

Diameter of Polyhedra: Limits of Abstraction

Friedrich Eisenbrand
Institute of Mathematics
EPFL
Lausanne, Switzerland
friedrich.eisenbrand@epfl.ch

Nicolai Hähnle
Institute of Mathematics
EPFL
Lausanne, Switzerland
nicolai.haehnle@epfl.ch

Thomas Rothvoß
Institute of Mathematics
EPFL
Lausanne, Switzerland
thomas.rothvoss@epfl.ch

ABSTRACT

We investigate the diameter of a natural abstraction of the 1-skeleton of polyhedra. Although this abstraction is simpler than other abstractions that were previously studied in the literature, the best 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 a superlinear lower bound.

Categories and Subject Descriptors

G.1.6 [Numerical Analysis]: Optimization—*linear programming*; G.2.1 [Discrete Mathematics]: Combinatorics

General Terms

Algorithms, Theory

Keywords

convex geometry, disjoint coverings, Hirsch conjecture, polyhedra

1. INTRODUCTION

One of the most prominent mysteries in convex geometry is the question whether the diameter of polyhedra is polynomial or not. If the largest diameter of a d -dimensional polyhedron with n facets is denoted by $\Delta_u(d, n)$, then the best known lower and upper bounds are $n - d + \lfloor d/5 \rfloor \leq \Delta_u(d, n) \leq n^{\log d + 1}$, shown by Klee and Walkup [13] and Kalai and Kleitman [11] respectively. 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 the 1-skeleton of non-degenerate polyhedra. Here, we call a polyhedron non-degenerate, if each vertex of the polyhedron is contained in exactly d facets. In the quest of bounding Δ_u one can restrict one's attention to non-degenerate polyhedra since, by perturbation, the polyhedron can be turned into a non-degenerate polyhedron,

whose diameter is at least as large as the one of the original polyhedron. Such abstractions have been studied in the literature for a long time [9, 2, 1]. The subject of our study is a natural *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. We observe simple proofs which show that the best known upper bounds do also hold for our simple base abstraction and we provide a superlinear lower bound on the diameter of these abstractions.

While this shows that 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 quest of proving linear upper bounds. 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.

Our base abstraction is a family $\mathcal{B}_{d,n}$ of connected graphs $G = (V, E)$. The vertices V of G are subsets of $[n] = \{1, \dots, n\}$ of cardinality 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 $u \cap v$.

The largest diameter of a graph in $\mathcal{B}_{d,n}$ is denoted by $D(d, n)$. We call d the *dimension* and n the *number of symbols* 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, there exists a path from an extreme point u to an extreme point v of P which does not leave the facets 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 $D(d, n) = \Omega(n^{3/2})$ if d is allowed to grow as a function of n . The non-trivial construction relies on disjoint covering designs and uses Hall's theorem which characterizes the existence of perfect matchings in a bipartite graph. At the same time the bound of Kalai and Kleitman [11], $D(d, n) \leq n^{\log d + 1}$, as well as the upper bound of Larman [14], $D(d, n) \leq 2^{d-1} \cdot n$, which is linear when the dimension is fixed, continue to hold in the base abstraction.

While the first bound is merely an adaption of the proof in [11], our proof of the second bound is much simpler than the one which was proved for polyhedra in [14].

We 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 uv is present if and only if $|u \cap v| = d - 1$.
- iii) Each subset $S \in \binom{[n]}{d-1}$ 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* is also true up to dimension 5 for abstract polytopes. Recall that the *d-step conjecture* states that the diameter of a d -dimensional polytope with $2d$ facets is bounded by d . Klee and Walkup [13] proved that the *d-step conjecture* is true if and only if the famous *Hirsch conjecture* is true. The Hirsch conjecture states that the diameter of a polytope is bounded by $n - d$. Klee and Walkup [13] were the first to prove that the *d-step conjecture* is true up to dimension 5. A big advantage of abstraction is that the validity of the *d-step conjecture* for abstract polytopes in fixed dimension d can be automatically checked with a computer. This does not seem to be so easy for the validity of the *d-step conjecture* of polytopes themselves. The situation for lower bounds on the diameter of abstract polytopes is as follows. Mani and Walkup [15] have provided an example of an abstract polytope with $d = 12$ and $n = 24$ whose diameter is larger than 12, see also [12]. Superlinear lower bounds for the diameter of abstract polytopes are not known.

