Electronic Colloquium on Computational Complexity, Report No. 46 (2009)

Transitive-Closure Spanners of the Hypercube and the Hypergrid

Arnab Bhattacharyya* Elena Grigorescu* Kyomin Jung* Sofya Raskhodnikova† David P. Woodruff‡

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. Transitive-closure spanners were introduced in [7] as a common abstraction for applications in access control, property testing and data structures.

In this work we study the number of edges in the sparsest 2-TC-spanners for the directed hypercube $\{0, 1\}^d$ and hypergrid $\{1, 2, \ldots, m\}^d$ with the usual partial order, $\preceq$, defined by: $x_1 \ldots x_d \preceq y_1 \ldots y_d$ if and only if $x_i \leq y_i$ for all $i \in \{1, ..., d\}$. We show that the number of edges in the sparsest 2-TC-spanner of the hypercube is $2^{cd+\Theta(\log d)}$, where $c \approx 1.1620$. We also present upper and lower bounds on the size of the sparsest 2-TC-spanner of the directed hypergrid. Our first pair of upper and lower bounds for the hypergrid is in terms of an expression with binomial coefficients. The bounds differ by a factor of $O(d^{2m})$ and, in particular, give tight (up to a $poly(d)$ factor) bounds for constant m. We also give a second set of bounds, which show that the number of edges in the sparsest 2-TC-spanner of the hypergrid is at most $m^d \log^d m$ and at least $\Omega\left(\max\left\{m^d \frac{\log^d m}{(2d \log\log m)^{d-1}}, (m-1)^d 2^{(c+\alpha-1)d}\right\}\right)$, where $c \approx 1.1620$, as above, and $\alpha > 0$ satisfies $1 + H_b(\alpha) < c$. The two sets of bounds are, in general, incomparable.

Our results rule out a class of approaches to monotonicity testing of functions of the form $f : \{0, 1\}^d \rightarrow R$ and, more generally, $f : \{1, 2, \ldots, m\}^d \rightarrow R$, where R is an arbitrary range. [7] showed that sparse 2-TC-spanners imply fast monotonicity testers, and used this connection to improve existing monotonicity testers for planar and other H-minor-free graphs. It left open the question, which was again raised at the 2008 Dagstuhl seminar on Sublinear Algorithms, of whether the 2-TC-spanner approach can improve monotonicity testers on the hypercube and hypergrid. We show that a fundamentally new approach is required.

*Massachusetts Institute of Technology, USA. Email:{abhatt,elena_g,kmjung}@mit.edu.
†Pennsylvania State University, USA. Email: sofya@cse.psu.edu.
‡IBM Almaden Research Center, USA. Email: dpwoodru@us.ibm.com.

ISSN 1433-8092

1 Introduction

Graph spanners were introduced in the context of distributed computing [29], and since then have found numerous applications, such as efficient routing [11, 12, 31, 32, 37], simulating synchronized protocols in unsynchronized networks [30], parallel and distributed algorithms for approximating shortest paths [9, 10, 15], and algorithms for distance oracles [5, 38].

Several variants on graph spanners have been defined. In this work, we focus on *transitive-closure* spanners that were introduced in [7] as a common abstraction for applications in access control, property testing and data structures.

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

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

Thus, a k-transitive-closure-spanner (or k-TC-spanner) is a graph with small diameter that preserves the connectivity of the original graph. In the applications above, the goal is to find the sparsest k-TC-spanner for a given k and G. The number of edges in the sparsest k-TC-spanner of G is denoted by $S_k(G)$.

Our Results In this work we investigate the size of the sparsest 2-TC-spanners for the directed hypercube and hypergrid. These graph families are natural network topologies, and other variants of spanners for the hypercube, hypergrid, and other restricted graph families have been extensively studied [14, 17, 24, 25, 26, 27, 28].

The *directed hypercube*, denoted $\mathcal{H}_d$, has the vertex set $\{0,1\}^d$ and the edge set $\{(x, y) : x_i \leq y_i \ \forall i \in \{1, \dots, d\}$ and $|x - y| = 1\}$ where $|x - y|$ represents the Hamming distance between the two strings.

The obvious bounds on $S_2(\mathcal{H}_d)$ are the number of edges in the d-dimensional hypercube, $2^{d-1}d$, and the number of edges in the transitive closure of $\mathcal{H}_d$, which is $3^d - 2^d$. (An edge in the transitive closure of $\mathcal{H}_d$ has 3 possibilities for each coordinate: both endpoints are 0, both endpoints are 1, or the first endpoint is 0 and the second is 1. This includes self-loops, so we have to subtract the number of vertices in $\mathcal{H}_d$ to get the desired quantity.) Thus, $2^{d-1}d \leq S_2(\mathcal{H}_d) \leq 3^d - 2^d$.

The following theorem, proved in Section 3, gives the size of the sparsest 2-TC-spanner of the hypercube up to a multiplicative term polylogarithmic in the size of the graph.

Theorem 1.1 (Hypercube)**.** *Let $S_2(\mathcal{H}_d)$ denote the number of edges in the sparsest 2-TC-spanner of $\mathcal{H}_d$. Then*

$$S_2(\mathcal{H}_d) = O(d^3 2^{cd}) \text{ and } \Omega(2^{cd}), \text{ where } c \approx 1.1620.$$

We prove the theorem by giving nearly matching upper and lower bounds on $S_2(\mathcal{H}_d)$ in terms of an expression with binomial coefficients, and later numerically estimating the value of the expression. We prove the upper bound in Theorem 1.1 by presenting a randomized construction of a 2-TC-spanner of the directed hypercube. Curiously, even though the upper and lower bounds above differ by a factor of $O(d^3)$, we can show that our construction yields a 2-TC-spanner of $\mathcal{H}_d$ of size within $O(d^2)$ of the optimal.

In Sections 4 and 5, we present upper and lower bounds on the size of the sparsest 2-TC-spanner of the hypergrid. The *directed hypergrid*, denoted $\mathcal{H}_{m,d}$, has vertex set $\{1, 2, \dots, m\}^d$ and edge set $\{(x, y) : \exists$ unique $i \in \{1, \dots, d\}$ such that $y_i - x_i = 1$ and for $j \neq i, y_j = x_j\}$. The straightforward bounds on the number of edges in a 2-TC-spanner of $\mathcal{H}_{m,d}$ in terms of the number of edges in the directed grid and in its transitive closure are $dm^{d-1}(m-1)$ and $\left(\frac{m^2+m}{2}\right)^d - m^d$, respectively.

In Section 4, we extend our analysis for the hypercube to give upper and lower bounds on $S_2(\mathcal{H}_{m,d})$ in terms of an expression with binomial coefficients (Theorem 4.1). The upper and lower bounds differ by a factor of $O(d^{2m})$ and, in particular, show that our randomized 2-TC-spanner construction is optimal up to a $poly(d)$ factor for constant m. The value of the combinatorial expression can be estimated numerically for small m. Specifically, $S_2(\mathcal{H}_{m,d}) = 2^{c_m d}\,\text{poly}(d)$, where $c_3 \approx 2.03$, $c_4 \approx 2.82$ and $c_5 \approx 3.24$, each significantly smaller than the exponents corresponding to the transitive closure sizes for the different m.

The following theorem, proved in Section 5, gives another set of explicit bounds on $S_2(\mathcal{H}_{m,d})$ which, in general, are incomparable to the bounds described above.

Theorem 1.2 (Hypergrid). *Let $S_2(\mathcal{H}_{m,d})$ denote the number of edges in the sparsest 2-TC-spanner of $\mathcal{H}_{m,d}$. Then for $m \geq 3$,*

$$S_2(\mathcal{H}_d) \leq m^d \log^d m \;\; \textit{and} \quad \geq \Omega\left(\max\left\{\frac{m^d \log^d m}{(2d \log\log m)^{d-1}}, (m-1)^d 2^{(c+\alpha-1)d}\right\}\right),$$

where c is the constant from Theorem 1.1 and $\alpha > 0$ satisfies $1 + H_b(\alpha) < c$.

We prove the upper bound in Theorem 1.2 by presenting a general construction of k-TC-spanners for graph products for arbitrary $k \geq 2$. The second term in the lower bound expression in Theorem 1.2 is derived from the lower bound for the hypercube. The first term in the lower bound expression is proved by a reduction of the 2-TC-spanner construction for $[m]^d$ to that for the $2 \times [m]^{d-1}$ grid and then directly analyzing the number of edges required for a 2-TC-spanner of $2 \times [m]^{d-1}$. This analysis is one of the more interesting combinatorial arguments in the paper. We show a tradeoff between the number of edges in the 2-TC-spanner of the $2 \times [m]^{d-1}$ grid that stay within the hyperplanes $\{1\} \times [m]^{d-1}$ and $\{2\} \times [m]^{d-1}$ versus the number of edges that cross from one hyperplane to the other. The proof proceeds in multiple stages; assuming an upper bound on the number of edges staying within the hyperplanes, each stage is shown to separately contribute a substantial number of edges crossing between the hyperplanes. The proof of this tradeoff lemma is already non-trivial for $d = 2$ and is presented first.

Motivation: TC-spanner method in monotonicity testing As shown in [7], TC-spanners have several applications. 2-TC-spanners for the hypercube and hypergrid are especially relevant for the application to monotonicity testing.

Testing monotonicity of functions [2, 6, 13, 16, 18, 19, 20, 22] is one of the most studied problems in property testing [21, 33]. Testing monotonicity is equivalent to several other testing problems [19]. Let V_n be a poset of n elements and $G_n = (V_n, E)$ be the relation graph, i.e., the Hasse diagram, for V_n. A function $f : V_n \to \mathbb{R}$ is called *monotone* if $f(x) \leq f(y)$ for all $(x, y) \in E$. We say f is ϵ-far from monotone if f has to be changed on at least an ϵ fraction of the domain to become monotone, that is, $\min_{\text{monotone } g} |\{x : f(x) \neq g(x)\}| \geq \epsilon n$. A monotonicity tester on G_n is an algorithm that, given an oracle for a function $f : V_n \to \mathbb{R}$, accepts if f is monotone but rejects with probability $\geq \frac{2}{3}$ if f is ϵ*-far* from monotone.

For instance, if G_n is a directed line, $\mathcal{H}_{n,1}$, the tester needs to determine whether the input sequence specified by f is sorted or ϵ-far from sorted. If G_n is a 2-dimensional grid, $\mathcal{H}_{m,2}$, (with vertex set $\{1, ..., m\} \times \{1, ..., m\}$ and edge set $\{(x, y) \mid x_1 = y_1 \text{ and } x_2 + 1 = y_2\} \cup \{(x, y) \mid x_1 + 1 = y_1 \text{ and } x_2 = y_2\}$), the goal is to determine whether the input matrix has non-decreasing rows and columns. Finally, if $G_n = \mathcal{H}_d$, one has to determine if the input function $f : \{0, 1\}^d \to R$ is monotone.

The optimal monotonicity tester for the directed line, proposed in [13], is based on the sparsest 2-TC-spanner for that graph. The following lemma from [7] proves that a sparse 2-TC-spanner for any partial order graph G_n implies an efficient monotonicity tester on G_n.

