

Steiner Transitive-Closure Spanners of Low-Dimensional Posets^{*,**}

Piotr Berman¹, Arnab Bhattacharyya², Elena Grigorescu³,
Sofya Raskhodnikova¹, David P. Woodruff⁴, and Grigory Yaroslavtsev¹

¹ Pennsylvania State University, USA

{berman,sofya,grigory}@cse.psu.edu

² Massachusetts Institute of Technology, USA

abhatt@mit.edu

³ Georgia Institute of Technology, USA

elena@cc.gatech.edu

⁴ IBM Almaden Research Center, USA

dpwoodru@us.ibm.com

Abstract. 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 $H = (V_H, E_H)$ such that (1) $V \subseteq V_H$ and (2) for all vertices $v, u \in V$, the distance from v to u in H is at most k if u is reachable from v in G , and ∞ otherwise. Motivated by applications to property reconstruction and access control hierarchies, we concentrate on Steiner TC-spanners of directed acyclic graphs or, equivalently, partially ordered sets. We study the relationship between the dimension of a poset and the size, denoted S_k , of its sparsest Steiner k -TC-spanner.

We present a nearly tight lower bound on S_2 for d -dimensional directed hypergrids. Our bound is derived from an explicit dual solution to a linear programming relaxation of the 2-TC-spanner problem. We also give an efficient construction of Steiner 2-TC-spanners, of size matching the lower bound, for all low-dimensional posets. Finally, we present a nearly tight lower bound on S_k for d -dimensional posets.

1 Introduction

Graph spanners were introduced in the context of distributed computing by Awerbuch [3] and Peleg and Schäffer [12], and since then have found numerous applications. Our focus is on transitive-closure spanners, introduced explicitly in [5], but studied prior to that in many different contexts (see references in [5]).

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 $H = (V, E_H)$ such

* E.G. is supported by NSF award CCR-0829672 and NSF award 1019343 to the Computing Research Association for the Computing Innovation Fellowship Program. S.R. and G.Y. are supported by NSF / CCF CAREER award 0845701. G.Y. is also supported by University Graduate Fellowship and College of Engineering Fellowship.

** Full version is available at <http://arxiv.org/abs/1011.6100>.

that: (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 $d_H(u, v) \leq k$ and if $d_G(u, v) = \infty$ then $d_H(u, v) = \infty$, where $d_G(u, v)$ denotes the distance from u to v in G . That is, a k -TC-spanner is a graph with a small diameter that preserves the connectivity of the original graph. The edges of the transitive closure of G , added to G to obtain a TC-spanner, are called *shortcuts* and the parameter k is called the *stretch*.

TC-spanners have numerous applications, and there has been a lot of work on finding sparse TC-spanners for specific graph families. See [13] for a survey. In some applications of TC-spanners, in particular, to access control hierarchies [2,8], 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 1.1 (Steiner TC-spanner). *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 $H = (V_H, E_H)$ such that: (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. For example, consider 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. Every ordinary 2-TC-spanner of this graph has $\Omega(n^2)$ edges. However, $K_{\frac{n}{2}, \frac{n}{2}}$ has a Steiner 2-TC-spanner with n edges: 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 factor of $\Theta(n)$ gap between the size of the sparsest Steiner 2-TC-spanner and the size of an ordinary 2-TC-spanner.

We focus on Steiner TC-spanners of directed *acyclic* graphs (DAGs) or, equivalently, partially ordered sets (posets). They represent the most interesting case in applications of TC-spanners. In addition, there is a reduction from constructing TC-spanners of graphs with cycles to constructing TC-spanners of DAGs, with a small loss in stretch ([13], Lemma 3.2), which also applies to Steiner TC-spanners.

The goal of this work is to understand the minimum number of edges needed to form a Steiner k -TC-spanner of a given graph G as a function of n , the number of nodes in G . More specifically, motivated by applications to access control hierarchies [2,8] and property reconstruction [4,11], described in Section 1.2, we study the relationship between the dimension of a poset and the size of its sparsest Steiner TC-spanner. The *dimension* of a poset G is the smallest d such that G can be embedded into a d -dimensional directed hypergrid via an order-preserving embedding. (See Definition 2.1). Atallah *et al.* [2], followed by De Santis *et al.* [8], use Steiner TC-spanners in key management schemes for access control hierarchies. They argue that many access control hierarchies are low-dimensional posets that come equipped with an embedding demonstrating low dimensionality. For this reason, we focus on the setting where the dimension d is small relative to the number of nodes n .

We also study the size of sparsest (Steiner) 2-TC-spanners of specific posets of dimension d , namely, d -dimensional directed hypergrids. Our lower bound on this

quantity improves the lower bound of [4] and nearly matches their upper bound. It implies that our construction of Steiner 2-TC-spanners of d -dimensional posets is optimal up to a constant factor for any constant number of dimensions. It also has direct implications for property reconstruction.

