

Transitive-Closure Spanners: A Survey^{*}

Sofya Raskhodnikova^{1**}

Pennsylvania State University, USA
sofya@cse.psu.edu

Abstract. We survey results on transitive-closure spanners and their applications. 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 different areas of computer science, and properties of these spanners have been rediscovered over the span of 20 years. The common task implicitly tackled in these diverse applications can be abstracted as the problem of constructing sparse TC-spanners.

In this article, we survey combinatorial bounds on the size of sparsest TC-spanners, and algorithms and inapproximability results for the problem of computing the sparsest TC-spanner of a given directed graph. We also describe multiple applications of TC-spanners, including property testing, property reconstruction, key management in access control hierarchies and data structures.

1 Introduction

A *spanner* is 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 Awerbuch [10] and Peleg and Schäffer [43] in the context of distributed computing, spanners for undirected graphs have found numerous applications, including efficient routing [22, 23, 45, 46, 55], simulating synchronized protocols in unsynchronized networks [44], parallel and distributed algorithms for approximating shortest paths [20, 21, 27], and algorithms for distance oracles [11, 56].

In the setting of directed graphs, three notions of spanners have been proposed: the direct generalization of the above definition [43], *roundtrip spanners* [23, 46] and *transitive-closure spanners* [15]. In this survey, we focus on the latter definition. It captures the notion that a spanner should have a small diameter but preserve the connectivity of the original graph. By diameter¹, we mean the

^{*} Parts of this survey are adapted from [15–17, 14].

^{**} Supported by National Science Foundation (NSF/CCF award 0729171 and NSF/CCF CAREER award 0845701).

¹ The definition of diameter used in this survey and other papers on transitive-closure spanners is nonstandard. The diameter is usually defined as the largest distance

largest distance between a pair (u, v) of nodes in a directed graph such that v is reachable from u .

Recall that the *transitive closure* of a graph $G = (V, E)$ is a graph (V, E_{TC}) where $(u, v) \in E_{TC}$ if and only if there is a directed path from u to v in G .

Definition 1.1 (TC-spanner [15]). *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 $d_H(u, v) \leq k$.

The edges from the transitive closure of G that are added to G to obtain a TC-spanner are called **shortcuts**, and the parameter k is called the **stretch**.

Notice that a k -TC-spanner of G is a directed k -spanner of the transitive closure of G . Nevertheless, TC-spanners are interesting in their own right due to the multiple TC-spanner-specific applications.

Before TC-spanners were introduced in [15], they were studied implicitly in access control, property testing, and data structures, and properties of these combinatorial objects have been discovered and rediscovered over the span of 20 years. In this work, we survey results on TC-spanners and their applications. We start by discussing a simple example of a TC-spanner of a directed line, which has been studied in different areas under different guises.

1.1 A Simple Example: TC-spanners of the Directed Line


Directed acyclic graphs (DAGs) represent the most interesting case in applications of TC-spanners. There is also a reduction from constructing TC-spanners of graphs with cycles to constructing TC-spanners of DAGs, with a small loss in stretch, which we present in Section 3.2. In this section, we illustrate the definition of TC-spanners by constructing sparse TC-spanners of the simplest DAG—the directed line. The directed line L_n consists of nodes² $[n]$ and edges $\{(i, i + 1) : 1 \leq i \leq n - 1\}$. As discussed in Section 3.1, TC-spanners of the line were implicitly studied in different contexts by multiple authors, many of whom discovered the optimal constructions we describe here.

An additional motivation for considering optimal TC-spanners of the directed line in a separate section is that it gives one of the simplest settings where inverse Ackermann’s functions (see Definitions 1.2) arise naturally—simple enough to explain in an undergraduate algorithms class.

It is easy to see that the transitive closure of L_n has $\binom{n}{2} = \Theta(n^2)$ edges. A TC-spanner, even of the smallest possible stretch—stretch 2, can be much sparser than the transitive closure.


between a pair of nodes in a graph, and is set to infinity if a graph contains a pair of nodes with no path from one to the other.

² We use $[n]$ to denote $\{1, 2, \dots, n\}$.



Lemma 1.1 (2-TC-spanner of the line). *For all $n \geq 3$, the directed line L_n has a 2-TC-spanner with at most $n \log n$ edges³.*

Proof. Our 2-TC-spanner, H , is a graph with vertex set $[n]$. We construct the edge set of H recursively. First, define the middle node $v_{mid} = \lceil \frac{n}{2} \rceil$. Use this node as a *hub*: namely, add edges (v, v_{mid}) for all nodes $v < v_{mid}$ and edges (v_{mid}, v) for all nodes $v > v_{mid}$. Then recurse on the two line segments resulting from removing v_{mid} from the current line. Proceed until each line segment contains exactly one node.




H is a 2-TC-spanner of the line L_n , since every pair of nodes $u, v \in [n]$ is connected by a path of length at most 2 via a hub. This happens in the stage of the recursion during which u and v are separated into different line segments, or one of these two nodes is removed.

To get the bound on the size of the 2-TC-spanner, observe that there are $\lceil \log n \rceil$ stages of the recursion. In each stage, every non-hub node connects to the hub in its current line segment, adding a total of at most n edges. Therefore, the constructed spanner has at most $n \log n$ edges. \square

The same idea can be extended to construct a 3-TC-spanner of the line graph:

Lemma 1.2 (3-TC-spanner of the line). *The directed line L_n has a 3-TC-spanner with $O(n \log \log n)$ edges.*

Proof. Again we construct the edge set of our 3-TC-spanner H recursively. For simplicity, assume that \sqrt{n} is an integer. Designate nodes which are the multiples of \sqrt{n} as hubs. Connect each non-hub node to the nearest hub before it and the nearest hub after it. More precisely, for each non-hub node v , let v_ℓ be v rounded down to the nearest multiple of \sqrt{n} and let $v_r = v_\ell + \sqrt{n}$. Add edge (v_ℓ, v) if $v_\ell \in [n]$ and edge (v, v_r) if $v_r \in [n]$.



³ Logarithms in this article are base 2 unless indicated otherwise.

Also, add edges between all hubs, orienting them from the smaller to the larger. Finally, remove the hubs from the current line and recurse on the \sqrt{n} resulting line segments. Proceed until each line segment contains exactly one node.

H is a 3-TC-spanner of the line L_n , since every pair of nodes $u, v \in [n]$ is connected by a path of length at most 3 via a pair of hubs. This happens in the stage of the recursion where u and v are separated into different line segments, or one of these two nodes is removed. For example, we add a path (u, u_r, v_ℓ, v) if u and v are not hubs.

Denote the number of edges in the spanner by $T(n)$. At the first stage of recursion, we add $\binom{\sqrt{n}}{2} \leq n$ edges to connect the hubs and at most 2 edges per non-hub node to connect non-hubs to hubs. Therefore, $T(n)$ satisfies the following recurrence:

$$T(n) \leq \begin{cases} 0 & \text{if } n \leq 1; \\ 3n + \sqrt{n} \cdot T(\sqrt{n}) & \text{if } n > 1. \end{cases}$$

The solution to this recurrence is $T(n) \leq 3n \log \log n$. □

This construction generalizes to TC-spanners of arbitrary constant stretch k , giving k -TC-spanners of size $O(n \cdot \lambda_k(n))$, where $\lambda_k(n)$ are very slowly growing functions of n , called the k^{th} -row inverse Ackermann functions.

Definition 1.2 (Inverse Ackermann functions⁴). *Let $\mathbb{R}^{\geq 0}$ be the set of non-negative real numbers. For every function $f : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$ satisfying $f(n) < n$ for all $n > 1$, define the function $f^*(n) : \mathbb{R}^{\geq 0} \rightarrow \mathbb{R}^{\geq 0}$ as:*

$$f^*(n) = \min\{k \in \mathbb{Z}^{\geq 0} : f^{(k)}(n) < 2\},$$

where $f^{(k)}$ denotes f composed with itself k times.

Define the k^{th} -row inverse Ackermann function $\lambda_k(n)$ as follows:

$$\lambda_0(n) = n/2, \lambda_1(n) = \sqrt{n}, \text{ and } \lambda_k(n) = \lambda_{k-2}^*(n) \text{ for } k \geq 2.$$

Intuitively, $f^*(n)$ represents the number of times f can be applied before the answer drops below 2. If $f(n) = n/2$ this number is $\Theta(\log n)$, if $f(n) = \sqrt{n}$ it is $\Theta(\log \log n)$, and if $f(n) = \log n$ it is $\log^* n$. Therefore, $\lambda_2(n) = \Theta(\log n)$, $\lambda_3(n) = \Theta(\log \log n)$ and $\lambda_4(n) = \Theta(\log^* n)$. Also note that $\lambda_k(n)$ is a non-decreasing function of n for all $k \geq 0$.

Lemma 1.3 (k -TC-spanner of the line). *The directed line L_n has a k -TC-spanner with at most $k \cdot n \cdot \lambda_k(n)$ edges.*

Proof. Lemmas 1.1 and 1.2 imply Lemma 1.3 for $k = 2$ and 3. (Recall that in the proof of Lemma 1.3, we gave an upper bound of $3n \log \log n$ on the size of the

⁴ We define these functions, following the spirit of the presentation by Seidel [50]. The prevalent definition (see, e.g., [5]) is complicated and yields asymptotically equivalent functions.

sparsest 3-TC-spanner of the line, even though we chose to hide the constant in the statement of the lemma.) We prove Lemma 1.3 by induction on k , using the constructions of 2- and 3-TC-spanners as base cases. Our induction hypothesis is that one can construct a $(k-2)$ -TC-spanner of the line L_n with at most $(k-2) \cdot n \cdot \lambda_{k-2}(n)$ edges. To construct a k -TC-spanner for $k > 3$, we proceed as for $k = 3$, but select even more nodes as hubs, connect them using an optimal $(k-2)$ -TC-spanner, add edges from each node to the nearest hub before and the nearest hub after it, and recurse. Then each pair of nodes (u, v) will be connected by a path that jumps from u to the smallest hub greater than u , then follows a path of length at most $k-2$ in the $(k-2)$ -TC-spanner on the hubs to reach the largest hub smaller than v , and uses one more edge to jump to v .

