Enumeration of interval graphs and d-representable complexes

Boris Bukh* R. Amzi Jeffs†

March 2022

Abstract

For each fixed $d \geq 1$, we obtain asymptotic estimates for the number of d-representable simplicial complexes on n vertices as a function of n. The case $d = 1$ corresponds to counting interval graphs, and we obtain new results in this well-studied case as well. Our results imply that the d-representable complexes comprise a vanishingly small fraction of d-collapsible complexes.

1 Introduction

Given a tuple $\mathcal{C} = (C_1, C_2, \ldots, C_n)$ of convex sets, one may record their intersection pattern using the *nerve complex*

$$\operatorname{nerve}(\mathcal{C}) \stackrel{\text{def}}{=} \Big\{\sigma \;\Big|\; \bigcap_{i\in\sigma} C_i \neq \emptyset\Big\}.$$

Observe that $\operatorname{nerve}(\mathcal{C})$ is a simplicial complex on the vertex set $[n] \stackrel{\text{def}}{=} \{1, 2, \ldots, n\}$.

A simplicial complex is called *d-representable* if it is the nerve of a tuple of convex sets in $\mathbb{R}^d$. Such a tuple is called a *d-representation* of Δ. Given a d-representation, one may fix a point p_σ in each nonempty region $\bigcap_{i\in\sigma} C_i$ and then replace each C_i by the convex hull of the p_σ it contains, obtaining a d-representation of the same complex in which each set is compact. Henceforth we will only consider d-representations consisting of compact convex sets.

The class of d-representable complexes enjoys many useful topological and combinatorial properties. In particular, Helly's theorem implies that a d-representable complex is completely determined by its d-skeleton. Indeed, if Δ is d-representable and $\sigma \subseteq [n]$ is a set of size $d + 2$ or larger, then σ is a face of Δ if and only if every subset of σ with size $d + 1$ is a face of Δ.

Perhaps the closest combinatorial analog of d-representability is *d-collapsibility*. A *free face* in a simplicial complex is a face that is contained in a unique facet. A *d-collapse* is the operation of deleting a free face with dimension $d - 1$ or less (and all faces that contain it). Finally, a simplicial complex Δ is called *d-collapsible* if there is a sequence of d-collapses from Δ to the empty complex.

In 1975 Wegner [18] proved that d-representable complexes are d-collapsible by ordering the regions corresponding to faces according to a generic linear function (see [15, Figure 3] for a modern sketch

*Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported in part by U.S. taxpayers through NSF CAREER grant DMS-1555149.

†Department of Mathematical Sciences, Carnegie Mellon University, Pittsburgh, PA 15213, USA. Supported by the National Science Foundation through Award No. 2103206.

of the proof). There are d-collapsible complexes that are not d-representable, for any $d \geq 1$, and there has been a variety of work exploring the properties of d-representable complexes, d-collapsible complexes and the related (even more general) class of d-Leray complexes [9, 6, 10, 14, 16]. Tancer's survey [15] provides a good overview of results in this area prior to 2013.

In this work we study, for fixed $d \geq 1$, the growth of the number of d-representable complexes on the vertex set $[n]$ as a function of n. We also consider the number of d-collapsible complexes. For each $d \geq 1$ we determine the rough rate of growth for these quantities and provide upper and lower bounds on the constants involved.

One-dimensional case (counting interval graphs). We first discuss the case $d = 1$, which is of particular interest. Since the 1-skeleton of a 1-representable complex determines the entire complex, 1-representable complexes are exactly the clique complexes of the well-studied class of *interval graphs*, which record the nonempty pairwise intersections of n closed intervals[1] on the real line. Thus the number of 1-representable complexes on the vertex set $[n]$ is equal to the number of (labeled) interval graphs on this vertex set.

There were several works [5, 8, 7, 20, 1] concerned with estimating the number of interval graphs. The currently sharpest asymptotic result is due to Gavoille and Paul [7] who proved that the number of interval graphs on n vertices is $\exp\big(2n \log n - O(n \log \log n)\big)$. The upper bound is straightforward, as an n-vertex interval graph is determined by the relative order of $2n$ interval endpoints in its 1-representation. The nontrivial part of their result is in the proof of the lower bound. Unaware of this result, weaker lower bounds were subsequently found by Yang and Pippenger [20] and Acan, Chakraborty, Jo and Satti [1]. A different approach was taken by Hanlon [8] who found a recursive expression for the generating function for the number of *unlabeled* interval graphs. In addition, there were several works [19, 11] on quick algorithmic enumeration of non-isomorphic interval graphs.

By relating the enumeration of interval graphs to that of *interval orders* and applying results of Brightwell and Keller [4], we find an improved asymptotic for the number of the interval graphs.

Theorem 1. *The number of interval graphs on the vertex set* $[n]$ *is* $e^{2n \log n - (2 + \log(\pi^2/6))n + O(\log n)}$.

Higher-dimensional case. Our results in dimension $d \geq 2$ are less precise than for $d = 1$.

Theorem 2. *The number* $f_d(n)$ *of* d*-representable complexes on the vertex set* $[n]$ *satisfies*

$$\tfrac{2}{d^d} n^d \log n - O(n^d) \leq \log f_d(n) \leq \tfrac{2}{(d-1)!} n^d \log n + O(n^d).$$

In particular, $f_d(n) = e^{\Theta(n^d \log n)}$.

It turns out that it is not difficult to estimate the number of d-collapsible complexes on the vertex set $[n]$. We obtain the following result.

