

Steiner Transitive-Closure Spanners of Low-Dimensional Posets*

Piotr Berman[†] Arnab Bhattacharyya[‡] Elena Grigorescu[§] Sofya Raskhodnikova[†]
David P. Woodruff[¶] Grigory Yaroslavtsev[†]

April 10, 2013

Abstract

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 . In some applications, the shortcut paths added to the graph in order to obtain small diameter can use Steiner vertices, that is, vertices not in the original graph G . The resulting spanner is called a *Steiner transitive-closure spanner* (Steiner TC-spanner).

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. In these applications, the goal is to find a sparsest Steiner k -TC-spanner of a poset G for a given k and G . The focus of this paper is 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.

We present a nearly tight lower bound on the size of Steiner 2-TC-spanners of d -dimensional directed hypergrids. It implies better lower bounds on the complexity of local reconstructors of monotone functions and functions with small Lipschitz constant. The lower bound is derived from an explicit dual solution to a linear programming relaxation of the Steiner 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 lower bound on the size of Steiner k -TC-spanners of d -dimensional posets. It shows that the best-known construction, due to De Santis *et al.*, cannot be improved significantly.

1 Introduction

Graph spanners were introduced in the context of distributed computing by Awerbuch [6] and Peleg and Schäffer [23], and since then have found numerous applications. Our focus is on transitive-closure spanners, introduced explicitly in [9], but studied prior to that in many different contexts [12, 11, 31, 3, 13, 26, 10, 27, 28, 15, 19, 4].

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 that: (1) E_H is a subset of the edges in the transitive

*A preliminary version of this paper appeared in ICALP 2011 [7].

[†]Pennsylvania State University, USA. {berman, sofya, grigory}@cse.psu.edu. 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.

[‡]Massachusetts Institute of Technology, USA. abhatt@mit.edu

[§]Georgia Institute of Technology, USA. elena@cc.gatech.edu. Supported in part by NSF award CCR-0829672 and NSF award 1019343 to the Computing Research Association for the Computing Innovation Fellowship Program.


[¶]IBM Almaden Research Center, USA. dpwoodru@us.ibm.com.

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 [24] for a survey. In some applications of TC-spanners, in particular, to access control hierarchies [4, 14], 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 ([24], 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 [4, 14] and property reconstruction [8, 21], 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.* [4], followed by De Santis *et al.* [14], 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 [8] 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. The focus on stretch $k = 2$ is motivated by this application.

Several classes of posets, in addition to low-dimensional hypergrids, are known to have small dimension. For example, if a *planar poset*, that is, a poset with a planar Hasse diagram, has *both* a minimum *and* a maximum element, it has dimension at most 2. A planar poset with *either* a minimum or a maximum element has dimension at most 3 [29]. (In general, however, a planar poset can have an arbitrary dimension [22]). One can also bound the dimension in terms of cardinality of the poset: every poset of cardinality $n \geq 4$ has dimension at most $n/2$ [20]. Also, if every element in an n -element poset has at most u points above it, then its dimension is at most $2(u+1) \log n + 1$ [18]. Thus, posets with low dimension occur quite naturally in a variety of settings.

1.1 Our Results

1.1.1 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 [8], 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. *Every (Steiner) 2-TC-spanner of $\mathcal{H}_{m,d}$ has $\Omega\left(\frac{m^d (\ln m - 1)^d}{(4\pi)^d}\right)$ 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 [9] 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 [8] 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 [8], 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.

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 [4, 14]. We note that the only poset of dimension 1 is the directed line $\mathcal{H}_{n,1}$.

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

Stretch k	Prior bounds on $S_k(G)$	Stretch k	Our bounds on $S_k(G)$
$2d - 1$	$O(n^2)$ [4]	2	$O(n \log^d n)$
$2d - 2 + t$ for $t \geq 2$	$O(n(\log^{d-1} n)\lambda_t(n))$ [4]		
$2d + O(\log^* n)$	$O(n \log^{d-1} n)$ [4]	≥ 3	$\Omega\left(n \left(\frac{\log n}{cd}\right)^d\right)$ for a fixed $c > 0$
3	$O(n \log^{d-1} n \log \log n)$ for fixed d [14]		
			$\Omega(n \log^{\lceil (d-1)/k \rceil} n)$ for fixed d

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

TC-spanners of directed lines were discovered under many different guises. They were studied implicitly in [31, 12, 11, 3, 13, 15, 5] and explicitly in [10, 28]. Chandra, Fortune and Lipton [12, 11] implicitly 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 function².