It remains to specify the number of hubs, which we will denote by h , and to analyze the size of the spanner. Let $f(n) = \lambda_{k-2}(n)$. We set $h = \frac{n}{f(n)} - 1$ and recurse on the segments of size $f(n)$. By the induction hypothesis, the size of the optimal $(k-2)$ -TC-spanner on the hubs is at most

$$(k-2) \cdot h \cdot \lambda_{k-2}(h) \leq (k-2) \cdot \frac{n}{f(n)} \cdot f\left(\frac{n}{f(n)}\right) \leq (k-2) \cdot \frac{n}{f(n)} \cdot f(n) \leq (k-2)n.$$

As before, each non-hub connects to at most 2 hubs. Therefore, the number of edges in constructed spanner, $T(n)$, satisfies the following recursion:

$$T(n) \leq \begin{cases} 0 & \text{if } n \leq 1; \\ kn + \frac{n}{f(n)} \cdot T(f(n)) & \text{if } n > 1. \end{cases}$$

The solution is $T(n) = k \cdot n \cdot f^*(n)$. This follows from the fact that $f^*(f(n)) = f^*(n) - 1$ for $n > 1$. Thus, $T(n) = k \cdot n \cdot \lambda_{k-2}^*(n) = k \cdot n \cdot \lambda_k(n)$. \square

As discussed in Section 3.1, the bound in Lemma 1.3 is tight when k is a constant.

1.2 A Brief Overview

TC-spanners were defined by Bhattacharyya *et al.* [15] as a common abstraction for several applications. Prior to that, Thorup [52] 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 on n nodes there are such k -TC-spanners with k polylogarithmic in n . He proved his conjecture for planar graphs [53], but later Hesse [37] gave a counterexample to Thorup's conjecture for general graphs. TC-spanners were also studied for directed trees: implicitly in [5, 9, 19, 25, 60] and explicitly in [18, 54]. The implicit results were interpreted as TC-spanner constructions in [15].

[15] presented several applications of TC-spanners: testing monotonicity of functions, key management in an access hierarchy and data structures for computing partial products in a semigroup. They also studied the computational problem of finding a sparsest k -TC-spanner for a given directed graph. They

presented algorithms and inapproximability results for this problem. Finally, they gave sparse TC-spanner constructions for new graph families.

Later, [16, 14] studied TC-spanners for the directed hypercube and hypergrid and presented an application of TC-spanners to property reconstruction. Steiner TC-spanners (see Definition 3.1) were formally introduced in [17], but studied before that in the context of access control hierarchies by [7] and [49]. Berman, Raskhodnikova and Ruan [13] improved algorithms presented in [15]. Finally, Jha and Raskhodnikova [39] pointed out the application of TC-spanners to testing if a function is Lipschitz.

1.3 Organization of This Survey

We start by introducing notation and basic graph-theoretic background in Section 2. In Section 3, we describe structural results on TC-spanners, that is, combinatorial bounds on their size. Structural results for specific graph families are surveyed in Section 3.1, while general structural results, applicable to all graphs, appear in Section 3.2. In Section 4, we survey results on the computational problem of finding a sparsest TC-spanner of a given directed graph and the more general problem of finding directed spanners. We describe approximation algorithms and hardness results for these problems. Finally, Section 5 presents multiple applications of TC-spanners, including property testing, property reconstruction, key management in access control hierarchies and data structures for computing partial product in a semigroup.

2 Preliminaries and Notation

We write $u \preceq_G v$ to denote that vertex v is reachable from vertex u in graph G . When the graph is clear from the context, we omit G . The *transitive closure* of a directed graph $G = (V, E)$, denoted $TC(G)$, is the directed graph (V, E') , where $E' = \{(u, v) : u \preceq_G v\}$. Vertices u and v are *comparable* if either $(u, v) \in TC(G)$ (that is, u is *below* v or, equivalently, *smaller than* v) or $(v, u) \in TC(G)$ (that is, u is *above* v or, equivalently, *larger than* v). This terminology and notation is usually used for partially-ordered sets (posets), which are equivalent to directed acyclic graphs, but can be also applied to general directed graphs.

A digraph G is *weakly connected* if replacing each directed edge in G with an undirected edge results in a connected undirected graph. A digraph is *strongly connected* if each vertex in the graph is reachable from every other vertex via a directed path. The graph of strongly connected components of a digraph G is the digraph obtained by contracting each strongly connected component into one vertex, while maintaining all the edges between these components.

A *transitive reduction* of G is a digraph G' with the fewest edges for which $TC(G') = TC(G)$. As shown by Aho *et al.* [2], a transitive reduction of a given graph can be computed efficiently via a greedy algorithm. The algorithm contracts each strongly connected component C to a vertex $v(C)$ to get a supergraph H , obtains a supergraph H' by greedily removing edges in H that do not change

its transitive closure, and finally uncontracts $v(C)$ to an arbitrary directed cycle on vertices in C , choosing a representative vertex of C to be incident to the edges incident to $v(C)$ in H' . Directed acyclic graphs have a unique transitive reduction. We say G is *transitively reduced* if G is equal to its own transitive reduction.

3 Overview of Structural Results on TC-spanners

For a directed graph G , we denote the number of edges in G by $|G|$ and the size of the sparsest k -TC-spanner of G by $S_k(G)$. (The size refers to the number of edges.) To put the following results in proper context, observe that if G has n vertices, $S_k(G) = O(n^2)$. Unlike in the undirected setting, where for every $k \geq 1$, all graphs on n vertices have $(2k - 1)$ -spanners with $O(n^{1+1/k})$ edges [6, 42, 56], sparsest TC-spanners (and hence, sparsest directed spanners) can have $\Omega(n^2)$ edges. An example of a graph with a sparsest 2-TC-spanner of size $\Omega(n^2)$ is the complete bipartite graph $K_{\frac{n}{2}, \frac{n}{2}}$ with $n/2$ vertices in each part and all edges directed from the first part to the second. Therefore, most constructions surveyed below are for TC-spanners of specific graph families. Nevertheless, there are several general results, described in Section 3.2.

3.1 TC-spanners of Specific Graph Families


TC-spanners of lines and trees. TC-spanners of the directed lines and directed trees were discovered under many different guises. They were studied implicitly in [5, 9, 19, 25, 60] and explicitly in [18, 54]. Alon and Schieber [5] implicitly gave tight bounds on $S_k(L_n)$. They showed that, for constant k , the size of the sparsest k -TC-spanner of the directed line is $\Theta(n \cdot \lambda_k(n))$, where $\lambda_k(n)$ is the k^{th} -row inverse Ackermann function (see Definition 1.2 and Lemma 1.3). [5] also showed that the smallest k for which $S_k(L_n) = O(n)$ is $O(\alpha(n))$, where $\alpha(n)$ is the inverse Ackermann function. (The inverse Ackermann function is defined by $\alpha(n) = \min\{k \in \mathbb{Z}^{\geq 0} : \lambda_{2^k}(n) \leq 3\}$.) Note that the size of any TC-spanner of L_n is at least $n - 1$, since all edges of the form $(i, i + 1)$ must be present in a TC-spanner to ensure the same connectivity as in L_n . [5, 19, 54] proved that sparsest k -TC-spanners of rooted directed trees asymptotically have the same number of edges as k -TC-spanners of the line.

TC-spanners of planar graphs. Thorup [52] considered a special case of TC-spanners of graphs G that have at most twice as many edges as G . In [53], he proved that all directed planar graphs G on n nodes have such TC-spanners with stretch polylogarithmic in n .

TC-spanners of graphs with small separators (H -minor-free graphs). A graph H is a *minor* of G if H is a (not necessarily induced) subgraph of a graph obtained from G by a sequence of edge contractions. A graph family \mathcal{F} is *minor-closed* if it contains every minor of every graph in \mathcal{F} . For a fixed graph H (e.g., K_5),

a minor-closed family \mathcal{F} is H -minor-free if $H \notin \mathcal{F}$. Examples of such families include planar graphs, bounded treewidth graphs, and bounded genus graphs, explicitly studied in applications in Section 5. Bhattacharyya *et al.* [15] gave an efficient construction of k -TC-spanners of H -minor-free graphs. For constant k , the size of the spanners is $O(n \cdot \log n \cdot \lambda_k(n))$, where $\lambda_k(\cdot)$ is the k^{th} -row inverse Ackermann function. This result allowed [15] to drastically improve monotonicity testers of Fischer *et al.* [33]. The application to monotonicity testing is described in Section 5.

The construction in [15] uses divide-and-conquer approach. A natural first attempt would be to use separators of Lipton and Tarjan [41]. Recall that an s -separator for a graph G on n nodes is a set of s nodes whose removal disconnects G into connected components of size at most $2n/3$. Observe that the proofs of Lemmas 1.1–1.3 implicitly use this approach for the special case of the line graph. There, at every stage graph separators play a role of hubs. To come up with efficient constructions for a wider family of graphs, Bhattacharyya *et al.* use the *path separators* for undirected H -minor free graphs due to Abraham and Gavoille [1]. An s -path-separator⁵ for a graph G on n nodes is a set of s paths whose removal disconnects G into connected components of size at most $2n/3$. For some graph families, path separators can be much smaller than ordinary separators. For example, planar graphs require ordinary separators of size $\Theta(\sqrt{n})$, but are 3-path separable [1]. For a simple case of a 2-dimensional $m \times m$ grid, a Lipton-Tarjan separator has size of $\Omega(m)$, but it is enough to remove m nodes on one path, say, a horizontal line that cuts through the middle, to separate it.




The path separators of Abraham and Gavoille were constructed for undirected graphs. Bhattacharyya *et al.* employ this construction on the undirected graph, resulting from ignoring the directions of the edges of the input directed graph. The resulting path separator for the original directed graph may be the union of many directed paths. Here we only explain the construction for the simple case when a separator for the graph (and a separator for every subgraph obtained

⁵ For a graph G to be s -path-separable one needs to be able to disconnect the graph by removing nodes on at most s paths from any minimum spanning tree of G . To keep our high-level overview simple, we do not get into details.

by removing a separator) consists of a small number of directed paths, as is the case, for instance, for a 2-dimensional grid $[m] \times [m]$ where all edges are directed towards vertices with larger coordinates. More precisely, we focus on the case when there exists an integer s , such that every graph obtained at any recursion stage has a separator with at most s directed paths, and moreover, this separator can be found efficiently. (See [15] for the general treatment.)

Even though we use a 2-dimensional grid as an example in this proof, TC-spanners of d -dimensional hypergrids are treated separately in the current section, after TC-spanners of H -minor-free graphs.

If a separator consists of a constant number of (say, at most s) directed paths, we can construct a k -TC-spanner of each path P in the separator as in the proof of Lemma 1.3. We also need to make sure that our TC-spanner contains short paths between all pairs of nodes that were using P to connect. To accomplish this, for each node u with a path to some node in a separator path P , let u' be the first node in P reachable from u . As we are constructing a k -TC-spanner of P , at each stage of the recursion, we add an edge from u to a hub h whenever we add an edge (u', h) . We deal symmetrically with each node v with a path from some node in P .



Now, if there is a path from u to v via some vertex in P , there is a path of length at most k in the spanner we are constructing. This is because u and v are connected to the same hubs as u' and v' and, as demonstrated in the proof of Lemma 1.3, u' and v' are connected by a path of length at most k via the hubs. Now, we can safely remove the paths in the separator and recurse on the resulting components. To distinguish this recursion from the recursion in the construction of the TC-spanners of the paths, we call it an *outer* recursion. Observe that at each stage of outer recursion we are adding no more edges per node than in the construction for the line—namely, $O(\lambda_k(n))$ edges. This results in $O(n \cdot \lambda_k(n))$ edges per stage. Since there are $O(\log n)$ stages of outer recursion, the constructed k -TC-spanner has size $O(n \cdot \log n \cdot \lambda_k(n))$.