Lemma 1.3 ([7]). *If a directed acyclic graph G_n has a 2-TC-spanner with $s(n)$ edges, then there exists a monotonicity tester on G_n that runs in time $O\left(\frac{s(n)}{\epsilon n}\right)$.*

This lemma led to significant improvements in monotonicity testers for several graph families, including planar graphs and, in general, H-minor-free graphs [7]. It left open the question, which was again raised at the 2008 Dagstuhl seminar on Sublinear Algorithms, of whether the 2-TC-spanner approach can improve monotonicity testers of functions of the form $f : \{0,1\}^d \to R$ and, more generally, $f : \{1,2,\dots,m\}^d \to R$, where R is an arbitrary range. Currently, the running time of the best tester for this problem is $O\left(\frac{d}{\epsilon}\log m \cdot \log|R|\right)$ [13], while the best known lower bound (for the hypercube with range $R = \{0,1\}$) is $\Omega(\log\log d)$ [19]. Even though for a fixed d, it is known that the optimal monotonicity tester for the grid runs in time $\Theta(\frac{\log m}{\epsilon})$ [22, 18], bridging the gap between the lower and upper bounds for arbitrary d has remained elusive. Lemma 1.3 showed that if a 2-TC-spanner of size $o(2^d d^2)$ for the hypercube or, more generally, a 2-TC-spanner of size $o(m^d d^2 \log^2 m)$ for the hypergrid were found, the monotonicity tester of [13] would be improved. Our $\Omega\left(m^2 \frac{\log^2 m}{\log\log m}\right)$ bound on the size of a 2-TC-spanner of the 2-dimensional grid (Theorems 1.2 and, specifically, 5.4) shows that the optimal monotonicity testers for constant-dimensional grids from [22] cannot be matched with the TC-spanner approach. Our lower bounds for the size of the sparsest 2-TC-spanners for the hypercube (Theorem 1.1) and the hypergrid (Theorem 1.2) rule out the TC-spanner approach for improving monotonicity testers on the hypercube and hypergrid. A fundamentally new approach is required.

Previous work on bounding S_k for other families of graphs Thorup [34] considered a special case of TC-spanners of graphs G that have at most twice as many edges as G, and conjectured that for all directed graphs G on n nodes there are such k-TC-spanners with k polylogarithmic in n. He proved his conjecture for planar graphs [35], but later Hesse [23] gave a counterexample to Thorup's conjecture for general graphs by constructing a family of graphs for which all $n^{\frac{1}{17}}$-TC-spanners need at least $n^{1+\Omega(1)}$ edges. TC-spanners were also studied for directed trees: implicitly in [3, 4, 8, 13, 39] and explicitly in [36]. The implicit results were interpreted as TC-spanner constructions in [7]. For the directed line, [3] (and later, [4]) expressed the size of the sparsest k-TC-spanner in terms of the inverse Ackermann function. The *Ackermann function* ([1]) is defined by: $A(1,j) = 2^j$, $A(i+1,0) = A(i,1)$, $A(i+1,j+1) = A(i, 2^{2^{A(i+1,j)}})$. The inverse Ackermann function is $\alpha(n) = \min\{i : A(i,1) \geq n\}$ and the i^{th}-row inverse is $\lambda_i(n) = \min\{j : A(i,j) \geq n\}$.

Lemma 1.4 ([3, 4, 7]). *Let $S_k(\mathcal{H}_{n,1})$ denote the number of edges in the sparsest k-TC-spanner of the directed line $\mathcal{H}_{n,1}$. Then $S_2(\mathcal{H}_{n,1}) = \Theta(n\log n)$, $S_3(\mathcal{H}_{n,1}) = \Theta(n\log\log n)$, $S_4(\mathcal{H}_{n,1}) = \Theta(n\log^* n)$ and, more generally, $S_k(\mathcal{H}_{n,1}) = \Theta(n\lambda_k(n))$ where $\lambda_k(n)$ is the inverse Ackermann function.*

[3, 8, 36] gave the same bound for directed trees on n nodes. [7] extended it to $O(n\log n \cdot \lambda_k(n))$ bound on S_k for H-minor-free graph families, which include planar graphs, bounded tree-width graphs, and bounded genus graphs.

2 Preliminaries

For a positive integer m, we denote $\{1,\dots,m\}$ by $[m]$. For $x \in \{0,1\}^d$, we use $|x|$ to denote the weight of x, that is, the number of non-zero coordinates in x. Level i in a hypercube contains all vertices of weight i. The partial order $\preceq$ on the hypergrid $\mathcal{H}_{m,d}$ is defined as follows: $x \preceq y$ for two vertices $x, y \in [m]^d$ iff $x_i \leq y_i$ for all $i \in [d]$. Vertices x and y are *comparable* if either y is *above* x (that is, $x \preceq y$) or y is *below* x (that is, $y \preceq x$).

We denote a path from v_1 to v_ℓ, consisting of edges $(v_1,v_2),(v_2,v_3),\dots,(v_{\ell-1},v_\ell)$ by $(v_1,\dots,v_\ell)$.

As usual, $\log$ denotes the logarithm base 2 and $\ln$ denotes the logarithm base e.

3 2-TC-spanners of the Hypercube

In this section we prove Theorem 1.1, namely, we analyze the size of the sparsest 2-TC-spanner of the d-dimensional hypercube $\mathcal{H}_d$. Lemma 3.1 presents the upper bound on $S_2(\mathcal{H}_d)$. Lemma 3.3 presents the lower bound. The upper and lower bounds differ only by a factor of $O(d^3)$, and are dominated by the same combinatorial expression. A numerical approximation to this expression is given in Lemma 3.4. Remark 3.1 at the end of the section explains why our randomized construction in Lemma 3.1 yields a 2-TC-spanner of $\mathcal{H}_d$ of size within $O(d^2)$ of the optimal. The missing material is deferred to Appendix A.

Lemma 3.1. *There is a* 2*-TC-spanner of* $\mathcal{H}_d$ *with* $O\left(d^3 \max_{i,j:i<j} \min_{k:i\le k\le j} \frac{\binom{d}{k}}{\binom{j-i}{k-i}} \max\left\{\binom{k}{i}, \binom{d-k}{d-j}\right\}\right)$ *edges.*

Proof. Consider the following probabilistic construction that connects all comparable vertices at levels i and j of $\mathcal{H}_d$ by paths of length at most 2:

Given levels $i, j \in \{0, 1, ..., d\}$, $i < j$,

1. Initialize the set $E_{i,j}$ to $\emptyset$.
2. Let $k_{i,j} = \operatorname{argmin}_{k:i\le k\le j} \left(\frac{\binom{d}{k}}{\binom{j-i}{k-i}} \max\left\{\binom{k}{i}, \binom{d-k}{d-j}\right\}\right)$.
3. Let $S_{i,j}$ be a set of $3d\frac{\binom{d}{k}}{\binom{j-i}{k-i}}$ vertices chosen uniformly at random from the set of $\binom{d}{k}$ vertices that are in weight level $k = k_{i,j}$.
4. For each vertex $v \in S_{i,j}$, set $E_{i,j}$ to $E_{i,j} \cup \{(x, v) : |x| = i \wedge x \prec v\} \cup \{(v, y) : |y| = j \wedge v \prec y\}$. That is, connect v to all comparable vertices in levels i and j.
5. Output $E_{i,j}$.

Claim 3.2. *For all* $0 \le i < j \le d$*, with probability at least* $\frac{1}{2}$*,* $E_{i,j}$ *contains a path of length at most 2 between any pair of vertices* (x, y) *such that* $x \prec y$*,* $|x| = i$*, and* $|y| = j$*.*

Proof. Consider any particular pair of vertices (x, y) such that $x \prec y$, $|x| = i$, and $|y| = j$. The number of vertices in level k that are greater than x and less than y is exactly $\binom{j-i}{k-i}$. So, the probability that $S_{i,j}$ does not contain such a vertex is: $\left(1 - \binom{j-i}{k-i}/\binom{d}{k}\right)^{3d\frac{\binom{d}{k}}{\binom{j-i}{k-i}}} \le e^{-3d}$. The number of comparable pairs (x, y) is $\binom{d}{i}\binom{d-i}{d-j}$. So, by the union bound, the probability that there exists an (x, y) such that no vertex $v \in S_{i,j}$ satisfies $x \prec v \prec y$ is at most $\binom{d}{i}\binom{d-i}{d-j}e^{-3d} \le 2^{2d}e^{-3d} < \frac{1}{2}$. $\square$

So, for every i and j, there exists a choice of $S_{i,j}$ such that comparable pairs from the two weight levels are connected by a path of length at most 2. Let $E^*_{i,j}$ be the set of edges returned by the algorithm when this $S_{i,j}$ is chosen. We set $E = \bigcup_{0\le i<j\le d} E^*_{i,j}$. By Claim 3.2, $(\{0, 1\}^d, E)$ is a 2-TC-spanner of $\mathcal{H}_d$.

Now, we show that the size of E is as claimed in the lemma statement. The main observation is that in step (4), for any specific $v \in S_{i,j}$, $|\{(x, v) : |x| = i \wedge x \prec v\} \cup \{(v, y) : |y| = j \wedge v \prec y\}|$ is exactly $\binom{k_{i,j}}{i} + \binom{d-k_{i,j}}{d-j}$. Therefore, for all $0 \le i < j \le d$,

$$|E^*_{i,j}| \le 3d \min_{k:i\le k\le j} \frac{\binom{d}{k}}{\binom{j-i}{k-i}} \left(\binom{k}{i} + \binom{d-k}{d-j}\right) \le 6d \min_{k:i\le k\le j} \frac{\binom{d}{k}}{\binom{j-i}{k-i}} \max\left\{\binom{k}{i}, \binom{d-k}{d-j}\right\}.$$

Since $|E| = \sum_{0\leq i<j\leq d} |E^*_{i,j}|$, where the sum has $O(d^2)$ terms, the claimed bound follows. ☐

Lemma 3.3. *Any 2-TC-spanner of $\mathcal{H}_d$ has $\Omega\left(\max_{i,j:i<j}\min_{k:i\leq k\leq j}\frac{\binom{d}{k}}{\binom{j-i}{k-i}}\max\left\{\binom{k}{i},\binom{d-k}{d-j}\right\}\right)$ edges.*

Proof. Let S be a 2-TC-spanner for $\mathcal{H}_d$. We will count the edges in S that occur on paths connecting two particular weight levels of $\mathcal{H}_d$. Let $P_{i,j}$ be the pairs $\{(v_1, v_2) : |v_1| = i, |v_2| = j, v_1 \prec v_2\}$. We will lower bound $e^*_{i,j}$, the number of edges in the paths of length at most 2 in S that connect the pairs $P_{i,j}$. Let $e_{k,\ell}$ denote the number of edges in S that connect vertices in level k to vertices in level ℓ. Then $e^*_{i,j} = e_{i,j} + \sum_{k=i+1}^{j-1}(e_{i,k} + e_{k,j})$.