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 for which all Steiner k -TC-spanners have the specified number of edges.

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 [30]. Also, as mentioned previously, in the application to access control hierarchies, such an embedding is usually given.) The Steiner vertices used in our construction for d -dimensional posets are necessary to obtain sparse TC-spanners, as manifested by the example presented after the proof of Theorem 2.2.

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 n for m^d and $(\ln n)/d$ for $\ln m$ in the theorem statement. This works for all n larger than some constant to the power d and 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.*

The restriction on n in the corollary is not an artifact of the proof, since for $d = \omega(\log n)$, the expression in the lower bound is $\omega(n^2)$. Thus, the corollary is false this range of parameters because the transitive closure of any graph has $O(n^2)$ 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 $n \log^{\Omega(d)} n$.

²The Ackermann function [1] is defined by: $A(1, j) = 2^j$, $A(i + 1, 0) = A(i, 1)$, $A(i + 1, j + 1) = A(i, 2^{A(i+1, j)})$. The inverse Ackermann function is $\alpha(n) = \min\{i : A(i, 1) \geq n\}$ and the i^{th} -row inverse is $\lambda_i(n) = \min\{j : A(i, j) \geq n\}$. Specifically, $\lambda_2(n) = \Theta(\log n)$, $\lambda_3(n) = \Theta(\log \log n)$ and $\lambda_4(n) = \Theta(\log^* n)$.

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 $\Omega(n \log^{\lceil (d-1)/k \rceil} n)$ 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 [8]. Instead, we construct an n -element poset embedded in $\mathcal{H}_{n,d}$ using the following randomized procedure: 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 [24]. We focus on two of them: property reconstruction, described in [8, 21], and key management for access control hierarchies, described in [4, 9, 14].

Property Reconstruction.

Property-preserving data reconstruction was introduced by Ailon, Chazelle, Comandur and Liu [2]. In this model, a reconstruction algorithm, called a *filter*, sits between a *client* and a *dataset*. A dataset is viewed as a function $f : \mathcal{D} \rightarrow \mathcal{R}$. A 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 a pre-specified *structural* property (e.g., be monotone or have a low Lipschitz constant) and differ from f as little as possible. Extending this notion, Saks and Seshadhri [25] defined *local* reconstruction. A filter is *local* if the output function g does not depend on the order of the queries. Local filters can be used for distributed computations [25] and in private data analysis [21]. A filter is *nonadaptive* if its lookups do not depend on the answers to previous lookups.

Our results on TC-spanners are relevant to reconstruction of two properties of functions: monotonicity and having a low Lipschitz constant. Reconstruction of monotone functions was considered in [2, 8, 25]. A function $f : [m]^d \rightarrow \mathbb{R}$ is called *monotone* if $f(x) \leq f(y)$ for all $(x, y) \in E(\mathcal{H}_{m,d})$. Reconstruction of functions with low Lipschitz constant was studied in [21]. A function $f : [m]^d \rightarrow \mathbb{R}$ has Lipschitz constant c if $|f(x) - f(y)| \leq c \cdot |x - y|_1$. In [8], 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 [21], 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 optimal for any constant d .

Key Management for Access Control Hierarchies.

Atallah *et al.* [4] 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 to enforcing the access hierarchy is to use a key management scheme of the following form [4, 14]. 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 $d_G(u, v)$. 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 . 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 our upper bound, this time is $O(d)$, which for small d (e.g., constant) is likely to be much less than $2g(n)$, 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 expensive in practice and in theory.³

2 Definitions and Observations

For integers $j \geq i$, 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 from a poset G to a poset G' is called an *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 $x' \preceq_{G'} y'$.

Definition 2.1 (Poset dimension, [17]). *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 $\mathcal{H}_{n,d}$.*

Dushnik and Miller [16] proved that 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 $O(dn \log n)$.*

Proof. Let G be a d -dimensional poset G with n elements. By Definition 2.1, it can be embedded into a hypergrid $\mathcal{H}_{m,d}$. If $m > n$, we can sort the elements by coordinates in each dimension and reassign the coordinate values to be in $[n]$, without changing their relative order, in $O(dn \log n)$ time. Then we can compute a linear extension L of G by merging d lists of points, each sorted by a different coordinate, in $O(dn)$ time. Finally, for each dimension $i \in [d]$, in $O(n)$ time, we can go through the list of elements sorted

³Any hash function which is secure against $\text{poly}(n)$ -time adversaries requires $g(n) \geq \text{polylog } n$ evaluation time under existing number-theoretic assumptions.

by the i th coordinate, and reassign the i th coordinates, so that all of them are distinct, resolving conflicts according to the order of L .

