect the cut pairs present in each G_j^i . Since $\sum_j |V(G_j^i)| \leq 2n$ and the number of G^i 's is $\Theta(1)$, the total number of edges added in each step of the outer recursion is $O(n \log n)$. The size of the remaining connected components halves after each graph fragmentation step. So, the outer recursion continues only $\log n$ times, making the total number of added edges $O(n \log^2 n)$. The construction results in a 2-TC-spanner because every pair of related vertices (u, v) is a cut pair at some level of the outer recursion. Then, u and v are both contained in some G^i . One can check that the above construction ensures that both u and v are adjacent to the same vertex on some separator dipath intersecting a dipath from u to v . The formal proof is below.

Proof of Lemma E.5. Let us first see why connecting every cut pair by a path of length at most 2, and recursing on smaller components produce a 2-TC-spanner for G . Indeed, if (u, v) is a cut pair, then the first step ensures a path of length at most 2 between them. If (u, v) is not a cut pair but there exist some dipaths from u to v , then u and v are in the same component C of $G \setminus S$, and there exists a dipath between them that lies entirely within this component. In this case, constructing a 2-TC-spanner on the subgraph induced by this C suffices to connect u and v by a path of length at most 2.

Let us now argue that this process connects every cut pair by a path of length at most 2. Consider some cut pair (u, v) . Let i be the smallest nonnegative integer such that every dipath from u to v intersects a path in $\cup_{i' \leq i} P_{i'}$. Therefore, there must be a dipath from u to v entirely contained in G^i , and by claim E.3, it follows that there is a j such that both u and v are in G_j^i . Suppose $P \in P_j^i$ is a separator dipath of length m (a power of 2) that intersects a dipath in G^i from u to v . Let $y_1 = \min_y \{y : u \rightsquigarrow_{G^i} p_y\}$ and $y_2 = \max_y \{y : p_y \rightsquigarrow_{G^i} v\}$. $y_1 \leq y_2$ because otherwise there cannot be a vertex on P that lies on a path from u to v . One possibility is that $y_1 = y_2 = m$ in which case the construction ensures that we add the edges (u, p_m) and (p_m, v) . Otherwise, there exists some $z \in \{1, 2, \dots, \log_2 m\}$ such that there is a unique $y \in \{1, 2, \dots, 2^z - 1\}$ for which $y \cdot m/2^z$ is in the interval $[y_1, y_2]$. Moreover, $u \rightsquigarrow p_{y_1} \rightsquigarrow p_{y \cdot m/2^z} \rightsquigarrow p_{y_2} \rightsquigarrow v$. Therefore, the construction above adds the edges $(u, p_{y \cdot m/2^z})$ and $(p_{y \cdot m/2^z}, v)$.

In connecting the cut pairs in G_j^i , we add at most $O(s |V(G_j^i)| \log |V(G_j^i)|)$ because there are at most $O(s)$ separator dipaths in G_j^i and for any separator dipath P , each vertex in G_j^i is connected to at most $2(\log_2 |V(P)| + 1) \leq O(\log |V(G_j^i)|)$ vertices on the path. Recall that there are only a constant number of G^i 's and $\sum_j |V(G_j^i)| \leq 2|V(G^i)|$. Thus, if $S(G)$ denotes the total number of edges in the constructed 2-TC-spanner for G , we have that:

$$S(G) \leq \sum_{C \text{ is a c.c. of } G \setminus S} S(C) + O \left(\max_i \sum_j O(|V(G_j^i)| \cdot \log |V(G_j^i)|) \right) \leq \sum_{C \text{ is a c.c. of } G \setminus S} S(C) + O(n \log n)$$

Since $|V(C)| \leq n/2$ for any connected component of $G \setminus S$, it follows that $S(G) = O(n \log^2 n)$.

If the strong path separators can be found in polynomial time, as is guaranteed, for example, in Theorem E.1, then it is clear that the above 2-TC-spanner can be constructed efficiently. \square

We now prove the part of Theorem E.2 concerning k -TC-spanners for general k . Again, assume G is transitively reduced and connected; now, we wish to construct H , a k -TC-spanner for G . We perform the same preprocessing as before in order to obtain induced subgraphs G^0, G^1, \dots and a corresponding s -path separator $S = P^0 \cup P^1 \cup \dots$. Define a cut pair (u, v) to be a pair of vertices in G such that $u \rightsquigarrow v$ and every directed path from u to v intersects a path in S . This time, our plan is to connect all cut pairs by a path of length at most k and then to recurse on each of the connected components that remain after removing the vertices in the paths of S . By the argument used earlier, this process produces a k -TC-spanner.

