

Diameter of Polyhedra: Limits of Abstraction

Friedrich Eisenbrand

Institute of Mathematics, EPFL, Lausanne, Switzerland
email: friedrich.eisenbrand@epfl.ch

Nicolai Hähnle

Institute of Mathematics, EPFL, Lausanne, Switzerland
email: nicolai.haehnle@epfl.ch

Alexander Razborov

University of Chicago
email: razborov@cs.uchicago.edu

Thomas Rothvoß

Institute of Mathematics, EPFL, Lausanne, Switzerland
email: thomas.rothvoss@epfl.ch

We investigate the diameter of a natural abstraction of the 1-skeleton of polyhedra. Even if this abstraction is more general than other abstractions previously studied in the literature, known upper bounds on the diameter of polyhedra continue to hold here. On the other hand, we show that this abstraction has its limits by providing an almost quadratic lower bound.

Key words: convex geometry; disjoint coverings; Hirsch conjecture; polyhedra

MSC2000 Subject Classification: Primary: 52B05; Secondary: 05B40

OR/MS subject classification: Primary: Linear Programming theory; Secondary: Combinatorics

History: Received: October 13, 2009; Revised: October 23, 2010.

1. Introduction. One of the most prominent mysteries in convex geometry is the question whether the diameter of polyhedra is polynomial in the number of its facets or not. If the largest diameter of a d -dimensional and possibly unbounded polyhedron with n facets is denoted by $\Delta_u(d, n)$, then the best known upper bound is $\Delta_u(d, n) \leq n^{1+\log d}$, shown by Kalai and Kleitman [13]. For a long time, the best known lower bound was $\Delta_u(d, n) \geq n - d + \lfloor d/5 \rfloor$, due to Klee and Walkup [15]. Recently, Santos [20] has given a lower bound of $\Delta_u(d, n) \geq (1 + \epsilon)(n - d)$, where d, ϵ are fixed and n is arbitrarily large. The gap which is left open here is huge, even after decades of intensive research on this problem.

Interestingly, the above upper bound holds also for simple *combinatorial abstractions of polyhedra* by which term we (loosely) mean a rigorously defined set of *purely combinatorial* properties of the polyhedra in question that are strong enough to allow non-trivial conclusions about its geometry. In the quest of bounding Δ_u one can restrict attention to non-degenerate polyhedra (we call a polyhedron *non-degenerate* if each vertex is contained in exactly d facets) since, by perturbation, any polyhedron can be turned into a non-degenerate polyhedron whose diameter is at least as large as the one of the original polyhedron. For this reason we also allow ourselves this simplifying assumption of non-degeneracy (all our results, though, perfectly hold without it).

Combinatorial abstractions have been studied in the literature for a long time [11, 2, 1]. The subject of this paper is a simple *base abstraction* which is defined by *one single feature*, common to all previously studied abstractions from which lower and upper bounds have been previously derived. As an extra evidence (besides simplicity) that our framework is quite natural, we give for it three different descriptions that all turn out to be pairwise equivalent.

Even if our abstraction is more general than previously considered ones, we nonetheless show that all known upper bounds do hold here with natural and simple proofs. On the other hand, we prove an almost quadratic lower bound on the diameter in this abstraction, and this constitutes the main concrete result of this paper.

While only one feature of the previously studied abstractions suffices to derive the best known upper bounds, our lower bound also shows the limits of this natural base abstraction for the purpose of proving *linear* upper bounds on the diameter. To prove such a bound, more features of the geometry of polyhedra will have to be understood and used than the single one that we identify here. Let us, however, note that a polynomial (or even quadratic!) upper bound in this framework still remains a possibility.

In the first description (see Section 2 for equivalent definitions), our *base abstraction* is a connected graph $G = (V, E)$. Here¹ $V \subseteq \binom{[n]}{d}$ and the edges E of G are such that the following connectivity condition holds:

- i) For each $u, v \in V$ there exists a path connecting u and v whose intermediate vertices all contain $u \cap v$.

Let $\mathcal{B}_{d,n}$ be the set of all graphs G with the above property; the largest diameter of a graph in $\mathcal{B}_{d,n}$ will be denoted by $D(d, n)$. We call d the *dimension* and n the *number of facets* of the abstraction.

Before we proceed, let us understand why this class contains the 1-skeletons of non-degenerate polyhedra in dimension d having n facets. In this setting, each vertex is uniquely determined by the d facets in which it is contained. If the facets are named $\{1, \dots, n\}$, then a vertex is uniquely determined by a d -element subset of $\{1, \dots, n\}$. Furthermore, for every pair of vertices u, v there exists a path which does not leave the minimal face in which both u and v are contained. This is reflected in condition i). Thus if $\Delta_u(d, n)$ is the maximum diameter of a non-degenerate polyhedron with n facets in dimension d , then $\Delta_u(d, n) \leq D(d, n)$ holds.

