d-collapsibility is NP-complete for $d \geq 4$

Martin Tancer*

November 28, 2009

Abstract

A simplicial complex is *d-collapsible* if it can be reduced to an empty complex by repeatedly removing (collapsing) a face of dimension at most $d-1$ that is contained in a unique maximal face. We prove that the algorithmic question whether a given simplicial complex is d-collapsible is NP-complete for $d \geq 4$ and polynomial time solvable for $d \leq 2$.

As an intermediate step, we prove that d-collapsibility can be recognized by the greedy algorithm for $d \leq 2$, but the greedy algorithm does not work for $d \geq 3$.

A simplicial complex is *d-representable* if it is the nerve of a collection of convex sets in $\mathbb{R}^d$. The main motivation for studying d-collapsible complexes is that every d-representable complex is d-collapsible. We also observe that known results imply that analogical algorithmic question for d-representable complexes is NP-hard for $d \geq 2$.

1 Introduction

Our task is to determine the computational complexity of recognition of *d-collapsible* simplicial complexes. These complexes were introduced by Wegner [Weg75] and studying them is motivated by Helly-type theorems, which we will discuss later. All the simplicial complexes[1] throughout the article are assumed to be finite.

*Department of Applied Mathematics and Institute for Theoretical Computer Science (supported by project 1M0545 of The Ministry of Education of the Czech Republic), Faculty of Mathematics and Physics, Charles University, Malostranské nám. 25, 118 00 Prague, Czech Republic. E-mail: tancer@kam.mff.cuni.cz

[1] We assume that the reader is familiar with simplicial complexes; introductory chapters of books like [Mat03, Hat01, Mun84] should provide a sufficient background. Unless stated otherwise, we work with abstract simplicial complexes, i.e., set systems K such that if $A \in \mathsf{K}$ and $B \subseteq A$ then $B \in \mathsf{K}$.

d-collapsible complexes. A face σ of a simplicial complex K is *collapsible* if it is contained in a unique maximal face of K (by "maximal face" we always mean "inclusion-maximal face"). Unless stated otherwise, we denote this maximal face by $\tau(\sigma)$. (We allow $\tau(\sigma) = \sigma$!) Moreover, if $\dim \sigma \leq d-1$, then σ is *d-collapsible.* By $[\sigma, \tau(\sigma)]$ we denote the set

$$\{\eta \in \mathsf{K} \mid \sigma \subseteq \eta \subseteq \tau(\sigma)\}$$

of faces of K that contain σ.

We assume that σ is d-collapsible and we say that the complex $\mathsf{K}' = \mathsf{K} \setminus [\sigma, \tau(\sigma)]$ arises from K by an *elementary d-collapse.* We denote this by

$$\mathsf{K} \to \mathsf{K}'.$$

If we want to stress σ we write

$$\mathsf{K}' = \mathsf{K}_\sigma.$$

A complex K *d-collapses* to a complex L, in symbols $\mathsf{K} \twoheadrightarrow \mathsf{L}$, if there is a sequence of elementary d-collapses

$$\mathsf{K} \to \mathsf{K}_2 \to \mathsf{K}_3 \to \cdots \to \mathsf{L}.$$

This sequence is called a *d-collapsing* (of K to L). Finally, a complex K is *d-collapsible* if $\mathsf{K} \twoheadrightarrow \emptyset$. An example of 2-collapsible complex is in Figure 1.

The computational complexity of d-collapsibility. How hard is it to decide whether a given simplicial complex is d-collapsible? We consider the computational complexity of this question (the size of an input is the number of faces of the complex in the question), regarding d as a fixed integer; we refer to it as d-COLLAPSIBILITY.

According to Lekkerkerker and Boland [LB62] (see also [Weg75]), 1-collapsible complexes are exactly clique complexes over *chordal graphs.* (A graph is chordal if it does not contain an induced cycle of size 4 or more.) 1-COLLAPSIBILITY is therefore polynomial time solvable. (Polynomiality of 1-COLLAPSIBILITY also follows from Theorem 1.2(i).)

The main result of this paper is the following:

Theorem 1.1. (i) *2-COLLAPSIBILITY is polynomial time solvable.*

(ii) *d-COLLAPSIBILITY is NP-complete for $d \geq 4$.*

Suppose that d is fixed. A *good face* is a d-collapsible face of K such that K_σ is d-collapsible; a *bad face* is d-collapsible face of K such that K_σ is not d-collapsible.

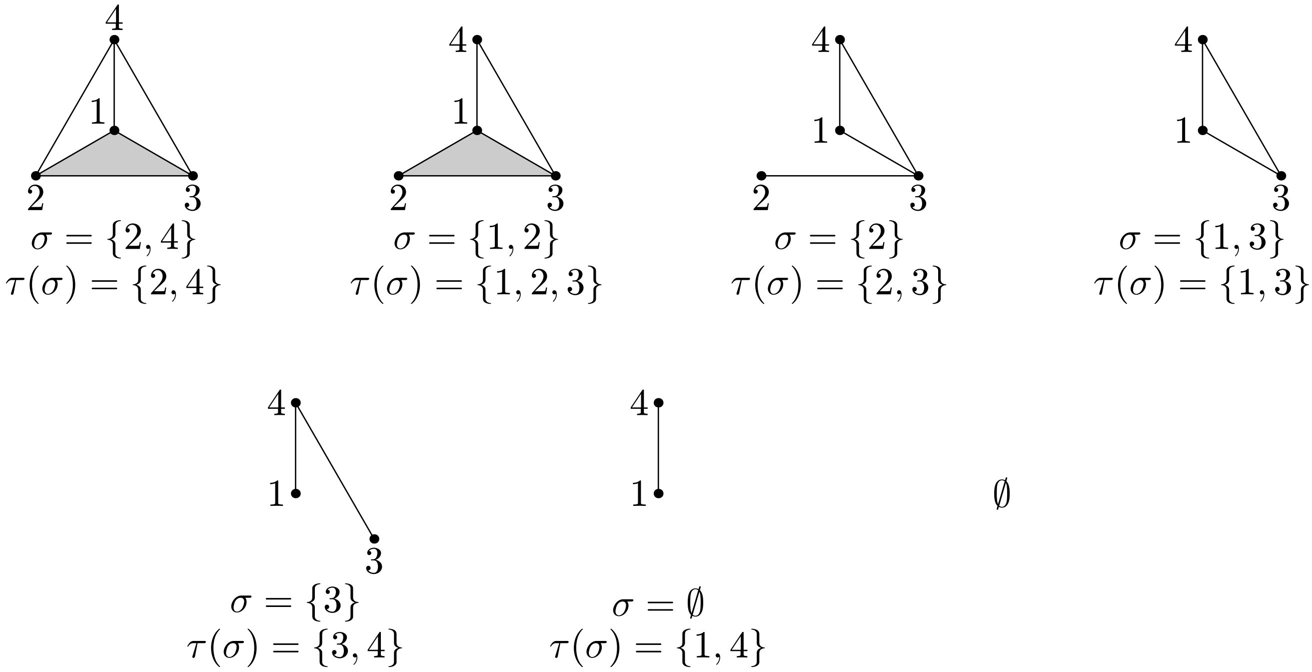

Figure 1: An example of 2-collapsing.

Now suppose that K is a d-collapsible complex. It is not immediately clear whether we can choose elementary d-collapses greedily in any order to d-collapse K, or whether there is a "bad sequence" of d-collapses such that the resulting complex is no longer d-collapsible. Therefore, we consider the following question: For which d there is a d-collapsible complex K such that it contains a bad face? The answer is:

Theorem 1.2. (i) *Let $d \leq 2$. Then every d-collapsible face of a d-collapsible complex is good.*

(ii) *Let $d \geq 3$. Then there exists a d-collapsible complex containing a bad d-collapsible face.*

Theorem 1.1(i) is a straightforward consequence of Theorem 1.2(i). Indeed, if we want to test whether a given complex is 2-collapsible, it is sufficient to greedily collapse d-collapsible faces. Theorem 1.2(i) implies that we finish with an empty complex if and only if the original complex is 2-collapsible.

Our construction for Theorem 1.2(ii) is an intermediate step to prove Theorem 1.1(ii).

***d*-representable complexes.** Famous Helly's theorem [Hel23] asserts that if $C_1, C_2, \ldots, C_n$ are convex sets in $\mathbb{R}^d$, $n \geq d+1$, and every $d+1$ of them have a common point, then $C_1 \cup C_2 \cup \cdots \cup C_n \neq \emptyset$. This theorem (and several other theorems in discrete geometry) deals with *intersection patterns* of convex sets in $\mathbb{R}^d$. It can be restated using the notion of *d-representable* complexes, which "record" the intersection patterns.

The *nerve* of a family $\mathcal{S} = \{S_1, S_2, \ldots, S_n\}$ of sets is the simplicial complex with vertex set $[n] = \{1, 2, \ldots, n\}$ and with the set $\sigma \subseteq [n]$ forming a face if $\bigcap_{i\in\sigma} S_i \neq \emptyset$. A simplicial complex is d-representable if it is isomorphic to the nerve of a family of convex sets in $\mathbb{R}^d$.

In this language, Helly theorem states that if a d-representable complex (with the vertex set V) contains all faces of dimension at most d, then it is already a *full simplex* 2^V. Beside Helly's theorem we also mention several other known results that can be formulated using d-representability. They include the fractional Helly theorem of Katchalski and Liu [KL79], the colorful Helly theorem of Lovász ([Lov74]; see also [Bár82]), the (p, q)-theorem of Alon and Kleitman [AK92], and the Helly type result of Amenta [Ame96].

***d*-Leray complexes.** Another related notion is a *d-Leray* simplicial complex, where K is d-Leray if every *induced subcomplex* of K (i.e. a subcomplex of the form $\mathsf{K}[X] = \{\sigma \cap X \mid \sigma \in \mathsf{K}\}$ for some subset X of the vertex set of K) has zero homology (over $\mathbb{Q}$) in dimensions d and larger. We will mention d-Leray complexes only briefly, thus the article should be accessible also for the reader not familiar with homology.

Relations among the preceding notions. Wegner [Weg75] proved that d-collapsible complexes are d-Leray and also that every d-representable complex is *d-collapsible.*

Many results on d-representable complexes can be generalized in terms of d-collapsible complexes, the results mentioned here even for d-Leray complexes.

For example, a topological generalization of Helly's theorem follows from Helly's own work [Hel30], a generalization of fractional Helly's theorem and (p, q)-theorem was done in [AKMM02], and a generalization of colorful Helly's theorem and Amenta's theorem was proved by Kalai and Meshulam [KM05], [KM08].

Dimensional gaps between collapsibility and representability were studied by Matoušek and this author [MT08]; an interesting variation on d-collapsibility was used by Matoušek in order to show that it is not easy to remove degeneracy in LP-type problems [Mat08].

Related complexity results. Similarly as d-COLLAPSIBILITY, we can also consider the computational complexity of d-REPRESENTABILITY and d-LERAY COMPLEX.

By a modification of result of Kratochvíl and Matoušek on string graphs ([KM89]; see also [Kra91]), one has that 2-REPRESENTABILITY is NP-hard. Moreover, this result also implies that d-REPRESENTABILITY is NP-hard for $d \geq 2$. Details are given in Appendix A.

Finally, d-LERAY COMPLEX is polynomial time solvable, since an equivalent characterization of d-Leray complexes is that it is sufficient to test whether the homology (of dimension greater or equal to d) of *links*[2] of faces of the complex in the question vanishes. These tests can be performed in a polynomial time; see [Mun84] (note that the k-th homology of a complex of dimension less than k is always zero; note also that the homology is over $\mathbb{Q}$, which simplifies the situation—counting homology for this case is indeed only a linear algebra).

We see that the notions of d-representable, d-collapsible and d-Leray complexes have a different behavior in the complexity questions (although they behave similarly for Helly-type theorems as discussed above).

Collapsibility in Whitehead's sense. Beside d-collapsibility, collapsibility in Whitehead's sense is much better known (called simply *collapsibility*). In the case of collapsibility, we allow only to collapse a face σ that is a proper subface of the unique maximal face containing σ. On the other hand, there is no restriction on dimension of σ.