Now we show how to connect cut pairs (u, v) with a path of length at most k . Do the following for every vertex $v \in V(G)$. Let $I = \{i : v \in V(G^i)\}$ and for each $i \in I$, let $J_i = \{j : v \in V(G_j^i)\}$. Do the following for each $i \in I$ and for each $j \in J_i$. Let P_j^i be the restriction of the paths in P^i to G_j^i . Break up the undirected paths into dipaths, increasing the size of P_j^i by a factor of at most 2. Do the following for each dipath $P \in P_j^i$. Let m be the length of P which visits vertices p_1, p_2, \dots, p_m in that order. Additionally, let $c(\ell)$ be a concave increasing function of ℓ , which satisfies $c(\ell) < \ell$ that we specify later. For simplicity of presentation, we omit all floors and ceilings. Let $c^*(\ell)$ denote the smallest z such that $c^z(\ell) = \Theta(1)$ where $c^z(\cdot)$ denotes the z th functional power of c . For each z such that $1 \leq z \leq c^*(m)$, add the following two edges to H : (i) an edge from v to $p_{y_1 \cdot c^z(m)}$ where $y_1 = \min_y \{1 \leq y < m/c^z(m) : v \rightsquigarrow p_{y \cdot c^z(m)} \text{ in } G\}$ and (ii) an edge to v from $p_{y_2 \cdot c^z(m)}$ where $y_2 = \max_y \{1 \leq y < m/c^z(m) : p_{y \cdot c^z(m)} \rightsquigarrow v \text{ in } G\}$. If any of the sets inside the min or max is empty, do not add the respective edge.

Finally, do the following for every dipath $P \in P_j^i$ that visits vertices p_1, p_2, \dots, p_m in that order.

CONNECT-ON(P)

1. Add to H the edges in a $(k - 2)$ -spanner of the induced subgraph of the transitive closure of G on $\{p_{c(m)}, p_{2c(m)}, \dots, p_m\}$.
2. Remove the points in the set $\{p_{c(m)}, p_{2c(m)}, \dots, p_m\}$ from P and run CONNECT-ON on each connected component of P that remains.

This completes our description of H . It is not too hard to see that H is indeed a k -TC-spanner, using reasoning as in the previous section. The only difference is that now, for a cut pair (u, v) , it could be that u and v are adjacent to different vertices on the separating dipath. But we have the guarantee by CONNECT-ON above that two path vertices in the same recursion level of CONNECT-ON have a path of length at most $k - 2$ between them. Hence, it follows that u and v have a path of length at most k between them.

Now, we bound the size of H . First, let us count the number of edges added in each step of the main recursion (that is, not counting the edges needed to connect pairs within components of size at most $n/2$). Denote by $\ell(n, k)$ the quantity $S_k(L_n)$, the size of the optimal k -TC-spanner for the directed line on n vertices. Let us count all the edges added that are incident to some separating dipath P of size m . Denote this quantity $f(m)$. By the definition of CONNECT-ON:

$$f(m) \leq O(n) + n/c(m)f(c(m)) + \ell(m/c(m), k - 2)$$

It can be seen that $f(m)$ is minimized when $\ell(m/c(m), k - 2) = O(m) = O(n)$. For example, for $k = 4$, $\ell(n, 2) = O(n \log n)$ and in this case, $c(m)$ should be chosen to be $\log m$ since $\frac{m}{\log m} \cdot \log \frac{m}{\log m} = O(m)$. In any case, once $c(m)$ is fixed, the solution to the above functional equation turns out to be: $f(m) \leq O(n \cdot c^*(n))$. Also, the number of edges added to the vertices not on the separating paths is also $O(n \cdot c^*(n))$. Making the same arguments as in the analysis of the 2-TC-spanner, we find that the number of edges added in total, counting all the paths and all the G^i 's is still $O(n \cdot c^*(n))$. Since there are $\log n$ levels of the recursion at the top level (at each level, the size of the largest component is decreased by a factor of 2), the total number of edges added to H is $O(n \cdot \log n \cdot c^*(n))$. Finally, we use the results of [5] about optimal k -TC-spanners of the directed line to conclude that $c^*(n) = O(\lambda_k(n))$. \square

F NP-Hardness of k -TC-spanner

Theorem D.1 breaks down for $k = \Omega(\log n)$. For these large values of k we have the following.