We say that a vertex v *covers* a pair of vertices (v_1, v_2) if S contains the edges (v_1, v) and (v, v_2) or, for the special case $v = v_1$, if S contains (v_1, v_2). Let $V^{(k)}_{i,j}$ be the set of vertices of weight k that cover pairs in $P_{i,j}$. Let α_k be the fraction of pairs in $P_{i,j}$ that are covered by a vertex in $V^{(k)}_{i,j}$. Since each pair in $P_{i,j}$ must be covered by a vertex in levels i to $j-1$, $\sum_{k=i}^{j-1}\alpha_k \geq 1$.

For any vertex $v \in V^{(k)}_{i,j}$, let in_v be the number of incoming edges from vertices of weight i incident to v and let out_v be the number of outgoing edges to vertices of weight j incident to v. For each $k \in \{i+1, ..., j-1\}$, since each vertex $v \in V^{(k)}_{i,j}$ covers $in_v \cdot out_v$ pairs,

$$\sum_{v\in V^{(k)}_{i,j}} in_v \cdot out_v \geq \alpha_k|P_{i,j}| = \alpha_k\binom{d}{i}\binom{d-i}{d-j}. \tag{1}$$

We upper bound $\sum_{v\in V^{(k)}_{i,j}} in_v \cdot out_v$ as a function of $e_{i,k} + e_{k,j}$, and then use Equation (1) to lower bound $e_{i,k} + e_{k,j}$.

For all $k \in \{i+1, ..., j-1\}$, variables in_v and out_v satisfy the following constraints:

$$\sum_{v\in V^{(k)}_{i,j}} in_v \leq e_{i,k}+e_{k,j},\quad \sum_{v\in V^{(k)}_{i,j}} out_v \leq e_{i,k}+e_{k,j},\quad in_v \leq \binom{k}{i}\forall v\in V^{(k)}_{i,j},\quad out_v \leq \binom{d-k}{d-j}\forall v\in V^{(k)}_{i,j}.$$

The last two constraints hold because in_v and out_v count the number of edges to a vertex of weight k from from vertices of weight i and from a vertex of weight k to vertices of weight j, respectively. We want to maximize $\sum_{v\in V^{(k)}_{i,j}} in_v \cdot out_v$ subject to the above constraints. Claim A.1, a technical statement proved in Appendix A, bounds the sum by proving that the maximum occurs when $in_v = \binom{k}{i}$ and $out_v = \binom{d-k}{d-j}$ for as many v as possible, subject to the remaining constraints. It gives us, for all $k \in \{i+1, ..., j-1\}$:

$$\sum_{v\in V^{(k)}_{i,j}} in_v \cdot out_v \leq 2(e_{i,k} + e_{k,j})\min\left\{\binom{k}{i},\binom{d-k}{d-j}\right\}.$$

Let $s_{i,k,j} = \frac{\binom{d}{i}\binom{d-i}{d-j}}{\min\left\{\binom{k}{i},\binom{d-k}{d-j}\right\}}$. From Equation (1), $e_{i,k} + e_{k,j} \geq \frac{1}{2}\alpha_k s_{i,k,j}$ for all $k \in \{i+1, ..., j-1\}$. Therefore,

$$e^*_{i,j} = e_{i,j} + \sum_{k=i+1}^{j-1}(e_{i,k} + e_{k,j}) \geq \alpha_i\binom{d}{i}\binom{d-i}{d-j} + \frac{1}{2}\sum_{k=i+1}^{j-1}\alpha_k s_{i,k,j} \geq \frac{1}{2}\sum_{k=i}^{j-1}\alpha_k s_{i,k,j} \geq \frac{1}{2}\min_{k:i\leq k\leq j} s_{i,k,j}$$

Since this holds for arbitrary i and j, the number of edges in the 2-TC-spanner $|S| \geq \frac{1}{2}\max_{i,j:i<j}\min_{k:i\leq k\leq j} s_{i,k,j}$.

Finally, a simple algebraic manipulation finishes the proof (see Claim A.2.) ☐

The following lemma completes the proof of Theorem 1.1.

Lemma 3.4. *Let* $s = \max_{i,j:i<j} \min_{k:i\le k\le j} \frac{\binom{d}{k}}{\binom{j-i}{k-i}} \max\left\{\binom{k}{i}, \binom{d-k}{d-j}\right\}$. *Then* $s = 2^{cd}$, *where* $c \approx 1.1620$.

Remark 3.1. We note that if the first maximum in the expression for s is replaced with the sum then Lemma 3.1 holds for $O(d\cdot s)$ instead of $O(d^3\cdot s)$ while Lemma 3.3 holds for $\Omega(d/s)$ instead of $\Omega(s)$. The proofs of these modified statements are similar. (We do not have an analogue of Lemma 3.4 for the modified expression for s.) Observe that the modified bounds differ by a factor of $O(d^2)$ instead of $O(d^3)$. This demonstrates that our randomized construction yields a 2-TC-spanner of $\mathcal{H}_d$ of size within $O(d^2)$ of the optimal. $\diamondsuit$

4 Tight Bounds for the Hypergrid in Terms of Combinatorial Expressions

In this section we generalize the arguments for the hypercube in Section 3 to the directed hypergrid $\mathcal{H}_{m,d}$. We obtain matching upper and lower bounds up to a d^{2m} factor in terms of an expression involving binomial coefficients (see Theorem 4.1). This expression can be evaluated numerically for small m, like in Lemma 3.4, to find the size of the sparsest 2-TC-spanner for $\mathcal{H}_{m,d}$ to within $\mathrm{poly}(d)$ factors.

Definition 4.1. *For the hypergrid* $\mathcal{H}_{m,d}$ *, define a* level *to be a set of vertices, indexed by vector* $\mathbf{i}\in[d]^m$ *with* $i_1+\cdots+i_m=d$, *that consists of vertices* $x=(x_1,\ldots,x_d)\in[m]^d$ *containing* i_1 *positions of value* 1, i_2 *positions of value* 2, …, *and* i_m *positions of value* m.

The number of vertices in level $\mathbf{i}=(i_1,i_2,\ldots,i_m)$ is the multinomial coefficient

$$\binom{d}{\mathbf{i}} = \binom{d}{i_1,...,i_d} = \binom{d}{i_1}\binom{d-i_1}{i_2}\binom{d-i_1-i_2}{i_3}\cdots\binom{d-\sum_{l=1}^{m-1} i_l}{i_m}.$$

Indeed, there are $\binom{d}{i_1}$ choices for the coordinates of value 1. For each such choice there are $\binom{d-i_1}{i_2}$ choices for the coordinates of value 2, and repeating this argument one obtains the above expression.

For levels $\mathbf{i},\mathbf{j}\in[d]^m$, say $\mathbf{j}$ *majorizes* $\mathbf{i}$, denoted $\mathbf{j}\succ\mathbf{i}$, if $\mathbf{j}$ contains a vertex which is above some vertex in $\mathbf{i}$, that is, if $\sum_{\ell=t}^{m} j_\ell \ge \sum_{\ell=t}^{m} i_\ell$ for all $t\in\{m,m-1,...,1\}$.

For $\mathbf{j}\succ\mathbf{i}$, the number of vertices y at level $\mathbf{i}$ comparable to a fixed vertex x at level $\mathbf{j}$ is

$$\mathcal{M}(\mathbf{i},\mathbf{j}) = \binom{j_m}{i_m}\binom{j_m+j_{m-1}-i_m}{i_{m-1}}\binom{j_m+j_{m-1}+j_{m-2}-i_m-i_{m-1}}{i_{m-2}}\cdots\binom{\sum_{l=1}^m j_l-\sum_{l=2}^m i_l}{i_1}.$$

Indeed, there are $\binom{j_m}{i_m}$ choices for the coordinates of value m in y. For each such choice, there are $\binom{j_m+j_{m-1}-i_m}{i_{m-1}}$ choices for the coordinates of value $m-1$ in y, and one can repeat this argument to obtain the claimed expression.

For $\mathbf{j}\succ\mathbf{i}$, the number of vertices y at level $\mathbf{j}$ comparable to a fixed vertex x at level $\mathbf{i}$ is

$$\mathcal{N}(\mathbf{i},\mathbf{j}) = \frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\binom{d}{\mathbf{i}}}.$$

Indeed, there are $\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}$ comparable pairs of vertices in levels $\mathbf{i}$ and $\mathbf{j}$, and level $\mathbf{i}$ contains $\binom{d}{\mathbf{i}}$ vertices. Since, by symmetry, each vertex in $\mathbf{i}$ is comparable to the same number of vertices in level $\mathbf{j}$, we get the desired expression.

Theorem 4.1. *Let* $\mathcal{B}(m,d) = \max_{\mathbf{i},\mathbf{j}:\mathbf{j}\succ\mathbf{i}} \min_{\mathbf{k}:\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}} \frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})} \max\{\mathcal{M}(\mathbf{i},\mathbf{k}),\mathcal{N}(\mathbf{k},\mathbf{j})\}$ *. Then the number of edges in the sparsest 2-TC-spanner of the directed hypergrid* $\mathcal{H}_{m,d}$ *is* $O\left(d^{2m}\mathcal{B}(m,d)\right)$ *and* $\Omega\left(\mathcal{B}(m,d)\right)$.

Theorem 4.1 follows from Lemmas B.1 and B.3 that appear in Appendix B.

5 Explicit Bounds for the Hypergrid

In this section we prove Theorem 1.2 that gives explicit bounds on the size of the sparsest 2-TC-spanners of $\mathcal{H}_{m,d}$. The bounds are stated separately in Corollary 5.2, Theorem 5.7 and Theorem 5.10. The upper bound in Corollary 5.2 is proved in Section 5.1. The lower bounds in Theorem 5.7 and Theorem 5.10 appear in Section 5.2.

5.1 Upper Bound

This section explains how to construct a TC-spanner of the Cartesian product of graphs G_1 and G_2 from TC-spanners of G_1 and G_2. Since the directed hypergrid is the Cartesian product of directed lines, and optimal TC-spanner constructions are known for the directed line, our construction yields sparse TC-spanners for the grid (Corollary 5.2). We start by defining two graph products: Cartesian and strong.

Definition 5.1 (Graph products)**.** *Given graphs* $G_1 = (V_1, E_1)$ *and* $G_2 = (V_2, E_2)$*, a product of* G_1 *and* G_2 *is a new graph* G *with vertex set* $V_1 \times V_2$*. For the* Cartesian graph product*, denoted by* $G_1 \times G_2$*, graph* G *contains an edge from* (u_1, u_2) *to* (v_1, v_2) *if and only if* $u_1 = v_1$ *and* $(u_2, v_2) \in E_2$*, or* $(u_1, v_1) \in E_1$ *and* $u_2 = v_2$*. For the* strong graph product*, denoted by* $G_1 \circ G_2$*, graph* G *contains an edge from* (u_1, u_2) *to* (v_1, v_2) *if and only if* $u_1 = v_1$ *and* $(u_2, v_2) \in E_2$*, or* $(u_1, v_1) \in E_1$ *and* $u_2 = v_2$*, or* $(u_1, v_1) \in E_1$ *and* $(u_2, v_2) \in E_2$*.*

For example, $\mathcal{H}_{m,2} = \mathcal{H}_{m,1} \times \mathcal{H}_{m,1}$ and TC$(\mathcal{H}_{m,2}) =$ TC$(\mathcal{H}_{m,1}) \circ TC(\mathcal{H}_{m,1})$, where TC$(G)$ denotes the transitive closure of G.