It remains to show that this transformation, call it f , does not change the partial order of G . Consider a pair of elements $x \prec y$ in G . If $x_i < y_i$ for some $i \in [d]$ then $f(x)_i < f(y)_i$ because we did not change the order of distinct coordinates. If $x_i = y_i$ for some $i \in [d]$ then $f(x)_i < f(y)_i$ since x precedes y in L . Therefore, $f(x) \prec f(y)$. Finally, consider incomparable elements x and y of G . Since they are incomparable, $x_i < y_i$ while $x_j > y_j$ for some $i, j \in [d]$. Then $f(x)_i < f(y)_i$ while $f(x)_j > f(y)_j$, and consequently, $f(x)$ and $f(y)$ are incomparable, as required. \square

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 $\mathcal{H}_{n,d}$, graph 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 $\text{lcp}(t_1, t_2) = p_i(t_1)$, 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 $(\text{lcp}(x_1, y_1), \dots, \text{lcp}(x_d, y_d))$. By construction, both edges (x, z) and (z, y) are in E_H . Point z can be found in $O(d)$ time, since $\text{lcp}(x_i, y_i)$ can be computed in $O(1)$ time, assuming $O(1)$ time bit operations on coordinates. \square

Note that the Steiner vertices used in our construction 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-spanner 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 Theorem 2.2, 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.

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.

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 $\mathcal{H}_{m,d}$.*

Proof. For each $s \in V_H - V$, we define $Prev(s) = \{x \in V : x \prec s\}$. If $Prev(s) = \emptyset$ then V_H is not minimal because H remains a Steiner k -TC-spanner of G when s is removed. We map each Steiner vertex s to $r(s)$, the replacement of s in $[m]^d$, whose i th coordinates for all $i \in [d]$ are $\max_{x \in Prev(s)} x_i$.

Consider an edge (x, y) in G . If $x, y \in V$ our embedding does not alter that edge. If $x \in V, y \in V_H - V$ then $x \in Prev(y)$ and $x \prec r(y)$ by the definition of r . If $x, y \in V_H - V$ then $Prev(x) \subseteq Prev(y)$ and the monotonicity of $\max(S)$ for sets implies $r(x) \preceq r(y)$. Finally, if $x \in V_H - V$ and $y \in V$ then for each $z \in Prev(x)$ and each $i \in [d]$, we have $z_i \leq y_i$ because $z \prec x \prec y$, and this implies $r(x) \preceq y$. \square

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 [9] 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 . (We need variables x_{vv} for notational convenience in the last part of the proof.) 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:

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

$$\begin{aligned} & y_{uv} - y'_{uvw} - y''_{uvw} \leq 0 && \forall u, v, w: u \preceq w \preceq v; \\ & y_{uv} \geq 0, y'_{uvw} \geq 0, y''_{uvw} \geq 0 && \forall u, v, w: u \preceq w \preceq v. \end{aligned} \quad (2)$$

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:

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

□

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

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

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

Proof. Below we denote $v - u$ by $x^0 = (x_1^0, \dots, x_d^0)$, a d -dimensional vector of ones $(1, \dots, 1)$ by $\vec{1}$ and $\prod_{i \in [d]} dx_i$ by dx .

$$\begin{aligned}
& \sum_{w: v \preceq w} \hat{y}'_{uvw} + \sum_{w: w \preceq u} \hat{y}''_{wuv} \\
&= \sum_{w: v \preceq w} \hat{y}_{uw} \frac{V(v-u)}{V(v-u) + V(w-v)} \\
&+ \sum_{w: w \preceq u} \hat{y}_{wv} \frac{V(v-u)}{V(u-w) + V(v-u)} \\
&= \sum_{w: v \preceq w} \frac{V(v-u)}{V(w-u)(V(v-u) + V(w-v))} \\
&+ \sum_{w: w \preceq u} \frac{V(v-u)}{V(v-w)(V(u-w) + V(v-u))} \\
&< 2 \sum_{x \in [0, m]^d} \frac{V(x^0)}{V(x^0 + x)(V(x^0) + V(x))} \\
&\leq 2^{2d+1} \sum_{x \in [1, m+1]^d} \frac{V(x^0)}{V(x^0 + x)(V(x^0) + V(x))} \tag{3}
\end{aligned}$$