Kalai [9] 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 [11] can also be proved in this setting. We will see later that the condition ii) is not necessary and that this bound, together with the linear bound in fixed dimension of Larman [14] 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 [8] and Matoušek, Sharir, and Welzl [16] 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 [6]. For unique sink orientations of cubes, which capture much of the interesting structure of these abstractions, Schurr and Szabó [18] 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 $O(n^3 e^{2\sqrt{n}})$ [7].

2. BASE ABSTRACTION AND CONNECTED LAYER FAMILIES

Let $G = (V, E)$ be a graph of our base abstraction $\mathcal{B}_{d,n}$. 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 subsets $\mathcal{L}_i \subseteq \binom{[n]}{d}$ for $i = 0, \dots, \ell$ with $\mathcal{L}_i = \{v \in V : \text{dist}(s, v) = i\}$. The sets \mathcal{L}_i satisfy the following conditions.

- a) *Disjointness*: $\forall i \neq j : \mathcal{L}_i \cap \mathcal{L}_j = \emptyset$.
- b) *Connectivity*: $\forall 0 \leq i < j < k \leq \ell$ 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 $\mathcal{L}_i \subseteq \binom{[n]}{d}, i = 0, \dots, \ell$ 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 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 lemma.

LEMMA 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.

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

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 $s \in S$ be a symbol of a vertex in a connected layer family. The *induction on s* is the following operation.

1. Remove all vertices from the connected layer family which do not contain s .
2. Remove s from all vertices.
3. Remove empty layers.

The next lemma follows directly from the definition of connected layer families.

LEMMA 2. *Consider a d -dimensional connected layer family with n symbols. Induction on 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 $\Delta_u(d, n) \leq n^{\log d+2}$ of [11] is up to today 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.

THEOREM 1. $D(d, n) \leq n^{1+\log d} - 1$.

PROOF. By Lemma 1 it is enough to show that the height 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. Then there exists a symbol $s \in \{1, \dots, n\}$ such that 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 $\ell_2 - \ell_1 - 1$. If we denote the maximum height of a d -dimensional connected layer family with n symbols by $h(d, n)$, then this yields the inequality

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

The bound is then proved by induction on n . Note that $h(1, n) = n$ and suppose 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

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

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 1. Notice that our bound on $D(d, n)$ is slightly better than the bound on $\Delta_u(d, n)$ in [11]. Kalai [10] pointed out that the $n^{\log d+2}$ bound can be improved to $n^{\log d+1}$ 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 [14]. 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 2. $D(d, n) \leq 2^{d-1} \cdot n - 1$.

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

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

Next we define a sequence s_i of symbols. The symbol s_1 is 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 one that reaches farthest among all symbols that are active in 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 $|i-j| \geq 2$.

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

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

\square




Figure 1: Illustration of the proof of Theorem 1




Figure 2: Illustration of the proof of Theorem 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 .

4. A LOWER BOUND ON THE DIAMETER OF THE BASE ABSTRACTION

Our goal is to construct a d -dimensional connected layered family with n symbols that has a large number of layers. The difficult condition to meet is connectivity b). Our first idea is to meet 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. 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 $n > k > r$ be natural numbers and let X be a ground set of size n . A k -element subset or k -subset of X is called a *block*. A collection C of blocks is called an (n, k, r) -*covering design* or simply a *covering* if every r -subset of X 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. Rödl [17] for example, proved a long standing conjecture of Erdős and Hanani [5] on 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. The next section is devoted to proving a lower bound on the number of disjoint $(n, r + 1, r)$ -covering designs. Here we make use of Hall's theorem on matchable bipartite graphs.