Lemma 5.1. *Let* G_1 *and* G_2 *be directed graphs with* k*-TC-spanners* S_1 *and* S_2*, respectively. Then* $S_1 \circ S_2$ *is a* k*-TC-spanner of* $G = G_1 \times G_2$*.*

Proof. Suppose (u, v) and (u', v') are comparable vertices in $G_1 \times G_2$. Then, by definition of the Cartesian product, $u \preceq u'$ in G_1 and $v \preceq v'$ in G_2. Let $(u_1, u_2, \ldots, u_\ell)$ be the shortest path in S_1 from $u = u_1$ to $u' = u_\ell$, and $(v_1, v_2, \ldots, v_t)$ the shortest path in S_2 from $v = v_1$ to $v' = v_t$. Assume w.t.o.g. that $l \leq t$. Then $((u_1, v_1), (u_2, v_2), \ldots, (u_\ell, v_\ell) \ldots, (u_\ell, v_t))$ is a path in $S_1 \circ S_2$ of length $t \leq k$, from (u, v) to (u', v'). Therefore, $S_1 \circ S_2$ is a k-TC-spanner of $G = G_1 \times G_2$. □

Lemma 5.1 together with previous results on the size of k-TC-spanners for the line $\mathcal{H}_{m,1}$, summarized in Lemma 1.4, imply an upper bound on the size of a k-TC-spanner of the directed hypergrid $\mathcal{H}_{m,d}$:

Corollary 5.2. *Let* $S_k(\mathcal{H}_{m,d})$ *denote the number of edges in the sparsest* k*-TC-spanner of the directed* d*-dimensional hypergrid* $\mathcal{H}_{m,d}$*. Then* $S_k(\mathcal{H}_{m,d}) = O(m^d \lambda_k(m)^d c^d)$ *for appropriate constant* c*.*

More precisely, $S_2(\mathcal{H}_{m,d}) \leq m^d \log^d m$ *for* $m \geq 3$*.*

Proof. Let S be a k-TC-spanner for the line $\mathcal{H}_{m,1}$. By Lemma 5.1, $S \circ \cdots \circ S$, where the strong graph product is applied d times, is a k-TC-spanner for the directed grid $\mathcal{H}_{m,d}$. By definition of the strong graph product, the number of edges in the resulting spanner is $(|E(S)| + m)^d - m^d$. Since the number of edges in the spanner, $|E(S)|$, is at least m, the main statement follows.

The more precise statement for $k = 2$ follows from Claim C.1 in Appendix C which gives a more careful analysis of the size of the sparsest 2-TC-spanner of the line: namely, $S_2(\mathcal{H}_{m,1}) \leq m \log m - m$ for $m \geq 3$. □

5.2 Lower Bounds

In this section we prove an explicit lower bound on the size of a 2-TC-spanner of the d-dimensional directed grid, stated in Theorem 1.2. Section 5.2.1 proves the first term in the lower bound expression for the special case of the 2-dimensional grid. Section 5.2.2 extends the proof to an arbitrary dimension. Section 5.2.3 proves the second term in the lower bound expression.

We start with an observation useful for all lower bounds in this section. It is tempting to think that a subgraph of a TC-spanner is itself a TC-spanner, however, in general, this is not the case. We observe that it is true for subgrids of a hypergrid that include all vertices between the lowest and the highest vertices in the subgrid.

Claim 5.3. *Let $x, y \in [m]^d$. Define $G_{x,y}$ to be the subgraph of $\mathcal{H}_{m,d}$ induced by the vertex set $\{z : x \preceq z \preceq y\}$. Every k-TC-spanner S of $\mathcal{H}_{m,d}$ must contain a k-TC-spanner of $G_{x,y}$.*

Proof. If a path (of length at most k) in S leaves $G_{x,y}$ it cannot return. □

5.2.1 Lower Bound for $d = 2$

In this section we prove a lower bound on the size of a 2-TC-spanner of the 2-dimensional directed grid, stated in Theorem 5.4. This is a special case of the lower bound in Theorem 1.2.

Theorem 5.4. *Any 2-TC-spanner of the 2-dimensional grid $\mathcal{H}_{m,2}$ has $\Omega\left(\frac{m^2 \log^2 m}{\log\log m}\right)$ edges.*

One way to prove the $\Omega(m \log m)$ lower bound on the size of a 2-TC-spanner for the directed line $\mathcal{H}_{m,1}$, stated in Lemma 1.4, is to observe that at least $\lfloor \frac{m}{2} \rfloor$ edges are cut when the line is halved: namely, at least one per vertex pair $(v, m - v + 1)$ for all $v \in \left[\lfloor \frac{m}{2} \rfloor\right]$. Continuing to halve the line recursively, we obtain the desired bound.

A natural extension of this approach to proving a lower for the grid is to recursively halve the grid along both dimensions, hoping that every such operations on an $m \times m$ grid cuts $\Omega(m^2 \log m)$ edges. This would imply that the size $S(m)$ of a 2-TC-spanner of the $m \times m$ grid satisfies the recurrence $S(m) = 4S(m/2) + \Omega(m^2 \log m)$; that is, $S(m) = \Omega(m^2 \log^2 m)$, matching the upper bound in Theorem 1.2.

An immediate problem with this approach is that in some 2-TC-spanners of the grid only $\Omega(m^2)$ edges connect vertices in different quarters. One example of such a 2-TC-spanner is the graph containing the transitive closure of each quarter and only at most $3m^2$ edges crossing from one quarter to another: namely, for each node u and each quarter q with vertices comparable to u, this graph contains an edge (u, v_q), where v_q is the smallest node in q comparable to u.

The TC-spanner in the example above is not optimal because it has too many edges inside the quarters. The first step in our proof of Theorem 5.4 is understanding the tradeoff between the number of edges *crossing* the cut and the number of edges *internal* to the subgrids, resulting from halving the grid along some dimension. The simplest manifestation of this tradeoff occurs when a $2 \times m$ grid is halved into two lines. (In the case of one line, there is no trade off: the $\Omega(m)$ bound on the number of crossing edges holds even if each half-line contains all edges its transitive closure.) Lemma 5.5 formulates the tradeoff for the two-line case, while taking into account only edges needed to connect comparable vertices on different lines by paths of length at most 2:

Lemma 5.5 (Two-Lines Lemma). *Let U be a graph with vertex set $[2] \times [m]$ that contains a path of length at most 2 from u to v for every $u \in \{1\} \times [m]$ and $v \in \{2\} \times [m]$, where $u \preceq v$. An edge (u, v) in U is called* internal *if $u_1 = v_1$, and* crossing *otherwise. If U contains at most $\frac{m \log^2 m}{32}$ internal edges, it must contain at least $\frac{m \log m}{16 \log \log m}$ crossing edges.*

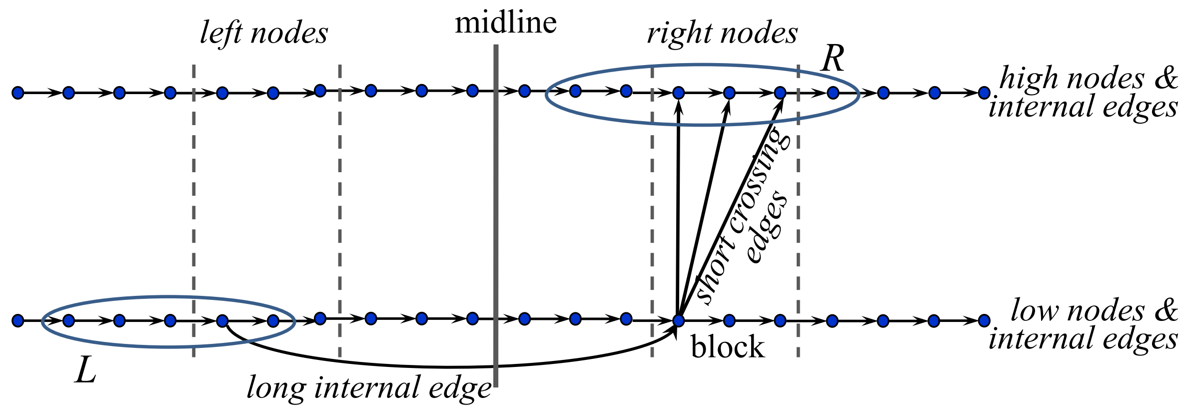

Figure 1: Illustration of the first stage in the proof of Lemma 5.5.

Note that if the number of internal edges is unrestricted, a 2-TC-spanner of $\mathcal{H}_{m,2}$ may have only m crossing edges.

Proof. The proof proceeds in $\frac{\log m}{2\log\log m}$ stages dealing with pairwise disjoint sets of crossing edges. In each stage, we show that U contains at least $\frac{m}{8}$ crossing edges in the prescribed set.

In the first stage, divide U into $\log^2 m$ blocks, each of length $\frac{m}{\log^2 m}$: namely, a node (v_1, v_2) is in block i if $v_2 \in \left[\frac{(i-1)\cdot m}{\log^2 m}+1, \frac{i\cdot m}{\log^2 m}\right]$. Call an edge *long* if it starts and ends in different blocks, and *short* otherwise. Assume, for contradiction, that U contains fewer than $\frac{m}{8}$ long crossing edges.

Call a node (v_1, v_2) *low* if $v_1 = 1$ (*high* if $v_1 = 2$), and *left* if $v_2 \in \left[\frac{m}{2}\right]$ (*right* otherwise). Also, call an edge (u, v) *low-internal* if $u_1 = v_1 = 1$ and *high-internal* if $u_1 = v_1 = 2$. Let L be the set of low left nodes that are not incident to long crossing edges. Similarly, let R be the set of high right nodes that are not incident to long crossing edges. Since there are fewer than $\frac{m}{8}$ long crossing edges, $|L| > \frac{m}{4}$ and $|R| > \frac{m}{4}$.

A node $u \in L$ can connect to a node $v \in R$ via a path of length at most 2 only by using a long internal edge. Observe that each long low-internal edge can be used by at most $\frac{m}{\log^2 m}$ such pairs (u, v): one low node u and high nodes v from one block. This is illustrated in Figure 1. Analogously, every long high-internal edge can be used by at most $\frac{m}{\log^2 m}$ such pairs. Since $|L|\cdot|R| > \frac{m^2}{16}$ pairs in $L \times R$ connect via paths of length at most 2, graph U contains more than $\frac{m^2}{16}\cdot\frac{\log^2 m}{m} = \frac{m\log^2 m}{16}$ long internal edges, which is a contradiction.

In each subsequent stage, call blocks used in the previous stage *megablocks*, and denote their length by B. Subdivide each megablock into $\log^2 m$ blocks of equal size. Call an edge *long* if it starts and ends in different blocks, but stays within one megablock. Assume, for contradiction, that U contains fewer than $\frac{m}{8}$ long crossing edges.