Malgouyres and Francés [MF08] proved that it is NP-complete to decide, whether a given 3-dimensional complex collapses to a given 1-dimensional complex. However, their construction does not apply for d-collapsibility. A key ingredient of their construction is that collapsibility distinguishes a Bing's house with thin walls and a Bing's house with a thick wall. However, they are not distinguishable from the point of view of d-collapsibility. They are both 3-collapsible, but none of them is 2-collapsible.

Technical issues. Throughout this paper we will use several technical lemmas about d-collapsibility. Since I think that the main ideas of the paper can be followed even without these lemmas I decided to put them separately to Section 5. The reader is encouraged to skip them for the first reading and look at them later for full details.

2 2-collapsibility

Here we prove Theorem 1.2(i).

The case $d = 1$ follows from the fact that d-collapsible complexes coincide with d-Leray ones ([LB62, Weg75]). Indeed, let K be a 1-collapsible complex and let σ be its 1-collapsible face. We have that K is 1-Leray, which implies that K_σ is 1-Leray (1-collapsing does not affect homology of dimensions 1 and

[2]A link of a face σ in a complex K is the complex $\{\eta \in \mathsf{K} \mid \eta \cup \sigma \in \mathsf{K}, \eta \cap \sigma = \emptyset\}$.

more). This implies that K_σ is 1-collapsible, i.e., σ is good. In fact, the case $d = 1$ can be also solved by a similar (simpler) discussion as the following case $d = 2$.

It remains to consider the case $d = 2$. Suppose that K is a 2-collapsible complex which, for contradiction, contains a bad 2-collapsible face $\sigma_B \in \mathsf{K}$. On the other hand, it also contains a good face σ_G since it is 2-collapsible. Moreover, we can, without loss of generality, suppose that K is the smallest complex (according to the number of faces) with these properties.

Claim 2.1. *Let σ be a good face of K and let σ' be a 2-collapsible face of K_σ. Then σ' is a good face of K_σ.*

Proof. The complex K_σ is 2-collapsible since σ is a good face of K. If σ' were a bad face of K_σ, then K_σ would be a smaller counterexample to Theorem 1.2(i) contradicting the choice of K. ☐

Recall that $\tau(\sigma)$ denotes the unique maximal superface of a collapsible face σ. Two collapsible faces σ and σ' are *independent* if $\tau(\sigma) \neq \tau(\sigma')$; otherwise, they are *dependent*. The symbol $\mathrm{St}(\sigma, \mathsf{K})$ denotes the (open) *star* of simplex σ in K, which consists of all superfaces of σ in K (including σ). We remark that $\mathrm{St}(\sigma, \mathsf{K}) = [\sigma, \tau(\sigma)]$ in case that σ is collapsible.

Claim 2.2. *Let $\sigma, \sigma' \in \mathsf{K}$ be independent 2-collapsible faces. Then σ is a 2-collapsible face of $\mathsf{K}_{\sigma'}$, σ' is a 2-collapsible face of K_σ, and $(\mathsf{K}_\sigma)_{\sigma'} = (\mathsf{K}_{\sigma'})_\sigma$.*

Proof. Since $\tau(\sigma) \neq \tau(\sigma')$, we have $\sigma \not\subseteq \tau(\sigma')$. Thus, $\mathrm{St}(\sigma, \mathsf{K}) = \mathrm{St}(\sigma, \mathsf{K}_{\sigma'})$, implying that $\tau(\sigma)$ is also a unique maximal face containing σ when considered in $\mathsf{K}_{\sigma'}$. It means that σ is a collapsible face of $\mathsf{K}_{\sigma'}$. Symmetrically, σ' is a collapsible face of K_σ. Finally,

$$(\mathsf{K}_\sigma)_{\sigma'} = (\mathsf{K}_{\sigma'})_\sigma = \mathsf{K} \setminus \{\eta \in \mathsf{K} \mid \sigma \subseteq \eta \text{ or } \sigma' \subseteq \eta\}\,.$$

☐

Claim 2.3. *Any two 2-collapsible faces of K are dependent.*

Proof. For contradiction, let σ, σ' be two independent 2-collapsible faces in K. First, suppose that one of them is good, say σ, and the second one, i.e. σ', is bad. The face σ' is a collapsible face of K_σ by Claim 2.2. Thus, $(\mathsf{K}_\sigma)_{\sigma'}$ is 2-collapsible by Claim 2.1. But $(\mathsf{K}_\sigma)_{\sigma'} = (\mathsf{K}_{\sigma'})_\sigma$ by Claim 2.2, which contradicts the assumption that σ' is a bad face.

Now suppose that σ and σ' are good faces. Then at least one of them is independent with σ_B, which yields the contradiction as in the previous case. Similarly, if both of σ and σ' are bad faces, then at least one of them is independent with σ_G. ☐

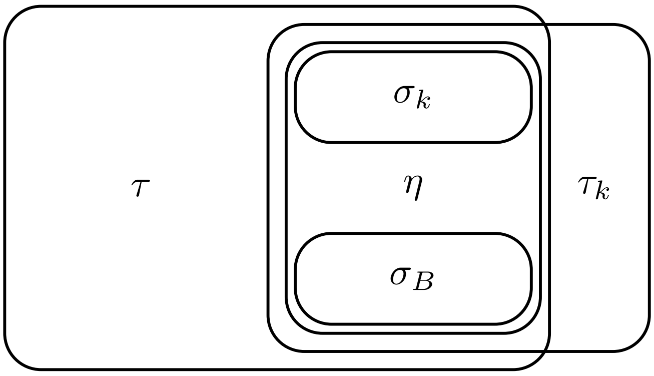

Figure 2: The simplices τ, τ_k and η.

Due to Claim 2.3 there exists a universal $\tau \in \mathsf{K}$ such that $\tau = \tau(\sigma)$ for every 2-collapsible $\sigma \in \mathsf{K}$. Let us remark that $\mathsf{K} \neq 2^\tau$ since σ_B is a bad face.

The following claim represents a key difference among 2-collapsibility and d-collapsibility for $d \geq 3$. It wouldn't be valid in case of d-collapsibility.

Claim 2.4. *Let σ be a good face and let σ' be a bad face. Then $\sigma \cap \sigma' = \emptyset$.*

Proof. It is easy to prove the claim in the case that either σ or σ' is a 0-face. Let us therefore consider the case that both σ and σ' are 1-faces. For contradiction suppose that $\sigma \cap \sigma' \neq \emptyset$, i.e. $\sigma = \{u, v\}$, $\sigma' = \{v, w\}$ for some mutually different $u, v, w \in \tau$. Then $\tau \setminus \{u\}$ is a unique maximal face in K_σ that contains σ', so $(\mathsf{K}_\sigma)_{\sigma'}$ exists. Similarly, $(\mathsf{K}_{\sigma'})_\sigma$ exists and the same argument as in the proof of Claim 2.2 yields $(\mathsf{K}_\sigma)_{\sigma'} = (\mathsf{K}_{\sigma'})_\sigma$. Similarly as in the proof of Claim 2.3, $(\mathsf{K}_\sigma)_{\sigma'}$ is 2-collapsible (due to Claim 2.1), but it contradicts the fact that σ' is a bad face. □

The complex K is 2-collapsible. Let $\mathsf{K} = \mathsf{K}_1 \to \mathsf{K}_2 \to \cdots \to \mathsf{K}_m = \emptyset$ be a 2-collapsing of K, where $\mathsf{K}_{i+1} = \mathsf{K}_i \setminus [\sigma_i, \tau_i]$. Clearly, $\tau_1 = \tau$. Let k be the minimal integer such that $\tau_k \not\subseteq \tau$. Such k exists since $\mathsf{K} \neq 2^\tau$.

Claim 2.5. *The face σ_k is a subset of τ, and it is not a 2-collapsible face of K.*

Proof. Suppose for contradiction that $\sigma_k \not\subseteq \tau$. Then $\mathrm{St}(\sigma_k, \mathsf{K}) = \mathrm{St}(\sigma_k, \mathsf{K}_i)$, since only subsets of τ are removed from K during the first i 2-collapses. It implies that σ_k is a 2-collapsible face of K contradicting the definition of τ.

It is not a 2-collapsible face of K since it is contained in τ and $\tau_k \not\subseteq \tau$. □

From the minimality of k, all the faces $\sigma_0, \sigma_1, \ldots, \sigma_{k-1}$ are good faces of K. Let $\eta = \sigma_k \cup \sigma_B$. See Figure 2. Claim 2.5 implies that $\eta \subseteq \tau$. By Claim 2.4 (and the fact that σ_k is not a good face—a consequence of Claim 2.5) the face η does not contain a good face. Thus, $\eta \in \mathsf{K}_k$. In particular $\eta \subseteq \tau_k$ since τ_k is a unique maximal face of K_k containing σ_k, hence $\sigma_B \subseteq \tau_k$. On the other hand, τ is a unique maximal face of $\mathsf{K} \supseteq \mathsf{K}_k$ containing σ_B, since σ_B is a 2-collapsible face, which implies $\tau_k \subseteq \tau$. It is a contradiction that $\tau_k \not\subseteq \tau$. □

3 d-collapsible complex with a bad d-collapse

In this section we prove Theorem 1.2(ii). Our proof relies on constructing d-dimensional d-collapsible complex C such that its first d-collapse is unique. We call this complex a *connecting gadget*. Precise properties of the connecting gadget are stated in Proposition 3.1.

Before stating the proposition we define the notion of *distant faces*. Suppose that K is a simplicial complex and let u, v be two of its vertices. By $\mathrm{dist}(u, v)$ we mean their distance in graph-theoretical sense in the 1-skeleton of K. We say that two faces $\omega, \eta \in \mathsf{K}$ are *distant* if $\mathrm{dist}(u, v) \geq 3$ for every $u \in \omega, v \in \eta$.

Proposition 3.1. *Let $d \geq 2$ and $t \geq 0$ be integers. There is a d-dimensional complex $\mathsf{C} = \mathsf{C}(\rho; \zeta_1, \ldots, \zeta_t)$ with the following properties:*

(i) *It contains $(d-1)$-dimensional faces $\rho, \zeta_1, \ldots, \zeta_t$ such that each two of them are distant faces.*

(ii) *Let $\mathsf{C}' = \mathsf{C}'(\rho; \zeta_1, \ldots, \zeta_t)$ be the subcomplex of C given by $\mathsf{C}' = 2^{\rho} \cup 2^{\zeta_1} \cup \cdots \cup 2^{\zeta_t}$. Then $\mathsf{C} \twoheadrightarrow (\mathsf{C}' \setminus \{\rho\})$. In particular, C is d-collapsible, since $(\mathsf{C}' \setminus \{\rho\})$ is d-collapsible.*

(iii) *The only d-collapsible face of C is the face ρ.*

(iv) *Suppose that d is a constant. Then the number of faces of C is $O(t)$.*

3.1 The complex $\mathsf{C}(\rho)$

We start our construction assuming $t = 0$; i.e., we construct the connecting gadget $\mathsf{C} = \mathsf{C}(\rho)$.