1.1 Our Results

Steiner 2-TC-spanners of Directed d -dimensional Grids. The *directed hypergrid*, denoted $\mathcal{H}_{m,d}$, has vertex set¹ $[m]^d$ and edge set $\{(x, y) : \exists \text{ unique } i \in [d] \text{ such that } y_i - x_i = 1 \text{ and if } j \neq i, y_j = x_j\}$. We observe (in Corollary 2.4) that for the grid $\mathcal{H}_{m,d}$, Steiner vertices do not help to create sparser k -TC-spanners. In [4], it was shown that for $m \geq 3$, sparsest (ordinary) 2-TC-spanners of $\mathcal{H}_{m,d}$ have size at most $m^d \log^d m$ and at least $\Omega\left(\frac{m^d \log^d m}{(2d \log \log m)^{d-1}}\right)$. They also give tight upper and lower bounds for the case of constant m and large d . Our first result is an improvement on the lower bound for the hypergrid for the case when m is significantly larger than d , i.e., the setting in the above applications.

Theorem 1.1. *All (Steiner) 2-TC-spanners of $\mathcal{H}_{m,d}$ have $\Omega\left(\frac{m^d (\ln m - 1)^d}{(4\pi)^d}\right)$ edges.*

The proof of Theorem 1.1 constructs a dual solution to a linear programming relaxation of the 2-TC-spanner problem. We consider a linear program (LP) for the sparsest 2-TC-spanner of $\mathcal{H}_{m,d}$. Our program is a special case of a more general LP for the sparsest directed k -spanner of an arbitrary graph G , used in [5] to obtain an approximation algorithm for that problem. We show that for our special case the integrality gap of this LP is small and, in particular, does not depend on n . Specifically, we find a solution to the dual LP by selecting initial values that have a combinatorial interpretation: they are expressed in terms of the *volume* of d -dimensional *boxes* contained in $\mathcal{H}_{m,d}$. For example, the dual variable corresponding to the constraint that enforces the existence of a length-2 path from u to v in the 2-TC-spanner is initially assigned a value inversely proportional to the number of nodes on the paths from u to v . The final sum of the constraints is bounded by an integral which, in turn, is bounded by an expression depending only on the dimension d .

We note that the best lower bound known previously [4] was proved by a long and sophisticated combinatorial argument that carefully balanced the number of edges that stay within different parts of the hypergrid and the number of edges that cross from one part to another. The recursion in the combinatorial argument is an inherent limitation of [4], resulting in suboptimal bounds even for constant d . In contrast, our linear programming argument can be thought of as assigning types to edges based on the volume of the boxes they define, and automatically balancing the number of edges of different types by selecting the correct coefficients for the constraints corresponding to those edges. It achieves an optimal bound for any constant number of dimensions.

¹ For a positive integer m , we denote $\{1, \dots, m\}$ by $[m]$.

Table 1. Steiner k -TC-spanner sizes for d -dimensional posets on n vertices for $d \geq 2$

Stretch k	Prior bounds on $S_k(G)$	Stretch k	Our bounds on $S_k(G)$
$2d - 1$	$O(n^2)$ [2]	2	$\Omega\left(n\left(\frac{\log n}{cd}\right)^d\right)$ for a fixed $c > 0$
$2d - 2 + t \ \forall t \geq 2$	$O(n(\log^{d-1} n)\lambda_t(n))$ [2]		
$2d + O(\log^* n)$	$O(n \log^{d-1} n)$ [2]		
3	$O(n \log^{d-1} n \log \log n)$ for fixed d [8]	≥ 3	$\Omega(n \log^{\lceil (d-1)/k \rceil} n)$ for fixed d

Steiner TC-spanners of General d -dimensional Posets. We continue the study of the number of edges in a sparsest Steiner k -TC-spanner of a poset as a function of its dimension, following [2,8]. We note that the only poset of dimension 1 is the directed line $\mathcal{H}_{n,1}$. TC-spanners of directed lines were discovered under many different guises. (See references in [5].) It was implicitly shown in [6,7] that, for constant k , the size of the sparsest k -TC-spanner of $\mathcal{H}_{n,1}$ is $\Theta(n \cdot \lambda_k(n))$, where $\lambda_k(n)$ is the k^{th} -row inverse Ackermann function.

Table 1 compares old and new results for $d \geq 2$. $S_k(G)$ denotes the number of edges in the sparsest Steiner k -TC-spanner of G . The upper bounds hold for all posets of dimension d . The lower bounds mean that there is an infinite family of d -dimensional posets with sparsest Steiner k -TC-spanners of the specified size.

Atallah *et al.* constructed Steiner k -TC-spanners with k proportional to d . De Santis *et al.* improved their construction for constant d . They achieved $O(3^{d-t} n t \log^{d-1} n \log \log n)$ edges for odd stretch $k = 2t + 1$, where $t \in [d]$. In particular, setting $t = 1$ gives $k = 3$ and $O(n \log^{d-1} n \log \log n)$ edges.