Call a node (v_1, v_2) *left* if it is in the left half of its megablock, that is, if $v_2 \le \frac{\ell+r}{2}$ whenever (v_1, v_2) is in a megablock $[2] \times \{\ell, \ldots, r\}$. (Call it *right* otherwise). Consider megablocks containing less than $\frac{B}{4}$ long crossing edges each. By an averaging argument, at least $\frac{m}{2B}$ megablocks are of this type. Within each such megablock more than $\frac{B}{4}$ low left nodes and more than $\frac{B}{4}$ high right nodes have no incident long crossing edges. By the argument from the first stage, each such megablock contributes more than $\frac{B^2}{16b}$ long internal edges, where $b = \frac{B}{\log^2 m}$ is the size of the blocks. Hence there must be more than $\frac{B^2}{16b}\cdot\frac{m}{2B} = \frac{m\log^2 m}{32}$ long internal edges, which is a contradiction to the fact that U contains at most $\frac{m\log^2 m}{32}$ internal edges.

We proceed to the next stage until each block is of length 1. Therefore, the number of stages, t, satisfies

$\frac{m}{\log^{2t} m} = 1$. That is, $t = \frac{\log m}{2 \log\log m}$, and each stage contributes $\frac{m}{8}$ new crossing edges, as desired. ☐

Next we generalize Lemma 5.5 to understand the tradeoff between the number of internal edges and crossing edges resulting from halving a 2-TC-spanner of an $2\ell \times m$ grid with the usual partial order.

Lemma 5.6. *Let S be a 2-TC-spanner of the directed $[2\ell] \times [m]$ grid. An edge (u, v) in S is called* internal *if $u_1, v_1 \in [\ell]$ or $u_1, v_1 \in \{\ell+1, \ldots, 2\ell\}$, and* crossing *otherwise. If S contains at most $\frac{\ell m \log^2 m}{64}$ internal edges, it must contain at least $\frac{\ell m \log m}{32 \log\log m}$ crossing edges.*

Proof. For each $i \in [\ell]$, we match the lines $\{i\} \times [m]$ and $\{2\ell - i + 1\} \times [m]$. Observe that a path of length at most 2 between the matched lines cannot use any edges with both endpoints in $\{i+1, \ldots, 2\ell - i\} \times [m]$. We modify S to ensure that there are no edges with only one endpoint in $\{i+1, \ldots, 2\ell - i\} \times [m]$ for all $i \in [\ell]$, and then apply Lemma 5.5 to the matched pairs of lines.

Call the $[\ell] \times [m]$ subgrid and all vertices and edges it contains *low*, and the remaining $\{\ell+1, \ldots, 2\ell\}$ subgrid and its vertices and edges *high*. Transform S into S' as follows: change each low internal edge (u, v) to $(u, (u_1, v_2))$, change each high internal edge (u, v) to $((v_1, u_2), v)$, and finally change each crossing edge $((i_1, j_1), (2\ell - i_2 + 1, j_2))$ to $((i, j_1), (2\ell - i + 1, j_2))$, where $i = \min(i_1, i_2)$. Intuitively, we are projecting the edges in S to be fully contained in one of the matched pairs of lines, while preserving whether the edge is internal or crossing. Crossing edges are projected onto the outer matched pair of lines chosen from the two pairs that contain the endpoints of a given edge.

Clearly, S' contains the same number of internal (crossing) edges as S. Observe that S' contains a path of length at most 2 from u to v for every comparable pair (u, v) where u is low, v is high, and u and v belong to the same pair of matched lines. Indeed, since S is a 2-TC-spanner, it contains either the edge (u, v) or a path (u, w, v). In the first case, S' also contains (u, v). In the second case, if (u, w) is a crossing edge S' contains $(u, (v_1, w_2), v)$, and if (u, w) is an internal edge S' contains $(u, (u_1, w_2), v)$. As claimed, each edge in S' belongs to one of the matched pairs of lines.

Finally, we apply Lemma 5.5. If S' contains at most $\frac{\ell m \log^2 m}{64}$ internal edges then at least half $\left(\text{i.e., } \frac{\ell}{2}\right)$ of the matched line pairs each contain at most $\frac{m \log^2 m}{32}$ internal edges. By Lemma 5.5, each of these pairs contributes at least $\frac{m \log m}{16 \log\log m}$ crossing edges. Thus S' must contain at least $\frac{\ell m \log m}{32 \log\log m}$ crossing edges. Since S contains as many crossing edges as S', the lemma follows. ☐

Now we prove Theorem 5.4 by recursively halving $\mathcal{H}_{m,2}$ along the horizontal dimension. Some resulting $\ell \times m$ subgrids may violate Lemma 5.6, but we can guarantee that the lemma holds for a constant fraction of the recursive steps for which $\ell \geq \sqrt{m}$. This is sufficient for obtaining the lower bound in the theorem.

Proof of Theorem 5.4. Assume m is a power of 2 for simplicity. For each step $i \in \{1, \ldots, \frac{1}{2}\log m\}$, partition $\mathcal{H}_{m,2}$ into the following 2^{i-1} equal-sized subgrids: $\{1, \ldots, l_i\} \times [m]$, $\{l_i + 1, \ldots, 2l_i\} \times [m]$, $\ldots$, $\{m - l_i + 1, \ldots, m\} \times [m]$ where $l_i = m/2^{i-1}$. For each of these subgrids, define internal and crossing edges as in Lemma 5.6. Now, suppose that there exists a step i such that at least half of the 2^{i-1} subgrids have $> \frac{l_i m \log^2 m}{64}$ internal edges. Since at a fixed i, the subgrids are disjoint, there are $2^{i-1}\Omega(l_i m \log^2 m) = \Omega(m^2 \log^2 m)$ edges in S, proving the theorem. On the other hand, suppose that for every $i \in \{1, \ldots, \frac{1}{2}\log m\}$, at least half of the 2^{i-1} subgrids have $\leq \frac{l_i m \log^2 m}{64}$ internal edges. Then, applying Lemma 5.6, the number of crossing edges in those subgrids is $\geq \frac{l_i m \log m}{32 \log\log m}$. Counting over all steps i and for all appropriate subgrids from those steps, the number of edges in S is bounded by $\Omega\left(m^2 \log m \frac{\log m}{\log\log m}\right) = \Omega\left(m^2 \frac{\log^2 m}{\log\log m}\right)$. ☐

5.2.2 Lower Bound for General d

This section generalizes the lower bound from the previous section to arbitrary d. The following theorem implies the first term in the lower bound expression in Theorem 1.2:

Theorem 5.7. *Any 2-TC-spanner of $\mathcal{H}_{m,d}$ has at least $\frac{m^d}{32}\frac{\log^d m}{(2d\log\log m)^{d-1}}$ edges.*

The main ingredient in the proof is the Two-Hyperplanes Lemma, an analogue of the Two-Lines Lemma (Lemma 5.5) for d dimensions. The main difficulty in extending the proof of the Two-Lines lemma to work for two hyperplanes is in generalizing the definitions of blocks and megablocks, so that, on one hand, each stage in the proof contributes a substantial number of crossing edges and, on the other hand, the crossing edges contributed in separate stages are pairwise disjoint.

Lemma 5.8 (Two-Hyperplanes Lemma)**.** *Let U be a graph with vertex set $[2]\times[m]^{d-1}$ that contains a path of length at most 2 from u to v for every $u\in\{1\}\times[m]^{d-1}$ and $v\in\{2\}\times[m]^{d-1}$, where $u\preceq v$. As in Lemma 5.5, an edge (u,v) in U is called* internal *if $u_1=v_1$, and* crossing *otherwise. Then, if U contains less than $\frac{m^{d-1}\log^d m}{(d-1)2^{2d+3}}$ internal edges, it must contain $\geq \frac{m^{d-1}}{8}\left(\frac{\log m}{2d\log\log m}\right)^{d-1}$ crossing edges.*

Proof. As for Lemma 5.5, the proof proceeds in several stages. The stages are indexed by $(d-1)$-tuples $\boldsymbol{i}$ in $\{0,1,\ldots,\frac{\log m}{d\log\log m}-1\}^{d-1}$. Then, the number of stages is $\left(\frac{\log m}{d\log\log m}\right)^{d-1}$. We show below that each stage contributes at least $\frac{m^{d-1}}{2^{d+2}}$ separate edges to the set of crossing edges, thus proving our lemma.

As in the proof of Lemma 5.5, at each stage vertices are partitioned into megablocks and blocks. In stage $\boldsymbol{i}=(i_1,\ldots,i_{d-1})$, we partition U into $(\log m)^{d(i_1+\cdots+i_{d-1})}$ equal-sized *megablocks* indexed by $\boldsymbol{b}=(b_1,\ldots,b_{d-1})$, where $b_j\in[\log^{d\cdot i_j} m]$ for all $j\in[d-1]$. A vertex v is in a megablock $\boldsymbol{b}$ if $v_{j+1}\in\left[(b_j-1)\frac{m}{\log^{di_j} m}+1, b_j\frac{m}{\log^{di_j} m}\right]$ for each $j\in[d-1]$. So, initially when $\boldsymbol{i}=\vec{0}$, there is only one megablock, and each time $\boldsymbol{i}$ increases by 1 in one coordinate, the volume of the megablocks shrinks by a factor of $\log^d m$.

Each megablock $\boldsymbol{b}$ is further partitioned into $(\log m)^{d(d-1)}$ equal-sized *blocks* indexed by $\boldsymbol{c}\in[\log^d m]^{d-1}$. A vertex v in a megablock $\boldsymbol{b}$ lies in block $\boldsymbol{c}$ if $(v-\boldsymbol{b}_{min})_{j+1}\in\left[(c_j-1)\frac{\ell_j}{\log^d m}+1, c_j\frac{\ell_j}{\log^d m}\right]$ for each $j\in[d-1]$, where $\boldsymbol{b}_{min}$ denotes the smallest vertex contained in megablock $\boldsymbol{b}$ and ℓ_j denotes the length of $\boldsymbol{b}$ in the the j'th dimension. Note that vertices $(1,v_2,\ldots,v_d)$ and $(2,v_2,\ldots,v_d)$ belong to the same (mega)block. At the last stage, each block contains only two vertices (differing by the first coordinate).

Next, we specify the set of crossing edges contributed at each stage. A crossing edge (u,v) in U is said to be *long* in stage $\boldsymbol{i}$ if:

(i) u and v lie in the same megablock, and

(ii) If u lies in block $(c_1,\ldots,c_{d-1})$ and v lies in block $(c'_1,\ldots,c'_{d-1})$, then $c_j<c'_j$ for all $j\in[d-1]$.

We claim that if $\boldsymbol{i}\neq\boldsymbol{i'}$, the sets of long crossing edges in stages $\boldsymbol{i}$ and $\boldsymbol{i'}$ are disjoint. To see this, let j be an index such that $i_j\neq i'_j$; suppose without loss of generality that $i_j<i'_j$. Then, the length of the megablocks in the j'th dimension for stage $\boldsymbol{i'}$ is at most the length of the blocks in the j'th dimension for stage $\boldsymbol{i}$. Hence, condition (ii) above implies that long crossing edges in stage $\boldsymbol{i}$ must have endpoints in different megablocks of stage $\boldsymbol{i'}$, and so violate condition (i) for being a long crossing edge in stage $\boldsymbol{i'}$.