Our main result is a super-linear lower bound on $D(d, n)$, namely $D(n/4, n) = \Omega(n^2 / \log n)$. The non-trivial construction relies on the notion of *disjoint covering designs* and in order to prove the existence of such designs with desired parameters we use the Lovász Local Lemma.

At the same time the bound of Kalai and Kleitman [13], $\Delta_u(d, n) \leq n^{1+\log d}$, as well as the upper bound of Larman [16], $\Delta_u(d, n) \leq 2^{d-1} \cdot n$, which is linear when the dimension is fixed, continue to hold for a base abstraction. While the first bound is merely an adaption of the proof in [13], our proof of the second bound is much simpler than the one which was proved for polyhedra in [16].

We strongly believe that the study of abstractions, asymptotic lower and upper bounds for those and the development of algorithms to compute bounds for fixed parameters d and n should receive more attention since they can help to understand the important features of the geometry of polyhedra that may help to improve the state-of-the-art of the diameter question.

Related abstractions Abstractions of polyhedra were already considered by Adler, Dantzig, and Murty [2] who studied *abstract polytopes*. Here, additionally to the condition i) of our base abstraction, the graph has to satisfy the following two conditions:

- ii) The edge (u, v) is present if and only if $|u \cap v| = d - 1$.
- iii) Each $e \in \binom{[n]}{d-1}$ is either contained in two vertices of G or it is not contained in any vertex of G .

Notice that this is an abstraction of non-degenerate d -dimensional *polytopes* with n facets, since condition iii) only holds for bounded polyhedra. Adler and Dantzig [3] showed that the diameter of abstract polytopes is bounded by $n - d$ if $n - d \leq 5$. This shows that the *d -step conjecture*² is also true up to dimension 5 for abstract polytopes. Klee and Walkup [15] proved that the *d -step conjecture* is true if and only if the famous *Hirsch conjecture*³ is true. Klee and Walkup [15] were the first to prove that the *d -step conjecture* is true up to dimension 5.

¹ $\binom{[n]}{d}$ is the family of all d -element subsets of $[n] = \{1, \dots, n\}$

²the *d -step conjecture* states that the diameter of a d -dimensional polytope with $2d$ facets is bounded by d

³the *Hirsch conjecture* states that the diameter of a d -dimensional polytope with n facets is bounded by $n - d$

A big advantage of any abstraction is that bounds on the diameter of abstract polytopes for fixed dimension d and number of facets n can be automatically checked with a computer. For example, Bremner and Schewe describe an automatic approach to check the d -step conjecture using a different abstraction based on oriented matroids [5]. They were able to verify this conjecture up to dimension 6. However, a recent construction of Santos [20] shows that the Hirsch conjecture and thus also the d -step conjecture are false in general.

The situation for lower bounds on the diameter of abstract polytopes in the setting of [2] is as follows. Mani and Walkup [17] have provided an example of an abstract polytope with $d = 12$ and $n = 24$ whose diameter is larger than 12, see also [14], and the construction of Santos [20] naturally applies to abstract polytopes as well. But superlinear lower bounds on the diameter of abstract polytopes are still not known.

Kalai [11] considered the abstraction in which, additionally to our base abstraction, only ii) has to hold. He called his abstraction *ultraconnected* set systems and showed that the upper bound [13] can also be proved in this setting. As demonstrated by our work, the condition ii) is not necessary and this bound, together with the linear bound in fixed dimension of Larman [16] also holds for the base abstraction which does not require condition ii).

We also want to mention recent progress in the study of abstractions of linear optimization problems. Kalai [10] and Matoušek, Sharir, and Welzl [18] were able to give subexponential upper bounds on the expected running time of randomized, purely combinatorial algorithms for linear programming. This prompted the study of various types of abstract optimization problems [8]. For unique sink orientations of cubes, which capture much of the interesting structure of these abstractions, Schurr and Szabó [21] proved a non-trivial lower bound of $\Omega(\frac{n^2}{\log n})$ for the running time of any deterministic algorithm, while the best known upper bound even in the acyclic case is still an expected running time of $O(n^3 e^{2\sqrt{n}})$ [9].

2. Base abstraction, connected layer families and interval evaluations. Let $G = (V, E) \in \mathcal{B}_{d,n}$ be a graph of our base abstraction. Recall that this means that $V \subseteq \binom{[n]}{d}$ and that the edges are such that the connectivity condition i) holds. Denote the length of a shortest path between two vertices u and v by $\text{dist}(u, v)$. Suppose that the diameter of G is the shortest path between the nodes s and t and suppose that $\text{dist}(s, t) = \ell$. If we label each vertex $v \in V$ with its distance to s , then we obtain non-empty subsets $\mathcal{L}_i \subseteq \binom{[n]}{d}$ for $i = 0, \dots, \ell$ with $\mathcal{L}_i = \{v \in V \mid \text{dist}(s, v) = i\}$. The sets \mathcal{L}_i satisfy the following conditions.