TC-spanners of hypergrids. The *directed hypergrid*, denoted $\mathcal{H}_{m,d}$, has vertex set $[m]^d$ and edge set $\{(x, y) : \exists i \in [d] \text{ such that } y_i - x_i = 1 \text{ and for } j \neq i, y_j = x_j\}$. For the special case $m = 2$, $\mathcal{H}_{2,d}$ is called a *hypercube* and is also denoted by \mathcal{H}_d .

2-TC-spanners of hypergrids are especially relevant for applications in property testing and property reconstruction. TC-spanners of hypergrids of general stretch k are used in the application to key management in an access hierarchy. The following results on TC-spanners of hypergrids are from [16, 14].

As a comparison point for bounds below, note that the obvious bounds on $S_2(\mathcal{H}_d)$ are the number of edges in the d -dimensional hypercube, $2^{d-1}d$, and the number of edges in the transitive closure of \mathcal{H}_d , which is $3^d - 2^d$. (An edge in the transitive closure of \mathcal{H}_d has 3 possibilities for each coordinate: both endpoints are 0, both endpoints are 1, or the first endpoint is 0 and the second is 1. This includes self-loops, so we subtract the number of vertices in \mathcal{H}_d to get the desired quantity.) Thus, $2^{d-1}d \leq S_2(\mathcal{H}_d) \leq 3^d - 2^d$. Similarly, the straightforward bounds on the number of edges in a 2-TC-spanner of $\mathcal{H}_{m,d}$ in terms of the number of edges in the directed grid and in its transitive closure are $dm^{d-1}(m-1)$ and $\left(\frac{m^2+m}{2}\right)^d - m^d$, respectively.

The following theorem gives upper and lower bounds on $S_2(\mathcal{H}_{m,d})$:

Theorem 3.1 (Hypergrid [16, 14]). *Let $S_2(\mathcal{H}_{m,d})$ denote the number of edges in the sparsest 2-TC-spanner of $\mathcal{H}_{m,d}$. Then for $m \geq 3$,*

$$S_2(\mathcal{H}_{m,d}) = \Omega\left(\frac{m^d \log^d m}{(2d \log \log m)^{d-1}}\right) \quad \text{and} \quad \leq m^d \log^d m.$$

The upper bound in Theorem 3.1 follows from a general construction of k -TC-spanners for graph products for arbitrary $k \geq 2$. The lower bound is proved by a reducing the 2-TC-spanner construction for $[m]^d$ to that for the $[2] \times [m]^{d-1}$ grid and then directly analyzing the number of edges required for a 2-TC-spanner of $[2] \times [m]^{d-1}$. The authors show a tradeoff between the number of edges in the 2-TC-spanner of the $[2] \times [m]^{d-1}$ grid that stay within the hyperplanes $\{1\} \times [m]^{d-1}$ and $\{2\} \times [m]^{d-1}$ versus the number of edges that cross from one hyperplane to the other. The proof proceeds in multiple stages. Assuming an upper bound on the number of edges staying within the hyperplanes, each stage is shown to separately contribute a substantial number of edges crossing between the hyperplanes.

Theorem 3.1 is most useful when m is large. When m is small, it is superseded by another set of bounds on $S_2(\mathcal{H}_{m,d})$, given in [16, 14], which are optimal up to a factor of d^{2m} . These bounds are formulated in terms of a complicated combinatorial expression, but value of this expression can be estimated numerically. Specifically, $S_2(\mathcal{H}_{m,d}) = 2^{c_m d} \text{poly}(d)$, where $c_2 \approx 1.1620$, $c_3 \approx 2.03$, $c_4 \approx 2.82$ and $c_5 \approx 3.24$, each significantly smaller than the exponents corresponding to the transitive closure sizes for the different m . More precisely, for the hypercube, $S_2(\mathcal{H}_d) = O(d^3 2^{c_2 d})$ and $\Omega(2^{c_2 d})$. The upper bound on $S_2(\mathcal{H}_d)$ is proved via a randomized construction of a 2-TC-spanner of the directed hypercube. Curiously, even though the upper and lower bounds above differ by a factor of $O(d^3)$, it is known that the randomized construction yields a 2-TC-spanner of \mathcal{H}_d of size within $O(d^2)$ of the optimal.


Steiner TC-spanners of d -dimensional posets. In some applications (in particular, to access control hierarchies [8, 9, 49, 7]), the shortcuts can use *Steiner* vertices, that is, vertices not in the original graph G . The resulting spanner is called a *Steiner TC-spanner*.

Definition 3.1 (Steiner TC-spanner [17]). *Given a directed graph $G = (V, E)$ and an integer $k \geq 1$, a **Steiner k -transitive-closure-spanner (Steiner k -TC-spanner)** of G is a directed graph $H = (V_H, E_H)$ satisfying:*

1. $V \subseteq V_H$;
2. for all vertices $u, v \in V$, if $d_G(u, v) < \infty$ then $d_H(u, v) \leq k$ and if $d_G(u, v) = \infty$ then $d_H(u, v) = \infty$.

Vertices in $V_H \setminus V$ are called *Steiner vertices*.

For some graphs, Steiner TC-spanners can be significantly sparser than ordinary TC-spanners. Before, our example of a graph with a 2-TC-spanner of size $\Omega(n^2)$ was a complete bipartite graph $K_{\frac{n}{2}, \frac{n}{2}}$ with $n/2$ vertices in each part and all edges directed from the first part to the second. This graph has a Steiner 2-TC-spanner of size n : it is enough to add one Steiner vertex v , edges to v from all nodes in the left part, and edges from v to all nodes in the right part. Thus, for $K_{\frac{n}{2}, \frac{n}{2}}$ there is a linear gap between the size of the sparsest Steiner 2-TC-spanner and the size of an ordinary 2-TC-spanner.



However, Bhattacharyya *et al.* [17] show that for directed hypergrids, Steiner vertices do not help: sparsest Steiner TC-spanners have the same size as TC-spanners with no Steiner vertices.

Lemma 3.1 ([17]). *If $\mathcal{H}_{m,d}$ has a Steiner k -TC spanner H , it also has a k -TC spanner of size $|H|$.*

Proof. We show how to replace one Steiner vertex in H with a grid vertex while keeping the same number of edges in the Steiner k -TC-spanner. This step can be repeated to remove all Steiner vertices.

Since $\mathcal{H}_{m,d}$ is acyclic, a cycle in H can contain at most one non-Steiner vertex, and therefore H will still remain a Steiner k -TC-spanner of $\mathcal{H}_{m,d}$ if this cycle is contracted to one vertex. Thus, we can assume without loss of generality that H is acyclic.

Let s be a Steiner vertex in H which does not have any other Steiner vertices below it. Let s' be the smallest vertex in $\mathcal{H}_{m,d}$ which is above all vertices v in $\mathcal{H}_{m,d}$ satisfying $v \preceq s$. (If there are no such v then s' is the grid vertex with all coordinates equal to 1.) Observe that s' always exists and is unique. Moreover,

every vertex in $\mathcal{H}_{m,d}$ that is above all such v is also above s' . By definition, s' is above all vertices in $\mathcal{H}_{m,d}$ which have a path to s . It is also below all vertices in $\mathcal{H}_{m,d}$ which are reachable from s . We replace all edges in H that have s as an endpoint with the corresponding edges with s' as an endpoint, and remove s from H . Every pair of vertices that was connected via a path of length at most k is still connected via the same path, with s replaced by s' if necessary. No new pair (u, v) of vertices in $\mathcal{H}_{m,d}$ got connected via s' since if u was below s and v was above s then $u \preceq s \preceq v$. The number of edges in H has not increased. \square

Atallah *et al.* [7], De Santis *et al.* [49] and Bhattacharyya *et al.* [17] study Steiner TC-spanners of directed *acyclic* graphs or, equivalently, partially ordered sets. Motivated by the application to access control hierarchies (described in Section 5), they focus on the relationship between the dimension of a poset and the size of its sparsest Steiner TC-spanner.

Definition 3.2 (Poset dimension). *The dimension of a poset G is the smallest d such that G can be embedded into a d -dimensional hypergrid $\mathcal{H}_{m,d}$ via an order-preserving embedding. A mapping from a poset G to a poset G' is called an order-preserving embedding if it respects the partial order, that is, all $x, y \in G$ are mapped to $x', y' \in G'$ such that $x \preceq_G y$ iff $x' \preceq_{G'} y'$.*

Each poset has a dimension. In particular, each poset with n elements can be embedded into a hypergrid $\mathcal{H}_{n,d}$, so that for all $i \in [d]$, the i th coordinates of images of all points are distinct.

Poset dimension is a fundamental and well-studied parameter in poset theory. For instance, Dilworth's famous chain partitioning theorem was originally intended as a lemma for proving a theorem about the dimension of distributive lattices [24]. A survey on poset dimension can be found in Trotter's monograph [57]. One important result for the discussion below, proved by Dushnik and Miller [26], is that for all m , the hypergrid $\mathcal{H}_{m,d}$ has dimension exactly d . Atallah *et al.* argue that many access control hierarchies are low-dimensional posets that come equipped with an embedding demonstrating low dimensionality.

Stretch k	Upper Bounds on $S_k(G)$	Lower Bounds on $S_k(G)$	Reference
$k = 2$	$O(n \log^d n)$	$\Omega\left(n \left(\frac{\log n}{\log \log n}\right)^d\right)$ for constant d	[17]
constant $k \geq 3$		$n \log^{\Omega(d)} n$ for constant d	[17]
$k = 2t + 1$ for $t \in [d]$	$O(3^{d-t} t \cdot n \log^{d-1} n \log \log n)$		[49]

Table 1. The size of the sparsest Steiner k -TC-spanners of d -dimensional posets on n vertices for $d \geq 2$

Observe that the only poset of dimension 1 is the directed line. Tight bounds for the size of (Steiner) TC-spanners of directed lines were discussed in the beginning of Section 3. Table 1 summarizes the best bounds for $d \geq 2$. The upper bounds hold for all posets of dimension d . The TC-spanners in the upper bounds can be constructed efficiently, given an explicit embedding of the poset into a d -dimensional grid. (Finding such an embedding is NP-hard [59].) Furthermore, paths of length at most k between all pairs of vertices in the resulting k -TC-spanners can be found efficiently. This is important for the application to access control hierarchies.