It remains to show that every stage contributes at least $\frac{m^{d-1}}{2^{d+2}}$ long crossing edges. For the sake of contradiction, suppose that the number of long crossing edges at some stage $\boldsymbol{i}$ is $<\frac{m^{d-1}}{2^{d+2}}$. Let $B=m^{d-1}/(\log m)^{d(i_1+\cdots+i_{d-1})}$ be the volume of the megablocks restricted to one of the two hyperplanes. By an averaging argument, at least $\frac{m^{d-1}}{2B}$ megablocks contain $<\frac{B}{2^{d+1}}$ long crossing edges (otherwise, there would be at least $\frac{m^{d-1}}{2^{d+2}}$ long crossing edges). But we show next that if a megablock contains $<\frac{B}{2^{d+1}}$ long crossing edges, then there are $\geq \frac{B \log^d m}{(d-1)2^{2d+2}}$ internal edges with both endpoints inside the megablock. This would imply that the total number of internal edges is $\geq \frac{m^{d-1}}{2B} \cdot \frac{B \log^d m}{(d-1)2^{2d+2}} = \frac{m^{d-1} \log^d m}{(d-1)2^{2d+3}}$, a contradiction.

Suppose then that a megablock contains $< \frac{B}{2^{d+1}}$ long crossing edges. Let Low be the set of vertices in the megablock with each coordinate at most the average value of that coordinate in the megablock, and $High$ the set of vertices with each coordinate greater than the average value of that coordinate. Then $|Low| \geq \frac{B}{2^d}, |High| \geq \frac{B}{2^d}$, and each vertex in Low is comparable to each vertex in $High$. By the bound on the number of long crossing edges, there must exist a set L of at least $\frac{B}{2^{d+1}}$ vertices in Low not incident to any long crossing edge, and a set R of at least $\frac{B}{2^{d+1}}$ vertices in $High$ not incident to any long crossing edges. L lies in the lower hyperplane, R in the upper hyperplane, and each vertex in L is comparable to each vertex in R. Call a crossing edge *short* if it satisfies condition (i), but violates condition (ii) above. A path in U of length at most 2 from a vertex in L to a vertex in R must consist of one internal edge and one short crossing edge. The number of short crossing edges incident to a given vertex v is at most $(d-1)\frac{B}{\log^d m}$, by counting, for each of the $d-1$ block indices, the number of vertices in the megablock that share the value of that block index with v. So, each internal edge helps connect at most $(d-1)\frac{B}{\log^d m}$ pairs of vertices. Since $\frac{B^2}{2^{2d+2}}$ pairs of vertices need to be connected by a path, there must exist at least $\frac{B^2}{2^{2d+2}} \cdot \frac{\log^d m}{(d-1)B} = \frac{B \log^d m}{(d-1)2^{2d+2}}$ internal edges. □

The analogue of Lemma 5.6 in d dimensions (Lemma 5.9) and the rest of the proof of Theorem 5.7 are straightforward generalizations of the 2-dimensional case.

Lemma 5.9. *Let S be a 2-TC-spanner of the directed $[2\ell]\times[m]^{d-1}$ grid. An edge (u, v) in S is called* internal *if $u_1, v_1 \in [\ell]$ or $u_1, v_1 \in \{\ell+1, \ldots, 2\ell\}$, and* crossing *otherwise. If S contains less than $\frac{\ell m^{d-1} \log^d m}{(d-1)2^{2d+3}}$ internal edges, it must contain at least $\geq \ell \frac{m^{d-1}}{8} \left(\frac{\log m}{2d \log\log m}\right)^{d-1}$ crossing edges.*

Proof sketch. We can generalize the proof of Lemma 5.6 in a straightforward way. For each $i \in [\ell]$, instead of matching the lines, we match the hyperplanes $\{i\} \times [m]^{d-1}$ and $\{2\ell - i + 1\} \times [m]^{d-1}$. □

Proof of Theorem 5.7. Assume m is a power of 2 for simplicity. For each step $i \in \{1, \ldots, \frac{1}{2}\log m\}$, partition $\mathcal{H}_{m,d}$ into the following 2^{i-1} equal-sized subgrids: $\{1, \ldots, l_i\} \times [m]^{d-1}$, $\{l_i + 1, \ldots, 2l_i\} \times [m]^{d-1}$, $\ldots$, $\{m - l_i + 1, \ldots, m\} \times [m]^d$ where $l_i = m/2^{i-1}$. For each of these subgrids, define internal and crossing edges as in Lemma 5.9. Now, suppose that there exists a step i such that at least half of the 2^{i-1} subgrids have $\geq \frac{l_i m^{d-1} \log^d m}{(d-1)2^{2d+3}}$ internal edges. Since at a fixed i, the subgrids are disjoint, there are at least $2^{i-2}\frac{l_i m^{d-1} \log^d m}{(d-1)2^{2d+3}} = \frac{m^d \log^d m}{(d-1)2^{2d+4}}$ edges in S, which is enough to prove the theorem. On the other hand, suppose that for every $i \in \{1, \ldots, \frac{1}{2}\log m\}$, at least half of the 2^{i-1} subgrids have $< \frac{l_i m^{d-1} \log^d m}{(d-1)2^{2d+3}}$ internal edges. Then, applying Lemma 5.9, the number of crossing edges in those subgrids is $\geq \frac{l_i m^{d-1}}{8} \left(\frac{\log m}{2d \log\log m}\right)^{d-1}$. Counting over all steps i and for all appropriate subgrids from those steps, the number of edges in S is lower-bounded by $\frac{\log m}{2} \cdot 2^{i-2} \cdot \frac{l_i m^{d-1}}{8} \left(\frac{\log m}{2d \log\log m}\right)^{d-1} = \frac{m^d}{32} \frac{\log^d m}{(2d \log\log m)^{d-1}}$. □

5.2.3 Lower Bound for Small m

Finally, we prove the lower bound on the size of the sparsest TC-spanner of the directed hypergrid, which builds up on the lower bound for the hypercube. This lower bound is especially relevant for small m.

Theorem 5.10. *Any 2-TC-spanner of $\mathcal{H}_{m,d}$ has at least $\Omega\left((m-1)^d 2^{(c+\alpha-1)d}\right)$ edges, where c is the constant from Theorem 1.1 and $\alpha > 0$ satisfies $1 + H_b(\alpha) < c$.*

Proof. Let S be a 2-TC-spanner of $\mathcal{H}_{m,d}$. For each $x \in [m-1]^d$, let U_x be the set $\{y : x \preceq y, |x-y|_\infty \leq 1\}$. By Claim 5.3, S must contain a 2-TC-spanner of U_x for each x. (Recall that in general, a subgraph of a TC-spanner need not be a TC-spanner.)

Call an edge of a 2-TC-spanner of $\mathcal{H}_d$ *long* if it connects x and y with $|x - y| \geq \alpha d$. Claim 5.11 implies that for each $x \in [m-1]^d$, S contains $\Omega(S_2(\mathcal{H}_d))$ long edges that belong to a TC-spanner of U_x, where $S_2(\mathcal{H}_d)$ is the number of edges in the sparsest 2-TC-spanner of the hypercube. Since endpoints of a long edge can agree on at most $(1-\alpha)d$ coordinates, each such edge belongs to at most $2^{(1-\alpha)d}$ subcubes U_x. Thus, S must contain $\Omega\left(\frac{(m-1)^d S_2(\mathcal{H}_d)}{2^{(1-\alpha)d}}\right)$ long edges. The claimed lower bound follows, by Theorem 1.1. □

Claim 5.11. *Let c be the constant from Lemma 3.4 and $\alpha > 0$ be a constant for which $1 + H_b(\alpha) < c$. Call an edge (x, y) in a 2-TC-spanner of $\mathcal{H}_d$* long *if $|x - y| \geq \alpha d$, where $|x - y|$ denotes the Hamming distance between x and y. Then every 2-TC-spanner of $\mathcal{H}_d$ must have $\Omega(S_2(\mathcal{H}_d))$ long edges.*

Proof. The number of pairs (x, y), where $x, y \in \{0,1\}^d$, $x \preceq y$ and $|x - y| < \alpha d$, is at most $2^d \cdot 2^{H_b(\alpha)d} = 2^{(1+H_b(\alpha))d}$, where $H_b(\cdot)$ is the binary entropy function. If $1 + H_b(\alpha) < c$, then every 2-TC-spanner of $\mathcal{H}_d$ must have $\Omega(S_2(\mathcal{H}_d))$ long edges. □

Remark 5.1. We note that the constant α can be improved slightly by adapting the proof of Lemma 3.3 to optimize the number of edges in the 2-TC-spanner of $\mathcal{H}_d$ between endpoints of Hamming distance at least αd, divided by $2^{(1-\alpha)d}$. The calculations are similar to those in Lemma 3.3. ♢

Theorem 1.2 that gives explicit bounds on the size of the sparsest 2-TC-spanners of $\mathcal{H}_{m,d}$ follows from Corollary 5.2, Theorem 5.7 and Theorem 5.10.

References

[1] W. Ackermann. Zum Hilbertshen aufbau der reelen zahlen. *Math. Ann.*, 99:118–133, 1928.

[2] N. Ailon and B. Chazelle. Information theory in property testing and monotonicity testing in higher dimension. *Inf. Comput.*, 204(11):1704–1717, 2006.

[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, K. B. Frikken, and M. Blanton. Dynamic and efficient key management for access hierarchies. In *ACM Conference on Computer and Communications Security*, pages 190–202, 2005.

[5] S. Baswana and S. Sen. Approximate distance oracles for unweighted graphs in expected $\tilde{O}(n^2)$ time. *ACM Transactions on Algorithms*, 2(4):557–577, 2006.

[6] T. Batu, R. Rubinfeld, and P. White. Fast approximate PCPs for multidimensional bin-packing problems. *Inf. Comput.*, 196(1):42–56, 2005.

[7] A. Bhattacharyya, E. Grigorescu, K. Jung, S. Raskhodnikova, and D. P. Woodruff. Transitive-closure spanners. In *Proceedings of the Twentieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 932–941, 2009.

[8] B. Chazelle. Computing on a free tree via complexity-preserving mappings. *Algorithmica*, 2:337–361, 1987.

[9] E. Cohen. Fast algorithms for constructing t-spanners and paths with stretch t. *SIAM J. Comput.*, 28(1):210–236, 1998.

[10] E. Cohen. Polylog-time and near-linear work approximation scheme for undirected shortest paths. *JACM*, 47(1):132–166, 2000.

[11] L. Cowen. Compact routing with minimum stretch. *J. Algorithms*, 38(1):170–183, 2001.

[12] L. Cowen and C. G. Wagner. Compact roundtrip routing in directed networks. *J. Algorithms*, 50(1):79–95, 2004.

[13] Y. Dodis, O. Goldreich, E. Lehman, S. Raskhodnikova, D. Ron, and A. Samorodnitsky. Improved testing algorithms for monotonicity. In *RANDOM*, pages 97–108, 1999.

[14] W. Duckworth and M. Zito. Sparse hypercube 3-spanners. *Discrete Applied Mathematics*, 103(1-3):289–295, 2000.

[15] M. Elkin. Computing almost shortest paths. In *PODC*, pages 53–62, 2001.

[16] F. Ergun, S. Kannan, S. R. Kumar, R. Rubinfeld, and M. Viswanathan. Spot-checkers. *JCSS*, 60(3):717–751, 2000.