We present the first construction of Steiner 2-TC-spanners for d -dimensional posets. In our construction, the spanners have $O(n \log^d n)$ edges, and the length-2 paths can be found in $O(d)$ time. This result is stated in Theorem 2.2 (in Section 2). Our construction, like all previous constructions, takes as part of the input an explicit embedding of the poset into a d -dimensional grid. (Finding such an embedding is NP-hard [15]. Also, as mentioned previously, in the application to access control hierarchies, such an embedding is usually given.) The Steiner vertices used in our construction are necessary to obtain sparse TC-spanners. An (easy) example that demonstrates this is deferred to the full version.

Theorem 1.1 implies that there is an absolute constant $c > 0$ for which our upper bound for $k = 2$ is tight within an $O((cd)^d)$ factor, showing that no drastic improvement in the upper bound is possible. To obtain a bound in terms of the number n of vertices and dimension d , substitute m^d with n and $\ln m$ with $(\ln n)/d$ in the theorem statement. This gives the following corollary.

Corollary 1.2 *There is an absolute constant $c > 0$ for which for all $d \geq 2$ and n larger than some constant to the power d , there exists a d -dimensional poset G on n vertices such that every Steiner 2-TC-spanner of G has $\Omega\left(n\left(\frac{\log n}{cd}\right)^d\right)$ edges.*

In addition, we prove a lower bound for all constant $k > 2$ and constant dimension d , which qualitatively matches known upper bounds. It shows that, in

particular, every Steiner 3-TC-spanner has size $\Omega(n \log n)$, and even with significantly larger constant stretch, every Steiner TC-spanner has size $n \log^{\Omega(d)} n$.

Theorem 1.3. *For all constant $d \geq 2$ and sufficiently large n , there exists a d -dimensional poset G on n vertices such that for all $k \geq 3$, every Steiner k -TC-spanner of G has $\Omega(n \log^{[(d-1)/k]} n)$ edges.*

This theorem (see Section 4) captures the dependence on d and greatly improves upon the previous $\Omega(n \log \log n)$ bound, which follows trivially from known lower bounds for 3-TC-spanners of a directed line.

The lower bound on the size of a Steiner k -TC-spanner for $k \geq 3$ is proved by the probabilistic method. We note that using the hypergrid as an example of a poset with large Steiner k -TC-spanners for $k > 2$ would yield a much weaker lower bound because $\mathcal{H}_{m,d}$ has a 3-TC-spanner of size $O((m \log \log m)^d)$ and, more generally, a k -TC-spanner of size $O((m \cdot \lambda_k(m))^d)$, where $\lambda_k(m)$ is the k^{th} -row inverse Ackermann function [4]. Instead, we construct an n -element poset embedded in $\mathcal{H}_{n,d}$ as follows: all poset elements differ on coordinates in dimension 1, and for each element, the remaining $d - 1$ coordinates are chosen uniformly at random from $[n]$. We consider a set of partitions of the underlying hypergrid into d -dimensional boxes, and carefully count the expected number of edges in a Steiner k -TC-spanner that cross box boundaries for each partition. We show that each edge is counted only a small number of times, proving that the expected number of edges in a Steiner k -TC-spanner is large. We conclude that some poset attains the expected number of edges.

Organization. We explain applications of Steiner TC-spanners in Section 1.2. Section 2 gives basic definitions and observations. In particular, our construction of sparse Steiner 2-TC-spanners for d -dimensional posets (the proof of Theorem 2.2) is presented there. Our lower bounds constitute the main technical contribution of this paper. The lower bound for the hypergrid for $k = 2$ (Theorem 1.1) is proved in Section 3. The lower bound for $k > 2$ (Theorem 1.3) is presented in Section 4.

1.2 Applications

Numerous applications of TC-spanners are surveyed in [13]. We focus on two of them: property reconstruction, described in [4,11], and key management for access control hierarchies, described in [2,5,8].

Property Reconstruction. A *local filter* [14] (see also a slightly modified definition in [4,11]) reconstructs an arbitrary function f to ensure that the reconstructed function g has the desired property, changing f only when necessary. A local filter is given a function f and a query x and, after looking up the value of f on a small number of points, it has to output $g(x)$ for some function g , which has the desired property and does not depend on x . If f has the property, g must be equal to f .

Our results on TC-spanners are relevant to reconstruction of two properties of functions: monotonicity, studied in [1,4,14] and having a low Lipschitz constant, studied in [11]. In [4], the authors proved that the existence of a local filter for monotonicity of functions with low lookup complexity implies the existence of a sparse 2-TC-spanner of $\mathcal{H}_{m,d}$. In [11], an analogous connection was drawn between local reconstruction of functions with low Lipschitz constant and 2-TC-spanners. Our improvement in the lower bound on the size of 2-TC-spanners of $\mathcal{H}_{m,d}$ directly translates into an improvement by the same factor in the lower bounds on lookup complexity of local nonadaptive filters for these two properties, showing they are nearly optimal for any constant d .