$$< 2^{2d+1} \int_{\mathbb{R}_+^d} \frac{V(x^0) dx}{V(x^0 + x)(V(x^0) + V(x))} \tag{4}$$

$$= 2^{2d+1} \int_{\mathbb{R}_+^d} \frac{V^2(x^0) dt}{V(t)V(x_0)(V(x^0) + \prod_i (t_i(x_i^0 + 1) + 1))} \tag{5}$$

$$< 2^{2d+1} \int_{\mathbb{R}_+^d} \frac{V(x^0) dt}{V(t)(V(x^0) + \prod_i t_i(x_i^0 + 1))}$$

$$= 2^{2d+1} \int_{\mathbb{R}_+^d} \frac{dt}{V(t)(\vec{1} + V(t-1))}.$$

The first two equalities above are obtained by plugging in values of \hat{y}' and \hat{y}'' from Steps 1 and 2 with appropriate indices. The first inequality is obtained by extending each sum to the whole subgrid. Here (3) holds because $\frac{1}{V(u)} \leq \frac{2^d}{V(u+1)}$ for all u , such that $u_i \geq 0$. In (4), the sum can be bounded from above by the integral because the summand is monotone in all variables. To get (5), we substitute x with t such that $x_i = t_i(x_i^0 + 1)$. Then $V(x_0 + x) = V(t)V(x_0)$, and $dx = V(x_0)dt$. To obtain the last inequality, we substitute $V(x^0)$ for $\prod_i (x_i^0 + 1)$.

Proposition 3.3. Let $I_d = \int_{\mathbb{R}_+^d} \frac{dt}{V(t)(\vec{1} + V(t-1))}$. Then $I_d \leq \frac{\pi^d}{2}$ for all d .

Proof. To bound the integral I_d , we first make a substitution $x_i = \frac{1-t_i}{1+t_i}$:

$$I_d = \int_{[-1, \dots, 1]^d} \frac{dx}{\prod_{1 \leq i \leq d} (1+x_i) + \prod_{1 \leq i \leq d} (1-x_i)}.$$

Then we bound the denominator using the inequality $a + b \geq 2\sqrt{ab}$ and get

$$I_d \leq \int_{[-1 \dots 1]^d} \frac{dx}{2 \sqrt{\prod_{1 \leq i \leq d} (1 + x_i) \times \prod_{1 \leq i \leq d} (1 - x_i)}} = \frac{J^d}{2},$$

where J denotes the following integral:

$$J = \int_{-1}^1 \frac{dx}{\sqrt{1 - x^2}} = \pi.$$

Therefore, $I_d \leq \frac{\pi^d}{2}$, as claimed. □

Lemma 3.2 follows from Proposition 3.3. □

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 Our Lower Bound for k -TC-spanners of d -dimensional Posets 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 $d \geq 2$.

To simplify the presentation, we first prove the special case of Lemma 4.1 for 2-dimensional posets in Section 4.1. The general case is proved in Section 4.2. 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. □

4.1 The Case of $d = 2$

This section proves a special case of Lemma 4.1 for 2-dimensional posets, which illustrates many of the ideas used in the proof of the general lemma. In both proofs, we assume that $\ell = \log n$ is an integer.

Lemma 4.2 (Special case of Lemma 4.1). $\mathbb{E}_{G \leftarrow \mathcal{G}_2} [S_k(G)] = \Omega(n \log n)$ for all $k \geq 3$ and $d = 2$.

Proof. 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 (Proposition 4.3) and bound the expected number of jumps from below (Proposition 4.4).

Proposition 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 $|E_H| \geq |\mathcal{J}|$.

Proof. To prove the statement, we establish an injective mapping from \mathcal{J} to E_H . By Lemma 2.3, we can assume that all Steiner vertices of H are embedded into $\mathcal{H}_{n,2}$. Given a jump (p_a, p_b) , let j be the index of the box-pair containing p_a in the bottom box and p_b in the top box. We define $e(a, b) \in E_H$ by following an arbitrary path from p_a to p_b in H . This path is contained in the box-pair $\mathbb{B}(i, j) \cup \mathbb{B}(i, j+1)$. We define $e(a, b)$ as the edge on that path that starts in $\mathbb{B}(i, j)$ and ends in $\mathbb{B}(i, j+1)$.