[17] G. Fertin, A. L. Liestman, T. C. Shermer, and L. Stacho. Edge-disjoint spanners in cartesian products of graphs. *Discrete Mathematics*, 296(2-3):167–186, 2005.

[18] E. Fischer. On the strength of comparisons in property testing. *Inf. Comput.*, 189(1):107–116, 2004.

[19] E. Fischer, E. Lehman, I. Newman, S. Raskhodnikova, R. Rubinfeld, and A. Samorodnitsky. Monotonicity testing over general poset domains. In *STOC*, pages 474–483, 2002.

[20] O. Goldreich, S. Goldwasser, E. Lehman, D. Ron, and A. Samorodnitsky. Testing monotonicity. *Combinatorica*, 20(3):301–337, 2000.

[21] O. Goldreich, S. Goldwasser, and D. Ron. Property testing and its connection to learning and approximation. *JACM*, 45(4):653–750, 1998.

[22] S. Halevy and E. Kushilevitz. Testing monotonicity over graph products. In *ICALP*, pages 721–732, 2004.

[23] W. Hesse. Directed graphs requiring large numbers of shortcuts. In *SODA*, pages 665–669, 2003.

[24] M.-C. Heydemann, J. G. Peters, and D. Sotteau. Spanners of hypercube-derived networks. *SIAM J. Discrete Math.*, 9(1):37–54, 1996.

[25] C. Laforest, A. L. Liestman, D. Peleg, T. C. Shermer, and D. Sotteau. Edge-disjoint spanners of complete graphs and complete digraphs. *Discrete Mathematics*, 203(1-3):133–159, 1999.

[26] A. Liestman and T. Shermer. Grid and hypercube spanners. *Simon Fraser University Technical Report*, 1991.

[27] A. L. Liestman and T. C. Shermer. Additive spanners for hypercubes. *Parallel Processing Letters*, 1:35–42, 1991.

[28] A. L. Liestman, T. C. Shermer, and C. R. Stolte. Degree-constrained spanners for multidimensional grids. *Discrete Applied Mathematics*, 68(1-2):119–144, 1996.

[29] D. Peleg and A. A. Schäffer. Graph spanners. *Journal of Graph Theory*, 13(1):99–116, 1989.

[30] D. Peleg and J. D. Ullman. An optimal synchronizer for the hypercube. *SIAM J. Comput.*, 18(4):740–747, 1989.

[31] D. Peleg and E. Upfal. A trade-off between space and efficiency for routing tables. *JACM*, 36(3):510–530, 1989.

[32] L. Roditty, M. Thorup, and U. Zwick. Roundtrip spanners and roundtrip routing in directed graphs. In *SODA*, pages 844–851, 2002.

[33] R. Rubinfeld and M. Sudan. Robust characterization of polynomials with applications to program testing. *SIAM Journal on Computing*, 25(2):252–271, 1996.

[34] M. Thorup. On shortcutting digraphs. In *WG*, pages 205–211, 1992.

[35] M. Thorup. Shortcutting planar digraphs. *Combinatorics, Probability & Computing*, 4:287–315, 1995.

[36] M. Thorup. Parallel shortcutting of rooted trees. *J. Algorithms*, 23(1):139–159, 1997.

[37] M. Thorup and U. Zwick. Compact routing schemes. In *ACM Symposium on Parallel Algorithms and Architectures*, pages 1–10, 2001.

[38] M. Thorup and U. Zwick. Approximate distance oracles. *JACM*, 52(1):1–24, 2005.

[39] A. C.-C. Yao. Space-time tradeoff for answering range queries (extended abstract). In *STOC*, pages 128–136, 1982.

A Missing Proofs from Section 3

The following claim was used in the proof of Lemma 3.3.

Claim A.1. *If the variables* $\{x_1, \ldots, x_n, y_1, \ldots, y_n\}$ *are subject to the constraints:*

$$\forall i \in [n] : 0 \leq x_i \leq \alpha \textit{ and } 0 \leq y_i \leq \beta, \quad \sum_{i=1}^{n} x_i \leq t, \quad \sum_{i=1}^{n} y_i \leq t,$$

then $\sum_{i=1}^{n} x_i y_i \leq 2t \cdot \min\{\alpha, \beta\}$.

Proof. Suppose $\sum x_i y_i$ is maximized when $x_i = \alpha_i$ and $y_i = \beta_i$ for all i. Also, suppose without loss of generality that the variables are indexed such that $\alpha_1\beta_1 \leq \alpha_2\beta_2 \leq \cdots \leq \alpha_n\beta_n$. Then $\alpha_i\beta_i > 0$ implies $\alpha_{i+1}\beta_{i+1} = \alpha\beta$. The reason is the following. Notice that either $\alpha_i \leq \alpha_{i+1}$ or $\beta_i \leq \beta_{i+1}$. (Otherwise, $\alpha_i\beta_i > \alpha_{i+1}\beta_{i+1}$.) Assume without loss of generality that $\beta_i \leq \beta_{i+1}$. If we set $\alpha_i' = \alpha_i + \alpha_{i+1} - \min\{\alpha_{i+1} + \alpha_i, \alpha\}$ and $\alpha_{i+1}' = \min\{\alpha_{i+1} + \alpha_i, \alpha\}$ then $\alpha_i'\beta_i + \alpha_{i+1}'\beta_{i+1} \geq \alpha_i\beta_i + \alpha_{i+1}\beta_{i+1}$ while all the constraints in the lemma statement are satisfied. So, we can replace α_i and α_{i+1} with α_i' and α_{i+1}' respectively. Then, from the above definitions, we see that $\alpha_i > 0$ implies $\alpha_{i+1} = \alpha$. Similarly, $\beta_i > 0$ implies $\beta_{i+1} = \beta$. Therefore, if $\alpha_i\beta_i > 0$ then $\alpha_{i+1}\beta_{i+1} = \alpha\beta$. In other words, $\alpha_i\beta_i = 0$ for $1 \leq i < m$ for some $m < n$, $\alpha_m\beta_m > 0$, and $\alpha_i\beta_i = \alpha\beta$ for $m < i \leq n$. Then, $\sum_{i=1}^{n} \alpha_i\beta_i \leq \left\lceil \frac{t}{\max\{\alpha,\beta\}} \right\rceil \alpha\beta \leq \frac{2t}{\max\{\alpha,\beta\}}\alpha\beta = 2t\min\{\alpha, \beta\}$. □

Claim A.2. $s_{i,k,j} = \frac{\binom{d}{k}}{\binom{j-i}{k-i}} \max\left\{\binom{k}{i}, \binom{d-k}{d-j}\right\}$.

Proof. Take the ratio of the two sides:

$$\frac{s_{i,k,j}}{\frac{\binom{d}{k}}{\binom{j-i}{k-i}}\max\left\{\binom{k}{i},\binom{d-k}{d-j}\right\}} = \frac{\binom{d}{i}\binom{d-i}{d-j}\binom{j-i}{k-i}}{\binom{d}{k}\binom{k}{i}\binom{d-k}{d-j}} = \frac{\binom{d}{i}\binom{d-i}{j-i}\binom{j-i}{k-i}}{\binom{d}{k}\binom{k}{i}\binom{d-k}{j-k}} = 1.$$

The first equality follows from the fact that $\max(x,y)\cdot\min(x,y) = x\cdot y$. The last equality can be proved either by expanding the binomial coefficients into factorials, or by realizing that both $\binom{d}{i}\binom{d-i}{j-i}\binom{j-i}{k-i}$ and $\binom{d}{k}\binom{k}{i}\binom{d-k}{j-k}$ count the number of ways i red balls, $j-k$ blue balls, and $k-i$ green balls can be placed into d slots, each of which can hold one ball at most. This completes the proof of the lemma. ☐

Proof of Lemma 3.4. We use the fact that $\binom{n}{cn} = 2^{(H_b(c)-o_n(1))n}$, where "$o_n(1)$" is a function of n that tends to zero as n tends to infinity, and $H_b(p) = -p\log p - (1-p)\log(1-p)$ is the binary entropy function. Substituting $i = \alpha d, j = \beta d$ and $k = \gamma d$ in the resulting expression for s, and taking the logarithm of both sides, we get

$$\log_2 s = \max_{0\le\alpha<\beta\le 1}\min_{\alpha\le\gamma\le\beta}\left[H_b(\gamma) - H_b\left(\frac{\gamma-\alpha}{\beta-\alpha}\right)(\beta-\alpha) + \max\left(H_b\left(\frac{\alpha}{\gamma}\right)\gamma, H_b\left(\frac{1-\beta}{1-\gamma}\right)(1-\gamma)\right)\right] d$$

In other words, $\log_2 s = cd$ where c is a constant. We can check numerically that $c \approx 1.1620$. ☐

B Missing Proofs from Section 4

Lemma B.1. *There is a* 2*-TC-spanner of* $\mathcal{H}_{m,d}$ *with*

$$O\left(d^{2m}\max_{\mathbf{i},\mathbf{j}:\mathbf{j}\succ\mathbf{i}}\min_{\mathbf{k}:\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}}\frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})}\max\{\mathcal{M}(\mathbf{i},\mathbf{k}),\mathcal{N}(\mathbf{k},\mathbf{j})\}\right) \text{ edges.}$$

Proof. Let $v \in \mathbf{i}$ denote that vertex v belongs to level $\mathbf{i}$. Consider the following probabilistic construction that connects comparable vertices at levels $\mathbf{i}$ and $\mathbf{j}$ of $\mathcal{H}_{m,d}$ by paths of length at most 2:

Given levels $\mathbf{i},\mathbf{j} \in [m]^d$, $\mathbf{j} \succ \mathbf{i}$,

1. Initialize the set $E_{\mathbf{i},\mathbf{j}}$ to $\emptyset$.
2. Let $\mathbf{k}_{\mathbf{i},\mathbf{j}} = \operatorname{argmin}_{\mathbf{k}:\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}}\left(\frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})}\max\{\mathcal{M}(\mathbf{i},\mathbf{k}),\mathcal{N}(\mathbf{k},\mathbf{j})\}\right)$.
3. Let $S_{\mathbf{i},\mathbf{j}}$ be a set of $d^m\frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})}$ vertices chosen uniformly at random from the set of $\binom{d}{\mathbf{k}}$ vertices that are in weight level $\mathbf{k} = \mathbf{k}_{\mathbf{i},\mathbf{j}}$.
4. For each vertex $v \in S_{\mathbf{i},\mathbf{j}}$, set $E_{\mathbf{i},\mathbf{j}}$ to $E_{\mathbf{i},\mathbf{j}} \cup \{(x,v) : x \in \mathbf{i} \wedge x \prec v\} \cup \{(v,y) : y \in \mathbf{j} \wedge v \prec y\}$. That is, connect v to all comparable vertices in levels $\mathbf{i}$ and $\mathbf{j}$.
5. Output $E_{\mathbf{i},\mathbf{j}}$.