Key Management for Access Control Hierarchies. Atallah *et al.* [2] 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 . In the approach from [2,8] to enforcing the access hierarchy, a user from a class u can compute cryptographic keys necessary to access a class v in time proportional to $d_G(u, v)$. To speed this up, Atallah *et al.* 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 . With this modification, the number of edges in H corresponds to the space complexity of the scheme, while the running time has two components: the time to find a path of length at most k from u to v in H and the time to compute the cryptographic keys. The second component is proportional to the stretch k of H . In our construction of Steiner 2-TC-spanners, the time to find length- k paths is $O(d)$. For small d , it is likely to be dominated by the second component which involves a (time-consuming) evaluation of a cryptographic hash function.

2 Definitions and Observations

For integers $j \geq i$, an interval $[i, j]$ refers to the set $\{i, i + 1, \dots, j\}$. Logarithms are always base 2, except for \ln which is the natural logarithm.

Each DAG $G = (V, E)$ is equivalent to a poset with elements V and partial order \preceq , where $x \preceq y$ if y is reachable from x in G . Elements x and y are *comparable* if $x \preceq y$ or $y \preceq x$, and *incomparable* otherwise. We write $x \prec y$ if $x \preceq y$ and $x \neq y$. The *hypergrid* $\mathcal{H}_{m,d}$ with dimension d and side length m was defined in the beginning of Section 1.1. Equivalently, it is the poset on elements $[m]^d$ with the *dominance order*, defined as follows: $x \preceq y$ for two elements $x, y \in [m]^d$ iff $x_i \leq y_i$ for all $i \in [d]$.

A mapping f from a poset G to a poset G' is called an *embedding* if it respects the partial order, that is, $f(x) \preceq_{G'} f(y)$ iff $x \preceq_G y$ for all $x, y \in G$.

Definition 2.1 ([10]). *Let G be a poset with n elements. The dimension of G is the smallest integer d such that G can be embedded into the hypergrid $\mathcal{H}_{n,d}$.*

As shown in [9], for any $m > 1$, the hypergrid $\mathcal{H}_{m,d}$ has dimension exactly d .

Fact 2.1. *Each d -dimensional poset G 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 elements are distinct. Moreover, such an embedding can be obtained from an arbitrary embedding of G into $\mathcal{H}_{n,d}$ in time $O(dn \log n)$.*

Sparse Steiner 2-TC-spanners for d -dimensional Posets. We give a simple construction of sparse Steiner 2-TC-spanners for d -dimensional posets. For constant d , it matches the lower bound from Section 3 up to a constant factor. Note that the construction itself works for arbitrary, not necessarily constant, d .

Theorem 2.2. *Each d -dimensional poset G on n elements has a Steiner 2-TC-spanner H of size $O(n \log^d n)$. Given an embedding of G into the hypergrid $\mathcal{H}_{n,d}$, H can be constructed in time $O(dn \log^d n)$. Moreover, for all $x, y \in G$, where $x \prec y$, one can find a path in H from x to y of length at most 2 in time $O(d)$.*

Proof. Consider an n -element poset G embedded into the hypergrid $\mathcal{H}_{n,d}$. Transform it, so that for all $i \in [d]$, the i th coordinates of images of all elements are distinct. (See Fact 2.1.) In this proof, assume that the hypergrid coordinates start with 0, i.e., its vertex set is $[0, n - 1]^d$. Let $\ell = \lceil \log n \rceil$ and $b(t)$ be the ℓ -bit binary representation of t , possibly with leading zeros. Let $p_i(t)$ denote the i -bit prefix of $b(t)$ followed by a single 1 and then $\ell - i - 1$ zeros. Let $lcp(t_1, t_2) = p_i(t_1)$, where i is the length of the longest common prefix of $b(t_1)$ and $b(t_2)$.

To construct a Steiner 2-TC-spanner (V_H, E_H) of G , we insert at most ℓ^d edges into E_H per each poset element. Consider a poset element with coordinates $x = (x_1, \dots, x_d)$ in the embedding. For each d -tuple $(i_1, \dots, i_d) \in [0, \ell - 1]^d$, let p be a hypergrid vertex whose coordinates have binary representations $(p_{i_1}(x_1), \dots, p_{i_d}(x_d))$. If $x \prec p$, we add an edge (x, p) to E_H ; otherwise, if $p \prec x$ we add an edge (p, x) to E_H . Note that only edges between comparable points are added to E_H .

Observe that for $d > (2 \log n)/(\log \log n)$, the theorem is trivial since then $n \log^d n > n^3$, and the transitive-closure of G has $O(n^2)$ edges and can be computed in $O(n^3)$ time. For smaller d , $\lceil \log n \rceil^d = O(\log^d n)$ and, consequently, E_H contains $O(n \log^d n)$ edges and can be constructed in $O(dn \log^d n)$ time, as described, if bit operations on coordinates can be performed in $O(1)$ time.