4.1 Disjoint covering designs

We call the maximum size of a family of disjoint (n, k, r) -coverings the *disjoint covering number*, denoted by $DC(n, k, r)$. We show the following theorem.

THEOREM 3. *Let $1 \leq r < n$. Then $DC(n, r + 1, r) \geq n - r(r + 2)$.*

An upper bound on $DC(n, r + 1, r)$ is $n - r$. This can be seen by picking an arbitrary r -set. This has to be covered by a block in every covering. The r -set can be completed to an $r + 1$ -set in $n - r$ different ways. Disjointness implies the upper bound. In the proof of Theorem 3 we use Hall's Theorem, see, e.g. [4], which we recall here.

THEOREM 4 (HALL'S THEOREM). *A bipartite graph $G = (V \cup W, E)$ has a matching which matches each node $v \in V$ if and only if $|N(S)| \geq |S|$ for each $S \subseteq V$, where $N(S)$ denotes the set of neighbors of S .*

PROOF OF THEOREM 3. We use the integers modulo n , \mathbb{Z}_n , as our ground set. For any subset $A \subseteq \mathbb{Z}_n$, define $\sum(A) := \sum_{a \in A} a$. The function $b : \binom{\mathbb{Z}_n}{r+1} \rightarrow \mathbb{Z}_n$, $b(B) := \sum(B)$ maps every block into a *bucket* $b^{-1}(\{i\})$, $i = 0, \dots, n - 1$. These buckets are disjoint sets of blocks, however they are not coverings. We make the following auxiliary claim.

Every r -element subset $A \subset \mathbb{Z}_n$ is covered in exactly $n - r$ buckets. Conversely, there exist exactly r buckets in which A is not covered by any block.

This claim is seen as follows: Let A be any r -element subset of \mathbb{Z}_n . Each block that covers A is of the form $A \cup \{x\}$ with $x \in \mathbb{Z}_n \setminus A$ and there are exactly $n - r$ such blocks.

Those blocks are mapped into the buckets with label $\hat{b}(x) = \sum(A) + x$. However, \hat{b} is injective, which implies that they are mapped into pairwise different buckets.

The buckets $0, \dots, n - r(r+2) - 1$ will be called *surviving* buckets, and the buckets numbered $n - r(r+2), \dots, n - 1$ will be called *supply* buckets. We will turn the surviving buckets into coverings by filling them with blocks from the supply buckets.

To this end, we construct the bipartite graph $G = (V \cup W, E)$, where

$$\begin{aligned} V &:= \{(A, j) \mid A \in \binom{\mathbb{Z}_n}{r} \text{ is not covered in surv. bucket } j\} \\ W &:= \{B \in \binom{\mathbb{Z}_n}{r+1} \mid B \text{ is contained in a supply bucket}\} \end{aligned}$$

There is an edge between $(A, j) \in V$ and $B \in W$ if $A \subset B$. We are done once we show that G satisfies the conditions of Hall's theorem, since we then add to each surviving bucket j the blocks which are matched to nodes of the form (A, j) . The so-constructed sets of blocks are disjoint $(n, r+1, r)$ coverings.

Let S be a subset of V . We will count the number e_S of edges in the subgraph induced by $S \cup N(S)$ to show that $|S| \leq |N(S)|$. Consider any $B \in N(S)$. There are $r+1$ different r -element subsets of B . By the auxiliary claim, each of these fails to be covered in exactly r buckets. Some of these buckets may be supply buckets. Nevertheless, this means that B has at most degree $r(r+1)$ in G . We thus have

$$e_S \leq r(r+1)|N(S)|$$

Now consider any $(A, j) \in S$. By the above claim, the set A is covered by $n - r$ buckets and thus by at least $r(r+1)$ supply buckets. This means that (A, j) has at least degree $r(r+1)$ in G . This shows that

$$e_S \geq r(r+1)|S|$$

and we conclude $|N(S)| \geq |S|$. \square

4.2 The lower bound construction

Let A and B be two disjoint sets of symbols. We will define two *components*, the *bridge* $\mathcal{B}(A)$ and *mesh* $\mathcal{M}(A, i; B, j)$. Both are connected layer families in their own right, but they satisfy additional conditions which allow us to stack them together.