Claim B.2. *For all* $\mathbf{i} \prec \mathbf{j}$*, with probability at least* $\frac{1}{2}$*,* $E_{\mathbf{i},\mathbf{j}}$ *contains a path of length at most 2 between any pair of vertices* (x,y) *such that* $x \prec y$*,* $x \in \mathbf{i}$*, and* $y \in \mathbf{j}$*.*

Proof. Fix x, y with $x \prec y$, and assume $x \in \mathbf{i}$, and $y \in \mathbf{j}$. We will first show that $Pr_{v \in \mathbf{k}}[x \prec v \prec y] \geq p$, where $p = \frac{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})}{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}$.

Toward that end, notice that there are $\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}$ pairs of comparable vertices (u, w) with $u \in \mathbf{i}, w \in \mathbf{j}$. Each vertex in $S_{\mathbf{i},\mathbf{j}}$ connects exactly $\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})$ pairs of nodes from levels $\mathbf{i}$ and $\mathbf{j}$. It is enough to show that for any such pair (u, w), the number of vertices at level $\mathbf{k}$ that are comparable to both u and v is independent of u, w, i.e., that number only depends on the levels $\mathbf{i}, \mathbf{k}, \mathbf{j}$ and it is therefore the same for all such pairs. To see that, for a vertex $u \in \mathbf{z}$, denote by $T_l(u)$ the set of positions of value l in u. Notice that $|T_l(u)| = z_l$. For $x \prec v \prec y$ it is the case that $T_m(x) \subseteq T_m(v) \subseteq T_m(y)$. Hence there are $\binom{j_m - i_m}{k_m - i_m}$ choices for the m-values in the vector v. Similarly, we must have $T_{m-1}(x) \subseteq T_{m-1}(v) \subseteq T_m(y) \cup T_{m-1}(y)$. Hence there are $\binom{j_m + j_{m-1} - k_m - i_{m-1}}{k_{m-1} - i_{m-1}}$ choices for the values $m-1$ in v. Repeating this process, we obtain that the number of possible v's does not depend on the particular choice of x and y.

Thus the probability that $S_{\mathbf{i},\mathbf{j}}$ does not contain such a vertex v with $x \prec v \prec y$ is $(1-p)^{d^m/p} \leq e^{-d^m}$.

The number of comparable pairs (x, y) is at most m^{2d}, and by the union bound, the probability that there exists (x, y) such that there is no $v \in S_{\mathbf{i},\mathbf{j}}$ with $x \prec v \prec y$ is at most $m^{2d}e^{-d^m} < 1/2$. □

So, for every $\mathbf{i}$ and $\mathbf{j}$, there exists a choice of $S_{\mathbf{i},\mathbf{j}}$ such that comparable pairs from the two weight levels are connected by a path of length at most 2. Let $E^*_{\mathbf{i},\mathbf{j}}$ be the set of edges returned by the algorithm when this $S_{\mathbf{i},\mathbf{j}}$ is chosen. We set $E = \bigcup_{i<j} E^*_{\mathbf{i},\mathbf{j}}$. By Claim B.2, $([m]^d, E)$ is a 2-TC-spanner of $\mathcal{H}_{m,d}$.

Now, we show that the size of E is as claimed in the lemma statement. The main observation is that in step (4), for any specific $v \in S_{\mathbf{i},\mathbf{j}}$, $|\{(x, v) : x \in \mathbf{i} \wedge x \prec v\} \cup \{(v, y) : y \in \mathbf{j} \wedge v \prec y\}|$ is exactly $\mathcal{M}(\mathbf{i},\mathbf{k}) + \mathcal{N}(\mathbf{k},\mathbf{j})$.

The claimed bound follows since $|E| = \sum_{\mathbf{j} \succ \mathbf{i}} |E^*_{\mathbf{i},\mathbf{j}}|$, where the sum has d^m terms.

□

Lemma B.3. *Any 2-TC-spanner of $\mathcal{H}_{m,d}$ has*

$$\Omega\left(\max_{\mathbf{i},\mathbf{j}:\mathbf{j}\succ\mathbf{i}} \min_{\mathbf{k}:\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}} \frac{\mathcal{M}(\mathbf{i},\mathbf{j})\binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k})\mathcal{N}(\mathbf{k},\mathbf{j})} \max\{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}\right) \textit{ edges}.$$

Proof. Let S be a 2-TC-spanner for $\mathcal{H}_{m,d}$. We will count the edges in S that occur on paths connecting two particular levels of $\mathcal{H}_{m,d}$. Let $P_{\mathbf{i},\mathbf{j}} = \{(v_1, v_2) : v_1 \in \mathbf{i}, v_2 \in \mathbf{j}, v_1 \prec v_2\}$. We will lower bound $e^*_{\mathbf{i},\mathbf{j}}$, the number of edges in the paths of length at most 2 in S, that connect the pairs $P_{\mathbf{i},\mathbf{j}}$. Notice that $|P(\mathbf{i},\mathbf{j})| = \binom{d}{\mathbf{j}}\mathcal{M}(\mathbf{i},\mathbf{j})$.

Let $e_{\mathbf{k},\boldsymbol{\ell}}$ denote the number of edges in S that connect vertices in level $\mathbf{k}$ to vertices in level $\boldsymbol{\ell}$. Then

$$e^*_{\mathbf{i},\mathbf{j}} = e_{\mathbf{i},\mathbf{j}} + \sum_{\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}} (e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}). \tag{2}$$

We say that a vertex v *covers* a pair of vertices (v_1, v_2) if S contains the edges (v_1, v) and (v, v_2) or, for the special case $v = v_1$, if S contains (v_1, v_2). Let $V^{(\mathbf{k})}_{\mathbf{i},\mathbf{j}}$ be the set of vertices in level $\mathbf{k}$ that cover pairs in $P_{\mathbf{i},\mathbf{j}}$. Let $\alpha_{\mathbf{k}}$ be the fraction of pairs in $P_{\mathbf{i},\mathbf{j}}$ that are covered by the vertices in $V^{(\mathbf{k})}_{\mathbf{i},\mathbf{j}}$. Since each pair in $P_{\mathbf{i},\mathbf{j}}$ must be covered by a vertex in levels $\mathbf{k}$ with $\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}$, we must have

$$\sum_{\mathbf{i}\prec\mathbf{k}\prec\mathbf{j}} \alpha_{\mathbf{k}} \geq 1. \tag{3}$$

For any vertex $v \in V^{(\mathbf{k})}_{\mathbf{i},\mathbf{j}}$, let in_v be the number of incoming edges from vertices of level $\mathbf{i}$ incident to v and let out_v be the number of outgoing edges to vertices of level j incident to v. For each level $\mathbf{k}$ with

$\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}$, since each vertex $v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}$ covers $in_v \cdot out_v$ pairs,

$$\sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} in_v \cdot out_v \geq \alpha_{\mathbf{k}} |P_{\mathbf{i},\mathbf{j}}| \geq \alpha_{\mathbf{k}} \mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}}. \tag{4}$$

We upper bound $\sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} in_v \cdot out_v$ as a function of $e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}$, and then use Equation (4) to lower bound $e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}$.

For all $\mathbf{k}$ with $\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}$, variables in_v and out_v satisfy the following constraints:

$$\sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} in_v \leq e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}, \quad \sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} out_v \leq e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}, \quad in_v \leq \mathcal{M}(\mathbf{i},\mathbf{k}) \, \forall v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}, \quad out_v \leq \mathcal{N}(\mathbf{k},\mathbf{j}) \, \forall v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}.$$

The last two constraints hold because in_v and out_v count the number of edges to a vertex of level $\mathbf{k}$ from from vertices of level $\mathbf{i}$ and from a vertex of level $\mathbf{k}$ to vertices of level $\mathbf{j}$, respectively. We want to maximize $\sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} in_v \cdot out_v$ subject to the above constraints. Claim A.1 bounds the sum by proving that the maximum occurs when $in_v = \mathcal{M}(\mathbf{i},\mathbf{k})$ and $out_v = \mathcal{N}(\mathbf{k},\mathbf{j})$ for as many v as possible, subject to the remaining constraints. It gives us, for all $\mathbf{k}$ with $\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}$,

$$\sum_{v \in V_{\mathbf{i},\mathbf{j}}^{(\mathbf{k})}} in_v \cdot out_v \leq 2(e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}) \min \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\} .$$

From Equation (4), $e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}} \geq \frac{1}{2} \alpha_{\mathbf{k}} \mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}} \frac{1}{\min \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}}$ for all $\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}$. Applying Equations (2) and (3), we get

$$\begin{aligned}
e^*_{\mathbf{i},\mathbf{j}} &= e_{\mathbf{i},\mathbf{j}} + \sum_{\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}} (e_{\mathbf{i},\mathbf{k}} + e_{\mathbf{k},\mathbf{j}}) \\
&\geq \frac{1}{2} \sum_{\mathbf{k}} \alpha_{\mathbf{k}} \frac{1}{\min \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}} \mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}} \\
&\geq \frac{1}{2} \min_{\mathbf{k}} \frac{1}{\min \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}} \mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}} \\
&= \frac{1}{2} \min_{\mathbf{k}} \frac{1}{\mathcal{M}(\mathbf{i},\mathbf{k}) \mathcal{N}(\mathbf{k},\mathbf{j})} \mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}} \max \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}.
\end{aligned}$$

Since this holds for arbitrary $\mathbf{i}$ and $\mathbf{j}$, the number of edges in the 2-TC-spanner is

$$|S| \geq \frac{1}{2} \max_{\mathbf{i},\mathbf{j}: \mathbf{i} \prec \mathbf{j}} \min_{\mathbf{i} \prec \mathbf{k} \prec \mathbf{j}} \frac{\mathcal{M}(\mathbf{i},\mathbf{j}) \binom{d}{\mathbf{j}}}{\mathcal{M}(\mathbf{i},\mathbf{k}) \mathcal{N}(\mathbf{k},\mathbf{j})} \max \{\mathcal{M}(\mathbf{i},\mathbf{k}), \mathcal{N}(\mathbf{k},\mathbf{j})\}.$$

□

C Missing Proofs from Section 5

Claim C.1. *For all $m \geq 3$, the directed line $\mathcal{H}_{m,1}$ has a 2-TC-spanner with at most $m \log m - m$ edges.*

Proof. Construct graph S on vertex set $[m]$ recursively. First, define the middle node $v_{mid} = \lceil \frac{m}{2} \rceil$. Add edges (v, v_{mid}) for all nodes $v < v_{mid}$ and edges (v_{mid}, v) for all nodes $v > v_{mid}$. Then recurse on the two line segments resulting from removing v_{mid} from the current line. Proceed until each line segment contains exactly one node. This construction is implicit in, e.g., [13].

S is a 2-TC-spanner for the line $\mathcal{H}_{m,1}$, since every pair of nodes $u, v \in [m]$ is connected by a path of length at most 2 via a middle node. This happens in the stage of the recursion where u and v are separated into different line segments, or one of these two nodes is removed.

There are $t = \lfloor \log m \rfloor$ stages of the recursion, and in each stage $i \in [t]$ each node that is not removed by the end of the this stage connects to the middle node in its current line segment. Since 2^{i-1} nodes are removed in the ith stage, exactly $m - (2^i - 1)$ edges are added in that stage. Thus, the total number of edges in S is $m \cdot t - (2^{t+1} - t - 2) \leq m \log m - m$. The last inequality holds for $m \geq 3$. □

ECCC ISSN 1433-8092
http://eccc.hpi-web.de/