The geometric realization of $\mathsf{C}(\rho)$. First, we describe the geometric realization, $\|\mathsf{C}\|$, of C. Let P be the d-dimensional *crosspolytope*, the convex hull

$$\mathrm{conv}\,\{\mathbf{e}_1, -\mathbf{e}_1, \ldots, \mathbf{e}_d, -\mathbf{e}_d\}$$

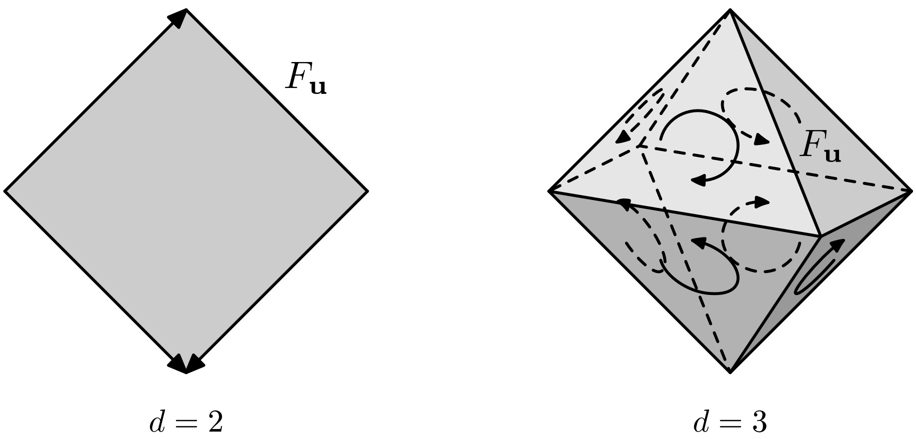

Figure 3: The space X. The arrows denote, which facets are identified.

of the vectors of the standard orthonormal basis and their negatives. It has 2^d *facets*

$$F_{\mathbf{s}} = \operatorname{conv}\{s_1\mathbf{e}_1, \ldots, s_d\mathbf{e}_d\},$$

where $\mathbf{s} = (s_i)_{i=1}^d \in \{-1,1\}^d$ ($\mathbf{s}$ for *sign*). We want to glue all facets together except the facet $F_{\mathbf{u}}$ where $\mathbf{u} = (1, \ldots, 1)$ (see Figure 3).

More precisely, let $\mathbf{s} \in \{-1,1\}^d \setminus \{\mathbf{u}\}$. Every $\mathbf{x} \in F_{\mathbf{s}}$ can be uniquely written as a convex combination $\mathbf{x} = \mathbf{x}_{\mathbf{a},\mathbf{s}} = a_1s_1\mathbf{e}_1 + \cdots + a_ds_d\mathbf{e}_d$ where $\mathbf{a} = (a_i)_{i=1}^d \in [0,1]^d$ and $\sum_{i=1}^d a_i = 1$. For every such fixed $\mathbf{a}$ we glue together the points in the set $\left\{\mathbf{x}_{\mathbf{s},\mathbf{a}} \,\middle|\, \mathbf{s} \in \{-1,1\}^d \setminus \{\mathbf{u}\}\right\}$; by X we denote the resulting space. We will construct C in such a way that X is a geometric realization of C.

Triangulations of the crosspolytope. We define two auxiliary triangulations of P—they are depicted in Figure 4. The simplicial complex J is the simplicial complex with vertex set $\{\mathbf{0}, \mathbf{e}_1, -\mathbf{e}_1, \ldots, \mathbf{e}_d, -\mathbf{e}_d\}$. The set of its faces is given by the maximal faces

$$\{\mathbf{0}, s_1\mathbf{e}_1, s_2\mathbf{e}_2, \ldots, s_d\mathbf{e}_d\} \text{ where } s_1, s_2, \ldots, s_d \in \{-1,1\}.$$

The complex J is a triangulation of P.

Let ϑ be the face $\{\mathbf{0}, \mathbf{e}_1, \ldots, \mathbf{e}_d\}$. The complex H is constructed by iterated *stellar subdivisions* starting with J and subdividing faces of $\mathsf{J} \setminus 2^{\vartheta}$ (first subdividing d-dimensional faces, then $(d-1)$-dimensional, etc.). Formally, H is a complex with the vertex set $(\mathsf{J} \setminus 2^{\vartheta}) \cup \vartheta$ and with faces of the form

$$\{\{\sigma_1, \ldots, \sigma_k\} \cup \tau\} \text{ where } \sigma_1 \supsetneq \cdots \supsetneq \sigma_k \supsetneq \tau; \sigma_1, \ldots, \sigma_k \in \mathsf{J}\setminus 2^{\vartheta}; \tau \subseteq \vartheta; k \in \mathbb{N}_0.$$

The construction of C. Informally, we obtain C from H by the same gluing as was used for constructing X from P.

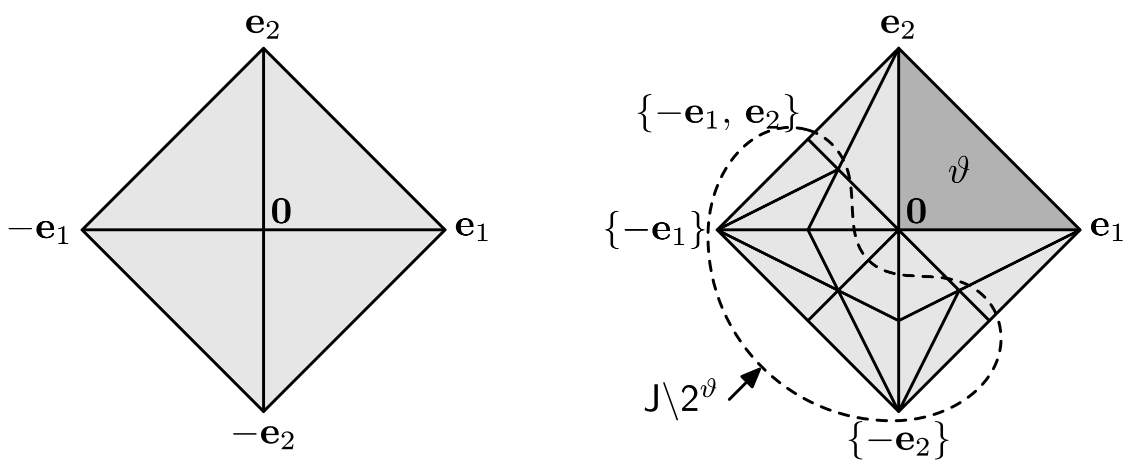

Figure 4: The triangulations J (left) and H (right) of P with $d = 2$.

Formally, let $\approx$ be an equivalence relation on $(\mathsf{J} \setminus 2^{\vartheta}) \cup \vartheta$ given by

$$\begin{array}{lcll} \mathbf{e}_i & \approx & \{-\mathbf{e}_i\} & \text{for } i \in [d], \\ \sigma_1 & \approx & \sigma_2 & \text{for } \sigma_1, \sigma_2 \in \mathsf{J} \setminus 2^{\vartheta}, \\ & & & \sigma_1 = \{s_1\mathbf{e}_{k_1}, \ldots, s_m\mathbf{e}_{k_m}\},\ \sigma_2 = \{s'_1\mathbf{e}_{k_1}, \ldots, s'_m\mathbf{e}_{k_m}\} \\ & & & \text{where } s_i, s'_i \in \{-1, 1\} \text{ and } 1 \le k_1 < \cdots < k_m \le d. \end{array}$$

For an equivalence relation $\equiv$ on a set X we define $\langle\equiv\rangle$ to be an equivalence relation on $\mathcal{Y} \subset 2^X$ inherited from $\equiv$; i.e., we have, for $Y_1, Y_2 \in \mathcal{Y}$, $Y_1\langle\equiv\rangle Y_2$ if and only if there is a bijection $f\colon Y_1 \to Y_2$ such that $f(y) \equiv y$ for every $y \in Y_1$.

We define $\mathsf{C} = \mathsf{H}/_{\langle\approx\rangle}$. One can prove that C is indeed a simplicial complex and also that $\|\mathsf{C}\|$ is homeomorphic to X (since the identification $\mathsf{C} = \mathsf{H}/_{\langle\approx\rangle}$ was chosen to follow the construction of X).

The faces of C are equivalence classes of $\langle\approx\rangle$. For notational convenience, we write only $[\sigma]$ instead of $[\sigma]_{\langle\approx\rangle}$ for such a class. By ρ we denote the face $[\{\mathbf{e}_1, \ldots, \mathbf{e}_d\}]$ of C.

3.2 The complex $\mathsf{C}(\rho; \zeta_1, \ldots, \zeta_t)$

Now we assume that $t \ge 1$ and we construct the complex $\mathsf{C}(\rho; \zeta_1, \ldots, \zeta_t)$, which is a refinement of $\mathsf{C}(\rho)$.

A suitable triangulation of a simplex. An example of the following construction is depicted in Figure 5. Let Δ be a d-dimensional (geometric) simplex with a set of vertices $V = \{\mathbf{v}_1, \ldots, \mathbf{v}_{d+1}\}$, let $\mathbf{b}$ be its barycentre, and let t be an integer. Next, we define

$$W = \left\{ \mathbf{w}_{i,j} \;\middle|\; \mathbf{w}_{i,j} = \frac{j}{3t}\mathbf{v}_i + \frac{3t - j}{3t}\mathbf{b}, i \in [d+1], j \in [3t] \right\}.$$

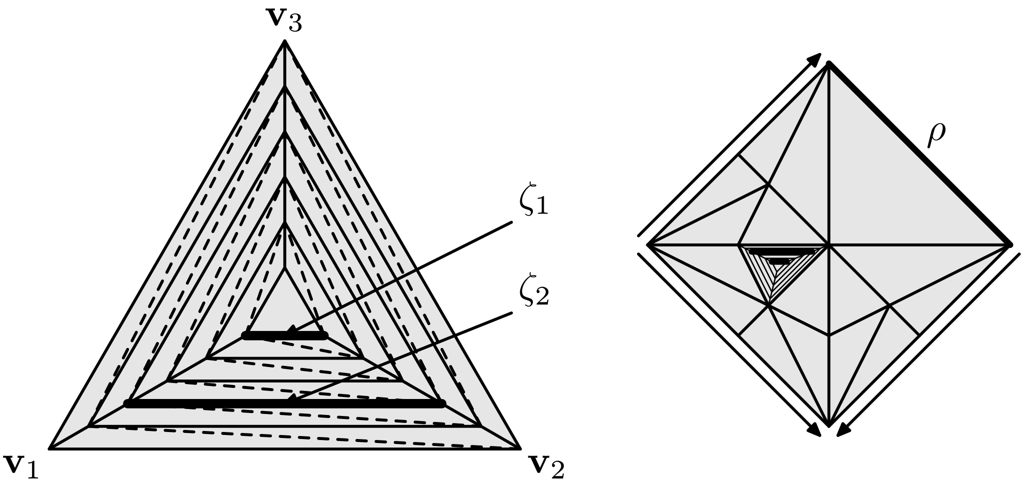

Figure 5: The complex $\mathsf{D}(\zeta_1, \zeta_2)$ (left) and $\mathsf{C}(\rho; \zeta_1, \zeta_2)$ (right), here $d = 2$.

Note that $V \subset W$. For $j \in [t]$, ζ_j is a $(d-1)$-face $\{\mathbf{w}_{1,3j-2}, \mathbf{w}_{2,3j-2}, \ldots, \mathbf{w}_{d,3j-2}\}$.

Now we define polyhedra $Q_1, \ldots, Q_{3t}$. The polyhedron Q_1 is the convex hull conv $\{\mathbf{w}_{1,1} \ldots, \mathbf{w}_{d+1,1}\}$. For $j \in [3t] \setminus \{1\}$ the polyhedron Q_j is the union of the convex hulls

$$\bigcup_{i \in [d+1]} \operatorname{conv}\{\mathbf{w}_{k,l} \mid k \in [d+1] \setminus \{i\}, l \in \{j-1, j\}\}.$$

The polyhedron Q_1 is a simplex. For $j > 1$, the polyhedra Q_j are isomorphic to the prisms $\partial\Delta^d \times [0,1]$, where Δ^d is a d-simplex. Each such prism admits a (standard) triangulation such that $\partial\Delta^d \times \{0\}$ and $\partial\Delta^d \times \{1\}$ are not subdivided (see [Mat03, Exercise 3, p. 12]).

By $\mathsf{D}(\zeta_1, \ldots, \zeta_t)$ we understand an abstract simplicial complex on a vertex set W, which comes from a triangulation of Δ obtained by first subdividing it into the polyhedra $Q_1, \ldots, Q_{3t}$ and subsequently triangulating these polyhedra as described above.

The definition of $\mathsf{C}(\rho; \zeta_1, \ldots, \zeta_t)$. Let ξ be a d-face of H such that $\|\xi\| \subset \operatorname{int}\|\mathsf{H}\|$. Suppose that the set V (from above) is the set of vertices of $\|\xi\|$. We define

$$\mathsf{C}(\rho; \zeta_1, \ldots, \zeta_t) = (\mathsf{C}(\rho) \setminus \{[\xi]\}) \cup \mathsf{D}(\zeta_1, \ldots, \zeta_t).$$

See Figure 5.