The lower bounds mean that there exists a poset of dimension d for which every Steiner k -TC-spanner has the specified number of edges. The lower bound for Steiner 2-TC-spanners holds for the hypergrid $\mathcal{H}_{m,d}$ and follows from the lower bound on $S_2(\mathcal{H}_{m,d})$ in Theorem 3.1 and the fact that Steiner vertices do not help for directed hypergrids (Lemma 3.1). The lower bound on the size of a Steiner k -TC-spanner for $k \geq 3$ holds for a poset obtained by a randomized construction.

Note that the Steiner vertices used in the constructions for d -dimensional posets are necessary to obtain sparse TC-spanners. Recall our example of a bipartite graph $K_{\frac{n}{2}, \frac{n}{2}}$ for which every 2-TC-spanners required $\Omega(n^2)$ edges. $K_{\frac{n}{2}, \frac{n}{2}}$ is a poset of dimension 2, and thus, by the upper bound in [17], has a Steiner 2-TC spanner of size $O(n \log^2 n)$. (As we mentioned before, for this graph there is an even better Steiner 2-TC spanner with $O(n)$ edges.) To see that $K_{\frac{n}{2}, \frac{n}{2}}$ is embeddable into a $[n] \times [n]$ grid, map each of the $n/2$ left vertices of $K_{\frac{n}{2}, \frac{n}{2}}$ to a distinct grid vertex in the set of incomparable vertices $\{(i, n/2 + 1 - i) : i \in [n/2]\}$, and similarly map each right vertex to a distinct vertex in the set $\{(n + 1 - i, i + n/2) : i \in [n/2]\}$. It is easy to see that this is a proper embedding.


3.2 General TC-spanner Constructions

Graphs that require a large number of shortcuts. We have seen in the beginning of Section 3 that, in general, TC-spanners can be large. However, in the example of $K_{\frac{n}{2}, \frac{n}{2}}$ we looked at, the graph itself was large and, in fact, we did not have to add any shortcuts to construct a 2-TC-spanner of that graph. Can one always construct a TC-spanner by adding a small number of edges to the original graph? Thorup [52] conjectured that all directed graphs G on n nodes have TC-spanners with stretch polylogarithmic in n and size at most $2|G|$. As mentioned before, he proved his conjecture for planar graphs [53], but later Hesse [37] gave a counterexample to Thorup's conjecture for general graphs. For all small $\epsilon > 0$, he constructed a family of graphs with $n^{1+\epsilon}$ edges for which all n^ϵ -TC-spanners require $\Omega(n^{2-\epsilon})$ edges.

2-TC-spanners from $2k$ -TC-spanners. Berman, Raskhodnikova and Ruan [13] show how to (efficiently) obtain a 2-TC-spanner of a graph G with diameter at most $2k$ by adding $O(n^{1-1/k} \cdot |G|)$ shortcuts. They prove that this relationship is nearly tight in the following sense: for every sufficiently small positive ϵ , there

are graphs with $2k$ -TC-spanners of size $n^{1+1/k}$ and no 2-TC-spanners of size less than $n^{2-\epsilon}$. These graphs are obtained by adjusting the parameters in the construction by Hesse mentioned above.

Moreover, as shown in [13], their upper bound is completely tight for the transformation from 3-TC-spanners to 2-TC-spanners: the number of added edges is asymptotically optimal, as evidenced by the 4-layered graph with m^2 nodes in layers 1 and 4 and m nodes in layers 2 and 3, where the edges are directed from smaller to larger layers and are formed as follows. There is a complete bipartite graph between layers 2 and 3. Each node in layer 2 is connected to m nodes in layer 1, and each node in layer 1 has outdegree 1. The edges between layers 3 and 4 are constructed in the same manner. The resulting graph has $3m^2$ edges and is a 3-TC-spanner. A 2-TC-spanner of this graph must connect m^4 pairs of vertices in layers 1 and 4 via paths of length at most 2. Each shortcut edge can be used by at most m such pairs. Therefore, at least m^3 shortcuts are required. Setting $n = 2m^2 + 2m$, we obtain a graph with a 3-TC-spanner of size $O(n)$, for which every 2-TC-spanner requires $\Omega(n^{3/2})$ edges.



TC-spanners with large stretch. Improving on the first result in this vein from [15], Berman, Raskhodnikova and Ruan [13] show that one can obtain a k -TC-spanner of any graph by adding $O(n^2/k^2)$ shortcut edges. This construction is efficient.

TC-spanners of Graphs with Cycles. Here we give a reduction from constructing TC-spanners of general directed graphs to constructing TC-spanners of directed acyclic graphs (DAGs).

Lemma 3.2. *Let G be a directed graph on n vertices, and G' be the graph of strongly connected components of G . Then $S_{k+2}(G) \leq S_k(G') + 2n$. Moreover, given a k -TC-spanner H' of G' , one can efficiently construct a $(k+2)$ -TC-spanner H of G with at most $|H'| + 2n$ edges.*

Proof. For each strongly connected component C of G , pick an arbitrary vertex v_C and call it a representative of C . To construct a $(k+2)$ -TC-spanner H of G from a k -TC-spanner H' of G' , first connect representatives of connected components to mimic the structure of H' : namely, add an edge (v_{C_1}, v_{C_2}) to H for every edge (C_1, C_2) in H' . Second, for every vertex u in the component C , where $u \neq v_C$, add edges (u, v_C) and (v_C, u) to H .

The resulting H has the same number of edges as H' plus at most 2 edges per vertex, added to connect each vertex to the representative of its strongly connected component. That is, $|H| \leq |H'| + 2n$. To see that H is a $(k+2)$ -TC-spanner of G , consider vertices u_1, u_2 in G , where u_2 is reachable from u_1 . Let v_1 and v_2 be the representatives of the components of u_1 and u_2 , respectively. Since H' is a k -TC-spanner of G' , there is a path of length at most k from the component of u_1 to the component of u_2 in H' . Therefore, H contains a path of length at most k from v_1 to v_2 . Since H also contains edges (u_1, v_1) and (v_2, u_2) , it contains a path of length at most $k+2$ from u_1 to u_2 . \square

4 Overview of Computational Results on Directed Spanners

The computational problem of finding the size of the sparsest k -TC-spanner of a given graph, called k -TC-SPANNER, was first considered in [15]. k -TC-SPANNER is a special case of a well-studied problem, called DIRECTED k -SPANNER, of finding the size of the sparsest k -spanner of a given (not necessarily transitively closed) directed graph [29, 28, 15, 13]. In this section, we survey approximation algorithms and inapproximability results for these two problems. All known algorithms on DIRECTED k -SPANNER also apply to two other variants, CLIENT/SERVER DIRECTED k -SPANNER and k -DIAMETER SPANNING SUBGRAPH, defined by Elkin and Peleg [29].

Problem	Stretch k	Approximability	Previous Work
DIRECTED k -SPANNER (and k -TC-SPANNER)	$k = 2$	$O(\log n)$ [29]	
	$k = 3$	$O(\sqrt{n} \cdot \log n)$ [13]	[28, 15]
	$k \geq 3$	$O(kn^{1-1/\lceil k/2 \rceil} \cdot \log n)$ [13]	[15]
k -TC-SPANNER only	$k \geq 3$	$O(n^{1-1/\lceil k/2 \rceil} \cdot \log n)$ [13]	[15]
	$k = \Omega\left(\frac{\log n}{\log \log n}\right)$	$O(n/k^2)$ [13]	[15]

Table 2. Summary: Algorithmic Results on DIRECTED k -SPANNER and k -TC-SPANNER

Algorithms for DIRECTED k -SPANNER and k -TC-SPANNER. All algorithms for DIRECTED k -SPANNER immediately yield algorithms for k -TC-SPANNER with the same approximation ratio because k -TC-SPANNER on input graph G is equivalent to DIRECTED k -SPANNER on input $\text{TC}(G)$. Table 2 summarizes the best known approximation algorithms for these problems for different stretch k . Elkin and Peleg [29] gave an $O(\log n)$ -approximation algorithm for DIRECTED 2-SPANNER. For $k = 3$, approximation algorithms were proposed in [28, 15, 13] with the best ratio, $O(\sqrt{n} \cdot \log n)$, due to [13]. In general, for $k > 3$, [13] prove an approximation ratio $O(kn^{1-1/\lceil k/2 \rceil} \cdot \log n)$, improving the first non-trivial

polynomial time algorithm for this problem, given in [15]. For the special case of k -TC-SPANNER, [13] give a slightly better ratio of $O(n^{1-1/\lceil k/2 \rceil} \cdot \log n)$. For large k , the best approximation ratio is $O(n/k^2)$, due to [13], again an improvement over the first non-trivial algorithm for this range of parameters, proposed in [15].

We briefly describe the two TC-spanner-specific approximation algorithms from [13]. They are based on the structural results mentioned in Section 3.2. The first algorithm runs the $O(\log n)$ -approximation algorithm from [29] for DIRECTED 2-SPANNER on the transitive closure of the input graph G . The analysis relies on the construction of 2-TC-spanners from k -TC-spanners, mentioned in Section 3.2. This construction proves that $S_2(G) \leq S_k(G) + O(n^{1-1/\lceil k/2 \rceil} \cdot TR(G))$, where $TR(G)$ denotes the size of a transitive reduction of G . (A transitive reduction was defined and discussed in Section 2.) Since the algorithm is guaranteed to output a 2-TC-spanner of size $O(\log n \cdot S_2(G)) = O(\log n \cdot S_k(G) + n^{1-1/\lceil k/2 \rceil} \log n \cdot TR(G))$, the result is an $O(n^{1-1/\lceil k/2 \rceil} \log n)$ -approximation. (Recall that $TR(G)$ is a lower bound on the size of a TC-spanner.)

The algorithm for large k is based on an efficient procedure that obtains a k -TC-spanner by adding $O(n^2/k^2)$ shortcut edges. It can be run on each weakly connected component separately. For a weakly connected component with n nodes, the size of a k -TC-spanner is at least $n - 1$, so the resulting graph is a $O(n/k^2)$ -approximation.

It is important to note that the algorithm for large k has a better approximation ratio than the corresponding hardness result for DIRECTED k -SPANNER. That is, k -TC-SPANNER is a strictly easier problem for this range of parameters.

Inapproximability of DIRECTED k -SPANNER. For completeness, we state the inapproximability results for DIRECTED k -SPANNER, even though they do not imply anything for k -TC-SPANNER. Kortsarz [40] showed that the $O(\log n)$ approximation ratio for DIRECTED 2-SPANNER cannot be improved unless $P=NP$. For all $\delta, \epsilon \in (0, 1)$ and $3 \leq k \leq n^{1-\delta}$, it is impossible to approximate DIRECTED k -SPANNER within a factor of $2^{\log^{1-\epsilon} n}$ in polynomial time, assuming $NP \not\subseteq DTIME(n^{\text{poly} \log n})$ [28, 30]. ($DTIME(f(n))$ denotes the class of languages decidable deterministically in time $f(n)$.) Thus, according to Arora and Lund's classification [38] of NP-hard problems, DIRECTED k -SPANNER is in class III, for $k \in [3, n^{1-\delta}]$. Moreover, [30] showed that proving that DIRECTED k -SPANNER is in class IV, that is, inapproximable within n^δ for some $\delta \in (0, 1)$, would resolve a long standing open question in complexity theory: namely, cause classes III and IV to collapse into a single class.