- a) *Disjointness*: for all $0 \leq i \neq j \leq \ell$, $\mathcal{L}_i \cap \mathcal{L}_j = \emptyset$;
- b) *Connectivity*: for all $0 \leq i < j < k \leq \ell$ and $u \in \mathcal{L}_i$, $v \in \mathcal{L}_k$ there is a $w \in \mathcal{L}_j$ such that $u \cap v \subseteq w$.

While condition a) clearly holds, let us argue why condition b) is also satisfied. Since we have the connectivity condition i) from our base abstraction, there exists a path from $u \in \mathcal{L}_i$ to $v \in \mathcal{L}_k$ whose intermediate vertices all contain the intersection $u \cap v$. These intermediate vertices have distance labels. Clearly, all distance labels between i and k must appear on this path which means in particular that the label j appears on this path. This shows that \mathcal{L}_j contains a vertex w containing $u \cap v$.

A sequence of non-empty sets $\mathcal{L}_i \subseteq \binom{[n]}{d}$, $i = 0, \dots, \ell$ which satisfies a) and b) is called a *connected layer family*, where the sets \mathcal{L}_i are referred to as *layers*. The elements of the ground set $\{1, \dots, n\}$ are the *symbols* of the connected layer family (they correspond to facets) and d is its *dimension*. The elements of each layer (subsets of $\{1, \dots, n\}$ of cardinality d) are again referred to as *vertices* of the layer. The *height* of this connected layer family is $\ell + 1$. We have argued above that a base abstraction of diameter ℓ naturally yields a connected layer family of height $\ell + 1$.

On the other hand, a connected layer family of height $\ell + 1$ yields a base abstraction of diameter ℓ by connecting all pairs of vertices u, v where $u \in \mathcal{L}_i$ and $v \in \mathcal{L}_{i+1}$ or $u \in \mathcal{L}_i$ and $v \in \mathcal{L}_i$. We therefore have the following result.

THEOREM 2.1 *The maximum diameter of a d -dimensional base abstraction with n symbols is the largest height of a d -dimensional connected layer family with n symbols minus one.*

The following is an example of a 2-dimensional connected layer family with $n = 6$ symbols and 7 layers. A set of symbols w is *active* on a layer \mathcal{L}_i if there exists a vertex of \mathcal{L}_i containing w . In our example,

we highlight the symbol 4 and, due to condition b) the layers on which 4 is active are consecutive. This holds for each symbol and thus the following example is a 2-dimensional connected layer family.

$$\begin{aligned}\mathcal{L}_0 &= \{\{1, 6\}\} \\ \mathcal{L}_1 &= \{\{1, 2\}, \{2, 6\}\} \\ \mathcal{L}_2 &= \{\{2, 5\}, \{1, 3\}, \{4, 6\}\} \\ \mathcal{L}_3 &= \{\{2, 4\}, \{1, 5\}, \{3, 6\}\} \\ \mathcal{L}_4 &= \{\{2, 3\}, \{1, 4\}, \{5, 6\}\} \\ \mathcal{L}_5 &= \{\{4, 5\}, \{3, 4\}\} \\ \mathcal{L}_6 &= \{\{3, 5\}\}\end{aligned}$$

Let us provide yet another visualization of our abstraction. An *integer interval* is a (possibly empty) subset $I \subseteq \mathbb{Z}$ of the form $\{x \in \mathbb{Z} \mid a \leq x \leq b\}$. An *interval evaluation on n symbols of dimension d* is a mapping ϕ that maps

$$\binom{[n]}{\leq d} := \{f \subseteq [n] \mid |f| \leq d\}$$

into the set Int of all integer intervals such that the following properties hold:

- Anti-monotonicity** $f \subseteq g \implies \phi(g) \subseteq \phi(f)$;
Continuousness $|f| < d \implies \phi(f) \subseteq \bigcup_{g \supset f} \phi(g)$;
Dimensionality restriction $|v| = d \implies |\phi(v)| \leq 1$.

The *height* of an interval evaluation ϕ is $|\phi(\emptyset)|$.

THEOREM 2.2 *The largest height of a d -dimensional connected layer family with n symbols is equal to the largest height of an interval evaluation $\phi : \binom{[n]}{\leq d} \rightarrow \text{Int}$.*

PROOF. Given a connected layer family $\mathcal{L}_0, \mathcal{L}_1, \dots, \mathcal{L}_\ell$, we let

$$\phi(f) := \{i \mid \exists v \in \mathcal{L}_i : v \supseteq f\}.$$

The fact that ϕ is an integer interval readily follows from the connectivity condition while all other properties in the definition of an interval evaluation are obvious.