Proof of Proposition 3.1. The claims (i), (iii) and (iv) follow straightforwardly from the construction. Regarding the claim (ii), informally, we first d-collapse the face ρ; after that we d-collapse the "interior" of C in order to collapse all d-dimensional faces except the faces that should remain in $\mathsf{C}' \setminus \{\rho\}$. Formally, we use Lemma 5.4. □

Gluing. As the name of connecting gadget suggests, we want to use it (in Section 4) for connecting several other complexes (gadgets). In particular, we want to have some notation for gluing this gadget. We introduce this notation here.

Suppose that $\sigma, \gamma_1, \dots, \gamma_t$ are already known $(d-1)$-dimensional faces of a given complex L. These faces are assumed to be distinct, but not necessarily disjoint. We start with the complex $\mathsf{K} = 2^\sigma \cup 2^{\gamma_1} \cup \cdots \cup 2^{\gamma_t}$. We take a new copy of $\mathsf{C}(\rho; \zeta_1, \dots, \zeta_t)$ and we perform identifications $\rho = \sigma$, $\zeta_1 = \gamma_1, \dots, \zeta_t = \gamma_t$. After these identifications, the complex $\mathsf{K} \cup \mathsf{C}(\rho; \zeta_1, \dots, \zeta_t)$ is denoted by $\mathsf{C}_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$. Note that C (before gluing) and $\mathsf{C}_{\text{glued}}$ are generally not isomorphic, since the gluing procedure can identify some faces of C. The complex $\mathsf{C}'_{\text{glued}} \subset \mathsf{C}_{\text{glued}}$ is defined analogically as $\mathsf{C}_{\text{glued}}$, using C' instead of C.[3]

We also have an analogy of Proposition 3.1(ii)–(iv).

Lemma 3.2. *Let* L, $\mathsf{C}_{\text{glued}} = \mathsf{C}_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$ *and* $\mathsf{C}'_{\text{glued}} = \mathsf{C}'_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$ *be the complexes from the previous paragraph. Then we have:*

(i) *If* σ *is a maximal face of* L, *then* $\mathsf{L} \cup \mathsf{C}_{\text{glued}} \twoheadrightarrow (\mathsf{L} \cup \mathsf{C}'_{\text{glued}}) \setminus \{\sigma\}$.

(ii) *The only* d*-collapsible face of* $\mathsf{C}_{\text{glued}}$ *is the face* σ.

(iii) *Suppose that* d *is a constant. Then the number of faces of* $\mathsf{C}_{\text{glued}}$ *is* $O(t)$.

Proof. The first claim follows from Lemma 5.6. The second claim follows from Proposition 3.1(i) and (iii). The last claim follows from Proposition 3.1(iv). □

3.3 The construction of a bad complex B

In this subsection, for $d \geq 3$, we construct a *bad* complex B, which is d-collapsible but it contains a bad face.

Let $S = \{p, q_1, \dots, q_{d-1}, r_1, \dots, r_d\}$ be a $2d$-element set. Consider the full simplex 2^S. We name its $(d-1)$-faces:

$\iota = \{p, q_1, \dots, q_{d-1}\}$ is an *initial* face,
$\lambda_i = \{q_1, \dots, q_{d-1}, r_i\}$ are *liberation* faces for $i \in [d]$,
$\sigma_B = \{r_1, \dots, r_d\}$, we will show that σ_B is a bad face.

The remaining $(d-1)$-faces are *attaching* faces; let us denote these faces by $\alpha_1, \dots, \alpha_t$.

We define B by

$$\mathsf{B} = 2^S \cup \mathsf{C}_{\text{glued}}(\iota; \alpha_1, \dots, \alpha_t).$$

[3]We recall that the complex C' was defined in the statement of Proposition 3.1.

Proof of Theorem 1.2(ii). We want to prove that B is d-collapsible, but it contains a bad d-collapsible face.

First, we observe that σ_B is a bad face. By Lemma 3.2(ii) and the inspection, the only d-collapsible faces of B are λ_i and σ_B for $i \in [d]$. After collapsing σ_B there is no d-collapsible face, implying that σ_B is a bad face.

In order to show d-collapsibility of B we need a few other definitions. The complex R is defined by

$$\mathsf{R} = \left\{\sigma \in 2^S \mid \text{if } \{q_1, \dots, q_{d-1}\} \subseteq \sigma \text{ then } \sigma \subseteq \iota \right\}.$$

We observe that $\mathsf{R} \setminus \{\iota\}$ is d-collapsible and also that $2^S \twoheadrightarrow \mathsf{R}$ by collapsing all liberation faces (in any order). In fact, the first observation is a special case of Lemma 4.1(ii) used for the NP-reduction.

Auxiliary complexes A, A' are defined in a similar way to B:

$$\begin{aligned} \mathsf{A} &= \mathsf{R} \cup \mathsf{C}_{\text{glued}}(\iota; \alpha_1, \dots, \alpha_t); \\ \mathsf{A}' &= (\mathsf{R} \cup \mathsf{C}'_{\text{glued}}(\iota; \alpha_1, \dots, \alpha_t)) \setminus \{\iota\}. \end{aligned}$$

We show d-collapsibility of B by the following sequence of d-collapses:

$$\mathsf{B} \twoheadrightarrow \mathsf{A} \twoheadrightarrow \mathsf{A}' = (\mathsf{R} \setminus \{\iota\}) \twoheadrightarrow \emptyset.$$

The fact that $\mathsf{B} \twoheadrightarrow \mathsf{A}$ is quite obvious—it is sufficient to d-collapse the liberation faces. More precisely, we use Lemma 5.3 with $\mathsf{K} = \mathsf{B}$, $\mathsf{K}' = 2^S$ and $\mathsf{L}' = \mathsf{R}$. The fact that $\mathsf{A} \twoheadrightarrow \mathsf{A}'$ follows from Lemma 3.2(i). We already observed that $\mathsf{R} \setminus \{\iota\} \twoheadrightarrow \emptyset$ when defining R. □

4 NP-completeness

Here we prove Theorem 1.1(ii). Throughout this section we assume that $d \geq 4$ is a fixed integer. We have that d-COLLAPSIBILITY is in NP, since if we are given a sequence of faces of dimension at most $d-1$, we can check in a polynomial time whether this sequence determine a d-collapsing of a given complex.

For NP-hardness, we reduce the problem 3-SAT to d-COLLAPSIBILITY. The problem 3-SAT is NP-complete according to Cook [Coo71]. Given a 3-CNF formula Φ, we construct a complex F that is d-collapsible if and only if Φ is satisfiable.

4.1 Simplicial gadgets

For purposes of the reduction we need several gadgets with similar properties. We call them *simplicial gadgets* since they consist of full simplices (on varying

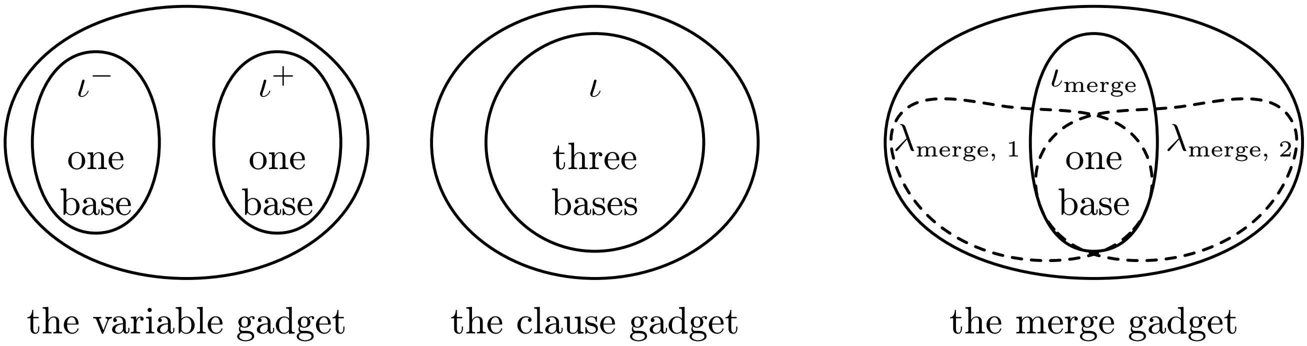

Figure 6: A schematic pictures of simplicial gadgets; the liberation faces of the merge gadget are distinguished.

number of vertices) with several distinguished $(d-1)$-faces. These gadgets generalize the complex 2^S defined in Subsection 3.3. Every simplicial gadget contains one or more $(d-1)$-dimensional pairwise disjoint *initial* faces. Every initial face ι contains several (possibly only one) distinguished $(d-2)$-faces called *bases* of ι. The *liberation* faces of the gadget are such $(d-1)$-faces λ that contain a base of some initial face ι, but $\lambda \neq \iota$. The remaining $(d-1)$-faces are *attaching* faces.

Now we define several concrete examples of simplicial gadgets.

The variable gadget. The *variable gadget* $\mathsf{V} = \mathsf{V}(\iota^+, \iota^-, \beta^+, \beta^-)$ is described by the following table:

vertices:	$p^+, q_1^+, \ldots, q_{d-1}^+, p^-, q_1^-, \ldots, q_{d-1}^-$;
initial faces:	$\iota^+ = \{p^+, q_1^+, \ldots, q_{d-1}^+\}$, $\iota^- = \{p^-, q_1^-, \ldots, q_{d-1}^-\}$;
bases:	$\beta^+ = \{q_1^+, \ldots, q_{d-1}^+\}$, $\beta^- = \{q_1^-, \ldots, q_{d-1}^-\}$.

The clause gadget. The *clause gadget* $\mathsf{G}(\iota, \lambda_1, \lambda_2, \lambda_3)$ is given by:

vertices:	$p_1, \ldots, p_d, q$;
initial face:	$\iota = \{p_1, p_2, \ldots, p_d\}$;
bases:	$\beta_1 = \iota \setminus \{p_1\}$, $\beta_2 = \iota \setminus \{p_2\}$. $\beta_3 = \iota \setminus \{p_3\}$.

Every base β_j is contained in exactly one liberation face $\lambda_j = \beta_j \cup \{q\}$.

The merge gadget. The *merge gadget* $\mathsf{M}(\iota_{\text{merge}}, \lambda_{\text{merge},1}, \lambda_{\text{merge},2})$ is given by:

vertices:	$p_1, \ldots, p_d, q, r$;
initial face:	$\iota_{\text{merge}} = \{p_1, p_2, \ldots, p_d\}$;
base:	$\iota_{\text{merge}} \setminus \{p_1\}$.

The merge gadget contains exactly two liberation faces, which we denote $\lambda_{\text{merge},1}$ and $\lambda_{\text{merge},2}$.

We close this subsection by proving a lemma about d-collapsings of simplicial gadgets.

Lemma 4.1. *Suppose that* S *is a simplicial gadget,* ι *is its initial face,* $\beta \subseteq \iota$ *is a base face, and* $\lambda_1, \dots, \lambda_t$ *are liberation faces containing* β*. Then d-collapsing of* $\lambda_1, \dots, \lambda_t$ *(even in any order) yields a complex* R *such that*

(i) ι *is a maximal face of* R*;*

(ii) $\mathsf{R} \setminus \{\iota\}$ *is d-collapsible;*

(iii) $\mathsf{R} \setminus \{\iota\} \twoheadrightarrow 2^{\iota'}$ *where* ι' *is an initial face different from* ι *(if exists).*

Proof. We prove each of the claims separately.

(i) Let V be a set of vertices of S and let $\lambda_{t+1} = \iota$. We (inductively) observe that d-collapsing of faces $\lambda_1, \dots, \lambda_k$ for $k \leq t$ yields a complex in which λ_{k+1} is contained in a unique maximal face $(V \setminus (\lambda_1 \cup \dots \cup \lambda_k)) \cup \beta$. This implies that R is well defined and also finishes the first claim, since

$$(V \setminus (\lambda_1 \cup \dots \cup \lambda_t)) \cup \beta = \iota.$$

We remark that the few details skipped here are exactly the same as in the proof of Lemma 5.1.