[1] There seems to be no consensus as to whether the empty set is an interval. However, the class of interval graphs is unaffected by the choice of a side in this question, for one may always replace each empty set in a 1-representation by an interval that intersects no other.

Theorem 3. *The number $g_d(n)$ of d-collapsible complexes on the vertex set $[n]$ satisfies*

$$\tfrac{1}{(d+1)^{d+1}} n^{d+1} + O(n^d) \le \log_2 g_d(n) \le \tfrac{1}{(d+1)!} n^{d+1} + O(n^d).$$

In particular, $g_d(n) = e^{\Theta(n^{d+1})}$.

This implies that there are far more d-collapsible complexes than d-representable complexes. That is surprising because nearly all published results for d-representable complexes in fact apply to d-collapsible complexes.

Another curious consequence of these estimates is that there are more d-representable complexes than $(d-1)$-collapsible complexes. That this is not obvious can be seen from the work of Tancer [14] who, for any d, described complexes that are 2-collapsible but not d-representable.

Thanks. We are grateful to the referee who made a number of useful suggestions. We are especially grateful to them for catching a subtle off-by-one mistake in the proof of Theorem 9.

2 Representable complexes in general: proof of Theorem 2

We prove Theorem 2 modulo the case $d = 1$, which is established in Section 3 by proving Theorem 1.

Upper bound. Let $(C_1, C_2, \dots, C_n)$ be a d-representation of a complex Δ on the vertex set $[n]$. Define $N \stackrel{\text{def}}{=} \binom{n}{d}$. Let π be a linear projection to the span of e_1 in $\mathbb{R}^d$, and for $\sigma \in \binom{[n]}{d}$ define $I_\sigma \stackrel{\text{def}}{=} \pi(\bigcap_{i\in\sigma} C_i)$. Observe that each I_σ is a (possibly empty) closed interval. We may regard the collection of intervals $\{I_\sigma \mid \sigma \in \binom{[n]}{d}\}$ as a 1-representation of a simplicial complex Γ on the vertex set $[N]$.

We claim that Γ completely determines the d-dimensional faces of Δ. Specifically, we claim that for any $\tau \in \binom{[n]}{d+1}$ we have $\tau \in \Delta$ if and only if the intervals $\{I_{\tau\setminus\{i\}} \mid i \in \tau\}$ share a common point. One direction is clear: if $p \in \bigcap_{i\in\tau} C_i$ then $\pi(p)$ lies in all the appropriate intervals. For the reverse inclusion, suppose that the intervals share a common point, and consider the fiber of π over this point. This fiber is a $(d-1)$-dimensional affine subspace of $\mathbb{R}^d$. Moreover, all d-fold intersections of the collection $\{C_i \mid i \in \tau\}$ contain a point in this fiber. Applying Helly's theorem inside the fiber (regarded as a copy of $\mathbb{R}^{d-1}$), we conclude that the fiber contains a point in $\bigcap_{i\in\tau} C_i$. Thus $\tau \in \Delta$ as desired.

Helly's theorem (now applied in $\mathbb{R}^d$) implies that the d-skeleton of Δ determines all of Δ. The faces dimension d and $d-1$ in Δ are determined by Γ, and the remainder of the d-skeleton is determined by choosing faces of dimension less than $d-1$. Thus we have

$$f_d(n) \le f_1(N) \cdot 2^{\binom{n}{<d}},$$

where $\binom{n}{<d} \stackrel{\text{def}}{=} \binom{n}{d-1} + \cdots + \binom{n}{0}$. Taking logarithms and using Theorem 1, we obtain

$$\log(f_d(n)) \le 2N \log N + O(N) + O(n^{d-1}) = \frac{2}{(d-1)!} n^d \log n + O(n^d).$$

Lower bound. Below, we shall use the notation $\Delta * w$ to denote the *cone over* Δ *with apex* w, which is the simplicial complex whose facets are obtained by adding a new vertex w to every facet of Δ.

Proposition 4. *Let V and W be disjoint sets. For each $w \in W$, let Δ_w be a $(d-1)$-representable complex on vertex set V. Then the simplicial complex*

$$\Delta \stackrel{\text{def}}{=} 2^V \cup \bigcup_{w \in W} (\Delta_w * w)$$

on vertex set $V \sqcup W$ is d-representable.

Proof. Pick a full-dimensional polytope P in $\mathbb{R}^d$ with at least $|W|$ many facets, and choose a point p in the interior of P. Let $\{C_w \mid w \in W\}$ be a collection of $(d-1)$-dimensional simplices, each contained in the interior of distinct facets of P. For each $w \in W$, let $\{A_{w,v} \mid v \in V\}$ be a $(d-1)$-representation of Δ_w contained in C_w. Lastly, for each $v \in V$, define

$$C_v \stackrel{\text{def}}{=} \operatorname{conv}\left(\{p\} \cup \bigcup_{w \in W} A_{w,v}\right).$$

We claim that the collection $\{C_v \mid v \in V\} \cup \{C_w \mid w \in W\}$ is a d-representation of Δ. Clearly all C_v share a common point, namely p, and all C_w are mutually disjoint from one another. Thus it suffices to argue for each $w \in W$ and $\sigma \subseteq V$ that C_w contains a point in $\bigcap_{v \in \sigma} C_v$ if and only if σ is a face of Δ_w. This follows by construction, since the intersection of $\{C_v \mid v \in V\}$ with C_w is exactly the $(d-1)$-representation $\{A_{w,v} \mid v \in V\}$ of Δ_w. □

Example 5. Figure 1 shows the construction in Proposition 4 for the case $V = \{1,2,3\}$, $W = \{4,5,6,7\}$ and $d = 2$. Here each vertex in W "picks out" a 1-representable complex on the vertex set $\{1,2,3\}$. Explicitly, Δ_4 is three isolated vertices, Δ_5 is a 2-simplex, Δ_6 is the 1-simplex on $\{2,3\}$, and Δ_7 has facets $\{1,3\}$ and $\{2,3\}$. In the figure each C_w is slightly thickened and shown in black, while C_1, C_2, and C_3 are slightly transparent.