In the opposite direction, given an interval evaluation ϕ , we may assume w.l.o.g. that $\phi(\emptyset) = \{0, 1, \dots, \ell\}$. Now we let

$$\mathcal{L}_i := \{v \in \binom{[n]}{d} \mid \phi(v) = \{i\}\}.$$

\mathcal{L}_i are obviously disjoint. Using continuousness, we prove by induction on $\nu = 0, \dots, d$ that $\phi(\emptyset) \subseteq \bigcup_{f \in \binom{[n]}{\nu}} \phi(f)$ and then anti-monotonicity implies that actually $\phi(\emptyset) = \bigcup_{f \in \binom{[n]}{\nu}} \phi(f)$. This, along with the dimensionality restriction, implies that \mathcal{L}_i are non-empty. Finally, if $0 \leq i < j < k \leq \ell$ and $u \in \mathcal{L}_i, v \in \mathcal{L}_k$ then by anti-monotonicity we conclude that $\phi(u \cap v) \supseteq [i, k]$, hence $j \in \phi(u \cap v)$. Arguing by induction as before, we find $w \supseteq u \cap v$ with $|w| = d$ and $\phi(w) = \{j\}$. This gives connectivity. \square

3. Upper bounds on the diameter of the base abstraction. Before we prove upper bounds on $D(d, n)$, we need an operation on connected layer families. This operation is motivated by the fact that the face of a polyhedron is again a polyhedron. Let $s \in \{1, 2, \dots, n\}$ be a symbol in a connected layer family. The *induction on s* is the following operation.

- (i) Remove all vertices from the connected layer family which do not contain s .
- (ii) Remove s from all vertices.
- (iii) Remove empty layers (and relabel non-empty labels starting from 0).

This operation looks particularly natural in the interval representation (Theorem 2.2): $\phi : \binom{[n]}{d} \rightarrow \text{Int}$ gives rise to the induced interval representation $\phi_s : \binom{[n] \setminus \{s\}}{d-1} \rightarrow \text{Int}$ defined simply by $\phi_s(f) := \phi(f \cup \{s\})$.

Either way, the next lemma follows directly from definitions.

LEMMA 3.1 *Given a d -dimensional connected layer family with n symbols, induction on any symbol s results in a $(d - 1)$ -dimensional connected layer family with $n - 1$ symbols.*

Induction on 4 of the layered family in our example above, results in the following connected layer family.

$$\begin{aligned}\mathcal{L}'_0 &= \{\{6\}\} \\ \mathcal{L}'_1 &= \{\{2\}\} \\ \mathcal{L}'_2 &= \{\{1\}\} \\ \mathcal{L}'_3 &= \{\{5\}, \{3\}\}\end{aligned}$$

The quasi-polynomial bound $\Delta_u(d, n) \leq n^{2+\log d}$ of [13] is up to now the best known bound on $\Delta_u(d, n)$. We prove this in the setting of our base abstraction by showing that this is also an upper bound on the height of a d -dimensional connected layer family with n symbols. All logarithms are with respect to base 2.




Figure 1: Illustration of the proof of Theorem 3.1

THEOREM 3.1 *The maximum diameter $D(d, n)$ of a d -dimensional base abstraction with n symbols is bounded by $n^{1+\log d} - 1$.*

PROOF. By Theorem 2.1 it is enough to show that the maximal height $h(d, n)$ of a d -dimensional connected layer family with n symbols is bounded by $n^{1+\log d}$. To this end, let $\mathcal{L}_0, \dots, \mathcal{L}_\ell$ be a connected layer family. Let $\ell_1 \geq -1$ be maximal such that the union of the vertices in $\mathcal{L}_0, \dots, \mathcal{L}_{\ell_1}$ contains at most $\lfloor n/2 \rfloor$ many symbols. Let $\ell_2 \leq \ell + 1$ be minimal such that the union of the vertices in $\mathcal{L}_{\ell_2}, \dots, \mathcal{L}_\ell$ contains at most $\lfloor n/2 \rfloor$ symbols. Since $|\mathcal{L}_0 \cup \dots \cup \mathcal{L}_{\ell_1+1}|, |\mathcal{L}_{\ell_2-1} \cup \dots \cup \mathcal{L}_\ell| > n/2$, there exists a symbol $s \in \{1, \dots, n\}$ belonging to both of these sets, and this s is active on all layers $\mathcal{L}_{\ell_1+1}, \dots, \mathcal{L}_{\ell_2-1}$, see Figure 1.

Now we observe that $\mathcal{L}_0, \dots, \mathcal{L}_{\ell_1}$ and $\mathcal{L}_{\ell_2}, \dots, \mathcal{L}_\ell$ are d -dimensional connected layer families with at most $\lfloor n/2 \rfloor$ symbols each. After inducing on the symbol s , which is active on all layers $\mathcal{L}_{\ell_1+1}, \dots, \mathcal{L}_{\ell_2-1}$ we obtain a $d - 1$ -dimensional connected layer family with $n - 1$ symbols of height at least $\ell_2 - \ell_1 - 1$. Thus, we get the recursion

$$h(d, n) \leq 2 \cdot h(d, \lfloor n/2 \rfloor) + h(d - 1, n - 1). \quad (1)$$

The bound is then proved by induction on n . Note that $h(1, n) = n$ and $h(d, n) = 0$ if $d > n$. Suppose now that $d, n \geq 2$. Applying (1) repeatedly, we obtain

$$\begin{aligned}h(d, n) &\leq 2 \cdot h(d, \lfloor n/2 \rfloor) + h(d - 1, n) \\ &\leq 2 \cdot \sum_{i=2}^d h(i, \lfloor n/2 \rfloor) + h(1, n).\end{aligned}$$