(ii) We observe that β is a maximal $(d-2)$-face of $\mathsf{R} \setminus \{\iota\}$ and $\mathsf{S}_\beta = \mathsf{R} \setminus \{\iota, \beta\}$, hence $\mathsf{R} \setminus \{\iota\} \to \mathsf{S}_\beta$. (We recall that K_σ denotes the resulting complex of an elementary d-collapse $\mathsf{K} \to \mathsf{K}_\sigma = \mathsf{K} \setminus [\sigma, \tau(\sigma)]$.) Next, $\mathsf{S}_\beta \twoheadrightarrow \mathsf{S}_\emptyset = \emptyset$ by Lemma 5.1.

(iii) Similarly as before we have $\mathsf{R} \setminus \{\iota\} \to \mathsf{S}_\beta$. Let v be a vertex of β, we have $\mathsf{S}_\beta \to \mathsf{S}_{\{v\}}$ by Lemma 5.1. The complex $\mathsf{S}_{\{v\}}$ is a full simplex (S with removed v), this complex even 1-collapse to $2^{\iota'}$ by collapsing vertices of $V \setminus (\iota' \cup \{v\})$ (in any order).

□

4.2 The reduction

Let the given 3-CNF formula be $\Phi = C^1 \wedge C^2 \wedge \dots \wedge C^n$, where each C^i is a clause with exactly three literals (we assume without loss of generality that every clause contains three different variables). Suppose that $x_1, \dots, x_m$ are variables appearing in the formula. For every such variable x_j we take a fresh copy of the variable gadget and we denote it by $\mathsf{V}_j = \mathsf{V}_j(\iota_j^+, \iota_j^-, \beta_j^+, \beta_j^-)$.

For every clause C^i containing variables x_{j_1}, x_{j_2} and x_{j_3} (in a positive or negative occurrence) we take a new copy of the clause gadget and we denote it by $\mathsf{G}^i = \mathsf{G}^i(\iota^i, \lambda^i_{j_1}, \lambda^i_{j_2}, \lambda^i_{j_3})$. Moreover, for C^i with $i \geq 2$, we also take a new copy of the merge gadget and we denote it $\mathsf{M}^i = \mathsf{M}^i(\iota^i_{\text{merge}}, \lambda^i_{\text{merge},1}, \lambda^i_{\text{merge},2})$.

Now we connect these simplicial gadgets by glued copies of the connecting gadget called *connections*.

Suppose that a variable x_j occurs positively in the clauses $C^{i_1}, \dots, C^{i_k}$. We construct the *positive occurrence connections* by setting

$$\mathsf{O}^+_j = \mathsf{C}_{\text{glued}}(\iota^+_j; \lambda^{i_1}_j, \dots, \lambda^{i_k}_j).$$

The *negative occurrence connections* O^-_j are constructed similarly (we use ι^-_j instead of ι^+_j; and we use clauses in which is x_j in negative occurrence).

The *merge connections* are defined by

$$\begin{array}{lcll} \mathsf{I}^1_1 & = & \mathsf{C}_{\text{glued}}(\iota^1; \lambda^2_{\text{merge},2}); & \\ \mathsf{I}^i_1 & = & \mathsf{C}_{\text{glued}}(\iota^i; \lambda^i_{\text{merge},1}) & \text{where } i \in \{2, \dots, n\}; \\ \mathsf{I}^i_2 & = & \mathsf{C}_{\text{glued}}(\iota^i_{\text{merge}}; \lambda^{i+1}_{\text{merge},2}) & \text{where } i \in \{2, \dots, n-1\}. \end{array}$$

For convenient notation we denote I^1_1 also by I^1_2.

Finally, the *tidy connection* is defined by

$$\mathsf{T} = \mathsf{C}_{\text{glued}}(\iota^n_{\text{merge}}; \alpha_1, \dots, \alpha_t)$$

where $\alpha_1, \dots, \alpha_t$ are attaching faces of all simplicial gadgets in the reduction, namely the variable gadgets V_j for $j \in [m]$, the clause gadgets G^i for $i \in [n]$, and the merge gadgets M^i for $i \in \{2, \dots, n\}$.

The complex F in the reduction is defined by

$$\mathsf{F} = \bigcup_{j=1}^{m} \mathsf{V}_j \cup \bigcup_{i=1}^{n} \mathsf{G}^i \cup \bigcup_{i=2}^{n} \mathsf{M}^i \cup \bigcup_{j=1}^{m} (\mathsf{O}^+_j \cup \mathsf{O}^-_j) \cup \bigcup_{i=1}^{n} \mathsf{I}^i_1 \cup \bigcup_{i=2}^{n-1} \mathsf{I}^i_2 \cup \mathsf{T}.$$

See Figure 7 for an example.

We observe that the number of faces of F is polynomial in the number of clauses in the formula (regarding d as a constant). Indeed, we see that the number of gadgets (simplicial gadgets and connections) is even linear in the number of variables. Each simplicial gadget has a constant size. Each connection has at most linear size due to Lemma 3.2(iii).

Collapsibility for satisfiable formulae. We assign each variable TRUE or FALSE so that the formula is satisfied. For every variable gadget V_j we proceed as follows. First, suppose that x_j is assigned TRUE. We d-collapse[4]

[4]Note that after d-collapsing a liberation face containing β^+_j the liberation faces containing β^-_j are no more d-collapsible (and vice versa). This will be a key property for showing that unsatisfiable formulae yield to non-collapsible complexes.

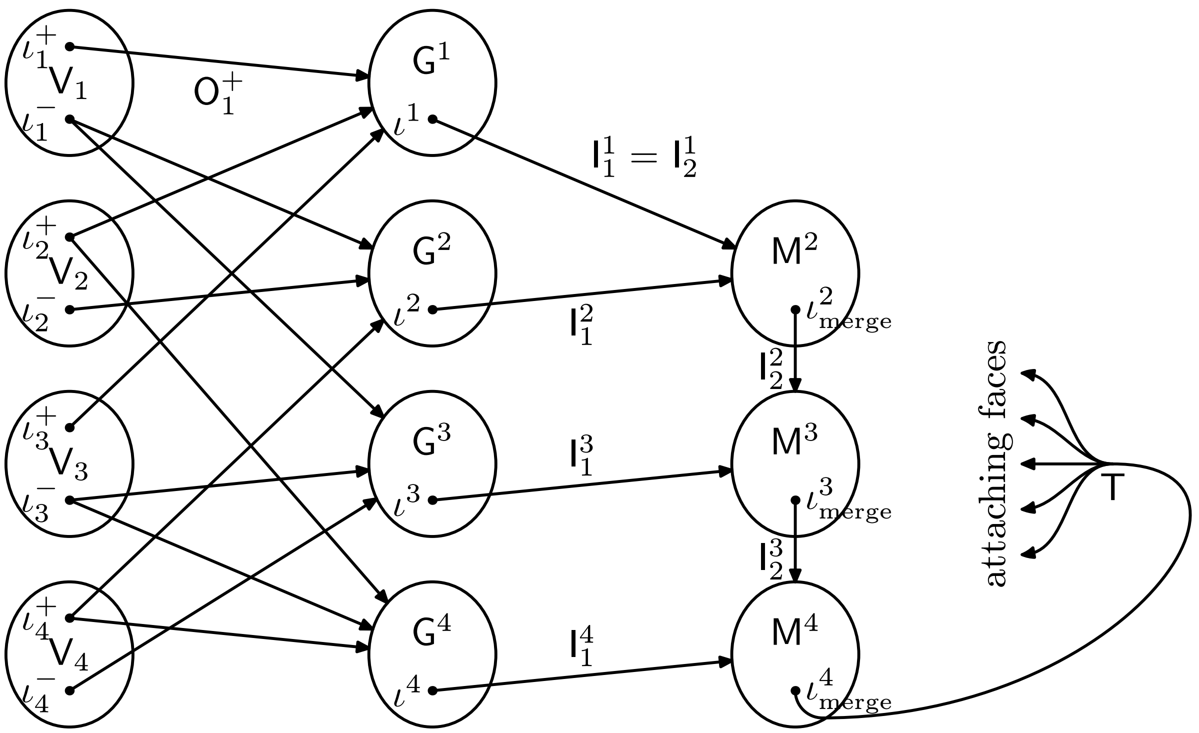

Figure 7: A schematic example of F for the formula $\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_4) \wedge (\neg x_1 \vee \neg x_3 \vee \neg x_4) \wedge (x_2 \vee \neg x_3 \vee x_4)$. Initial faces are drawn as points. (Multi)arrows denote connections. Each (multi)arrow points from the unique d-collapsible face of the connection to simplicial gadgets that are attached to the connection by some of its liberation faces.

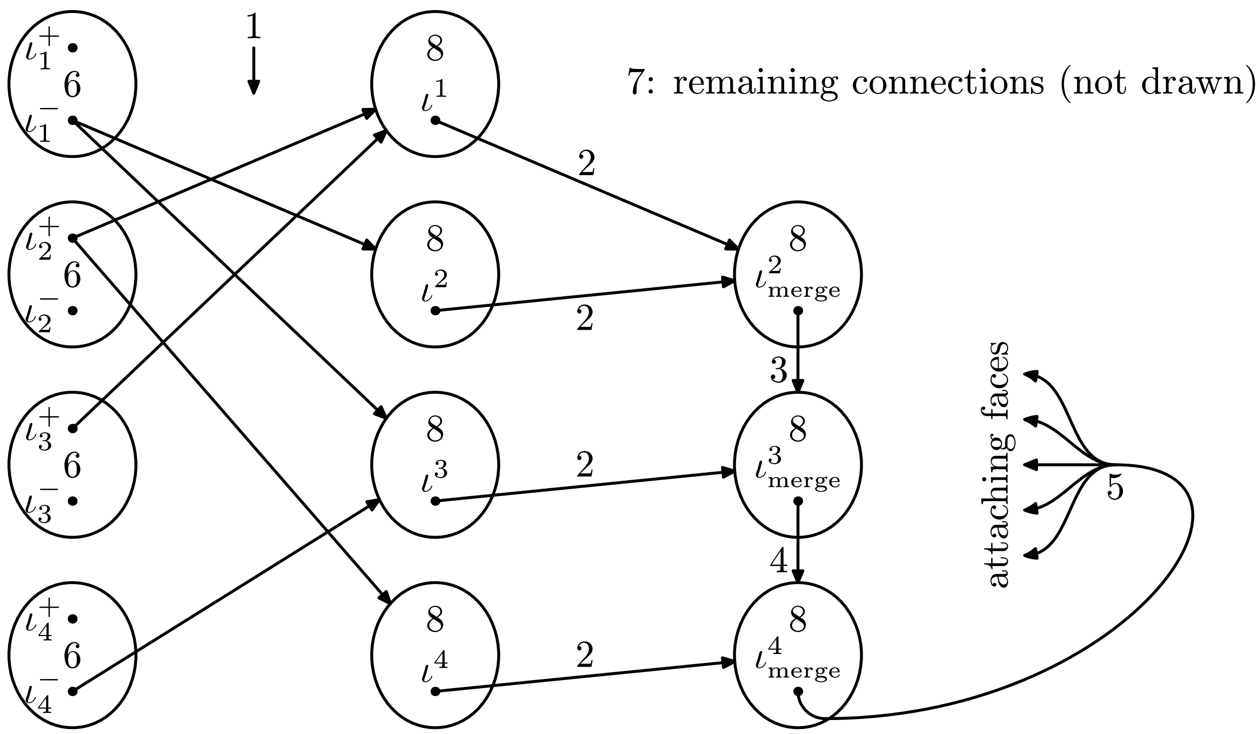

Figure 8: d-collapsing of F for the Φ from Figure 7 assigned (FALSE, TRUE, TRUE, FALSE). The numbers denote the order in which the parts of F vanish.

the liberation faces containing β_j^+ (see Lemma 4.1(i)), after that ι_j^+ is d-collapsible, and we d-collapse O_j^+ (following Lemma 3.2(i) in the same way as in the proof of Theorem 1.2(ii)). Similarly, we d-collapse O_j^- if x_j is assigned FALSE.

We use several times Lemma 4.1(i) and Lemma 3.2(i) in the following paragraphs. The use is very similar is in the previous one, thus we do not mention these lemmas again.