Inapproximability of k -TC-SPANNER. Table 3 summarizes inapproximability results for k -TC-SPANNER for different values of k . For constant k , the hardness results are the same as for DIRECTED k -SPANNER, even though the reductions are much more technically involved. Observe that a stronger inapproximability result for $k > 2$ would imply the same inapproximability for DIRECTED- k -SPANNER and, as shown in [30], collapse classes III and IV in Arora and Lund's classification. For nonconstant k for which there exists a sufficiently small $\gamma > 0$ such that $k \leq n^{1-\gamma}$, we know that the problem is NP-hard, but not much beyond


Stretch k	Inapproximability	Assumption	Notes
$k = 2$	$\Omega(\log n)$	$P \neq NP$	Matches the upper bound
constant $k \geq 3$	$\Omega(2^{\log^{1-\epsilon} n})$ $\forall \epsilon \in (0, 1)$	$NP \not\subseteq DTIME(n^{\text{polylog } n})$	Improvement implies breakthrough
$k \leq n^{1-\delta}$ $\forall \delta \in (0, 1)$	$\Omega(1 + \delta)$	$P \neq NP$	

Table 3. Summary of Hardness Results on k -TC-SPANNER; all results are from [15]

that. This contrasts sharply with the known hardness of DIRECTED k -SPANNER, but, as mentioned previously, k -TC-SPANNER is known to be strictly easier for some (but not all) k in that range.


The $2^{\log^{1-\epsilon} n}$ -inapproximability of k -TC-SPANNER for constant $k \geq 3$ in [15] matches the inapproximability of DIRECTED k -SPANNER for the same stretch in [30]. As is the case for DIRECTED k -SPANNER, the reduction is from a problem called MIN-REP, whose inapproximability DIRECTED k -SPANNER inherits. However, as illustrated in [15], all known hard instances for DIRECTED k -SPANNER cannot imply anything better than $\Omega(1)$ -hardness for k -TC-SPANNER. Intuitively, inapproximability of k -TC-SPANNER is harder to prove than inapproximability of DIRECTED k -SPANNER because an instance of k -TC-SPANNER must be transitively-closed, and thus, have more “shortcut” routes between pairs of vertices. The construction of hard instances of k -TC-SPANNER in [15] uses so-called *generalized butterfly* and *broom* graphs. The paths in these graphs are well-structured, making it possible to analyze many different routes in the transitive closure of a hard instance.

The reduction from MIN-REP to k -TC-SPANNER in [15] is quite involved. We briefly describe some of the ideas behind the reduction. An instance of MIN-REP is a bipartite graph G , where each part consists of n nodes partitioned into r clusters of size n/r . The clusters in the left part are called $\mathcal{A}_1, \dots, \mathcal{A}_r$ and the clusters in the right part are $\mathcal{B}_1, \dots, \mathcal{B}_r$.




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. Elkin and Peleg [28] showed that MIN-REP is $2^{\log^{1-\epsilon} n}$ -inapproximable.

We now describe *generalized butterfly* and *broom* graphs used in the reduction. *Generalized butterflies* were defined by Woodruff [58]. Each node in a generalized butterfly has k coordinates: (a_1, \dots, a_{k-1}, i) , where $a_1, \dots, a_{k-1} \in [d]$ and $i \in [k]$. There is an edge from node (a_1, \dots, a_{k-1}, i) to node $(b_1, \dots, b_{k-1}, i+1)$ iff for all $j \neq i, a_j = b_j$.

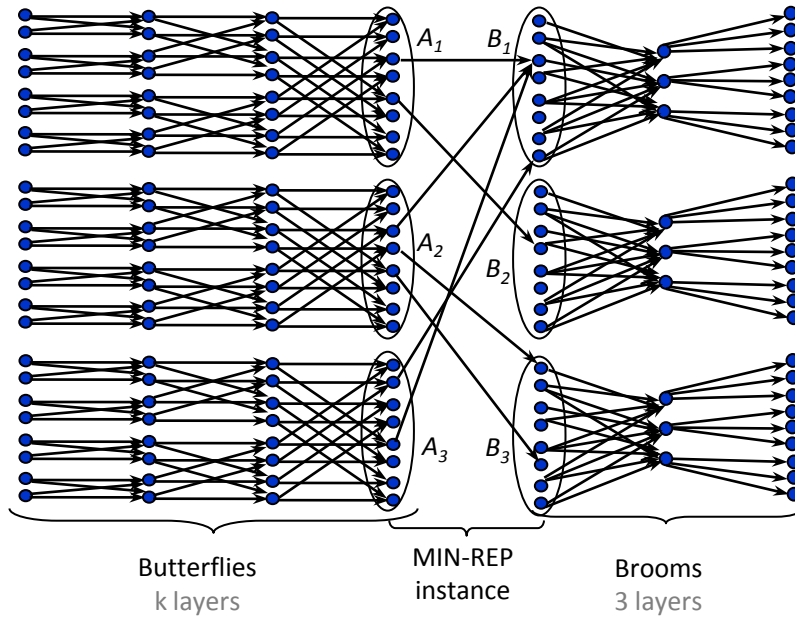


Since there are d possibilities for b_i , each node has outdegree d . Similarly, each node has indegree d . It is easy to see that there is a unique shortest path of length $k - 1$ from any node in layer 1 to any node in layer k . Moreover, any shortcut is on at most d^{k-3} such paths because if it connects layer i to layer $i + \ell$ (where $\ell \geq 2$) it fixes all but $i - 1$ coordinates of the first node and all but $k - (i + \ell)$ coordinates of the second. Thus, at least d^{k+1} shortcuts are needed to reduce the diameter from $k - 1$ to $k - 2$.

A *broom* is a 3-layer graph, where the two leftmost layers form a bipartite clique, and the right layer consists of degree-1 nodes, attached to nodes in the middle layer. Each node in the first and second layer has outdegree d . All edges are directed from left to right.



Given an instance of MIN-REP, we construct an instance G of k -TC-SPANNER 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 vertices in layer k of the butterfly. The parameter d is determined by the size of \mathcal{A}_i and k . (We can add isolated vertices to each cluster of the given MIN-REP instance to ensure that $|\mathcal{A}_i|$ is a $(k - 1)$ st power.) Next, each \mathcal{B}_j is identified with the leftmost layer of a disjoint broom graph. All edges of G are directed towards the rightmost nodes of the brooms. The resulting graph has diameter $k + 2$.



A k -TC-spanner H of G is formed as follows. Let O be a minimum rep-cover of the underlying MIN-REP instance. For each butterfly, include all shortcuts from layer $k - 2$ to comparable vertices in layer k which are also in O . In addition, include all shortcuts from vertices in layer $k + 1$ which are also in O to all comparable nodes in the last layer. Since O is a rep-cover, H is a k -TC-spanner. The size of H is $|G| + d^2|O|$ because for each vertex in O we add shortcuts to d^2 vertices (in layer $k - 2$ for vertices in the left clusters of MIN-REP, and in layer $k + 3$ for vertices in the right clusters).

If H were optimal, then approximating its size would approximate a minimum rep-cover of the original MIN-REP instance within the same factor. To ensure that H is optimal, [15] carefully modify the original MIN-REP instance and only then apply the reduction we described.

5 Applications of TC-spanners

We describe four types of applications that use sparse TC-spanners: property testing, property reconstruction, key management in an access hierarchy and data structures for computing partial products in a semigroup. For property testing, we give two applications: to testing monotonicity of functions and to testing if a function is Lipschitz. All these applications, with the exception of testing Lipschitz functions and property reconstruction, were pointed out and described in [15]. The application to Lipschitz functions is from [39]. The application to property reconstruction is from [14].

5.1 Applications to Property Testing

We start by describing the application to testing monotonicity of functions. We also point out the limitations of TC-spanner techniques and related open questions in the area.

Monotonicity testing. Monotonicity of functions [31, 35, 25, 12, 32, 33, 36, 4, 15, 16, 51] is one of the most widely studied properties in property testing [34, 47]. Fischer *et al.* [33] 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. For instance, if G_n is a directed line L_n , the tester needs to determine whether the input sequence, specified by f , is sorted or ϵ -far from sorted. If G_n is a 2-dimensional grid $\mathcal{H}_{m,2}$, the goal is to determine whether the input matrix has non-decreasing rows and columns. The optimal monotonicity tester for the directed line L_n , proposed by Dodis *et al.* [25], is based on the sparsest 2-TC-spanner for that graph. Implicit in the proof of Proposition 9 in [25] is a lemma relating the complexity of a monotonicity tester for L_n to the size of a 2-TC-spanner of L_n . Bhattacharyya *et al.* [15] generalized this lemma by observing that a sparse 2-TC-spanner for any partial order graph G_n implies an efficient monotonicity tester on G_n .

Lemma 5.1 ([15]). *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 $O\left(\frac{s(n)}{\epsilon n}\right)$.*

Proof. The tester selects $\frac{8s(n)}{\epsilon n}$ 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 some selected edge (x, y) is *violated* by f , that is, $f(x) > f(y)$.

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 $\frac{\epsilon n}{4}$ edges in H . Then each selected edge is violated with probability $\frac{\epsilon n}{4s(n)}$, and the lemma follows by elementary probability theory.

Denote the transitive closure of G_n by $TC(G_n)$. 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 $TC(G_n)$. This implies ([33], 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 at least $\frac{\epsilon n}{2}$ violated edges in $TC(G_n)$ [25]. We will establish a map from the set of edges in M to the set of violated edges in H , so that each violated edge in H is the image of at most 2 edges in M . 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, at most 2 edges in M can be mapped to one violated edge in H . Thus, the 2-TC-spanner H has $\geq \frac{\epsilon n}{4}$ 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_{2n} is a 3-TC-spanner of 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.

As demonstrated by Lemma 5.1, all the 2-TC-spanner constructions yield monotonicity testers for functions defined on the corresponding posets. This lemma led to significant improvements in monotonicity testers for several graph families, including planar graphs and, in general, H -minor-free graphs [15]. Indeed, [15] achieve testers with $O(\log^2 n/\epsilon)$ queries for H -minor-free graphs using their construction of sparse 2-TC-spanners for this graph family, whereas the previous tester, due to Fischer *et al.* [33], worked only for planar graphs and required $\Theta(\sqrt{n}/\epsilon)$ queries.