By induction, this is bounded by

$$\begin{aligned}h(d, n) &\leq 2(d - 1)(2d)^{\log n - 1} + n \\ &= (2d)^{\log n - 1} (2(d - 1) + n / ((2d)^{\log n - 1})) \\ &\leq (2d)^{\log n}\end{aligned}$$

In the last inequality we have used $d \geq 2$ and thus $(2d)^{\log n - 1} \geq n^2/4$. Since $n \geq 2$ one can conclude $n / ((2d)^{\log n - 1}) \leq 4/n \leq 2$. \square

REMARK 3.1 Notice that our bound on $D(d, n)$ is slightly better than the bound on $\Delta_u(d, n)$ in [13]. Kalai [12] pointed out that the $n^{2+\log d}$ bound can be improved to $n^{1+\log d}$ which matches the upper bound for the diameter of base abstractions that we provide above.

3.1 A linear bound in fixed dimension. Next we provide a linear upper bound on $D(d, n)$ in the case where the dimension d is fixed. The original proof for polyhedra is due to Larman [16]. We would like to point out that the proof in our setting is much simpler than the original one. Our constant however is slightly worse, because our base case of induction is weaker. Larman’s bound on the diameter of polyhedra is $\Delta_u(d, n) \leq 2^{d-3}n$.

THEOREM 3.2 The maximum diameter $D(d, n)$ of a d -dimensional base abstraction with n symbols is bounded by $2^{d-1} \cdot n - 1$.

PROOF. Let $F = (\mathcal{L}_0, \dots, \mathcal{L}_\ell)$ be a connected layer family with n symbols of dimension d . We prove the claim by induction on d . For $d = 1$, one has at most n vertices which implies that the height $h(1, n)$ is bounded by n .

For a symbol s , let $[L(s), U(s)] \subseteq \{0, \dots, \ell\}$ be the interval which corresponds to the layers on which s is active, i.e.,

$$\begin{aligned} L(s) &= \min\{i \mid \exists u \in \mathcal{L}_i : s \in u\} \\ U(s) &= \max\{i \mid \exists u \in \mathcal{L}_i : s \in u\} \end{aligned}$$

(in the interval representation, $[L(s), U(s)]$ is simply $\phi(\{s\})$).

Next we define a sequence s_i of symbols. The symbol s_1 is the one whose interval of active layers contains the starting layer \mathcal{L}_0 and reaches farthest among those whose interval starts at 0. In other words, $s_1 = \operatorname{argmax}_{s \in \{1, \dots, n\}} \{U(s) \mid L(s) = 0\}$. If s_1, \dots, s_j are given and $U(s_j) < \ell$, the symbol s_{j+1} is the one that reaches farthest among all symbols that are active in the layer $U(s_j) + 1$. In other words

$$s_{j+1} = \operatorname{argmax}_{s \in \{1, \dots, n\}} \{U(s) \mid L(s) \leq U(s_j) + 1\}.$$

The starting points of these intervals hash the connected layer family $F = (\mathcal{L}_0, \dots, \mathcal{L}_\ell)$ into connected layer families $F_1 = (\mathcal{L}_{L(s_1)}, \dots, \mathcal{L}_{U(s_1)})$ and $F_i = (\mathcal{L}_{U(s_{i-1})+1}, \dots, \mathcal{L}_{U(s_i)})$ for $i = 2, \dots, k$, see Figure 2. The important observation due to our construction is the following:

the symbols in F_i and F_j are disjoint if $|i - j| \geq 2$

(otherwise, any symbol in their intersection should have been chosen instead of $s_{\min(i,j)+1}$). Let n_i denote the number of symbols in F_i . The above observation implies $\sum_{i=1}^k n_i \leq 2 \cdot n$.

Since the symbol s_i is active on each layer of F_i , induction on s_i leaves the height unchanged. This implies that

$$\begin{aligned} \operatorname{height}(F) &= \sum_{i=1}^k \operatorname{height}(F_i) \leq \sum_{i=1}^k h(d-1, n_i) \\ &\leq \sum_{i=1}^k 2^{d-2} n_i \leq 2^{d-1} n \end{aligned}$$

□

4. A lower bound on the diameter of the base abstraction. Our goal is to construct a d -dimensional connected layer family with n symbols that has a large number of layers. The difficult condition to meet is connectivity b). Our first idea is to satisfy this condition by enforcing that each $(d-1)$ -subset of the symbols is contained in a vertex of *each layer* of the connected layer family (and then we boost this construction). If this holds, then b) is clearly satisfied. How many layers can a d -dimensional connected layer family have that satisfies the property above? This question is related to the question of *covering designs*, a classical topic in combinatorics.