After d-collapsings described above, we have that every clause gadget G^i contains at least one liberation face that is d-collapsible, since we have chosen such an assignment that the formula is satisfied. We d-collapse this liberation face and after that the face ι^i is d-collapsible. We continue with d-collapsing the merge gadgets I_1^i for $i \in [n]$.

The next step is that we gradually d-collapse the merge gadgets I_2^i for $i \in \{2, \dots, n-1\}$. For this, we have that both liberation faces of I_2^2 are d-collapsible, we d-collapse them and we have that ι^2_{merge} is d-collapsible. We d-collapse I_2^2 and now we continue with the same procedure with I_2^3, the I_2^4, etc.

Finally, we d-collapse the tidy gadget. The d-collapsing of tidy gadget makes all the attaching faces of simplicial gadgets d-collapsible. After this “tidying up” we can d-collapse all variable gadgets (using Lemma 4.1(iii)), then all remaining connections, and at the end all remaining simplicial gadgets due to Lemma 4.1(ii).

Non-collapsibility for unsatisfiable formulae. We suppose that Φ is unsatisfiable and we prove that F is not d-collapsible.

For contradiction, we suppose that F is d-collapsible. Let

$$\mathsf{F} = \mathsf{F}_1 \to \mathsf{F}_2 \to \cdots \to \emptyset$$

be a d-collapsing of F. We call it *our* d-collapsing. For a technical reason, according to Lemma 5.2, we can assume that first $(d-1)$-dimensional faces are collapsed and after that faces of less dimensions are removed.

Let us fix a subcomplex F_ℓ in our d-collapsing. Let N be a connection (one of that forming F) and let $\mathsf{N}_\ell = \mathsf{F}_\ell \cap \mathsf{N}$. We say that N is *activated* in F_ℓ if N_ℓ is a proper subcomplex of N.

The connection N is defined as $\mathsf{C}_{\text{glued}}(\sigma; \gamma_1, \ldots, \gamma_s)$ for some $(d-1)$-faces $\sigma, \gamma_1, \ldots, \gamma_s$ of simplicial gadgets in F. We remark that Lemma 3.2(ii) implies that if N is activated in F_ℓ then $\sigma \notin \mathsf{F}_\ell$.

We also prove the following lemma about activated connections.

Lemma 4.2. *Let F_ℓ be a complex from our d-collapsing such that T is not activated in F_ℓ. Then we have:*

(i) *Let $j \in [m]$. If the positive occurrence connection O_j^+ is activated in F_ℓ, then the negative occurrence connection O_j^- is not activated in F_ℓ (and vice versa).*

(ii) *Let $i \in [n]$. If the merge connection I_1^i is activated in F_ℓ, then at least one of the three occurrence connections attached to G^i is activated in F_ℓ.*

(iii) *Let $i \in \{2, \ldots, n-1\}$. If the merge connection I_2^i is activated in F_ℓ, then the merge connections I_1^i and I_2^{i-1} are activated in F_ℓ.*

Proof. Let us consider first $\ell - 1$ d-collapses of our d-collapsing

$$\mathsf{F} = \mathsf{F}_1 \to \mathsf{F}_2 \to \cdots \to \mathsf{F}_\ell,$$

where $\mathsf{F}_{k+1} = \mathsf{F}_k \setminus [\sigma_k, \tau_k]$ for $k \in [\ell - 1]$. According to assumption on our d-collapsing, we have that $\sigma_1, \ldots, \sigma_{\ell-1}$ are $(d-1)$-dimensional (since T is not activated in F_ℓ yet).

Now we prove each of the claims separately.

(i) For a contradiction we suppose that both O_j^+ and O_j^- are activated in F_ℓ.

We consider the variable gadget V_j. We say that an index $k \in [\ell - 1]$ is *relevant* if $\sigma_k \in \mathsf{V}_j$. We observe that if k is a relevant index then σ_k is a liberation face or an initial face of V_j, because attaching faces are contained in T.

By *positive* face we mean either the initial face ι_j^+ or a liberation face containing β_j^+. A *negative* face is defined similarly. Let k^+ (respectively k^-) be the smallest relevant index such that σ_{k^+} is a positive face (respectively negative face). These indexes have to exist, since both O_j^+ and O_j^- are activated in F_ℓ. Without loss of generality $k^+ < k^-$.

We show that σ_{k^-} is not a d-collapsible face of $\mathsf{F}_{k^- -1}$, thus we get a contradiction. Indeed, let $S = \sigma_{k^+} \setminus \sigma_{k^-}$. We have $|S| \geq 2$, since $d \geq 4$ (here we crucially use this assumption). Let $s \in S$. Then we have $\sigma_{k^-} \cup \{s\} \in \mathsf{F}_{k^- -1}$, because $\sigma_{k^-} \cup \{s\}$ does not contain a positive subface (it does not contain β_j^+, since $|\sigma_{k^-} \cap \beta_j^+| \leq 1$, but $|\beta_j^+| \geq 3$). On the other hand $\sigma_{k^-} \cup S \notin \mathsf{K}_{k^- -1}$, since it contains σ_{k^+}. I.e., σ_{k^-} is not in a unique maximal face.

(ii) We again define relevant index; this time $k \in [\ell - 1]$ is *relevant* if $\sigma_k \in \mathsf{G}^i$. We consider the smallest relevant index k'. Again we have that $\sigma_{k'}$ is either the initial face ι^i or a liberation face of G^i. In fact, $\sigma_{k'}$ cannot be ι^i: by Lemma 3.2(ii) we would have that $\mathsf{I}_1^i \subseteq \mathsf{F}_{k'-1}$ and also $\mathsf{G}^i \subseteq \mathsf{F}_{k'-1}$ form minimality of k', which would contradict that $\sigma_{k'}$ is a collapsible face of $\mathsf{F}_{k'-1}$. Thus $\sigma_{k'}$ is a liberation face of G^i. This implies, again by Lemma 3.2(ii), that at least one of the occurrence gadgets attached to liberation faces is activated even in $\mathsf{F}_{k'-1}$.

(iii) By a similar discussion as in previous step we have that at least one of the liberation faces $\lambda^i_{\text{merge},1}$ and $\lambda^i_{\text{merge},2}$ of M^i have to be d-collapsed before d-collapsing ι^i_{merge}. However, we observe that d-collapsing only one of these faces is still insufficient for possibility of d-collapsing ι^i_{merge}. Hence both of the liberation faces have to be d-collapsed, which implies that both the gadgets I_1^i and I_2^{i-1} are activated in F_ℓ.

□

We prove of also an analogy of Lemma 4.2 for the tidy gadget. We have to modify the assumptions, that is why we use a separate lemma. The proof is essentially same as the proof of Lemma 4.2(iii), therefore we omit it.

Lemma 4.3. *Let ℓ be the largest index such that T is not activated in F_ℓ, then the merge connections I_1^n and I_2^{n-1} are activated in F_ℓ.*

□

Now we can quickly finish the proof of non-collapsibility. Let ℓ be the integer from Lemma 4.3. By this lemma and repeatedly using Lemma 4.2(iii) we have that all merge connections are activated in F_ℓ. By Lemma 4.2(ii), for every clause gadget G^i there is an occurrence gadget attached to G^i, which is activated in F_ℓ. Finally, Lemma 4.2(i) implies that for every variable x_j at most one of the occurrence gadgets O_j^+, O_j^- is activated in F_ℓ. Let us assign x_j TRUE if it is O_j^+ and FALSE otherwise. This is satisfying assignment, since for every G^i at least one occurrence gadget attached to it is activated in F_ℓ. This contradicts the fact that Φ is unsatisfiable.

□

5 Technical properties of d-collapsing

In this section, we prove several auxiliary lemmas on d-collapsibility used throughout the paper.

5.1 d-collapsing faces of dimension strictly less than $d-1$

Lemma 5.1. *Let K be a complex, d an integer, and σ a d-collapsible face (in particular, $\dim\sigma \leq d-1$). Let $\sigma' \supseteq \sigma$ be a face of K of dimension at most $d-1$. Then σ' is d-collapsible and $\mathsf{K}_{\sigma'} \twoheadrightarrow \mathsf{K}_\sigma$.*

Proof. We assume that $\sigma \neq \sigma'$ otherwise the proof is trivial.

First, we observe that $\tau(\sigma)$ is a unique maximal face containing σ'. Indeed, $\sigma' \subseteq \tau(\sigma)$, since $\tau(\sigma)$ is the unique maximal face containing σ, and also if $\eta \supseteq \sigma'$, then $\eta \supseteq \sigma$, which implies $\eta \subseteq \tau(\sigma)$. Hence we have that σ' is d-collapsible.

Let v_1 be a vertex of $\sigma' \setminus \sigma$. It is sufficient to prove that $\mathsf{K}_{\sigma'} \twoheadrightarrow \mathsf{K}_{\sigma' \setminus \{v_1\}}$ and proceed by induction. Thus, for simplicity of notation, we can assume that $\sigma' = \sigma \cup \{v_1\}$.

Let $v_2, \dots, v_t$ be vertices of $\tau(\sigma) \setminus \sigma'$. By η_i we denote the face $\sigma \cup \{v_i\}$ for $i \in [t]$. (In particular, $\sigma' = \eta_1$.) For $i \in [t+1]$ we define a complex K_i by the formula

$$\mathsf{K}_i = \{\eta \in \mathsf{K} \mid \eta \not\supseteq \eta_1, \dots, \eta \not\supseteq \eta_{i-1}\} = \{\eta \in \mathsf{K} \mid \text{if } \eta \supseteq \sigma \text{ then } v_j \notin \eta \text{ for } j < i\}.$$

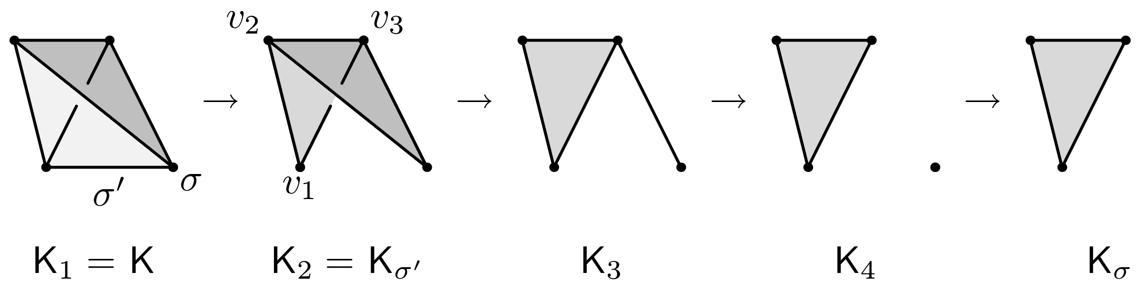

Figure 9: An example of 2-collapsing $\mathsf{K} \to \mathsf{K}_{\sigma'} \twoheadrightarrow \mathsf{K}_\sigma$.

From these descriptions we have that η_i is a d-collapsible face of K_i contained in a unique maximal face $\tau_i = \tau(\sigma) \setminus \{v_1, \dots, v_{i-1}\}$. Moreover $(\mathsf{K}_i)_{\eta_i} = \mathsf{K}_{i+1}$. Thus, we have a d-collapsing

$$\mathsf{K} = \mathsf{K}_1 \to \mathsf{K}_2 \to \cdots \to \mathsf{K}_t.$$

See Figure 9 for an example.

To finish the proof it remains to observe that $\mathsf{K}_2 = \mathsf{K}_{\sigma'}$ and $\mathsf{K}_{t+1} = \mathsf{K}_\sigma \cup \{\sigma\}$, hence $\mathsf{K}_{t+1} \to \mathsf{K}_\sigma$. □

As a corollary, we obtain the following lemma.