It remains to show that the mapping $e(a, b)$ is injective. Consider an edge (u, v) of H with $u = (u_1, u_2)$ and $v = (v_1, v_2)$. Observe that there is a unique box-pair $\mathbb{B}(i, j) \cup \mathbb{B}(i, j+1)$ such that $v \in \mathbb{B}(i, j)$ and $u \in \mathbb{B}(i, j+1)$. (Indices i and j can be determined by finding the number of the form $j2^{\ell-i}$ in the interval $[u_2, v_2 - 1]$, such that $\ell - i$ is maximized.) At most one jump (a, b) satisfies $p_a \in \mathbb{B}(i, j)$, $p_b \in \mathbb{B}(i, j+1)$ and $a \leq u_1 \leq v_1 \leq b$, since the intervals $[a, b]$ are disjoint for all jumps in a box pair. \square




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

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

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 . This is well defined since box-pairs with respect to $\mathbb{BP}(i)$ partition $[n]^2$. 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

Propositions 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

4.2 The Case of Constant d

This section proves Lemma 4.1, a generalization of Lemma 4.2 and the main building block in the proof of Theorem 1.3.

Proof of Lemma 4.1. Generalizing the proof for $d = 2$ to arbitrary constant d , we consider ℓ^{d-1} partitions of $[n]^d$ into boxes, where $\ell = \log n$. As before, we assume ℓ is an integer. In this proof, let $\ell' = \lfloor \ell/(d-1) \rfloor$ and $d' = \lceil (d-1)/k \rceil$.

Definition 4.3 (Box partition). *Given vectors $\vec{i} = (i_1, \dots, i_{d-1}) \in [\ell']^{d-1}$ and $\vec{j} = (j_1, \dots, j_{d-1}) \in [2^{i_1}] \times \dots \times [2^{i_{d-1}}]$, the box $\mathbb{B}(\vec{i}, \vec{j})$ is $[n] \times I_{j_1}^{i_1} \times \dots \times I_{j_{d-1}}^{i_{d-1}}$, and the box partition $\mathbb{BP}(\vec{i})$ is a partition of $[n]^d$ that contains boxes $\mathbb{B}(\vec{i}, \vec{j})$ for all eligible \vec{j} .*

Next we generalize the definition of the set of jumps \mathcal{J} . We denote $(d-1)$ -dimensional vectors $(0, \dots, 0)$ and $(1, \dots, 1)$ by $\vec{0}$ and $\vec{1}$, respectively. We say that a vector \vec{j} is *odd* if all of its coordinates are odd. Now we form box-pairs from boxes $\mathbb{B}(\vec{i}, \vec{j})$ and $\mathbb{B}(\vec{i}, \vec{j} + \vec{1})$ for odd vectors \vec{j} . Analogously to the 2-dimensional case, we call $\mathbb{B}(\vec{i}, \vec{j})$ the *bottom box* and $\mathbb{B}(\vec{i}, \vec{j} + \vec{1})$ the *top box* in a box-pair.

Definition 4.4 (Jumps). *Given a poset G , embedded into $\mathcal{H}_{n,d}$, and an index vector $\vec{i} \in [\ell']^{d-1}$, a jump generated by the box partition $\mathbb{BP}(\vec{i})$ is a pair (p_a, p_b) of elements of G , such that for some odd vector \vec{j} , the following holds: $p_a \in \mathbb{B}(\vec{i}, \vec{j})$, $p_b \in \mathbb{B}(\vec{i}, \vec{j} + \vec{1})$, but $p_c \notin \mathbb{B}(\vec{i}, \vec{j}) \cup \mathbb{B}(\vec{i}, \vec{j} + \vec{1})$ for all $c \in (a, b)$. The set of jumps generated by all partitions $\mathbb{BP}(\vec{i})$ for $\vec{i} \in [\ell']^{d-1}$ is denoted by \mathcal{J} .*