Let X be a ground set of size n . Extending previously used notation, by $\binom{X}{d}$ we denote the family of all its d -element subsets. Fix natural numbers $r < k < n$. An element $b \in \binom{X}{k}$ will be called a *block*.




Figure 2: Illustration of the proof of Theorem 3.2. Suppose the x -axis denotes level indices. The black lines denote intervals $[L(s_i), U(s_i)]$ where s_i is active, while gray lines denotes the layers contained in families F_i .

A collection C of blocks is called an (n, k, r) -covering design or simply a *covering* if every $a \in \binom{X}{r}$ is contained in at least one of the blocks in C . Well studied is the smallest size (number of blocks) of an (n, k, r) -covering. Rdl [19] for example, proved a long standing conjecture of Erdős and Hanani [7] on the asymptotic size of covering designs for fixed k and r .

Now the layers (covering designs) have to be disjoint from each other. This means that we need *disjoint families* of $(n, d, d - 1)$ -covering designs. We want as many disjoint covering designs as possible.

The question of how many disjoint $(n, d, d - 1)$ -coverings exist has, to the best of our knowledge, not been studied before. Since every $a \in \binom{X}{d-1}$ can be covered by precisely $(n - d + 1)$ blocks $b \in \binom{X}{d}$, $(n - d + 1)$ is an obvious upper bound on their number. The next section is devoted to proving an almost matching $\Omega(n/\log n)$ lower bound based upon a simple application of Lovász Local Lemma.

4.1 Disjoint covering designs. We denote the maximum size of a family of disjoint (n, k, r) -coverings by $DC(n, k, r)$, and, given our motivations, we are mostly interested in the case $k = r + 1$. As we already observed, $D(n, r + 1, r) \leq n - r$; on the other hand we have the following theorem.

THEOREM 4.1 *Let $1 \leq r < n$. Then $DC(n, r + 1, r) \geq \lfloor \frac{n-r}{3 \ln n} \rfloor$.*

PROOF. Set

$$\ell := \lfloor (n - r)/(3 \ln n) \rfloor.$$

Pick a random coloring $\chi : \binom{X}{r+1} \rightarrow [\ell]$ of $\binom{X}{r+1}$ into ℓ colors, and let

$$C_i := \{b \in \binom{X}{r+1} \mid \chi(b) = i\}.$$

It suffices to show that the event “every C_i is an $(n, r + 1, r)$ -covering design” has a non-zero probability.

Re-phrasing it differently, for any $a \in \binom{X}{r}$ denote by E_a the event that the set

$$\text{Sh}(a) := \{b \in \binom{X}{r+1} \mid b \supseteq a\}$$

of cardinality $(n - r)$ does *not* intersect at least one C_i (i.e., this particular a violates the covering condition for some $i \in [\ell]$). Then the event we are interested in is complementary to $\bigvee_{a \in \binom{X}{r}} E_a$, and, thus, we need to prove that

$$\mathbf{P} \left[\bigvee_{a \in \binom{X}{r}} E_a \right] < 1. \tag{2}$$

The probability p of any individual E_a is easy to estimate:

$$p := \mathbf{P}[E_a] \leq \ell \cdot \mathbf{P}[\text{Sh}(a) \cap C_i = \emptyset] = \ell \cdot \frac{(\ell - 1)^{n-r}}{\ell^{n-r}} \leq \ell \cdot e^{-\frac{n-r}{\ell}} \leq n^{-2}, \tag{3}$$

where the last inequality holds due to our choice of ℓ . On the other hand, E_a is mutually independent of all events $E_{a'}$ with $|a' \cap a| < r - 1$ (since $|a' \cap a| < r - 1$ implies $\text{Sh}(a') \cap \text{Sh}(a) = \emptyset$). As $|\{a' \in \binom{X}{r} \mid$

$|a' \cap a| = r - 1$ is $r(n - r) \leq \frac{n^2}{4}$, (2) follows from our bound (3) on p by Lovász Local Lemma (see e.g. [4, Corollary 5.1.2]). \square

4.2 Connected layer families. Let A and B be two disjoint sets of symbols with $|A| = |B| = m$, and let $0 < i, j < m$. We first define the main building block of our construction, the *mesh* $\mathcal{M}(A, i; B, j)$. This is a connected layer family of dimension $(i + j)$ with the set of symbols $A \cup B$, but it also satisfies additional conditions that will allow us to stack different meshes together.