Proposition 6. *For any $0 \le m \le n$ we have $f_d(n) \ge f_{d-1}(m)^{n-m}$.*

Proof. We will create $f_{d-1}(m)^{n-m}$ many distinct d-representable complexes on the vertex set $[n]$. Define V to be the first m elements of $[n]$, and let W be the last $n-m$ elements of $[n]$. For each $w \in W$, let Δ_w be a $(d-1)$-representable complex on vertex set V. Observe that we may recover Δ_w from the d-representable complex Δ formed in Proposition 4 (in particular, Δ_w is the link of the vertex w in this complex). Thus every collection $\{\Delta_w \mid w \in W\}$ of $(d-1)$-representable complexes determines a unique d-representable complex on the vertex set $[n]$. We have $f_{d-1}(m)$ choices for each Δ_w, and since we have $n-m$ many vertices in W the bound follows. □

Corollary 7. *For every $d \ge 1$ and sufficiently large n we have*

$$\log f_d(n) \ge \frac{2}{d^d} n^d \log n + O(n^d).$$

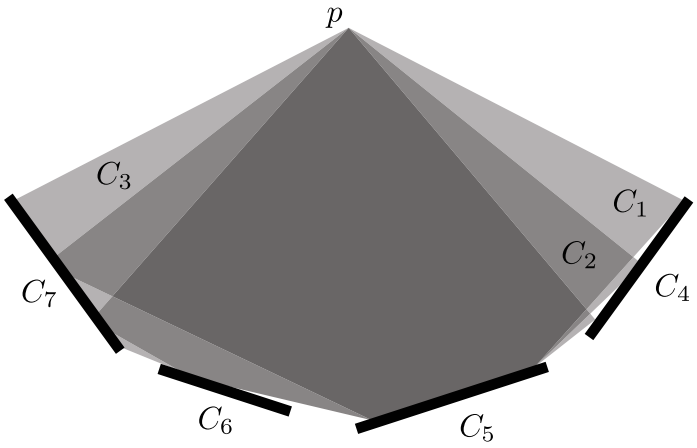

Figure 1: The construction used in Proposition 4.

Proof. We proceed by induction on d. The base case $d = 1$ follows from Theorem 1, which we will establish in Section 3. Since the function $f_d(n)$ is increasing in n, it suffices to consider the case when n is a multiple of d. For the inductive step, let $m = \frac{d-1}{d}n$. Using Proposition 6 and the inductive hypothesis, we compute

$$\begin{aligned}
\log f_d(n) &\geq \log f_{d-1}(m)^{n-m} \\
&= \tfrac{n}{d} \log f_{d-1}(\tfrac{d-1}{d}n) \\
&\geq \tfrac{n}{d} \left(\tfrac{2}{(d-1)^{d-1}} (\tfrac{d-1}{d}n)^{d-1} \log(\tfrac{d-1}{d}n) + O\big(n^{d-1}\big) \right) \\
&= \tfrac{n}{d} \cdot \tfrac{2}{(d-1)^{d-1}} \cdot \tfrac{(d-1)^{d-1}}{d^{d-1}} \cdot n^{d-1} \log n + O(n^d) \\
&= \tfrac{2}{d^d} n^d \log n + O(n^d).
\end{aligned}$$

□

Better lower bound in dimension 2. The choice of the constant $\frac{d-1}{d}$ in the definition of m above is optimal, and so $2/d^d$ is the best constant that can be deduced from the construction in Proposition 6. However, in Proposition 8 below we will use a slightly different construction to obtain the constant $\frac{3}{4}$ when $d = 2$, which is better than the constant $\frac{2}{2^2} = \frac{1}{2}$ from Corollary 7. One can use the inductive argument in Corollary 7 to prove the lower bound of $\frac{3}{d^d} n^d \log n + O(n)$, for $d \geq 2$, by using Proposition 8 as the base case. It would be interesting to determine the constant β_2 such that $\log f_2(n) = \beta_2 n^2 \log n + O(n)$.

Proposition 8. *For sufficiently large n we have $f_2(n) \geq f_1(n/4)^{3n/2}$. In particular, it follows that $\log f_2(n) \geq \frac{3}{4} n^2 \log n + O(n^2)$ for sufficiently large n.*

Proof. Since $f_2(n)$ is increasing in n, it suffices to prove the result when n is a multiple of 4, say $n = 4m$. We wish to construct $f_1(n/4)^{3n/2} = f_1(m)^{6m}$ distinct 2-representable complexes. Our