Theorem F.1. For any $k < n^{1-\epsilon}$ for any $\epsilon > 0$, it is NP-hard to approximate the size of the sparsest k -TC-spanner within a factor of $1 + \gamma$, for some $\gamma = \Omega\left(\frac{1}{k}\right)$.

Proof of Theorem F.1. We use a reduction from 3NODECOVER to show that, unless $P = NP$, k -TC-SPANNER cannot be approximated within a factor of $1 + \Omega\left(\frac{1}{k}\right)$. That is, for constant k , the problem is APX-hard. An instance of 3-Node Cover consists of a collection D of subsets of a universe X . Each subset contains at most 3 elements, and each element of X is contained in at most 2 subsets. The goal is to output a minimum size subcollection $M \subseteq D$ whose union is X . We need the following result in [11]:

Lemma F.2. 3NODECOVER is NP-hard to approximate within a factor of $1 + c$, for some constant $c > 0$.

We now give a reduction from 3NODECOVER to k -TC-SPANNER. For a given instance R of 3NODECOVER we construct the following graph G . Let V_1 be the set of vertices representing each set $d \in D$. Let V_2 be the set of vertices representing each element $t \in X$. Draw a directed edge from each vertex in V_1 corresponding to $d \in D$ to the vertices in V_2 corresponding to elements of d . Add an extra vertex a and draw directed edges from a to every element of V_1 . For each vertex $v \in V_1$ add $k - 1$ new vertices v_1, v_2, \dots, v_{k-1} and connect them via a directed path of length k passing through $a, v_1, v_2, \dots, v_{k-1}, v$ in the given order. Call this path $P(v)$.

Let OPT_S be the size of a minimum k -TC-spanner on G and OPT_{3NC} be the size of the solution to the initial instance $R = (D, X)$ of 3NODECOVER. Let $|D| = n$.

Claim F.3. $OPT_S = OPT_{3NC} + kn + \sum_{d \in D} |d|$.

Proof. We show that there is a sparsest spanner H that contains only edges from V_1 to V_2 , from vertex a to some vertices in V_1 and the edges on the paths $P(v)$, for all $v \in V_1$.

Note that all edges from V_1 to V_2 need to be included in a sparsest spanner, since each of them forms a unique directed path connecting its endpoints. There are $\sum_{d \in D} |d|$ such edges. Similarly, for each vertex $v \in V_1$, all the edges on the path $P(v)$ need to be included in H . In total, there are kn such edges.

For $v \in V_1$, suppose that H contains an edge (v_i, t) , for some $v_i \in P(v)$ and $t \in V_2$. We claim that such an edge (v_i, t) can be removed and substituted with an edge (a, u) , where $u \in V_1$ s.t. u is adjacent to t . Indeed, all vertices on $P(v)$ except for a are already at distance $\leq k$ from t in H . Thus, to reach t from a it is enough to include in H a directed edge from a to any $u \in V_1$ that is adjacent to t .

Similarly, an edge e between vertices on a path $P(v)$ in H can be replaced by an edge (a, v) . Indeed, such an edge e can only be useful to connect a to some vertices in V_2 , via a path that passes through v .

Therefore, among the edges from a to V_1 a sparsest spanner need only contain the edges that connect a to a minimum set of vertices in V_1 that cover V_2 . Thus, there are exactly OPT_{3NC} such edges. \square

Suppose that there exists $0 < \gamma$ and an algorithm \mathcal{A} that computes the size of a sparsest k -TC-spanner within $1 + \gamma$. Namely, \mathcal{A} outputs s , such that $OPT_S \leq s \leq (1 + \gamma)OPT_S$. We show that $\gamma \geq \frac{c}{19+6k}$, where c is the constant from Lemma F.2.

Each set d contained in an optimal solution to R covers at most 3 elements of the universe X . Therefore $|X| \leq 3OPT_{3NC}$. Any element of X is contained in at most 2 sets of D , and therefore, $|X| \geq \frac{n}{2}$. This implies that $n \leq 6OPT_{3NC}$. Let $s' = s - (k + 3)n$. Then

$$\begin{aligned} OPT_{3NC} \leq s' &\leq OPT_{3NC} + \gamma(OPT_{3NC} + kn + 3n) \\ &\leq OPT_{3NC} + \gamma(OPT_{3NC} + 6kOPT_{3NC} + 18OPT_{3NC}) \\ &= OPT_{3NC}(1 + \gamma(19 + 6k)). \end{aligned}$$

Finally, using F.2, it follows that $\gamma \geq \frac{c}{19+6k}$. Thus $\gamma = \Omega\left(\frac{1}{k}\right)$. \square