For all pairs of poset elements $x = (x_1, \dots, x_d)$ and $y = (y_1, \dots, y_d)$, such that $x \prec y$, there is an intermediate point z with coordinates whose binary representations are $(lcp(x_1, y_1), \dots, lcp(x_d, y_d))$. By construction, both edges (x, z) and (z, y) are in E_H . Point z can be found in $O(d)$ time, since $lcp(x_i, y_i)$ can be computed in $O(1)$ time, assuming $O(1)$ time bit operations on coordinates. \square

Equivalence of Steiner and non-Steiner TC-spanners for Hypergrids. Our lower bound on the size of 2-TC-spanners for d -dimensional posets of size n is obtained by proving a lower bound on the size of the Steiner 2-TC-spanner of $\mathcal{H}_{m,d}$ where $m = n^{1/d}$. The following lemma, used in Section 4, implies Corollary 2.4 that shows that sparsest Steiner and non-Steiner 2-TC-spanners of $\mathcal{H}_{m,d}$ have the same size. The proof of the lemma is deferred to the full version.

Lemma 2.3 *Let G be a poset on elements $V \subseteq [m]^d$ with the dominance order and $H = (V_H, E_H)$ be a Steiner k -TC-spanner of G with minimal V_H . Then H can be embedded into $\mathcal{H}_{m,d}$.*

Corollary 2.4 *If $\mathcal{H}_{m,d}$ has a Steiner k -TC-spanner H , it also has a k -TC-spanner with the same number of nodes and at most the same number of edges.*

3 Lower Bound for 2-TC-spanners of the Hypergrid

In this section, we prove Theorem 1.1 that gives a nearly tight lower bound on the size of (Steiner) 2-TC-spanners of the hypergrids $\mathcal{H}_{m,d}$. By Corollary 2.4, we only have to consider non-Steiner TC-spanners.

Proof (of Theorem 1.1). We start by introducing an LP for the sparsest 2-TC-spanner of an arbitrary graph. Our lower bound on the size of a 2-TC-spanner of $\mathcal{H}_{m,d}$ is obtained by finding a feasible solution to the dual program, which, by definition, gives a lower bound on the objective function of the primal.

An Integer LP for Sparsest 2-TC-spanner. For each graph $G = (V, E)$, we can find the size of a sparsest 2-TC-spanner by solving the following $\{0,1\}$ -LP, a special case of an LP from [5] for directed k -spanners. For all vertices $u, v \in V$ satisfying $u \preceq v$, we introduce variables $x_{uv} \in \{0, 1\}$. For $u \neq v$, they correspond to potential edges in a 2-TC-spanner H of G . For all vertices $u, v, w \in V$ satisfying $u \preceq w \preceq v$, we introduce auxiliary variables $x'_{uvw} \in \{0, 1\}$, corresponding to potential paths of length at most 2 in H . The $\{0,1\}$ -LP is as follows:

$$\begin{aligned}
 &\text{minimize} && \sum_{u,v: u \preceq v} x_{uv} \\
 &\text{subject to} && x_{uw} - x'_{uvw} \geq 0, x_{wv} - x'_{uvw} \geq 0 && \forall u, v, w: u \preceq w \preceq v; \\
 &&& \sum_{w: u \preceq w \preceq v} x'_{uvw} \geq 1 && \forall u, v: u \preceq v.
 \end{aligned}$$

Given a solution to the LP, we can construct a 2-TC-spanner $H = (V, E_H)$ of G of size not exceeding the value of the objective function by including (u, v) in E_H iff the corresponding variable $x_{uv} = 1$ and $u \neq v$. In the other direction, given a 2-TC-spanner $H = (V, E_H)$ of G , we can find a feasible solution of the LP with the value of the objective function not exceeding $|E_H| + |V|$. Let $E'_H = E_H \cup L$, where L is the set of loops (v, v) for all $v \in V$. Then we set $x_{uv} = 1$ iff $(u, v) \in E'_H$ and $x'_{uvw} = 1$ iff both $(u, w) \in E'_H$ and $(w, v) \in E'_H$. Therefore, the size of a sparsest 2-TC-spanner of G and the optimal value of the objective function of the LP differ by at most $|V|$. They are asymptotically equivalent because $|V| = O(|E_H|)$ for every weakly connected graph G .