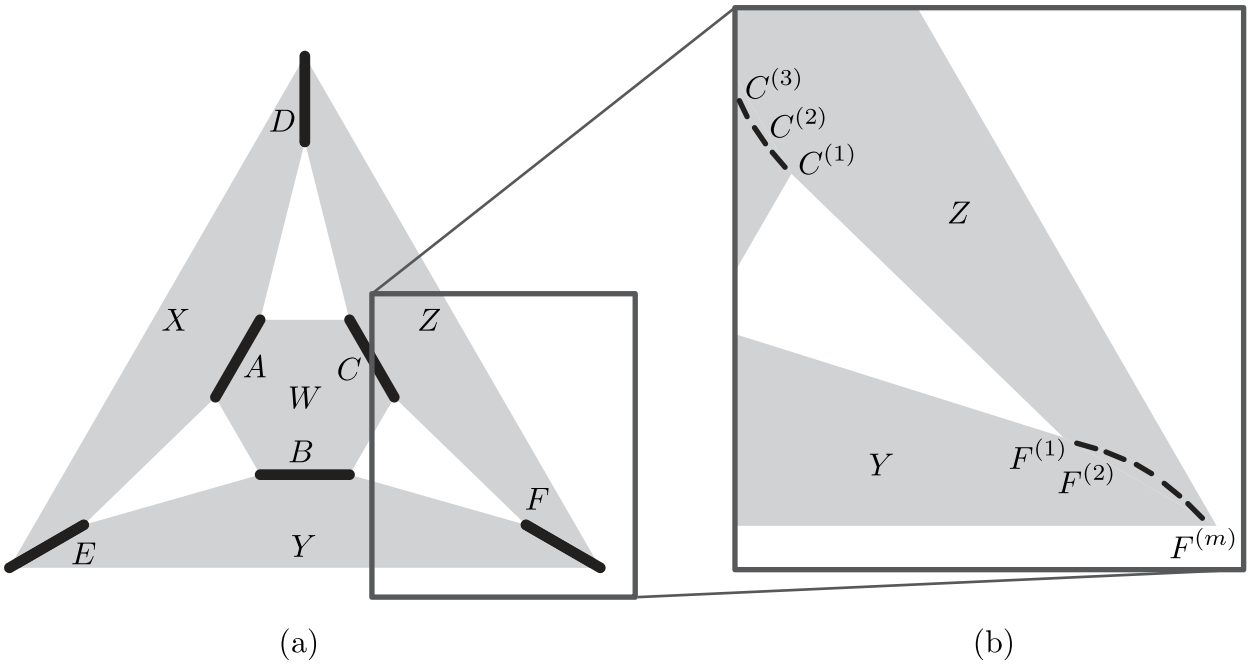

Figure 2: (a) The convex regions W, X, Y and Z in $\mathbb{R}^2$ used in the proof of Proposition 8. (b) The small secants chosen near each arc.

starting point is an arrangement of closed convex sets W, X, Y and Z in the plane, shown in Figure 2. These sets have disjoint interiors, and any pair shares a line segment along the boundary. These line segments are labeled A, B, C, D, E and F in the figure.

Consider replacing each segment by a smooth arc, oriented so that the arc bends closer to the corresponding label in the figure. Choose m arbitrarily small disjoint line segments $A^{(1)}, A^{(2)}, \ldots, A^{(m)}$ whose endpoints lie on the arc A. Note that these segments are in convex position, in the sense that they form a subset of the edges of a convex polygon. Similarly, choose m-many small disjoint line segments in convex position with endpoints along B, C, D, E and F. We thus obtain $6m$ line segments in total. Since these line segments are secants to the arcs A through F, respectively, we call them simply *secants*.

By making our arcs sufficiently close to the original line segments and making the secants sufficiently small, we can guarantee the following properties:

(i) For any $i \in [m]$, the convex hull of $A^{(i)} \cup B^{(i)} \cup C^{(i)}$ does not contain a point in any secant except for $A^{(i)}$, $B^{(i)}$ and $C^{(i)}$.

(ii) For any $i \in [m]$, the secant $D^{(i)}$ together with all $A^{(j)}$s and all $E^{(k)}$s forms a collection of $2m+1$ disjoint line segments in convex position.

(iii) Condition (ii) holds, but with (D, A, E) replaced by the triple (E, B, F). Similarly, it also holds with (D, A, E) replaced by (F, C, D).

Now, for every $i \in [m]$, fix a tuple $(\Delta_{A,i}, \Delta_{B,i}, \ldots, \Delta_{F,i})$ of 1-representable complexes on the vertex set $[m]$. Place a 1-representation of $\Delta_{A,i}$ inside the secant $A^{(i)}$. Let $A_1^{(i)}, \ldots, A_m^{(i)} \subset A^{(i)}$ be the line segments in the representation of $\Delta_{A,i}$. Likewise, place 1-representations of all other complexes inside their corresponding secants, denoting them by $B_j^{(i)}, C_j^{(i)}$ and so forth. We obtain $6m$ 1-representations arranged in the plane, with m-many near each of the segments A, B, C, D, E and F.

We now define a 2-representation of a complex Δ on the vertex set

$$\{w_1, \ldots, w_m, x_1, \ldots, x_m, y_1, \ldots, y_m, z_1, \ldots, z_m\},$$

which can be regarded as a relabeling of $[n]$. For any $i \in [m]$, define $C_{w_i} \stackrel{\text{def}}{=} \text{conv}(A^{(i)} \cup B^{(i)} \cup C^{(i)})$. Moreover, define convex sets

$$C_{x_i} \stackrel{\text{def}}{=} \text{conv}\Big(D^{(i)} \cup \bigcup_{j=1}^{m} A_i^{(j)} \cup \bigcup_{j=1}^{m} E_i^{(j)}\Big),$$
$$C_{y_i} \stackrel{\text{def}}{=} \text{conv}\Big(E^{(i)} \cup \bigcup_{j=1}^{m} B_i^{(j)} \cup \bigcup_{j=1}^{m} F_i^{(j)}\Big),$$
$$C_{z_i} \stackrel{\text{def}}{=} \text{conv}\Big(F^{(i)} \cup \bigcup_{j=1}^{m} C_i^{(j)} \cup \bigcup_{j=1}^{m} D_i^{(j)}\Big).$$