Let A be a set of symbols, $|A| = m$ and let $C_0, \dots, C_{\ell-1}$ be a family of disjoint $(m, d, d-1)$ -coverings of the ground set A with $\ell = m - d^2 + 1$ by Theorem 3. Define the connected layer family $\mathcal{B}(A)$ such that $C_0, \dots, C_{\ell-1}$ are the layers of $\mathcal{B}(A)$.

LEMMA 3. $\mathcal{B}(A)$ is a d -dimensional connected layer family with $m - d^2 + 1$ layers using only symbols from A . Furthermore, all $(d-1)$ -subsets of A are active on all layers of $\mathcal{B}(A)$.

PROOF. Condition a) for connected layer families is obviously true. Also, all $(d-1)$ -subsets of A are covered by each C_j , so all $(d-1)$ -subsets are active on all layers. This is sufficient for condition b) to hold. \square

Next we come to the definition of the mesh. Let A and B be disjoint sets of symbols with $|A| = |B| = m$ and $0 < i, j < d$ with $i + j = d$. Let $\mathcal{A} = \{C_0^A, \dots, C_{\ell-1}^A\}$ be a family of disjoint $(m, i, i-1)$ -coverings of A and $\mathcal{B} = \{C_0^B, \dots, C_{\ell-1}^B\}$

a family of disjoint $(m, j, j-1)$ -coverings of B . Let $q := \max\{i, j\}$. By Theorem 3, we can obtain families such that $\ell = m - q^2 + 1$. We can 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

$$\mathcal{L}_k := \{M \in C_\alpha^A \otimes C_\beta^B \mid \alpha + \beta = k\}$$

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

LEMMA 4. The mesh $\mathcal{M}(A, i; B, j)$ is a d -dimensional connected layer family with $m - q^2 + 1$ layers 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 the 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 M . Since \mathcal{A} is complete, $M \cap A$ is contained in a block of C_α^A for some α . Furthermore, $M \cap B$ is covered by every C_β^B , $0 \leq \beta < \ell$. Therefore, M is covered by every $C_\alpha^A \otimes C_{k-\alpha}^B$, $0 \leq k < \ell$, and thus M 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

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

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

The sequence of $d-1$ mesh components is used to interpolate between the two sets of symbols.

LEMMA 5. The sequence of layers in the 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 of $d-1$ or fewer symbols are active in contiguous subsequences of layers. Lemma 3 and Lemma 4 imply that all components have the property that whenever a set of $d-1$ or fewer symbols is active anywhere in the component, it is active in all layers of that component, so we only need to show that every such set is active in a contiguous sequence of components.

Consider an i -element subset of A , where $0 < i < d$. This set is clearly active in the bridge $\mathcal{B}(A)$. It is also active in the subsequent mesh components $\mathcal{M}(A, d-1; B, 1), \dots, \mathcal{M}(A, i; B, d-i)$. It cannot be active in any of the later components, since vertices in those components contain fewer than i symbols from A . A symmetric argument applies to i -element subsets of B .