Next we generalize the definitions of location $\lambda_i(p_a)$ and relative position $\rho_i(p_a)$. Unlike in the 2-dimensional case, now some boxes (and, consequently, some points) do not belong to box-pairs. For each odd \vec{j} , we group boxes $\mathbb{B}(\vec{i}, \vec{j} + \vec{\alpha})$ for all $\vec{\alpha} \in \{0, 1\}^{d-1}$ into a *megabox*. We call \vec{j} the *index vector* of the

megabox, and refer to α as the *relative position* of a box in the megabox. Observe that megaboxes with respect to $\mathbb{BP}(\vec{i})$ partition $[n]^d$. Given \vec{i} , let $\lambda_{\vec{i}}(p_a)$ be the index vector \vec{j} of the megabox of p_a with respect to $\mathbb{BP}(\vec{i})$, and let $\rho_{\vec{i}}(p_a)$ be the relative position vector $\vec{\alpha}$ of the box of p_a in the megabox. In other words, to obtain $\lambda_{\vec{i}}(p_a)$, we take the index \vec{j} of the box $\mathbb{B}(\vec{i}, \vec{j})$ containing p_a , and round its coordinates down to the nearest odd numbers. Then $\rho_{\vec{i}}(p_a) = \vec{j} - \lambda_{\vec{i}}(p_a)$, where \vec{j} is the index of the box $\mathbb{B}(\vec{i}, \vec{j})$ containing p_a .

Proposition 4.5. *Let G be a poset, embedded into $\mathcal{H}_{n,d}$, and $H = (V_H, E_H)$ be a Steiner k -TC-spanner of G . Then $|E_H| = \Omega(|\mathcal{J}|/\ell^{d-1-d'})$.*

Proof. To prove the statement, we establish a mapping from \mathcal{J} to E_H that takes $O(\ell^{d-1-d'})$ jumps to one edge. By Lemma 2.3, we can assume that all Steiner vertices of H are embedded into $\mathcal{H}_{n,d}$.

First, we describe how to map a jump (p_a, p_b) to an edge $e(a, b) \in E_H$. Each such jump is generated by a box partition $\mathbb{BP}(\vec{i})$ for some \vec{i} . We follow an arbitrary path of length at most k in H from p_a to p_b , say, $(p_a = u_0, \dots, u_k = p_b)$, and let $e(a, b)$ be an edge (u_{c-1}, u_c) with the maximum (over all $c \in [k]$) Hamming distance between $\rho_{\vec{i}}(u_{c-1})$ and $\rho_{\vec{i}}(u_c)$. Note that the maximum distance is at least d' because $\rho_{\vec{i}}(u_0) = \vec{0}$ and $\rho_{\vec{i}}(u_k) = \vec{1}$. That is, for $(u, v) = e(a, b)$, the difference $\rho_{\vec{i}}(v) - \rho_{\vec{i}}(u)$ is a vector in $\{0, 1\}^{d-1}$ with at least d' ones. In addition, the edge $e(a, b)$ belongs to the megabox of p_a and p_b .

Now we count the jumps (p_a, p_b) mapped to a specific edge (u, v) . First, we find all such jumps generated by a single box partition $\mathbb{BP}(\vec{i})$. Observe that, for such a jump, p_a and p_b belong to the same megabox as u and v , i.e., $\lambda_{\vec{i}}(u)$. Moreover, interval $[a, b]$ contains $[u_1, v_1]$. Since intervals $[a, b]$ are disjoint for all jumps in a megabox, this uniquely determines $[a, b]$. Hence, there is at most one such jump.

It remains to count box partitions $\mathbb{BP}(\vec{i})$ which can generate a jump mapped to a specific edge (u, v) . Recall that $\rho_{\vec{i}}(v) - \rho_{\vec{i}}(u)$ must be a vector in $\{0, 1\}^{d-1}$ with at least d' ones. There are less than 2^{d-1} such vectors. Consider one of these vectors, say, $\vec{\gamma}$. If for some $t \in [d-1]$, $\gamma_t = 1$ then i_t is uniquely determined by the largest power of 2 that divides a number in $[u_t, v_t - 1]$. When $\gamma_t = 0$, there are at most ℓ' possible values of i_t because $\vec{i} \in [\ell']^{d-1}$. Since d is a constant, there are at most $2^{d-1}(\ell')^{d-1-d'} = O(\ell^{d-1-d'})$ possible vectors \vec{i} , such that $\mathbb{BP}(\vec{i})$ could have generated a jump (p_a, p_b) .