Now, by properties (i) and (ii) above, the sets $C_{x_1} \cap C_{w_i}, \ldots, C_{x_m} \cap C_{w_i}$ form a 1-representation $\Delta_{A,i}$. Thus we may recover $\Delta_{A,i}$ from Δ: it is the link of w_i, restricted to the vertex set $\{x_1, \ldots, x_m\}$. Similar reasoning allows us to recover $\Delta_{B,i}, \ldots, \Delta_{F,i}$ (for example, $\Delta_{F,i}$ is the link of z_i restricted to the vertex set $\{y_1, \ldots, y_m\}$).

Thus every choice of m-many 6-tuples $(\Delta_{A,i}, \Delta_{B,i}, \ldots, \Delta_{F,i})$ of 1-representable complexes on the vertex set $[m]$ yields a unique 2-representable complex on vertex set $[n]$. There are $f_1(m)^{6m}$ ways to choose m-many such 6-tuples, and so the first bound follows. The bound $\log f_2(n) \geq \frac{3}{4} n^2 \log n + O(n)$ follows by taking logarithms and applying Theorem 1. □

3 Interval graphs and interval orders: Proof of Theorem 1

To prove Theorem 1 we will relate the number of interval graphs to the number of interval orders. Given a tuple $(I_1, \ldots, I_n)$ of closed intervals, its *interval order* is the partial order on $[n]$ where $i < j$ if and only if I_i is completely to the left of I_j. Observe that one may recover the interval graph associated to a tuple from the interval order. Indeed, i and j are incomparable if and only if their intervals overlap, and so the interval graph is the incomparability graph of the interval order.

Surprisingly, it is possible to go in the opposite direction: to turn an interval graph into an interval order. That leads to to the following result, which shows that $f_1(n)$, the number of interval graphs, is very close to the number of interval orders.

Theorem 9 (Sandwich theorem)**.** *Let $g(n)$ denote the number of interval orders on $[n]$. Then*

$$g(n-2) \leq f_1(n) \leq g(n).$$

The upper bound in Theorem 9 is immediate. For the lower bound, we construct an injective map

$$\{\text{Interval orders on } [n-2]\} \longrightarrow \{\text{Interval graphs on } [n]\}.$$

Compressed representations. To start, we fix a convenient representation for each interval order on $[n-2]$. Namely, a representation is *compressed* if the following three conditions hold:

(i) Consecutive left endpoints (respectively, right endpoints) are the same.

(ii) The endpoints take values in $[2n-4]$.

(iii) The subset of integers in $[2n-4]$ that are endpoints of some interval form a consecutive sequence starting at 1.

Note that we can easily obtain a compressed representation from any representation by first sliding consecutive endpoints of the same type together, then applying a monotone rescaling and shift to achieve (ii) and (iii), none of which will change the interval order being represented. We will regard a compressed representation as a pair of functions $L\colon [n-2] \to [2n-4]$ and $R\colon [n-2] \to [2n-4]$, where $L(i)$ (respectively, $R(i)$) is the coordinate of the left (respectively, right) endpoint of the i-th interval. In a compressed representation, condition (iii) guarantees that left endpoints only occur at odd values, while right endpoints only occur at even values. With this in mind, it will also be convenient to regard a compressed representation as a pair of lists $\mathcal{A} = (A_1, A_2, \ldots, A_{n-2})$ and $\mathcal{B} = (B_1, B_2, \ldots, B_{n-2})$ where

$$A_i \stackrel{\text{def}}{=} \{j \in [n-2] \mid L(j) = 2i-1\}, \text{ and}$$
$$B_i \stackrel{\text{def}}{=} \{j \in [n-2] \mid R(j) = 2i\}.$$

In other words, $\mathcal{A}$ and $\mathcal{B}$ record the different endpoints which appear at the various coordinates in $[2n-4]$. Note that we may determine (L, R) from $(\mathcal{A}, \mathcal{B})$ and vice versa. Also note that because of condition (iii), if A_i is empty, then so are B_i and A_{i+1}, and if B_i is empty, then so are A_{i+1} and B_{i+1}.

The injective map. Once we have fixed a compressed representation for every interval order on $[n-2]$, we are ready to encode interval orders into interval graphs. Let (L, R) be a compressed representation of an interval order on $[n-2]$, and let $m \stackrel{\text{def}}{=} \max\{R(i) \mid i \in [n-2]\}$. We define new intervals $J_1, J_2, \ldots, J_n$ where

$$J_i \stackrel{\text{def}}{=} [L(i), R(i)+2] \quad \text{for } i \in [n-2],$$
$$J_{n-1} \stackrel{\text{def}}{=} [0, 1], \text{ and}$$
$$J_n \stackrel{\text{def}}{=} [m+1, m+2].$$

Finally, let G be the interval graph represented by $\mathcal{J} = (J_1, J_2, \ldots, J_n)$. We claim that we can recover (L, R) from G. In fact, we will recover $(\mathcal{A}, \mathcal{B})$ iteratively.

For any $i \in [n-2]$, let G_i be the graph obtained from G by deleting $n-1$, and all vertices in $\bigcup_{j=1}^{i-1} B_j$. In particular, G_1 is the result of deleting only the vertex $n-1$. Geometrically, G_i is the interval graph obtained by deleting all intervals in $\mathcal{J}$ which lie strictly to the left of $2i+1$, except possibly J_n.