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

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

Now consider an $(A, i; B, j)$ -set with $i > 0$, $j > 0$, and $i + j < d$. This set is only active in mesh components. To be more precise, it is active in $\mathcal{M}(A, i'; B, j')$ if and only if $i \leq i'$ and $j \leq j'$. These mesh components form a contiguous sequence. \square

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

THEOREM 5. $D(\lfloor \sqrt{n}/2 \rfloor, n) = \Omega(n^{3/2})$

PROOF. In the previously described construction, every component contributes $m - d^2 + 1$ layers. Counting the number of components, we can compute the number of layers ℓ :

$$\ell = (d + 1)(m - d^2 + 1)$$

By setting $m = 2d^2$, we obtain $\ell \geq d^3 + 1 = n^{3/2}/8 + 1$, where $n = 2m = 4d^2$ is the total number of symbols. \square

5. FINAL REMARKS

There are many interesting questions related to abstractions which deserve further inspection. One question for example 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. Maybe the question whether $D(d, n)$ is polynomial in d and n has a simpler answer than the question whether $\Delta_u(d, n)$ is polynomial or not.

Another question concerns the number of disjoint $(n, r + 1, r)$ -covering designs $DC(n, r + 1, r)$. We show that $DC(n, r + 1, r) \geq n - r(r + 2)$ whereas the best upper bound $DC(n, r + 1, r) \leq n - r$ is the result of a simple counting argument.

6. REFERENCES

- [1] Ilan Adler. Lower bounds for maximum diameters of polytopes. *Math. Programming Stud.*, pages 11–19, 1974. Pivoting and extensions.
- [2] Ilan Adler, George Dantzig, and Katta Murty. Existence of A -avoiding paths in abstract polytopes. *Math. Programming Stud.*, (1):41–42, 1974. Pivoting and extensions.
- [3] Ilan Adler and George B. Dantzig. Maximum diameter of abstract polytopes. *Mathematical Programming Study*, (1):20–40, 1974. Pivoting and extensions.
- [4] Reinhard Diestel. *Graph theory*. Springer-Verlag, New York, 2 edition, 2000.
- [5] Paul Erdős and Haim Hanani. On a limit theorem in combinatorial analysis. *Publ. Math. Debrecen*, 10:10–13, 1963.
- [6] Bernd Gärtner. A subexponential algorithm for abstract optimization problems. *SIAM Journal on Computing*, 24(5):1018–1035, 1995.
- [7] Bernd Gärtner. The random-facet simplex algorithm on combinatorial cubes. *Random Structures & Algorithms*, 20(3):353–381, 2002. Probabilistic methods in combinatorial optimization.
- [8] Gil Kalai. A subexponential randomized simplex algorithm (extended abstract). In *STOC*, pages 475–482. ACM, 1992.
- [9] Gil Kalai. Upper bounds for the diameter and height of graphs of convex polyhedra. *Discrete & Computational Geometry*, 8:363–372, 1992.
- [10] Gil Kalai. Linear programming, the simplex algorithm and simple polytopes. *Mathematical Programming*, 79(1-3, Ser. B):217–233, 1997. Lectures on mathematical programming (ismp97) (Lausanne, 1997).
- [11] 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.
- [12] Victor Klee and Peter Kleinschmidt. The d -step conjecture and its relatives. *Math. Oper. Res.*, 12(4):718–755, 1987.
- [13] Victor Klee and David W. Walkup. The d -step conjecture for polyhedra of dimension $d < 6$. *Acta Math.* 133, pages 53–78, 1967.
- [14] David G. Larman. Paths on polytopes. *Proceedings London Math. Soc.*, pages 161–178, 1970.
- [15] Peter Mani and David W. Walkup. A 3-sphere counterexample to the W_0 -path conjecture. *Math. Oper. Res.*, 5(4):595–598, 1980.
- [16] Jiří Matoušek, Micha Sharir, and Emo Welzl. A subexponential bound for linear programming. *Algorithmica*, 16(4/5):498–516, October/November 1996.
- [17] Vojtěch Rödl. On a packing and covering problem. *European Journal of Combinatorics*, (5):69–78, 1985.
- [18] 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. An International Journal of Mathematics and Computer Science*, 31(4):627–642, 2004.