A Fractional Relaxation of the Dual LP. Every feasible solution of the following fractional relaxation of the dual LP gives a lower bound on the optimal value of the objective function of the primal:

$$\begin{aligned}
 &\text{maximize} && \sum_{u,v: u \preceq v} y_{uv} \\
 &\text{subject to} && \sum_{w: v \preceq w} y'_{uvw} + \sum_{w: w \preceq u} y''_{uvw} \leq 1 && \forall u, v: u \preceq v; && (1) \\
 &&& y_{uv} - y'_{uvw} - y''_{uvw} \leq 0 && \forall u, v, w: u \preceq w \preceq v; && (2) \\
 &&& y_{uv} \geq 0, y'_{uvw} \geq 0, y''_{uvw} \geq 0 && \forall u, v, w: u \preceq w \preceq v.
 \end{aligned}$$

Finding a Feasible Solution for the Dual. When the graph G is a hypergrid $\mathcal{H}_{n,d}$, we can find a feasible solution of the dual, which gives a lower bound on the objective function of the primal. To do that, we perform the following three steps. First, we choose initial values \hat{y}_{uv} for the variables y_{uv} of the dual program and, in Lemma 3.1, give a lower bound on the resulting value of the objective function of the primal program. Second, we choose initial values \hat{y}'_{uvw} and \hat{y}''_{uvw} for variables y'_{uvw} and y''_{uvw} so that (2) holds. Finally, in Lemma 3.2, we give an upper bound on the left-hand side of (1) for all $u \preceq v$. Our bound is a constant larger than 1 and independent of n . We obtain a feasible solution to the dual by dividing the initial values of the variables (and, consequently, the value of the objective function) by this constant.

Step 1. For a vector $x = (x_1, \dots, x_d) \in [0, m - 1]^d$, let the *volume* $V(x)$ denote $\prod_{i \in [d]} (x_i + 1)$. This corresponds to the number of hypergrid points inside a d -dimensional box with corners u and v , where $v - u = x$. We start building a solution to the dual by setting $\hat{y}_{uv} = \frac{1}{V(v-u)}$ for all $u \preceq v$. This gives the value of the objective function of the dual program, according to the following lemma.

Lemma 3.1 $\sum_{u,v: u \preceq v} \hat{y}_{uv} > m^d (\ln m - 1)^d.$

Proof. Substituting $1/(V(v - u))$ for \hat{y}_{uv} , we get:

$$\begin{aligned}
 \sum_{u,v: u \preceq v} \hat{y}_{uv} &= \sum_{u,v: u \preceq v} \frac{1}{V(v - u)} = \sum_{l \in [m]^d} \prod_{i \in [d]} \frac{m - l_i + 1}{l_i} = \left(\sum_{l \in [m]} \frac{m - l + 1}{l} \right)^d \\
 &> ((m + 1) \ln(m + 1) - m)^d > m^d (\ln m - 1)^d. \quad \square
 \end{aligned}$$

Step 2. The values of \hat{y}'_{uvw} and \hat{y}''_{uvw} are set as follows to satisfy (2) tightly (without any slack):

$$\hat{y}'_{uvw} = \hat{y}_{uw} \frac{V(v-u)}{V(v-u) + V(w-v)}, \quad \hat{y}''_{uvw} = \hat{y}_{uv} - \hat{y}'_{uvw} = \hat{y}_{uw} \frac{V(w-v)}{V(v-u) + V(w-v)}.$$

Step 3. The initial values \hat{y}'_{uvw} and \hat{y}''_{uvw} do not necessarily satisfy (1). The following lemma, whose proof is deferred to the full version, gives an upper bound on the left-hand side of all constraints in (1).

Lemma 3.2 *For all $u \preceq v$, $\sum_{w: v \preceq w} \hat{y}'_{uvw} + \sum_{w: w \preceq u} \hat{y}''_{wuv} \leq (4\pi)^d$.*

Finally, we obtain a feasible solution by dividing initial values \hat{y}_{uv} , \hat{y}'_{uvw} and \hat{y}''_{uvw} by the upper bound $(4\pi)^d$ from Lemma 3.2. Then Lemma 3.1 gives the desired bound on the value of the objective function:

$$\sum_{u,v: u \preceq v} \frac{\hat{y}_{uv}}{(4\pi)^d} > m^d \left(\frac{\ln m - 1}{4\pi} \right)^d.$$

This concludes the proof of Theorem 1.1. □

4 Lower Bound for k -TC-Spanners for $k > 2$

In this section, we prove Theorem 1.3 that gives a lower bound on the size of Steiner k -TC-spanners of d -dimensional posets for $k > 2$ and $d \geq 2$.

Proof (of Theorem 1.3). Unlike in the previous section, the poset which attains the lower bound is constructed probabilistically, not explicitly.

We consider n -element posets G embedded in the hypergrid $\mathcal{H}_{n,d}$, where the partial order is given by the dominance order $x \preceq y$ on $\mathcal{H}_{n,d}$. The elements of G are points $p_1, p_2, \dots, p_n \in [n]^d$, where the first coordinate of each p_a is a . (By Fact 2.1, each d -dimensional poset with n elements can be embedded into $\mathcal{H}_{n,d}$, so that the first coordinates of all points are distinct.) Let \mathcal{G}_d be a distribution on such posets G , where the last $d - 1$ coordinates of each point p_a are chosen uniformly and independently from $[n]$.