We briefly discuss the limitations of the TC-spanner method for constructing monotonicity testers. The lower bounds in [16, 14] on the size of the sparsest 2-TC-spanners for the hypercube and the hypergrid (described in Theorem 3.1) rule out the TC-spanner approach for improving monotonicity testers on the hypercube and hypergrid. Currently, the running time of the best tester for monotonicity of functions of the form $f : \{0, 1\}^d \rightarrow R$ and, more generally, $f : [m]^d \rightarrow R$, where R is an arbitrary range, is $O(\frac{d}{\epsilon} \log m \cdot \log |R|)$ [25]. The best known lower bound (for the hypercube with range $R = \{0, 1\}$) is $\Omega(\log \log d)$ [33]. (There are better bounds for restricted classes of tests in [33] and [51].) Even

though for a fixed d , it is known that the optimal monotonicity tester for the grid runs in time $\Theta(\frac{\log m}{\epsilon})$ [36, 32], bridging the gap between the lower and upper bounds for arbitrary d has remained elusive. Lemma 5.1 showed that if a 2-TC-spanner of size $o(2^d d^2)$ for the hypercube or, more generally, a 2-TC-spanner of size $o(m^d d^2 \log^2 m)$ for the hypergrid were found, the monotonicity tester of [25] would be improved. In the light of the lower bounds for the hypercube and the hypergrid, a fundamentally new approach is required.

Testing if a function is Lipschitz. In the important special case when G_n is the directed line, Lemma 5.1 yields an optimal tester for whether a function of the form $f : [n] \rightarrow R$ is monotone or, equivalently, of whether a list of n elements is sorted, that runs in time $O(\log n/\epsilon)$. (There is another optimal tester for this problem that was discovered first [31].) Jha and Raskhodnikova [39] observe that the test and analysis in Lemma 5.1 apply to any property of a list of numbers if (a) it can be expressed in terms of pairs of list elements and (b) it is transitive: namely, for all $x \prec y \prec z$, whenever (x, y) and (y, z) are not *violated*, (x, z) is also not violated. In particular, it applies to testing whether a function of the form $f : [n] \rightarrow \mathbb{R}$ is Lipschitz. A function $f : \mathcal{D} \rightarrow \mathcal{R}$ is called *Lipschitz* if $\text{dist}_{\mathcal{R}}(f(x), f(y)) \leq \text{dist}_{\mathcal{D}}(x, y)$ for all x, y in \mathcal{D} , where $\text{dist}_{\mathcal{R}}$ and $\text{dist}_{\mathcal{D}}$ denote the distance functions on the range and domain of f , respectively. Testing the Lipschitz property has applications to programs with noisy inputs and to data privacy.

Note that the Lipschitz property was defined in terms of pairs of domain elements. Consider a function $f : [n] \rightarrow \mathbb{R}$, where the domain and range are equipped with distance functions $\text{dist}_{\mathcal{D}}(x, y) = |x - y|$ and $\text{dist}_{\mathcal{R}}(f(x), f(y)) = |f(y) - f(x)|$. We say a pair (x, y) is *violated* if $|f(y) - f(x)| > |y - x|$. Then if (x, y) and (y, z) are not violated, it implies that neither is (x, z) . Thus, the requirements (a) and (b) above hold and, using their observation, Jha and Raskhodnikova get a $O(\log n/\epsilon)$ Lipschitz test for functions of the form $f : [n] \rightarrow \mathbb{R}$ via the optimal 2-TC-spanner construction of the line.

5.2 Application to Property Reconstruction

Property-preserving data reconstruction was introduced by Ailon, Chazelle, Co-mandur and Liu [3]. In this model, a reconstruction algorithm, called a *filter*, sits between a *client* and a *dataset*. A dataset is viewed as a function $f : \mathcal{D} \rightarrow \mathcal{R}$. Client accesses the dataset using *queries* of the form $x \in \mathcal{D}$ to the filter. The filter *looks up* a small number of values in the dataset and outputs $g(x)$, where g must satisfy some fixed *structural* property \mathcal{P} . Extending this notion, Saks and Seshadhri [48] defined *local* reconstruction. A filter is *local* if it allows for a local (or distributed) implementation: namely, if the output function g does not depend on the order of the queries.

Definition 5.1 (Local filter). A local filter for reconstructing property \mathcal{P} is an algorithm A that has oracle access to a function $f : \mathcal{D} \rightarrow \mathcal{R}$, and to an auxiliary random string ρ (the “random seed”), and takes as input $x \in \mathcal{D}$. For fixed f

and ρ , A runs deterministically on input x to produce an output $A_{f,\rho}(x) \in \mathcal{R}$. As x varies over the domain \mathcal{D} , this defines a function $g : \mathcal{D} \rightarrow \mathcal{R}$, where $g(x) = A_{f,\rho}(x)$. (Note that a local filter has no internal state to store previously made queries.) The filter must satisfy the following conditions:

- For each f and ρ , the function g output by the filter satisfies \mathcal{P} .
- If f satisfies \mathcal{P} , then g is identical to f with probability at least $1 - \delta$, for some $\delta \leq 1/3$. The parameter δ is called error probability.

In answering query $x \in \mathcal{D}$, the filter A may ask for values of f at domain points of its choice using its oracle access to f . Each such access made to the oracle is called a *lookup* to distinguish it from the client query x . A local filter is *non-adaptive* if the set of domain points that the filter looks up to answer an input query x does not depend on answers given by the oracle.

Saks and Seshadhri also required that g must be sufficiently close to f : *With high probability (over the choice of ρ), $\text{Dist}(g, f) \leq B(n) \cdot \text{Dist}(f, \mathcal{P})$, where $B(n)$ is called the error blow-up. ($\text{Dist}(g, f)$ is the number of points in the domain on which f and g differ. $\text{Dist}(f, \mathcal{P})$ is $\min_{g \in \mathcal{P}} \text{Dist}(g, f)$.)* If a local filter satisfies this condition along with Definition 5.1, we call it *distance-respecting*.

Local Monotonicity Reconstructors. The first property considered in the reconstruction [3] and local reconstruction [48] models was monotonicity of functions. (See Section 5.1 for a definition.) A (distance-respecting) filter for monotonicity can be used, for example, when a program will run correctly only if its input is sorted. Then, instead of accessing the input directly, the program can access it via a filter, which will ensure that the program always sees a sorted input, making small corrections when necessary. A local filter can be implemented in a distributed manner with an additional guarantee that every program run on the same not-quite-sorted input will see the same corrected version. This can be done by supplying the same random string to each copy of the filter.

To motivate monotonicity reconstructors for hypergrids, consider the scenario of rolling admissions: An admissions office assigns d scores to each application, such as the applicant’s GPA, SAT results, essay quality, etc. Based on these scores, some complicated (third-party) algorithm outputs the probability that a given applicant should be accepted. The admissions office wants to make sure “on the fly” that strictly better applicants are given higher probability, that is, probabilities are *monotone* in scores. A hypergrid monotonicity filter may be used here. And, as before, if the filter is local, it can be implemented in a distributed manner, guaranteeing the same results for all filters running in parallel.

Saks and Seshadhri [48] give a *distance-respecting* local monotonicity filter for the directed hypergrid, $\mathcal{H}_{m,d}$, that makes $(\log m)^{O(d)}$ lookups per query. No non-trivial monotonicity filter for the hypercube \mathcal{H}_d (performing $o(2^d)$ lookups per query) is known. One of the monotonicity filters in [3] is a local filter for the directed line $\mathcal{H}_{m,1}$ with $O(\log m)$ lookups per query (but a worse error blow up than in [48]). As observed in [48], this upper bound is tight. Notably, all known local filters for monotonicity property are *non-adaptive*. A lower bound of $2^{\alpha d}$,

on the number of lookups per query for a *distance-respecting* local monotonicity filter on \mathcal{H}_d with *error blow-up* $2^{\beta d}$, where α, β are sufficiently small constants, appeared in [48].

[14] show how to construct sparse 2-TC-spanners from local monotonicity reconstructors with low lookup complexity. These constructions, in conjunction with lower bounds on the size of 2-TC-spanners of the hypergrid and hypercube, described in Section 3.1, imply lower bounds on lookup complexity of local monotonicity reconstructors with arbitrary blow-up. The transformations from non-adaptive and adaptive reconstructors are stated in Theorems 5.1 and 5.2, respectively.

Theorem 5.1 (Transformation from non-adaptive Local Monotonicity Reconstructors to 2-TC-spanners, [14]). *Let $G_n = (V_n, E)$ be a poset on n nodes. Suppose there is a non-adaptive local monotonicity reconstructor A for G_n that looks up at most $\ell(n)$ values to answer any query $x \in V_n$ and has error probability at most δ . Then there is a 2-TC-Spanner of G_n with at most $O(n\ell(n) \cdot \lceil \log n / \log(1/\delta) \rceil)$ edges.*

Proof. Let A be a local reconstructor given by the statement of the theorem. Let \mathcal{F} be the set of pairs (x, y) with x, y in V_n such that $x \prec y$. Then, \mathcal{F} is of size at most $\binom{n}{2}$. Given $(x, y) \in \mathcal{F}$, let $\text{cube}(x, y)$ be the set $\{z \in V_n : x \preceq z \preceq y\}$. Define function $f^{(x,y)}(v)$ to be 1 on all $v \succeq x$ and 0 everywhere else. Also, define function $f^{(\overline{x}, \overline{y})}(v)$, which is identical to $f^{(x,y)}(v)$ for all $v \notin \text{cube}(x, y)$ and 0 for $v \in \text{cube}(x, y)$. Both, $f^{(x,y)}$ and $f^{(\overline{x}, \overline{y})}$, are monotone functions for all $(x, y) \in \mathcal{F}$. Let A_ρ be the deterministic algorithm which runs A with the random seed fixed to ρ . We say a string ρ is *good* for $(x, y) \in \mathcal{F}$ if filter A_ρ on input $f^{(x,y)}$ returns $g = f^{(x,y)}$ and on input $f^{(\overline{x}, \overline{y})}$ returns $g = f^{(\overline{x}, \overline{y})}$.

Now we show that there exists a set S of size $s \leq \lceil 2 \log n / \log(1/2\delta) \rceil$, consisting of strings used as random seeds by A , such that for every $(x, y) \in \mathcal{F}$ some string $\rho \in S$ is good for (x, y) . We choose S by picking strings used as random seeds uniformly and independently at random. Since A has error probability at most δ , we know that for every monotone f , with probability at least $1 - \delta$ (with respect to the choice of ρ), the function $A_{f,\rho}$ is identical to f . Then, for fixed $(x, y) \in \mathcal{F}$ and uniformly random ρ ,

$$\begin{aligned} \Pr[\rho \text{ is not good for } (x, y)] &\leq \Pr[A_\rho \text{ on input } f^{(x,y)} \text{ fails to output } f^{(x,y)}] \\ &\quad + \Pr[A_\rho \text{ on input } f^{(\overline{x}, \overline{y})} \text{ fails to output } f^{(\overline{x}, \overline{y})}] \leq 2\delta. \end{aligned}$$

Since strings in S are chosen independently, $\Pr[\text{no } \rho \in S \text{ is good for } (x, y)] \leq (2 \cdot \delta)^s$, which, for $s = \lceil 2 \log n / \log(1/2\delta) \rceil$, is at most $1/n^2 < 1/|\mathcal{F}|$. By a union bound over \mathcal{F} ,

$$\Pr[\text{for some } (x, y) \in \mathcal{F}, \text{ no } \rho \in S \text{ is good for } (x, y)] < 1.$$

Thus, there exists a set S with required properties.