Fix a family $\mathcal{A} = \{C_0^A, \dots, C_{\ell-1}^A\}$ of disjoint $(m, i, i - 1)$ -coverings with the ground set A and, likewise, let $\mathcal{B} = \{C_0^B, \dots, C_{\ell-1}^B\}$ be a family of disjoint $(m, j, j - 1)$ -coverings on B . We assume that

$$\ell = \min\{DC(m, i + 1, i), DC(m, j + 1, j)\}, \quad (4)$$

and we can also assume without loss of generality that \mathcal{A} and \mathcal{B} are *complete* in the sense that $\bigcup_{\alpha=0}^{\ell-1} C_\alpha^A = \binom{A}{i}$ and $\bigcup_{\alpha=0}^{\ell-1} C_\alpha^B = \binom{B}{j}$. Now for $k \in \{0, \dots, \ell - 1\}$, define layers

$$\mathcal{L}_k := \bigcup_{\alpha+\beta=k} (C_\alpha^A \otimes C_\beta^B),$$

where addition is modulo ℓ and \otimes is defined by $C_\alpha^A \otimes C_\beta^B := \{f \dot{\cup} g \mid f \in C_\alpha^A, g \in C_\beta^B\}$. That is, vertices are formed by combining i symbols from A with j symbols from B , and we call a set of this form an $(A, i; B, j)$ -set. The layers $\mathcal{L}_0, \dots, \mathcal{L}_{\ell-1}$ (arbitrarily ordered) form the *mesh* $\mathcal{M}(A, i; B, j)$, see Figure 3.

$$\begin{aligned} \mathcal{L}_0 &= C_0^A \otimes C_0^B \cup C_1^A \otimes C_{\ell-1}^B \cup C_2^A \otimes C_{\ell-2}^B \cup \dots \cup C_{\ell-1}^A \otimes C_1^B \\ \mathcal{L}_1 &= C_0^A \otimes C_1^B \cup C_1^A \otimes C_0^B \cup C_2^A \otimes C_{\ell-1}^B \cup \dots \cup C_{\ell-1}^A \otimes C_2^B \\ \mathcal{L}_2 &= C_0^A \otimes C_2^B \cup C_1^A \otimes C_1^B \cup C_2^A \otimes C_0^B \cup \dots \cup C_{\ell-1}^A \otimes C_3^B \\ &\vdots \\ \mathcal{L}_{\ell-1} &= C_0^A \otimes C_{\ell-1}^B \cup C_1^A \otimes C_{\ell-2}^B \cup C_2^A \otimes C_{\ell-3}^B \cup \dots \cup C_{\ell-1}^A \otimes C_0^B \end{aligned}$$

Figure 3: Illustration of the mesh $\mathcal{M}(A, i; B, j)$

LEMMA 4.1 *The mesh $\mathcal{M}(A, i; B, j)$ is a $(i + j)$ -dimensional connected layer family whose vertices are $(A, i; B, j)$ -sets. Furthermore, all proper subsets of each $(A, i; B, j)$ -set are active on all layers.*

PROOF. Since C_α^A and C_β^B are pairwise disjoint, each $(A, i; B, j)$ -set appears at most once during the construction. Thus condition a) holds.

Consider an $(A, i; B, j - 1)$ -set f . Since \mathcal{A} is complete, $f \cap A$ is contained in a block of C_α^A for some α . Furthermore, $f \cap B$ is covered by every C_β^B , $0 \leq \beta < \ell$. Therefore, f is covered by every $C_\alpha^A \otimes C_{k-\alpha}^B$, $0 \leq k < \ell$, and thus f is active on every layer. An analogous argument applies to $(A, i - 1; B, j)$ -sets. This shows that all proper subsets of $(A, i; B, j)$ -sets are active on all layers. In particular, condition b) of the definition of connected layer families holds. \square

Fix now the dimension d and stack $d - 1$ meshes of dimension d together in the following order to form the layers of the final construction:

$$\begin{aligned} &\mathcal{M}(A, d - 1; B, 1) \\ &\mathcal{M}(A, d - 2; B, 2) \\ &\dots \\ &\mathcal{M}(A, 1; B, d - 1). \end{aligned}$$

We take first all layers of the first mesh, then append all layers of the second mesh, and so on, so that the total number of layers obtained is the sum of the number of layers of all $d - 1$ meshes that are used in the construction.

LEMMA 4.2 *The sequence of layers in the order described above is a d -dimensional connected layer family with $2m$ symbols.*

PROOF. One can easily check that each d -subset of $A \cup B$ appears at most once as a vertex.

To verify condition b), one has to check that all sets $f \in \binom{A \cup B}{\leq d-1}$ are active in contiguous subsequences of layers. And this is immediate from the following description based upon Lemma 4.1:

f is active on any given layer of the mesh $\mathcal{M}(A, i; B, j)$ if and only if $i \geq |f \cap A|$ and $j \geq |f \cap B|$.

(in other words, in the interval representation $\phi(f)$ consists of the meshes $\mathcal{M}(A, |f \cap A|; B, d - |f \cap A|)$, $\mathcal{M}(A, |f \cap A| + 1; B, d - |f \cap A| - 1), \dots, \mathcal{M}(A, d - |f \cap B|; B, d - |f \cap B|)$). \square

We are now ready to prove our main result, an almost quadratic lower bound on the largest diameter $D(d, n)$ of our base abstraction of dimension d with n symbols.