Therefore, $O(\ell^{d-1-d'})$ jumps map to the same edge of E_H and, consequently, $|E_H| = \Omega(|\mathcal{J}|/\ell^{d-1-d'})$. \square

Proposition 4.6. *When a poset G is drawn from the distribution \mathcal{G}_d , the expected size of \mathcal{J} is $\Omega(\ell^{d-1}n)$.*

Proof. We first analyze the expected number of jumps generated by the partition $\mathbb{BP}(i)$ for a specific i . Under the distribution \mathcal{G}_d , the values $\rho_{\vec{i}}(p_a)$ are independent and uniform over $\{0, 1\}^{d-1}$ for all $a \in [n]$. Let P be the set of elements p_a in all the box-pairs, i.e., elements with $\rho_{\vec{i}}(p_a)$ equal to $\vec{0}$ and $\vec{1}$. The expected size of P is $n/2^{d-2}$. We group together elements p_a of P that have equal values of $\lambda_{\vec{i}}(p_a)$, and sort elements within groups in increasing order of their first coordinate a .

Observe that random variables $\rho_{\vec{i}}(p_a)$ within each group are uniform and independent because random variables $\lambda_{\vec{i}}(p_a)$ and $\rho_{\vec{i}}(p_a)$ are independent for all $a \in [n]$. Now, if we list $\rho_{\vec{i}}(p_a)$ in the sorted order for all elements in a particular group, we get a sequence of $\vec{0}$ s and $\vec{1}$ s. Two consecutive entries correspond to a jump iff they are $\vec{0}\vec{1}$. The last position in a group cannot correspond to the beginning of a jump. The expected number of positions that can correspond to the beginning of a jump in all groups is $n/2^{d-2}$ minus the expected number of groups. Let $m(\vec{i})$ denote the number of megaboxes with respect to a box partition $\mathbb{BP}(\vec{i})$. The number of groups is at most $m(\vec{i})$. On every position in the reordered sequence that is not the final position in its group, the expected number of jumps started is $1/4$, so the expected number of jumps is at least $(n/2^{d-2} - m(\vec{i}))/4 = n/2^d - m(\vec{i})/4$.

The number of megaboxes in all box partitions is

$$\sum_{\vec{i} \in [\ell']^{d-1}} m(\vec{i}) = \sum_{\vec{i} \in [\ell']^{d-1}} \prod_{t=1}^{d-1} 2^{i_t-1} = \left(\sum_{i_1=1}^{\ell'} 2^{i_1-1} \right)^{d-1} < 2^{\ell'(d-1)} \leq 2^\ell = n.$$

Therefore, the expected number of jumps generated by all box partitions is at least

$$(\ell')^{d-1} n / 2^d - \frac{1}{4} \sum_{\vec{i} \in [\ell']^{d-1}} m(\vec{i}) \geq (\ell')^{d-1} n / 2^d - n/4 = \Omega(\ell^{d-1} n).$$

The last equality holds because d is constant. □

By linearity of expectation, Propositions 4.5 and 4.6 imply that the expected number of edges in a Steiner TC-spanner H of G under the distribution \mathcal{G}_d is

$$\Omega\left(\mathbb{E}_{G \leftarrow \mathcal{G}_d}(|\mathcal{J}|) / \ell^{d-1-d'}\right) = \Omega(\ell^{d-1} n / \ell^{d-1-d'}) = \Omega(n \ell^{d'}) = \Omega(n \log^{\lceil (d-1)/k \rceil} n).$$