We construct our 2-TC-spanner $H = (V_n, E_H)$ of G_n using set S described above. Let $\mathcal{N}_\rho(x)$ be the set consisting of x and all vertices looked up by A_ρ

on query x . For each string $\rho \in S$ and each vertex $x \in V_n$, connect x to all comparable vertices in $\mathcal{N}_\rho(x)$ (other than itself) and orient these edges according to their direction in G_n .

We prove H is a 2-TC-Spanner as follows. Suppose not, *i.e.*, there exists $(x, y) \in \mathcal{F}$ with no path of length at most 2 in H from x to y . Consider $\rho \in S$ which is *good* for (x, y) . Define function h by setting $h(v) = f^{(x,y)}(v)$ for all $v \notin \text{cube}(x, y)$. Then $h(v) = f^{(\bar{x}, \bar{y})}(v)$ for all $v \notin \text{cube}(x, y)$, by definition of $f^{(\bar{x}, \bar{y})}$. For all $v \in \text{cube}(x, y)$, set $h(v)$ to 1 for $v \in \mathcal{N}_\rho(x)$ and to 0 for $v \in \mathcal{N}_\rho(y)$. All unassigned points are set to 0. By the assumption above, $\mathcal{N}_\rho(x) \cap \mathcal{N}_\rho(y)$ does not contain any points in $\text{cube}(x, y)$. Therefore, h is well-defined. Since, ρ is *good* for (x, y) and h is identical to $f^{(x,y)}$ for all look ups made on query x , $A_\rho(x) = h(x) = 1$. Similarly, $A_\rho(y) = h(y) = 0$. But $x \prec y$, so $A_{h,\rho}(v)$ is not monotone. Contradiction.

The number of edges in H is at most

$$\sum_{x \in V_n, \rho \in S} |\mathcal{N}_\rho(x)| \leq n \cdot \ell \cdot s \leq n\ell \cdot \lceil 2 \log n / \log(1/2\delta) \rceil. \quad \square$$

The next theorem applies even to *adaptive* local monotonicity reconstructors. It takes into account how many lookups on query x are points incomparable to x . In particular, if there are no such lookups, then the constructed 2-TC-spanner is of the same size as in Lemma 5.1.

Theorem 5.2 (Transformation from adaptive Local Monotonicity Reconstructors to 2-TC-spanners, [14]). *Let $G_n = (V_n, E)$ be a poset on n nodes. Suppose there is a (possibly adaptive) local monotonicity reconstructor A for G_n that, for any query $x \in V_n$, looks up at most $\ell_1(n)$ vertices comparable to x and at most $\ell_2(n)$ vertices incomparable to x , and has error probability at most δ . Then there is a 2-TC-Spanner of G_n with at most $O(n\ell_1(n) \cdot 2^{\ell_2(n)} \lceil \log n / \log(1/\delta) \rceil)$ edges.*

Proof. Define \mathcal{F} , $f^{(x,y)}$, $f^{(\bar{x}, \bar{y})}$, A_ρ and S as in the proof of Theorem 5.1. As before, for each $x \in V_n$, we define sets $\mathcal{N}_\rho(x)$, and construct the 2-TC-Spanner H by connecting each x to comparable points in $\mathcal{N}(x)$ for all $\rho \in S$ and orienting the edges according to G_n . However, now $\mathcal{N}_\rho(x)$ is a union of several sets $\mathcal{N}_\rho^{b,w}(x)$, indexed by $b \in \{0, 1\}$ and $w \in \{0, 1\}^{\ell_2(n)}$. (In addition, $\mathcal{N}_\rho(x)$ contains x .) For each $x \in V_n$, $b \in \{0, 1\}$ and $w \in \{0, 1\}^{\ell_2(n)}$, let $\mathcal{N}_\rho^{b,w}(x) \subseteq V_n$ be the set of lookups performed by A_ρ on query x , assuming that the oracle answers all lookups as follows. When a lookup y is comparable to x , answer 0 if $y \prec x$, b if $y = x$, 1 if $x \prec y$. Otherwise, if y is the i 'th lookup made to an incomparable point for some $i \in [\ell_2]$, answer $w[i]$. Recall that we set $\mathcal{N}_\rho(x)$ to be the union of $\mathcal{N}_\rho^{b,w}$ for all $b \in \{0, 1\}$ and all $w \in \{0, 1\}^{\ell_2(n)}$. This completes the description of $\mathcal{N}_\rho(x)$ and construction of H .

The argument that H is a 2-TC-spanner proceeds similarly to that in the proof of Theorem 5.1. The caveat is that an adaptive local filter might choose lookups based on the answers to previous lookups. The constructed function h sets all points comparable to x to 0 if they are below x and 1 if they are above x .

However, points incomparable to x might be comparable to y and might be set to 0 or 1, depending on whether they are above or below y . Since we included sets of points queried under all these possibilities in $\mathcal{N}_\rho(x)$, we can now conclude that $A_\rho(x) = h(x) = 1$. The same applies for y . So, $A_{h,\rho}$ outputs a non-monotone function, witnessed the pair (x, y) . Contradiction.

We proceed to bound the number of edges E_H in H . For each $\rho \in S$, $x \in V_n$, $b \in \{0, 1\}$, and $w \in \{0, 1\}^{\ell_2(n)}$, the number of vertices in $N_{b,w}^\rho(x)$ comparable to x is at most $\ell_1(n)$. Therefore,

$$|E_H| \leq \ell_1(n) \cdot 2 \cdot 2^{\ell_2(n)} \cdot |S| \leq O\left(n \cdot \ell_1(n) \cdot 2^{\ell_2(n)} \lceil \log n / \log(1/\delta) \rceil\right). \quad \square$$

In Theorems 5.1 and 5.2 when δ is sufficiently small the bounds on the 2-TC-Spanner size become $O(n\ell(n))$ and $O(n\ell_1(n) \cdot 2^{\ell_2(n)})$, respectively. As pointed out earlier, all known monotonicity reconstructors are non-adaptive. It is an open question whether it is possible to give a transformation from adaptive local monotonicity reconstructors to 2-TC-spanners without incurring an exponential dependence on the number of lookups made to points incomparable to the query point. It is not known if this dependence is an artifact of the proof or an indication that lookups to incomparable points might be helpful for adaptive local monotonicity reconstructors.

Theorems 5.1 and Theorem 5.2 imply the following lower bounds on the lookup complexity of local monotonicity reconstructors. These lower bounds hold for any error-blow up.

Corollary 5.1 ([14]). *Consider a nonadaptive local monotonicity filter with constant error probability δ . If the filter is for functions $f : \mathcal{H}_{m,d} \rightarrow \mathbb{R}$, it must perform $\Omega\left(\frac{\log^{d-1} m}{d^d(2 \log \log m)^{d-1}}\right)$ lookups per query. If the filter is for functions $f : \mathcal{H}_d \rightarrow \mathbb{R}$, it must perform $\Omega(2^{\alpha d}/d)$ lookups per query, where $\alpha \geq 0.1620$.*

Corollary 5.2 ([14]). *Consider an (adaptive) local monotonicity filter with constant error probability δ , that for every query $x \in V_n$, looks up at most ℓ_2 vertices incomparable to x . If the filter is for functions $f : \mathcal{H}_{m,d} \rightarrow \mathbb{R}$, it must perform $\Omega\left(\frac{\log^{d-1} m}{2^{\ell_2} d^d (2 \log \log m)^{d-1}}\right)$ lookups to vertices comparable to x per query x . If the filter is for functions $f : \mathcal{H}_d \rightarrow \mathbb{R}$, it must perform $\Omega(2^{\alpha d - \ell_2}/d)$ comparable lookups, where $\alpha \geq 0.1620$.*

Prior to [14], no lower bounds for monotonicity reconstructors on $\mathcal{H}_{m,d}$ with dependence on both m and d were known. Unlike the bound in [48], the TC-spanner-based lower bounds hold for any error blow-up. These bounds are tight for reconstructors that are either non-adaptive or perform the number of incomparable lookups that is polylogarithmic in the number of points in the domain. Specifically, for the hypergrid $\mathcal{H}_{m,d}$ of constant dimension d , the number of lookups is $(\log m)^{\Theta(d)}$, and for the hypercube \mathcal{H}_d , it is $2^{\Theta(d)}$ for any error blow-up.

5.3 Application to Key Management in Access Control Hierarchies

Atallah *et al.* [9] used sparse Steiner TC-spanners to construct efficient key management schemes for access control hierarchies. An *access hierarchy* is a partially ordered set G of access classes. Each user is entitled to access a certain class and all classes reachable from the corresponding node in G . One approach for devising a cryptographic protocol that enforces the access hierarchy is to have the users follow a *key management scheme* [8, 9, 49, 7, 17]. Here, each edge (i, j) has an associated public key $P(i, j)$, and each node i , an associated secret key k_i . Only users with the secret key for a node have the required permissions for the associated access class. The public and secret keys are designed so that there is an efficient algorithm A which takes k_i and $P(i, j)$ and generates k_j , but for each (i, j) in G , it is computationally hard to generate k_j without knowledge of k_i . Thus, a user can efficiently generate the required keys to access a descendant class, but not other classes. The number of runs of algorithm A needed to generate a secret key k_v from a secret key k_u is equal to the number of edges on the shortest path from u to v in G . To speed this up, Atallah *et al.* [7] suggest adding edges and nodes to G to increase connectivity. To preserve the access hierarchy represented by G , the new graph H must be a Steiner TC-spanner of G . The number of edges in H corresponds to the space complexity of the scheme, while the stretch k of the spanner corresponds to the time complexity.

We note that the time to find the path from u to v is also important in this application. In the upper bounds from [17] listed in Table 1, this time is $O(d)$, which for, say, constant d is likely to be much less than $2g(n)$ or $3g(n)$, where $g(n)$ is the time to run algorithm A . This is because algorithm A involves the evaluation of a cryptographic hash function, which is very expensive: any hash function secure against poly(n)-time adversaries requires $g(n) \geq \text{polylog } n$ evaluation time under existing number-theoretic assumptions.