Proposition 10. *Fix $i \in [n-3]$, and let A be the set of vertices in $\bigcup_{j=1}^{i} A_j$ which appear in G_i. For each $v \in A$, let $N(v)$ be the closed neighborhood of v in G_i. If n is not the only vertex in G_i, then*

(i) *A is not empty,*

(ii) *the various $N(v)$ for $v \in A$ are totally ordered by containment, so there is a unique inclusion-minimal such neighborhood N,*

(iii) *B_i is the set of vertices $v \in A$ with $N(v) = N$, and*

(iv) *A_{i+1} is equal to $N \setminus (A \cup \{n\})$.*

Proof. To prove item (i), observe that since G_i contains a vertex not equal to n, the set B_i must be nonempty. Any $v \in B_i$ is a vertex of G_i, and has $L(v) < R(v) = 2i$, which implies $v \in A$.

For item (ii), observe that A is exactly the set of $v \in [n-2]$ with $L(v) \leq 2i-1$ and $R(v) \geq 2i$. Given $v \in A$, the right endpoint of J_v is $R(v)+2$. Among the various vertices $u \in [n-2]$ that appear in G_i, we see that J_v intersects J_u if and only if $R(v)+2 > L(u)$. This happens if and only if $u \in A$ or $u \in A_j$ for some j satisfying $i+1 \leq j \leq (R(v)+2)/2$. Consequently, for each $v \in A$ the set $N(v)$ is exactly $A \cup \bigcup_{j=i+1}^{j=(R(v)+2)/2} A_j$, where we use the convention $A_{(m+2)/2} \stackrel{\text{def}}{=} \{n\}$ (due to intersection with J_n). Thus the various neighborhoods of vertices in A are totally ordered by containment, with strictly larger neighborhoods arising from strictly larger values of $R(v)$. This proves (ii).

From the above observations, we have $N(v) = N$ if and only if $R(v) = 2i$, that is, if and only if $v \in B_i$. At the beginning of the proof we observed that if $v \in B_i$ then $v \in A$. This proves (iii). Finally, the formula above implies that $N = A \cup A_{i+1}$, plus possibly the vertex n. This union is disjoint, so $N \setminus (A \cup \{n\}) = A_{i+1}$, proving (iv). □

The proposition above implies that if we know G, $A_1, \ldots, A_i$, and $B_1, \ldots, B_{i-1}$ (and hence also G_i), then we may compute A_{i+1} and B_i. Note that A_1 is equal to the neighborhood of $n-1$ in G. Thus, given only G, we may first compute A_1, and then apply the proposition repeatedly to obtain $A_1, \ldots, A_{n-2}$ and $B_1, \ldots, B_{n-3}$. Finally, B_{n-2} will simply consist of the vertices missing from all of the already-computed B_i. Thus we recover $(\mathcal{A}, \mathcal{B})$ from G, and hence recover our original interval order. This proves that our map from interval orders to interval graphs is injective, as desired, concluding the proof of Theorem 9.

Proof of Theorem 1. Brightwell and Keller [4, Theorem 6], building upon earlier works of Zagier [21] and Bousquet-Mélou, Claesson, Dukes and Kitaev [3], proved that the number $g(n)$ of interval orders on $[n]$ satisfies

$$g(n) \sim (n!)^2 \sqrt{n} \left(\frac{6}{\pi^2}\right)^n \left(E_0 + \frac{E_1}{n} + \frac{E_2}{n^2} + \cdots\right)$$

where the E_i are constants. By applying Stirling's approximation and simplifying, we obtain

$$g(n) = e^{2n \log n - (2 + \log(\pi^2/6))n + O(\log n)}.$$

Replacing n by $n-2$, we see that the same asymptotic holds for $g(n-2)$. Thus by Theorem 9 we obtain the same asymptotic equality for $f_1(n)$, proving Theorem 1.

4 Collapsible complexes: proof of Theorem 3

In the proof of Theorem 2 we constructed many d-representable complexes by starting with the case $d = 1$ and using induction on d. For the inductive step, we merged several $(d-1)$-representable complexes on the same vertex set into a single d-representable complex. It turns out that the very same operation turns a collection of $(d-1)$-collapsible complexes into a single d-collapsible complex.

The following is the direct analogue of Proposition 4 for collapsible complexes.

Proposition 11. *Let V and W be disjoint sets. For each $w \in W$, let Δ_w be a $(d-1)$-collapsible complex on vertex set V. Then the simplicial complex*

$$\Delta \stackrel{\text{def}}{=} 2^V \cup \bigcup_{w \in W} (\Delta_w * w)$$

on the vertex set $V \sqcup W$ is d-collapsible.

Proof. For each w, fix a $(d-1)$-collapsing sequence of Δ_w. By adding w to the free faces in this sequence, we obtain a sequence of d-collapses from $\Delta_w * w$ to Δ_w. Since the only faces of Δ that contain w are those in $\Delta_w * w$, we may perform this sequence of d-collapses in Δ to remove all the faces containing w. After doing this for every $w \in W$, we are left with the simplex 2^V, which is d-collapsible. Thus Δ is d-collapsible. □

We use the preceding result to construct many d-collapsible complexes by starting with many 1-collapsible complexes.