Lemma 5.2. *Suppose that* K *is a d-collapsible complex. Then there is a d-collapsing of* K *such that first only $(d-1)$-dimensional faces are collapsed and after that faces of dimensions less then $(d-1)$ are removed.*

Proof. Suppose that we are given a d-collapsing of K. Suppose that in some step we d-collapse a face σ that is not maximal and its dimension is less than $d-1$. Let us denote this step by $\mathsf{K}' \to \mathsf{K}'_\sigma$. Let $\sigma' \supseteq \sigma$ be such face of K' that either $\dim \sigma' = d-1$ or σ' is a maximal face. Then we replace this step by d-collapsing $\mathsf{K}' \to \mathsf{K}'_{\sigma'} \twoheadrightarrow \mathsf{K}_\sigma$.

We repeat this procedure until every d-collapsed face is either of dimension $d-1$ or maximal. We observe that this procedure can be repeated only finitely many times, since in every replacement we increase the number of elementary d-collapses in the d-collapsing, whence this number is bounded by the number of faces of K.

Finally, we observe that if we first remove a maximal face of dimension less than $d-1$ and then we d-collapse a $(d-1)$-dimensional face, we can swap these steps with the same result. □

5.2 d-collapsing to a subcomplex

Suppose that K is a simplicial complex, K' is a subcomplex of it, which d-collapses to a subcomplex L'. If certain conditions are satisfied, then we can

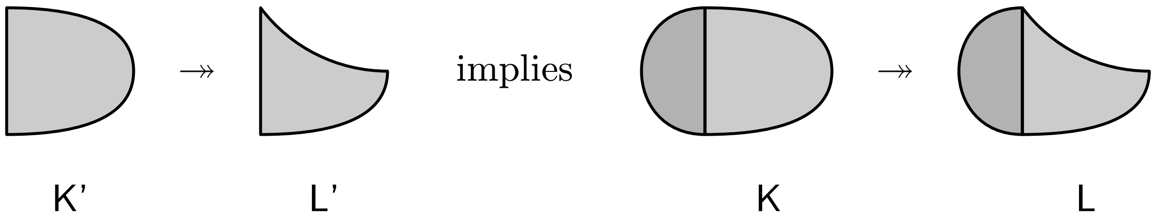

Figure 10: Complexes K, K', L and L' from the statement of Lemma 5.3.

perform d-collapsing $\mathsf{K}' \twoheadrightarrow \mathsf{L}'$ in whole K; see Figure 10 for an illustration. The precise statement is given in the following lemma.

Lemma 5.3 (d-collapsing a subcomplex)**.** *Let K be a simplicial complex, K' a subcomplex of K, and L' a subcomplex of K'. Assume that if $\sigma \in \mathsf{K}' \setminus \mathsf{L}'$, $\eta \in \mathsf{K}$, and $\eta \supseteq \sigma$, then $\eta \in \mathsf{K}' \setminus \mathsf{L}'$. Moreover assume that $\mathsf{K}' \twoheadrightarrow \mathsf{L}'$. Then $\mathsf{L} = (\mathsf{K} \setminus \mathsf{K}') \cup \mathsf{L}'$ is a simplicial complex and $\mathsf{K} \twoheadrightarrow \mathsf{L}$.*

Proof. It is straightforward to check that L is a simplicial complex using the equivalence

$$\eta \in \mathsf{L} \text{ if and only if } \eta \in \mathsf{K} \text{ and } \eta \notin \mathsf{K}' \setminus \mathsf{L}'.$$

In order to show $\mathsf{K} \twoheadrightarrow \mathsf{L}$, it is sufficient to show the following (and proceed by induction over elementary d-collapses):

Suppose that σ' is a d-collapsible face of K' such that $\mathsf{K}'_{\sigma'} \supseteq \mathsf{L}'$. Then we have

1. *σ' is a d-collapsible face of K.*

2. *If $\sigma \in \mathsf{K}'_{\sigma'} \setminus \mathsf{L}'$, $\eta \in \mathsf{K}_{\sigma'}$ and $\eta \supseteq \sigma$, then $\eta \in \mathsf{K}'_{\sigma'} \setminus \mathsf{L}'$.*

3. *$\mathsf{L} = (\mathsf{K}_{\sigma'} \setminus \mathsf{K}'_{\sigma'}) \cup \mathsf{L}'$.*

We prove the claims separately:

1. We know that $\sigma' \notin \mathsf{L}'$, since $\mathsf{K}'_{\sigma'} \supseteq \mathsf{L}'$. Thus, $\sigma' \in \mathsf{K}' \setminus \mathsf{L}'$. If $\eta' \in \mathsf{K}$ and $\eta' \supseteq \sigma'$, then, by the assumption of the lemma, $\eta' \in \mathsf{K}' \setminus \mathsf{L}' \subseteq \mathsf{K}'$. In particular, the maximal faces in K' containing σ' coincide with the maximal faces in K containing σ'. It means that σ' is a d-collapsible face of K.

2. We have $\mathsf{K}'_{\sigma'} \setminus \mathsf{L}' \subseteq \mathsf{K}' \setminus \mathsf{L}'$ and $\mathsf{K}_{\sigma'} \subseteq \mathsf{K}$. Thus the assumption of the lemma implies that $\eta \in \mathsf{K}' \setminus \mathsf{L}'$. Next we have $\mathsf{K}_{\sigma'} \cap \mathsf{K}' = \mathsf{K}'_{\sigma'}$, since the maximal faces in K' containing σ' coincide with the maximal faces in K containing σ'. We conclude that $\eta \in \mathsf{K}'_{\sigma'} \setminus \mathsf{L}'$.

3. One can check that $\mathsf{K} \setminus \mathsf{K}' = \mathsf{K}_{\sigma'} \setminus \mathsf{K}'_{\sigma'}$.

□

Suppose that $\mathcal{F}$ is a set system. For an integer k we define the graph $G_k(\mathcal{F}) = (V(G_k), E(G_k))$ as follows:

$$\begin{aligned} V(G_k) &= \{F \in \mathsf{K} \mid |F| = k+1 \text{ (i.e. } \dim F = k \text{ if } F \text{ is regarded as a face)}\}; \\ E(G_k) &= \{\{F, F'\} \mid F, F' \in V(G_k), F \cap F' \in \mathcal{F} \text{ and } |F \cap F'| = k\}. \end{aligned}$$

Lemma 5.4 (d-collapsing a d-dimensional complex)**.** *Suppose that* K *is a* d*-dimensional complex,* L *is its subcomplex and the following conditions are satisfied:*

- $\mathsf{K} \setminus \mathsf{L}$ *contains a* d*-collapsible face* σ *such that* $\tau(\sigma) \in \mathsf{K} \setminus \mathsf{L}$*,*
- $G_d(\mathsf{K} \setminus \mathsf{L})$ *is connected,*
- *For every* $(d-1)$*-face* $\eta \in \mathsf{K} \setminus \mathsf{L}$ *there are at most two* d*-faces in* $\mathsf{K} \setminus \mathsf{L}$ *containing* η*.*

Then $\mathsf{K} \twoheadrightarrow \mathsf{L}$*.*

Proof. See figure 11 when following the proof. Let $\tau_0 = \tau(\sigma)$, $\tau_1, \ldots, \tau_j$ be an order of vertices of $G_d(\mathsf{K} \setminus \mathsf{L})$ such that for every $i \in [j]$ the vertex τ_i has a neighbor $\tau_{n(i)}$ with $n(i) < i$. Such an order exists by the second condition. Let $\sigma_i = \tau_i \cap \tau_{n(i)}$.

Consider the following sequence of elementary d-collapses

$$\begin{aligned} \mathsf{K} &\to \mathsf{K}_0 = \mathsf{K}_\sigma, \\ \mathsf{K}_{i-1} &\to \mathsf{K}_i = (\mathsf{K}_{i-1})_{\sigma_i} \text{ for } i \in [j]. \end{aligned}$$

This sequence is indeed a sequence of elementary d-collapses, since $\tau_{n(i)} \notin \mathsf{K}_{i-1}$, thus τ_i is a unique maximal face containing σ_i in K_{i-1} by the third condition. Moreover, $\sigma_i \in \mathsf{K} \setminus \mathsf{L}$. Thus, K_j is a supercomplex of L.

The set system $\mathsf{K}_j \setminus \mathsf{L}$ contains only faces of dimensions $d-1$ or less. Hence $\mathsf{K}_j \twoheadrightarrow \mathsf{L}$ by removing faces, which establishes the claim. □

5.3 Gluing distant faces

Let k be an integer. Suppose that K is a simplicial complex and let $\omega = \{u_1, \ldots, u_{k+1}\}$, $\eta = \{v_1, \ldots, v_{k+1}\}$ be two k-faces of K. By

$$\mathsf{K}(\omega = \eta)$$

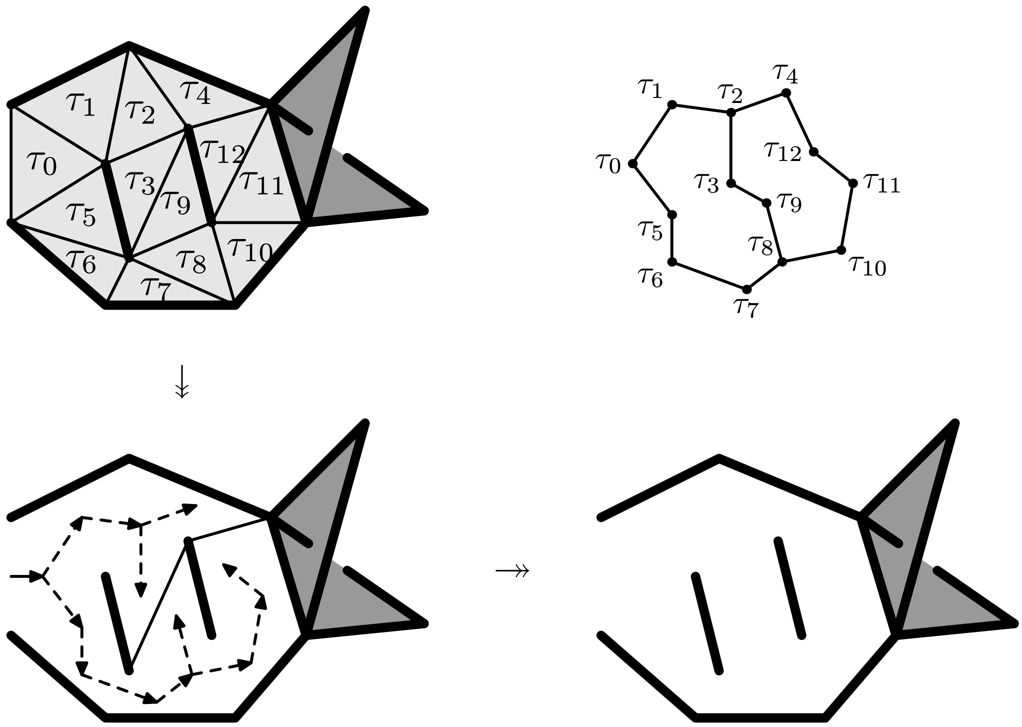

Figure 11: In top right picture there are complexes K and L from Lemma 5.4; L is thick and dark. In top left picture there is the graph $G_2(\mathsf{K}\backslash\mathsf{L})$. Collapsing $\mathsf{K} \twoheadrightarrow \mathsf{L}$ is in bottom pictures.

we mean the resulting complex under the identification $u_1 = v_1, \ldots, u_{k+1} = v_{k+1}$ (note that this complex is not unique—it depends on the order of vertices in ω and η; however, the order of vertices is not important for our purposes).

In a similar spirit, we define

$$\mathsf{K}(\omega_1 = \eta_1, \ldots, \omega_t = \eta_t)$$

for k-faces $\omega_1, \ldots, \omega_t, \eta_1, \ldots, \eta_t$.

Lemma 5.5 (Collapsing glued complex)**.** *Suppose that ω and η are two distant faces in a simplicial complex K. Let L be a subcomplex of K such that $\omega, \eta \in \mathsf{L}$. Suppose that K d-collapses to L. Then $\mathsf{K}(\omega = \eta)$ d-collapses to $\mathsf{L}(\omega = \eta)$.*

Proof. Let $\mathsf{K} \to \mathsf{K}_2 \to \mathsf{K}_3 \to \cdots \to \mathsf{L}$ be a d-collapsing of K to L. Our task is to show that

$$\mathsf{K}(\omega = \eta) \to \mathsf{K}_2(\omega = \eta) \to \mathsf{K}_3(\omega = \eta) \to \cdots \to \mathsf{L}(\omega = \eta)$$

is a d-collapsing of $\mathsf{K}(\omega \simeq \eta)$ to $\mathsf{L}(\omega \simeq \eta)$.