5.4 Application to Computing Partial Products in a Semigroup

Yao [60] 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) to be able to compute partial products $s_i \circ s_{i+1} \circ \dots \circ s_j$ for all $i, j \in [n]$ with at most k queries to a small database of pre-computed partial products. Examples of a semigroup (S, \circ) include (\mathbb{R}, \min) , the space of real d -dimensional vectors with operation $(x_1, \dots, x_d) \circ (y_1, \dots, y_d) = (\min(x_1, y_1), \dots, \min(x_d, y_d))$, and the space of real $d \times d$ matrices equipped with the multiplication operation.

Bhattacharyya *et al.* [15] point out that the problem of computing partial products in a semigroup reduces to finding a sparsest k -TC-spanner for a directed line L_{n+1} . If the database stores a product $s_u \circ \dots \circ s_v$ for each k -TC-spanner edge $(u, v + 1)$, every product $s_i \circ \dots \circ s_j$ can be computed by multiplying the products corresponding to the edges on a path of length at most k from i to $j + 1$ in the k -TC-spanner for L_{n+1} .

Chazelle [19] and Alon and Schieber [5] 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 us to compute the product of elements associated with all vertices on the path from i to j , for all vertices i, j in T . As before, only k queries to the data structure are allowed for each product computation. The generalized problem reduces to finding a sparsest k -TC-spanner for a 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 . To compute the product corresponding to a path from u to v in T , we consider 2 cases: (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.

Acknowledgment

The author would like to thank Oded Goldreich for persuading her to write this survey and Adam Smith, Ramesh T.K. and Piotr Berman for useful comments.

References

1. Abraham, I., Gavaille, C.: Object location using path separators. In: PODC. pp. 188–197 (2006)
2. Aho, A.V., Garey, M.R., Ullman, J.D.: The transitive reduction of a directed graph. *SIAM J. Comput.* 1(2), 131–137 (1972)
3. Ailon, N., Chazelle, B., Comandur, S., Liu, D.: Property-preserving data reconstruction. *Algorithmica* 51(2), 160–182 (2008)
4. Ailon, N., Chazelle, B.: Information theory in property testing and monotonicity testing in higher dimension. *Inf. Comput.* 204(11), 1704–1717 (2006)
5. Alon, N., Schieber, B.: Optimal preprocessing for answering on-line product queries. Tech. Rep. 71/87, Tel-Aviv University (1987)
6. Althöfer, I., Das, G., Dobkin, D., Joseph, D., Soares, J.: On sparse spanners of weighted graphs. *Discrete & Computational Geometry* 9(1), 81–100 (1993)
7. Atallah, M.J., Blanton, M., Fazio, N., Frikken, K.B.: Dynamic and efficient key management for access hierarchies. *ACM Trans. Inf. Syst. Secur.* 12(3), 1–43 (2009)
8. Atallah, M.J., Blanton, M., Frikken, K.B.: Key management for non-tree access hierarchies. In: SACMAT. pp. 11–18 (2006)
9. Atallah, M.J., Frikken, K.B., Blanton, M.: Dynamic and efficient key management for access hierarchies. In: ACM Conference on Computer and Communications Security. pp. 190–202 (2005)

10. Awerbuch, B.: Communication-time trade-offs in network synchronization. In: PODC. pp. 272–276 (1985)
11. Baswana, S., Sen, S.: Approximate distance oracles for unweighted graphs in expected $\tilde{O}(n^2)$ time. *ACM Transactions on Algorithms* 2(4), 557–577 (2006)
12. Batu, T., Rubinfeld, R., White, P.: Fast approximate PCPs for multidimensional bin-packing problems. *Inf. Comput.* 196(1), 42–56 (2005)
13. Berman, P., Raskhodnikova, S., Ruan, G.: Finding sparser directed spanners (2010), manuscript
14. Bhattacharyya, A., Grigorescu, E., Jha, M., Jung, K., Raskhodnikova, S., Woodruff, D.: Lower bounds for local monotonicity reconstruction from transitive-closure spanners. In: *RANDOM* (2010)
15. Bhattacharyya, A., Grigorescu, E., Jung, K., Raskhodnikova, S., Woodruff, D.: Transitive-closure spanners. In: *SODA*. pp. 932–941 (2009)
16. Bhattacharyya, A., Grigorescu, E., Jung, K., Raskhodnikova, S., Woodruff, D.: Transitive-closure spanners of the hypercube and the hypergrid (2009), eCCC Report TR09-046
17. Bhattacharyya, A., Grigorescu, E., Raskhodnikova, S., Woodruff, D.: Steiner transitive-closure spanners of d -dimensional posets (2010), manuscript
18. Bodlaender, H.L., Tel, G., Santoro, N.: Tradeoffs in non-reversing diameter. *Nordic Journal of Computing* 1(1), 111 – 134 (1994)
19. Chazelle, B.: Computing on a free tree via complexity-preserving mappings. *Algorithmica* 2, 337–361 (1987)
20. Cohen, E.: Fast algorithms for constructing t -spanners and paths with stretch t . *SIAM J. Comput.* 28(1), 210–236 (1998)
21. Cohen, E.: Polylog-time and near-linear work approximation scheme for undirected shortest paths. *JACM* 47(1), 132–166 (2000)
22. Cowen, L.: Compact routing with minimum stretch. *J. Algorithms* 38(1), 170–183 (2001)
23. Cowen, L., Wagner, C.G.: Compact roundtrip routing in directed networks. *J. Algorithms* 50(1), 79–95 (2004)
24. Dilworth, R.P.: A decomposition theorem for partially ordered sets. *The Annals of Mathematics, Second Series* 51(1), 161–166 (1950)
25. Dodis, Y., Goldreich, O., Lehman, E., Raskhodnikova, S., Ron, D., Samorodnitsky, A.: Improved testing algorithms for monotonicity. In: *RANDOM*. pp. 97–108 (1999)
26. Dushnik, B., Miller, E.: Concerning similarity transformations of linearly ordered sets. *Bulletin Amer. Math. Soc.* 46, 322–326 (1940)
27. Elkin, M.: Computing almost shortest paths. In: *PODC*. pp. 53–62 (2001)
28. Elkin, M., Peleg, D.: Strong inapproximability of the basic k -spanner problem. In: *ICALP*. pp. 636–647 (2000)
29. Elkin, M., Peleg, D.: The client-server 2-spanner problem with applications to network design. In: *SIROCCO*. pp. 117–132 (2001)
30. Elkin, M., Peleg, D.: The hardness of approximating spanner problems. *Theory Comput. Syst.* 41(4), 691–729 (2007)
31. Ergun, F., Kannan, S., Kumar, S.R., Rubinfeld, R., Viswanathan, M.: Spot-checkers. *JCSS* 60(3), 717–751 (2000)
32. Fischer, E.: On the strength of comparisons in property testing. *Inf. Comput.* 189(1), 107–116 (2004)
33. Fischer, E., Lehman, E., Newman, I., Raskhodnikova, S., Rubinfeld, R., Samorodnitsky, A.: Monotonicity testing over general poset domains. In: *STOC*. pp. 474–483 (2002)

34. Goldreich, O., Goldwasser, S., Ron, D.: Property testing and its connection to learning and approximation. *JACM* 45(4), 653–750 (1998)
35. Goldreich, O., Goldwasser, S., Lehman, E., Ron, D., Samorodnitsky, A.: Testing monotonicity. *Combinatorica* 20(3), 301–337 (2000)
36. Halevy, S., Kushilevitz, E.: Testing monotonicity over graph products. In: *ICALP*. pp. 721–732 (2004)
37. Hesse, W.: Directed graphs requiring large numbers of shortcuts. In: *SODA*. pp. 665–669 (2003)
38. Hochbaum, D. (ed.): *Approximation Algorithms for NP-hard Problems*. PWS Publishing Company, Boston (1997)
39. Jha, M., Raskhodnikova, S.: Testing and reconstruction of Lipschitz functions with applications to data privacy (2010), manuscript
40. Kortsarz, G.: On the hardness of approximating spanners. *Algorithmica* 30(3), 432–450 (2001)
41. Lipton, R.J., Tarjan, R.E.: A separator theorem for planar graphs. *SIAM Journal on Applied Mathematics* 36(2), 177–189 (1979), <http://www.jstor.org/stable/2100927>
42. Peleg, D.: *Distributed computing: a locality-sensitive approach*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA (2000)
43. Peleg, D., Schäffer, A.A.: Graph spanners. *Journal of Graph Theory* 13(1), 99–116 (1989)
44. Peleg, D., Ullman, J.D.: An optimal synchronizer for the hypercube. *SIAM J. Comput.* 18(4), 740–747 (1989)
45. Peleg, D., Upfal, E.: A trade-off between space and efficiency for routing tables. *JACM* 36(3), 510–530 (1989)
46. Roditty, L., Thorup, M., Zwick, U.: Roundtrip spanners and roundtrip routing in directed graphs. In: *SODA*. pp. 844–851 (2002)
47. Rubinfeld, R., Sudan, M.: Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing* 25(2), 252–271 (1996)
48. Saks, M.E., Seshadhri, C.: Parallel monotonicity reconstruction. In: *Proceedings of the 19th Annual Symposium on Discrete Algorithms (SODA)*. pp. 962–971 (2008)
49. Santis, A.D., Ferrara, A.L., Masucci, B.: Efficient provably-secure hierarchical key assignment schemes. In: *MFCS*. pp. 371–382 (2007)
50. Seidel, R.: Understanding the inverse Ackermann function (2006), available at <http://cgi.di.uoa.gr/~ewcg06/invited/Seidel.pdf>
51. Soriano, D.G., Matsliah, A., Chakraborty, S., Briet, J.: Monotonicity testing and shortest-path routing on the cube (2010), eCCC Report TR10-048
52. Thorup, M.: On shortcutting digraphs. In: *Graph-Theoretic Concepts in Computer Science*. vol. 657, pp. 205–211 (1993)
53. Thorup, M.: Shortcutting planar digraphs. *Combinatorics, Probability & Computing* 4, 287–315 (1995)
54. Thorup, M.: Parallel shortcutting of rooted trees. *J. Algorithms* 23(1), 139–159 (1997)
55. Thorup, M., Zwick, U.: Compact routing schemes. In: *ACM Symposium on Parallel Algorithms and Architectures*. pp. 1–10 (2001), [cite-seer.ist.psu.edu/thorup01compact.html](http://citeseer.ist.psu.edu/thorup01compact.html)
56. Thorup, M., Zwick, U.: Approximate distance oracles. *JACM* 52(1), 1–24 (2005)
57. Trotter, W. (ed.): *Combinatorics and Partially Ordered Sets: Dimension Theory*. Johns Hopkins University Press, Baltimore, MD (1992)
58. Woodruff, D.P.: Lower bounds for additive spanners, emulators, and more. In: *FOCS*. pp. 389–398 (2006)

59. Yannakakis, M.: The complexity of the partial order dimension problem. *JMAA* 3(3), 351–358 (1982)
60. Yao, A.C.C.: Space-time tradeoff for answering range queries (extended abstract). In: *STOC*. pp. 128–136 (1982)