THEOREM 4.2 $D(n/4, n) = \Omega(n^2/\log n)$.

PROOF. By Theorem 4.1 and (4) with $m = n/2$, in the previously described construction every mesh contributes $\Omega(n/\log n)$ layers, and altogether there are $n/2 - 1$ meshes. \square

5. Final remarks. There are many interesting questions related to abstractions which deserve further inspection.

First and foremost, is $D(d, n)$ bounded by a polynomial in n ? We note that even if we have proved an almost quadratic lower bound on $D(d, n)$, it appears as if the ideas underlying our construction completely break apart beyond that point.

Another interesting question is whether the addition of one or two of the conditions ii) or iii) strengthens the base abstraction, or whether the diameters of the corresponding abstractions are related via polynomial factors. In the latter case, the diameter of any abstraction would be polynomial if and only if this was the case for the base abstraction.

Finally, can we remove the annoying logarithmic factor in our bound on the number of disjoint $(n, r + 1, r)$ -covering designs in Theorem 4.1?

Acknowledgments. We would like to thank an anonymous referee for many useful remarks. Friedrich Eisenbrand and Nicolai Hähnle were supported by the Swiss National Science Foundation (SNF). Part of this work was done while Alexander Razborov was with the Steklov Mathematical Institute, supported by the Russian Foundation for Basic Research, and with the Toyota Technological Institute at Chicago. Thomas Rothvoß was supported by the German Research Foundation (DFG) within the Priority Program 1307 “Algorithm Engineering”.

References

- [1] Ilan Adler. Lower bounds for maximum diameters of polytopes. *Mathematical Programming Study*, (1):11–19, 1974.
- [2] Ilan Adler, George Dantzig, and Katta Murty. Existence of A -avoiding paths in abstract polytopes. *Mathematical Programming Study*, (1):41–42, 1974.
- [3] Ilan Adler and George B. Dantzig. Maximum diameter of abstract polytopes. *Mathematical Programming Study*, (1):20–40, 1974.
- [4] Noga Alon and Joel Spencer. *The probabilistic method*. Wiley-Interscience, 3 edition, 2008.
- [5] David Bremner and Lars Schewe. Edge-Graph Diameter Bounds for Convex Polytopes with Few Facets. *arXiv:0809.0915v3 [math.CO]*, 2009.
- [6] Reinhard Diestel. *Graph theory*. Springer-Verlag, New York, 2 edition, 2000.
- [7] Paul Erdős and Haim Hanani. On a limit theorem in combinatorial analysis. *Publicationes Mathematicae Debrecen*, 10:10–13, 1963.
- [8] Bernd Gärtner. A subexponential algorithm for abstract optimization problems. *SIAM Journal on Computing*, 24(5):1018–1035, 1995.
- [9] Bernd Gärtner. The random-facet simplex algorithm on combinatorial cubes. *Random Structures & Algorithms*, 20(3):353–381, 2002.
- [10] Gil Kalai. A subexponential randomized simplex algorithm (extended abstract). In *STOC*, pages 475–482. ACM, 1992.
- [11] Gil Kalai. Upper bounds for the diameter and height of graphs of convex polyhedra. *Discrete & Computational Geometry*, 8:363–372, 1992.

- [12] Gil Kalai. Linear programming, the simplex algorithm and simple polytopes. *Mathematical Programming*, 79(1-3, Ser. B):217–233, 1997.
- [13] Gil Kalai and Daniel J. Kleitman. A quasi-polynomial bound for the diameter of graphs of polyhedra. *Bulletin of the American Mathematical Society*, 26:315–316, 1992.
- [14] Victor Klee and Peter Kleinschmidt. The d -step conjecture and its relatives. *Mathematics of Operations Research*, 12(4):718–755, 1987.
- [15] Victor Klee and David W. Walkup. The d -step conjecture for polyhedra of dimension $d < 6$. *Acta Mathematica* 133, pages 53–78, 1967.
- [16] David G. Larman. Paths on polytopes. *Proceedings of the London Mathematical Society*, pages 161–178, 1970.
- [17] Peter Mani and David W. Walkup. A 3-sphere counterexample to the W_v -path conjecture. *Mathematics of Operations Research*, 5(4):595–598, 1980.
- [18] Jiří Matoušek, Micha Sharir, and Emo Welzl. A subexponential bound for linear programming. *Algorithmica*, 16(4/5):498–516, 1996.
- [19] Vojtěch Rödl. On a packing and covering problem. *European Journal of Combinatorics*, (5):69–78, 1985.
- [20] Francisco Santos, A counterexample to the Hirsch conjecture. *arXiv:1006.2814v1 [math.CO]*, 2010.
- [21] Ingo Schurr and Tibor Szabó. Finding the sink takes some time: an almost quadratic lower bound for finding the sink of unique sink oriented cubes. *Discrete & Computational Geometry*, 31(4):627–642, 2004.