Recall that $S_k(G)$ denotes the size of the sparsest Steiner k -TC-spanner of poset G . The following lemma gives a lower bound on the expected size of a Steiner k -TC-spanner of a poset drawn from \mathcal{G}_d .

Lemma 4.1 $\mathbb{E}_{G \leftarrow \mathcal{G}_d} [S_k(G)] = \Omega(n \log^{\lceil \frac{d-1}{k} \rceil} n)$ for all $k \geq 3$ and constant $d \geq 2$.

In this extended abstract, we only prove the special case of Lemma 4.1 for 2-dimensional posets (Lemma 4.2). The general case is deferred to the full version. Since Lemma 4.1 implies the existence of a poset G , for which every Steiner k -TC-spanner has $\Omega(n \log^{\lceil (d-1)/k \rceil} n)$ edges, Theorem 1.3 follows. □

The Case of $d = 2$. Next we prove a special case of Lemma 4.1 for 2-dimensional posets, which illustrates many ideas used in the proof of Lemma 4.1.

Lemma 4.2 $\mathbb{E}_{G \leftarrow \mathcal{G}_2} [S_k(G)] = \Omega(n \log n)$ for all $k \geq 3$ and $d = 2$.

Proof. We can assume that $\ell = \log n$ is an integer. To analyze the expected number of edges in a Steiner TC-spanner H of G , we consider ℓ partitions of $[n]^2$ into horizontal strips. We call strips *boxes* for compatibility with the case of general d .

Definition 4.1 (Box partition). For each $i \in [\ell]$, define sets of equal size that partition $[n]$ into 2^i intervals: the j th such set, for $j \in [2^i]$, is $I_j^i = [(j - 1)2^{\ell-i} + 1, j2^{\ell-i}]$. Given $i \in [\ell]$, and $j \in [2^i]$, the box $\mathbb{B}(i, j)$ is $[n] \times I_j^i$ and the box partition $\mathbb{BP}(i)$ is a partition of $[n]^2$ that contains boxes $\mathbb{B}(i, j)$ for all $j \in [2^i]$.

For each odd j , we group boxes $\mathbb{B}(i, j)$ and $\mathbb{B}(i, j + 1)$ into a *box-pair*. We call j the *index* of the box-pair and refer to $\mathbb{B}(i, j)$ and $\mathbb{B}(i, j + 1)$ as the *bottom* and the *top* box in the box-pair. Recall that a poset G consists of elements $p_1, p_2, \dots, p_n \in [n]^2$, where the first coordinate of each p_a is a . We analyze the expected number of edges in a Steiner TC-spanner H of G that cross from bottom to top boxes in all box-pairs. To do that, we identify pairs of poset elements (p_a, p_b) , called *jumps*, that force such edges to appear. By Lemma 2.3, we can assume that all Steiner vertices of H are embedded into $\mathcal{H}_{n,2}$. Therefore, if p_a is in the bottom box and p_b is in the top box of the same box-pair then H must contain an edge from the bottom to the top box. To ensure that we count such an edge just once, we consider only p_a and p_b for which no other point p_c with $c \in (a, b)$ is contained in this box pair. Next we define *jumps* formally. This concept is also illustrated in Figure 1.

Definition 4.2 (Jumps). Given a poset G , embedded into $\mathcal{H}_{n,2}$, and an index $i \in [\ell]$, a jump generated by the box partition $\mathbb{BP}(i)$ is a pair (p_a, p_b) of elements of G , such that for some odd $j \in [2^i]$, the following holds: $p_a \in \mathbb{B}(i, j)$, $p_b \in \mathbb{B}(i, j + 1)$, but $p_c \notin \mathbb{B}(i, j) \cup \mathbb{B}(i, j + 1)$ for all $c \in (a, b)$. The set of jumps generated by all partitions $\mathbb{BP}(i)$ for $i \in [\ell]$ is denoted by \mathcal{J} .

Next we establish that the number of jumps in a poset G is a lower bound on the number of edges in a Steiner TC-spanner of G (Claim 4.3) and bound the expected number of jumps from below (Claim 4.4).

Claim 4.3. Let G be a poset, embedded into $\mathcal{H}_{n,2}$, and $H = (V_H, E_H)$ be a Steiner k -TC-spanner of G . Then $|E_H| \geq |\mathcal{J}|$.

To prove the claim, we establish an injective mapping from \mathcal{J} to E_H . The proof is deferred to the full version.

Claim 4.4. When a poset G is drawn from the distribution \mathcal{G}_2 , the expected size of \mathcal{J} is at least $n(\ell - 1)/4$.