This concludes the proof of Lemma 4.1. □

References

- [1] W. Ackermann. Zum hilbertschen Aufbau der reellen Zahlen. *Math. Ann.*, 99:118–133, 1928.
- [2] N. Ailon, B. Chazelle, S. Comandur, and D. Liu. Property-preserving data reconstruction. *Algorithmica*, 51(2):160–182, 2008.
- [3] N. Alon and B. Schieber. Optimal preprocessing for answering on-line product queries. Technical Report 71/87, Tel-Aviv University, 1987.
- [4] M. J. Atallah, M. Blanton, N. Fazio, and K. B. Frikken. Dynamic and efficient key management for access hierarchies. *ACM Trans. Inf. Syst. Secur.*, 12(3), 2009.
- [5] M. J. Atallah, K. B. Frikken, and M. Blanton. Dynamic and efficient key management for access hierarchies. In V. Atluri, C. Meadows, and A. Juels, editors, *ACM Conference on Computer and Communications Security*, pages 190–202. ACM, 2005.
- [6] B. Awerbuch. Communication-time trade-offs in network synchronization. In *PODC*, pages 272–276, 1985.
- [7] P. Berman, A. Bhattacharyya, E. Grigorescu, S. Raskhodnikova, D. P. Woodruff, and G. Yaroslavtsev. Steiner transitive-closure spanners of low-dimensional posets. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (1)*, volume 6755 of *Lecture Notes in Computer Science*, pages 760–772. Springer, 2011.
- [8] A. Bhattacharyya, E. Grigorescu, M. Jha, K. Jung, S. Raskhodnikova, and D. P. Woodruff. Lower bounds for local monotonicity reconstruction from transitive-closure spanners. *SIAM J. Discrete Math.*, 26(2):618–646, 2012.
- [9] A. Bhattacharyya, E. Grigorescu, K. Jung, S. Raskhodnikova, and D. P. Woodruff. Transitive-closure spanners. *SIAM J. Comput.*, 41(6):1380–1425, 2012.
- [10] H. L. Bodlaender, G. Tel, and N. Santoro. Trade-offs in non-reversing diameter. *Nordic J. of Computing*, 1(1):111–134, 1994.
- [11] A. K. Chandra, S. Fortune, and R. J. Lipton. Lower bounds for constant depth circuits for prefix problems. In J. Díaz, editor, *ICALP*, volume 154 of *LNCS*, pages 109–117. Springer, 1983.
- [12] A. K. Chandra, S. Fortune, and R. J. Lipton. Unbounded fan-in circuits and associative functions. *J. Comput. Syst. Sci.*, 30(2):222–234, 1985.

- [13] B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.
- [14] A. De Santis, A. L. Ferrara, and B. Masucci. New constructions for provably-secure time-bound hierarchical key assignment schemes. *Theor. Comput. Sci.*, 407(1-3):213–230, 2008.
- [15] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In *RANDOM*, pages 97–108, 1999.
- [16] B. Dushnik and E. Miller. Concerning similarity transformations of linearly ordered sets. *Bulletin Amer. Math. Soc.*, 46:322–326, 1940.
- [17] B. Dushnik and E. W. Miller. Partially ordered sets. *Amer. J. Math.*, 63:600–610, 1941.
- [18] Z. Füredi and J. Kahn. On the dimension of ordered sets of bounded degree. *Order*, 3:15–20, 1986.
- [19] W. Hesse. Directed graphs requiring large numbers of shortcuts. In *SODA*, pages 665–669. ACM/SIAM, 2003.
- [20] T. Hiraguchi. On the dimension of partially ordered sets. *Science Reports Kanazawa University*, 1:77–94, 1951.
- [21] M. Jha and S. Raskhodnikova. Testing and reconstruction of Lipschitz functions with applications to data privacy. In R. Ostrovsky, editor, *FOCS*, pages 433–442. IEEE, 2011.
- [22] D. Kelly. On the dimension of partially ordered sets. *Discrete Mathematics*, 35(1-3):135–156, 1981.
- [23] D. Peleg and A. A. Schäffer. Graph spanners. *J. Graph Theory*, 13(1):99–116, 1989.
- [24] S. Raskhodnikova. Transitive-closure spanners: A survey. In O. Goldreich, editor, *Property Testing*, volume 6390 of *LNCS*, pages 167–196. Springer, 2010.
- [25] M. E. Saks and C. Seshadhri. Local monotonicity reconstruction. *SIAM J. Comput.*, 39(7):2897–2926, 2010.
- [26] M. Thorup. On shortcutting digraphs. In E. W. Mayr, editor, *WG*, volume 657 of *LNCS*, pages 205–211. Springer, 1992.
- [27] M. Thorup. Shortcutting planar digraphs. *Combinatorics, Probability & Computing*, 4:287–315, 1995.
- [28] M. Thorup. Parallel shortcutting of rooted trees. *J. Algorithms*, 23(1):139–159, 1997.
- [29] W. T. Trotter and J. Moore. The dimension of planar posets. *Journal of Combinatorial Theory, Series B*, 22:54–67, 1977.
- [30] M. Yannakakis. The complexity of the partial order dimension problem. *SIAM Journal on Matrix Analysis and Applications*, 3(3):351–358, 1982.
- [31] A. C.-C. Yao. Space-time tradeoff for answering range queries (extended abstract). In H. R. Lewis, B. B. Simons, W. A. Burkhard, and L. H. Landweber, editors, *STOC*, pages 128–136. ACM, 1982.