Theorem 3 (restated). *The number $g_d(n)$ of d-collapsible complexes on the vertex set $[n]$ satisfies*

$$\tfrac{1}{(d+1)^{d+1}} n^{d+1} + O(n^d) \le \log_2 g_d(n) \le \tfrac{1}{(d+1)!} n^{d+1} + O(n^d).$$

In particular, $g_d(n) = e^{\Theta(n^{d+1})}$.

Proof of the lower bound. Since $g_d(n)$ is increasing in n, it suffices to prove that

$$\log_2 g_d(n) \ge \frac{1}{(d+1)^{d+1}} n^{d+1} \qquad \text{when } n \text{ is divisible by } d+1. \tag{1}$$

We prove this by induction on d. When $d = 1$, $g_d(n)$ is the number of labeled chordal graphs on the vertex set $[n]$, since 1-collapsible complexes are exactly the clique complexes of chordal graphs. Recall that a *split graph* is a graph that can be partitioned into a clique and an independent set (with the edges between being arbitrary). Split graphs are chordal, and one may construct $2^{n^2/4}$ split graphs on the vertex set $[n]$ by forming a clique on the first half of the vertices and adding arbitrary edges between the first and second halves. This proves (1) when $d = 1$.

Assume $d \ge 2$. Let V be the first $\frac{d}{d+1} n$ vertices, and let W be the remaining $\frac{1}{d+1} n$ vertices. For each vertex $w \in W$, let Δ_w be a $(d-1)$-collapsible complex on vertex set V. Observe that we may recover Δ_w from the d-collapsible simplicial complex Δ defined in Theorem 3, namely Δ_w is the link of w in this complex. Thus we obtain a unique d-collapsible complex on the vertex set $[n]$ for every choice of the various Δ_w. This implies that $g_d(n) \ge g_{d-1}(\frac{d}{d+1} n)^{n/(d+1)}$. Taking logarithms and applying the inductive hypothesis, we obtain the lower bound as follows:

$$\log_2 g_d(n) \ge \log_2 \left(g_{d-1}(\tfrac{d}{d+1} n)^{n/(d+1)} \right) \ge \tfrac{n}{d+1} \cdot \tfrac{1}{d^d} \tfrac{d^d}{(d+1)^d} n^d = \tfrac{1}{(d+1)^{d+1}} n^{d+1}.$$ □

Proof of the upper bound. It is well-known (see for example the discussion in [10]) that d-collapsible complexes satisfy Helly's theorem, and so are determined by their d-skeleta. Since the d-skeleton of an n-vertex complex can be specified by which of the $N \stackrel{\text{def}}{=} \binom{n}{d+1} + \binom{n}{d} + \cdots + \binom{n}{0}$ possible faces belong to it, it follows the number of d-collapsible complexes on $[n]$ is at most 2^N. □

5 Problems and remarks

- Bender, Richmond and Wormald [2] proved that almost all chordal graphs are split, and so the lower bound on the number of 1-collapsible complexes in Theorem 3 is tight. We suspect that the lower bound remains tight also in higher dimensions, i.e., that $\log_2 g_d(n) = \frac{1}{(d+1)^{d+1}} n^{d+1} + O(n^d)$ for all $d \geq 1$.

- The lower bound in the Sandwich Theorem (Theorem 9 in Section 3) can be improved to $\frac{n(n-1)}{2} g(n-2) \leq f_1(n)$. To obtain this bound one must first note that the interval graph we construct has an essentially unique representation by intervals, up to reflection about a common point. This follows from the fact that the interval graph has no "buried subgraphs" in the sense of Hanlon [8]. Thus even when we regard the graph as unlabeled, this unique representation allows us to recognize the vertices n and $n-1$ as the outermost intervals, but not to tell these two vertices apart. Hence, by choosing which two elements of $[n]$ label these vertices, we can upgrade the injective map in the proof of Sandwich Theorem to a map

$$\{\text{Interval orders on } [n-2]\} \times \{(i,j) \in [n]^2 \mid i \neq j\} \longrightarrow \{\text{Interval graphs on } [n]\}$$

  that is at most 2-to-1.

- Our results easily extend to the *unlabeled* case. Indeed, the proof of Sandwich Theorem yields an injective map

$$\{\text{Unlabeled interval orders on } [n-2]\} \longrightarrow \{\text{Doubly-rooted interval graphs on } [n]\}.$$

  So, writing $\overline{f_d}(n)$ and $\overline{g}(n)$ for the unlabeled versions of $f_d(n)$ and $g(n)$, we obtain the inequality $\frac{1}{n(n-1)} \overline{g}(n-2) \leq \overline{f_1}(n) \leq \overline{g}(n)$. As above, the factor of $\frac{1}{n(n-1)}$ can be upgraded to $\frac{1}{2}$ by appealing to Hanlon's work [8]. Either way, from the enumeration of unlabeled interval orders [4, Theorem 1] we deduce that $\overline{f_1}(n) = e^{n \log n - (1 + \log(\pi^2/6))n + O(\log n)}$.

  Since the number of labeled and the number of unlabeled d-representable complexes differ by a factor of at most $n!$, and $f_d(n) = e^{\Theta(n^d \log n)}$, it follows that $\overline{f_d}(n) = e^{\Theta(n^d \log n)}$ for $d \geq 2$ as well. Similarly, the number of unlabeled d-collapsible complexes is $e^{\Theta(n^{d+1})}$ for $d \geq 1$, as per Theorem 3.