Fig. 1. Box partition $\mathbb{BP}(2)$ and jumps it generates

Proof. We first find the expected number of jumps generated by the partition $\mathbb{BP}(i)$ for a specific i . Let $\lambda_i(p_a)$ be the index j of the box-pair $\mathbb{B}(i, j) \cup \mathbb{B}(i, j+1)$ that contains p_a . Let $\rho_i(p_a)$ be 0 if p_a is in the bottom box of that box pair, and 1 otherwise. One can think of $\lambda_i(p_a)$ as the location of p_a , and of $\rho_i(p_a)$ as its relative position within a box-pair. Importantly, when G is drawn from \mathcal{G}_2 , that is, the second coordinates of points p_a for all $a \in [n]$ are chosen uniformly and independently from $[n]$, then random variables $\rho_i(p_a)$ are independent and uniform over $\{0, 1\}$ for all $a \in [n]$.

We group together points p_a that have equal values of $\lambda_i(p_a)$, and sort points within groups in increasing order of their first coordinate a . Since there are 2^{i-1} box-pairs, the number of groups is at most 2^{i-1} . Observe that random variables $\rho_i(p_a)$ within each group are uniform and independent because random variables $\lambda_i(p_a)$ and $\rho_i(p_a)$ are independent for all $a \in [n]$. Now, if we list $\rho_i(p_a)$ in the sorted order for all points in a particular group, we get a sequence of 0s and 1s. Two consecutive entries correspond to a jump iff they are 01. The last position in a group cannot correspond to the beginning of a jump. The number of positions that can correspond to the beginning of a jump in all groups is n minus the number of groups, which gives at least $n - 2^{i-1}$. For each such position, the probability that it starts a jump (i.e., the probability of 01) is $1/4$. Thus, the expected number of jumps generated by the partition $\mathbb{BP}(i)$ is at least $(n - 2^{i-1})/4$.

Summing over all $i \in [\ell]$, we get the expected number of jumps in all partitions: $(n\ell - \sum_{i=1}^{\ell} 2^{i-1})/4 > n(\ell - 1)/4 = \Omega(n \log n)$. \square

Claims 4.3 and 4.4 imply that, for a poset G drawn from \mathcal{G}_2 , the expected number of edges in a Steiner TC-spanner of G is $\Omega(n \log n)$, concluding the proof of Lemma 4.2. \square

References

1. Ailon, N., Chazelle, B., Comandur, S., Liu, D.: Property-preserving data reconstruction. *Algorithmica* 51(2), 160–182 (2008)
2. 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) (2009)
3. Awerbuch, B.: Communication-time trade-offs in network synchronization. In: *PODC*, pp. 272–276 (1985)
4. Bhattacharyya, A., Grigorescu, E., Jha, M., Jung, K., Raskhodnikova, S., Woodruff, D.P.: Lower bounds for local monotonicity reconstruction from transitive-closure spanners. In: Serna, M., Shaltiel, R., Jansen, K., Rolim, J. (eds.) *APPROX 2010, LNCS*, vol. 6302, pp. 448–461. Springer, Heidelberg (2010)
5. Bhattacharyya, A., Grigorescu, E., Jung, K., Raskhodnikova, S., Woodruff, D.P.: Transitive-closure spanners. In: Mathieu, C. (ed.) *SODA*, pp. 932–941. SIAM, Philadelphia (2009)
6. Chandra, A.K., Fortune, S., Lipton, R.J.: Lower bounds for constant depth circuits for prefix problems. In: Díaz, J. (ed.) *ICALP 1983. LNCS*, vol. 154, pp. 109–117. Springer, Heidelberg (1983)

7. Chandra, A.K., Fortune, S., Lipton, R.J.: Unbounded fan-in circuits and associative functions. *J. Comput. Syst. Sci.* 30(2), 222–234 (1985)
8. De Santis, A., Ferrara, A.L., Masucci, B.: New constructions for provably-secure time-bound hierarchical key assignment schemes. *Theor. Comput. Sci.* 407(1-3), 213–230 (2008)
9. Dushnik, B., Miller, E.: Concerning similarity transformations of linearly ordered sets. *Bulletin Amer. Math. Soc.* 46, 322–326 (1940)
10. Dushnik, B., Miller, E.W.: Partially ordered sets. *Amer. J. Math.* 63, 600–610 (1941)
11. Jha, M., Raskhodnikova, S.: Testing and reconstruction of Lipschitz functions with applications to data privacy. *Electronic Colloquium on Computational Complexity (ECCC)* TR11-057 (2011)
12. Peleg, D., Schäffer, A.A.: Graph spanners. *J. Graph Theory* 13(1), 99–116 (1989)
13. Raskhodnikova, S.: Transitive-closure spanners: A survey. In: Goldreich, O. (ed.) *Property Testing*. LNCS, vol. 6390, pp. 167–196. Springer, Heidelberg (2010)
14. Saks, M.E., Seshadhri, C.: Local monotonicity reconstruction. *SIAM J. Comput.* 39(7), 2897–2926 (2010)
15. Yannakakis, M.: The complexity of the partial order dimension problem. *SIAM Journal on Matrix Analysis and Applications* 3(3), 351–358 (1982)