- It is natural to consider, instead of the entire nerve, a skeleton thereof. Let $f_{d,k}(n)$ be the number of k-dimensional simplicial complexes that are k-skeletons of some d-representable complex on the vertex set $[n]$. This paper treated the case $k = d$. The case $(k,d) = (1,2)$ was treated in [12], where it is shown that almost almost every intersection graph of convex sets in the plane can be partitioned into 4 parts such that 3 of them induce a clique and the 4th one splits into

two cliques with no edge running between them. The case $d \geq 2k+1$ is trivial because every k-complex is a k-skeleton of some d-representable complex [17, 13]. We do not know how $f_{d,k}$ behaves for any other pair (k, d).

- The constructions of large families of d-representable complexes (Propositions 4 and 8) suggest the following problem. Say that a pair of convex sets C and C' are *tangent to each other* if there is a point $p \in \partial C \cap \partial C'$ such that both C and C' are smooth at p, have the same tangent hyperplane at p, but lie on the opposite sides of the hyperplane. Both Propositions 4 and 8 implicitly construct families of (not necessarily distinct) convex sets with many tangencies. For example, Proposition 8 constructs a family of n convex sets in the plane whose tangency graph is the complete 4-partite graph with $\approx \frac{3}{4}\binom{n}{2}$ edges. We do not know if the maximum number of tangencies among n convex sets in $\mathbb{R}^d$ is $\frac{d+1}{d+2}\binom{n}{2} + o(n^2)$.

References

[1] Hüseyin Acan, Sankardeep Chakraborty, Seungbum Jo, and Srinivasa Rao Satti. Succinct encodings for families of interval graphs. *Algorithmica*, 83(3):776–794, 2021. arXiv:1902.09228.

[2] E. A. Bender, L. B. Richmond, and N. C. Wormald. Almost all chordal graphs split. *J. Austral. Math. Soc. Ser. A*, 38(2):214–221, 1985.

[3] Mireille Bousquet-Mélou, Anders Claesson, Mark Dukes, and Sergey Kitaev. $(2+2)$-free posets, ascent sequences and pattern avoiding permutations. *J. Combin. Theory Ser. A*, 117(7):884–909, 2010. arXiv:0806.0666.

[4] Graham Brightwell and Mitchel T. Keller. Asymptotic enumeration of labelled interval orders. arXiv:1111.6766, November 2011.

[5] Joel E. Cohen, János Komlós, and Thomas Mueller. The probability of an interval graph, and why it matters. In *Relations between combinatorics and other parts of mathematics (Proc. Sympos. Pure Math., Ohio State Univ., Columbus, Ohio, 1978)*, Proc. Sympos. Pure Math., XXXIV, pages 97–115 (loose errata). Amer. Math. Soc., Providence, R.I., 1979.

[6] Jürgen Eckhoff. An upper-bound theorem for families of convex sets. *Geometriae Dedicata*, 19:217–227, 1985.

[7] Cyril Gavoille and Christophe Paul. Optimal distance labeling for interval graphs and related graph families. *SIAM J. Discrete Math.*, 22(3):1239–1258, 2008.

[8] Phil Hanlon. Counting interval graphs. *Trans. Amer. Math. Soc.*, 272(2):383–426, 1982.

[9] Gil Kalai. Intersection patterns of convex sets. *Israel J. Math.*, 48(2-3):161–174, 1984.

[10] Jiří Matoušek and Martin Tancer. Dimension gaps between representability and collapsibility. *Discrete Comput. Geom.*, 42(4):631–639, 2009. arXiv:0803.3520.

[11] Patryk Mikos. Efficient enumeration of non-isomorphic interval graphs. *Discrete Math. Theor. Comput. Sci.*, 23(1):Paper No. 2, 23, 2021. arXiv:1906.04094.

[12] János Pach, Bruce Reed, and Yelena Yuditsky. Almost all string graphs are intersection graphs of plane convex sets. *Discrete Comput. Geom.*, 63(4):888–917, 2020.

[13] G. Ya. Perel'man. Realization of abstract k-skeletons as k-skeletons of intersections of convex polyhedra in $\mathbf{R}^{2k-1}$. In *Geometric questions in the theory of functions and sets*, pages 129–131. Kalinin. Gos. Univ., Kalinin, 1985.

[14] Martin Tancer. Non-representability of finite projective planes by convex sets. *Proceedings of the American Mathematical Society*, 138(9):3285–3291, 2010.

[15] Martin Tancer. Intersection patterns of convex sets via simplicial complexes: a survey. In *Thirty essays on geometric graph theory*, pages 521–540. Springer, New York, 2013. arXiv:1102.0417.

[16] Martin Tancer. Recognition of collapsible complexes is NP-complete. *Discrete and Computational Geometry*, 55:21–38, 2016. arXiv:1211.6254.

[17] Gerd Wegner. *Eigenschaften der Nerven homologisch-einfacher Familien im* $\mathbb{R}^n$. PhD thesis, Universität Göttingen, 1967.

[18] Gerd Wegner. d-collapsing and nerves of families of convex sets. *Archiv der Mathematik*, 26:317–321, 1975.

[19] Kazuaki Yamazaki, Toshiki Saitoh, Masashi Kiyomi, and Ryuhei Uehara. Enumeration of nonisomorphic interval graphs and nonisomorphic permutation graphs. *Theoret. Comput. Sci.*, 806:310–322, 2020.

[20] Joyce C. Yang and Nicholas Pippenger. On the enumeration of interval graphs. *Proc. Amer. Math. Soc. Ser. B*, 4:1–3, 2017. arXiv:1609.02479.

[21] Don Zagier. Vassiliev invariants and a strange identity related to the Dedekind eta-function. *Topology*, 40(5):945–960, 2001.