It is sufficient to show $\mathsf{K}(\omega = \eta) \to \mathsf{K}_2(\omega = \eta)$ and proceed by induction.

For purposes of this proof, we distinguish faces before gluing $\omega = \eta$ by Greek letters, say σ, σ', and after gluing by Greek letters in brackets, say $[\sigma], [\sigma']$. E.g., we have $\omega \neq \eta$, but $[\omega] = [\eta]$.

Suppose that $\mathsf{K}_2 = \mathsf{K}_\sigma$ for a d-collapsible face σ. We want to show that $[\tau(\sigma)]$ is the unique maximal face containing $[\sigma]$. By the distance condition, we can without loss of generality assume that $\sigma \cap \eta = \emptyset$ (otherwise we swap ω and η). Suppose $[\sigma'] \supseteq [\sigma]$. Now we show that $\sigma' \supseteq \sigma$: if $\sigma \cap \omega = \emptyset$ we have it immediately (since vertices of σ are not identified with any other); if $\sigma \cap \omega \neq \emptyset$ then $\sigma' \cap \eta = \emptyset$ due to the distance condition, which implies $\sigma' \supseteq \sigma$. Hence $\tau(\sigma) \supseteq \sigma'$, and $[\tau(\sigma)] \supseteq [\sigma']$. Thus $[\tau(\sigma)]$ is the unique maximal face containing $[\sigma]$. ☐

Lemma 5.6 (Collapsing of connecting gadget)**.** *Let t be an integer. Let $\overline{\mathsf{L}}$ be a complex with distinct d-dimensional faces σ, $\gamma_1, \dots, \gamma_t$ such that σ is a maximal face of $\overline{\mathsf{L}}$. Let $\mathsf{C} = \mathsf{C}(\rho, \zeta_1, \dots, \zeta_t)$ and $\mathsf{C}' = \mathsf{C}'(\rho, \zeta_1, \dots, \zeta_t)$ be complexes defined in Section 3.*

Then the complex $(\overline{\mathsf{L}} \dot\cup \mathsf{C})(\sigma = \rho, \zeta_1 = \gamma_1, \dots, \zeta_t = \gamma_t)$ d-collapses to the complex $(\overline{\mathsf{L}} \dot\cup \mathsf{C}')(\sigma = \rho, \zeta_1 = \varphi_1, \dots, \zeta_t = \gamma_t) \setminus \{\sigma\}$.

Proof. First, we observe that

$$(\overline{\mathsf{L}} \dot\cup \mathsf{C})(\sigma = \rho) \twoheadrightarrow (\overline{\mathsf{L}} \dot\cup \mathsf{C}')(\sigma = \rho) \setminus \{\sigma\}.$$

This follows from Lemma 5.3 by setting $\mathsf{K} = (\overline{\mathsf{L}} \dot\cup \mathsf{C})(\sigma = \rho)$, $\mathsf{K}' = \mathsf{C}$, $\mathsf{L}' = \mathsf{C}' \setminus \{\sigma\}$, and then $\mathsf{L} = (\overline{\mathsf{L}} \dot\cup \mathsf{C}')(\sigma = \rho) \setminus \{\sigma\}$. Assumptions of the lemma are satisfied by Proposition 3.1(ii) and the inspection.

Now it is sufficient to iterate Lemma 5.5, assumptions are satisfied by Proposition 3.1(i). ☐

6 Acknowledgement

I would like to thank to Jiří Matoušek for discussing the contents and reading the preliminary version of this paper.

References

[AK92] N. Alon and D. Kleitman. Piercing convex sets and the Hadwiger Debrunner (p, q)-problem. *Adv. Math.*, 96(1):103–112, 1992.

[AKMM02] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam. Transversal numbers for hypergraphs arising in geometry. *Adv. Appl. Math*, 130:2509–2514, 2002.

[Ame96] N. Amenta. A short proof of an interesting Helly-type theorem. *Discrete Comput. Geom*, 15:423–427, 1996.

[Bár82] I. Bárány. A generalization of Carathéodory's theorem. *Discrete Math.*, 40:141–152, 1982.

[Coo71] S. A. Cook. The complexity of theorem proving procedures. In *Proc. 3rd Ann. ACM Symp. on Theory of Computing*, pages 151–158, 1971.

[Hat01] A. Hatcher. *Algebraic Topology.* Cambridge University Press, Cambridge, 2001.

[Hel23] E. Helly. Über mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresbericht Deustch. Math. Verein.*, 32:175–176, 1923.

[Hel30] E. Helly. Über Systeme von abgeschlossenen Mengen mit gemeinschaftlichen Punkten. *Monaths. Math. und Physik*, 37:281–302, 1930.

[KL79] M. Katchalski and A. Liu. A problem of geometry in R^n. *Proc. Amer. Math. Soc.*, 75:284–288, 1979.

[KM89] J. Kratochvíl and J. Matoušek. NP-hardness results for intersection graphs. *Coment. Math. Univ. Carolin.*, 30:761–773, 1989.

[KM94] J. Kratochvíl and J. Matoušek. Intersection graphs of segments. *J. Comb. Theory Ser. B*, 62(2):289–315, 1994.

[KM05] G. Kalai and R. Meshulam. A topological colorful Helly theorem. *Adv. Math.*, 191(2):305–311, 2005.

[KM08] G. Kalai and R. Meshulam. Leray numbers of projections and a topological Helly type theorem. *J. Topology*, 1(3):551–556, 2008.

[Kra91] J. Kratochvíl. String graphs II. Recognizing string graphs is NP-hard. *J. Comb. Theory Ser. B*, 52:67–78, 1991.

[Kra94] J. Kratochvíl. A special planar satisfiability problem and a consequence of its NP-completeness. *Discrete Appl. Math.*, 50(3):297–302, 1994.

[LB62] C. G. Lekkerkerker and J. C. Boland. Representation of a finite graph by a set of intervals on the real line. *Fund. Math.*, 51:45–64, 1962.

[Lov74] L. Lovász. Problem 206. *Matematikai Lapok*, 25:181, 1974.

[Mat03] J. Matoušek. *Using the Borsuk-Ulam Theorem.* Springer, Berlin etc., 2003.

[Mat08] J. Matoušek. Removing degeneracy in LP-type problems revisited. *Discrete and Comp. Geometry*, 2008. To appear.

[MF08] R. Malgouyres and A. R. Francés. Determining whether a simplicial 3-complex collapses to a 1-complex is NP-complete. *DGCI*, pages 177–188, 2008.

[MT08] J. Matoušek and M. Tancer. Dimension gaps between representability and collapsibility. *Discrete and Comp. Geometry*, 2008. To appear.

[Mun84] J. R. Munkres. *Elements of Algebraic Topology.* Addison - Wesley, 1984.

[Weg75] G. Wegner. d-collapsing and nerves of families of convex sets. *Arch. Math.*, 26:317–321, 1975.

A The complexity of d-representability

In this appendix we prove that d-REPRESENTABILITY is NP-hard for $d \geq 2$.

Intersection graphs. Let $\mathcal{F}$ be a set system. The *intersection graph* $I(\mathcal{F})$ of $\mathcal{F}$ is defined as the (simple) graph such that the set of its vertices is the set $\mathcal{F}$ and the set of its edges is the set $\{\{F, F'\} \mid F, F' \in \mathcal{F}, F \neq F', F \cap F' \neq \emptyset\}$. Alternatively, $I(\mathcal{F})$ is the 1-skeleton of the nerve of $\mathcal{F}$.

A *string graph* is a graph, which is isomorphic to an intersection graph of finite collection of curves in the plane. By STR we denote the set of all string graphs. By CON we denote the class of intersection graphs of finite collections of convex sets in the plane and by SEG we denote the class of intersection graphs of finite collections of segments in the plane. Finally, by SEG(≤ 2) we denote the class of intersection graphs of finite collections of segments in the plane such that no three segments share a common point.

Suppose that G is a string graph. A system $\mathcal{C}$ of curves in the plane such that G is isomorphic $I(\mathcal{C})$ is called an STR-*representation* of G. Similar definitions apply for another classes. Similarly, but for simplicial complexes, suppose that K is a d-representable simplicial complex. A system $\mathcal{C}$

of convex sets in $\mathbb{R}^d$ such that K is isomorphic to the nerve of $\mathcal{C}$ is called a *d-representation* of K.

We have STR $\supseteq$ CON $\supseteq$ SEG (actually, it is known that the inclusions are strict). Furthermore, suppose that we are given a graph $G \in$ SEG. By Kratochvíl and Matoušek [KM94, Lemma 4.1], there is a SEG-representation of G such that no two parallel segments of this representation intersect. By a small perturbation, we can even assume that no three segments of this representation share a common point. Hence SEG = SEG(≤ 2).

NP-hardness of 2-representability. Kratochvíl and Matoušek [KM89] prove that for the classes mentioned above (i.e. STR, CON and SEG), it is NP-hard to recognize whether a given graph belongs to the given class. For this they reduce planar 3-connected 3-satisfiability (P3C3SAT) to this problem (see [Kra94] for the proof of NP-completeness of P3C3SAT and another background). More precisely (see [KM89, the proof of Prop. 2]), given a formula Φ of P3C3SAT they construct a graph $G(\Phi)$ such that $G(\Phi) \in$ SEG if the formula is satisfiable, but $G(\Phi) \notin$ STR if the formula is unsatisfiable. Moreover, we already know that this yields $G(\Phi) \in$ SEG(≤ 2) for satisfiable formulae.

Let us consider $G(\Phi)$ as a 1-dimensional simplicial complex. We will derive that $G(\Phi)$ is 2-representable if and only if Φ is satisfiable.

If we are given a 2-representation of $G(\Phi)$ it is also a CON-representation of $G(\Phi)$, since $G(\Phi)$ is 1-dimensional. Hence $G(\Phi)$ is not 2-representable for unsatisfiable formulae.

On the other hand, a SEG(≤ 2)-representation of $G(\Phi)$ is also a 2-representation of $G(\Phi)$. Thus $G(\Phi)$ is 2-representable for satisfiable formulae.

In summary, we have that 2-REPRESENTABILITY is NP-hard.

***d*-representability of suspension.** Let K be a simplicial complex and let a and b be two new vertices. By *suspension* of K we mean the simplicial complex

$$\operatorname{susp} \mathsf{K} = \mathsf{K} \cup \{\{a\} \cup \sigma \mid \sigma \in \mathsf{K}\} \cup \{\{b\} \cup \sigma \mid \sigma \in \mathsf{K}\}.$$

Lemma A.1. *Let d be an integer. A simplicial complex K is $(d-1)$-representable if and only if the suspension $\operatorname{susp} \mathsf{K}$ is d-representable.*

Proof. First, we suppose that K is $(d-1)$-representable and we show that $\operatorname{susp} \mathsf{K}$ is d-representable. Let $K_1, \ldots, K_t \subseteq \mathbb{R}^{d-1}$ be convex set from a $(d-1)$-representation of K. Let $K(a)$ and $K(b)$ be hyperplanes $\mathbb{R}^{d-1} \times \{0\}$ and $\mathbb{R}^{d-1} \times \{1\}$ in $\mathbb{R}^d$. It is easy to see, that the nerve of the family

$$\{K_1 \times [0,1], \ldots, K_t \times [0,1], K(a), K(b)\}$$

of convex sets in $\mathbb{R}^d$ is isomorphic to $\operatorname{susp} \mathsf{K}$.

For the reverse implication, we suppose that $\operatorname{susp} \mathsf{K}$ is d-representable and we show that K is $(d-1)$-representable. Suppose that $K(a), K(b), K_1 \dots, K_t$ is a d-representation of $\operatorname{susp} \mathsf{K}$ ($K(a)$ corresponds to a and $K(b)$ corresponds to b). We have that $\{a, b\} \notin \operatorname{susp} \mathsf{K}$, thus there is a hyperplane $H \subseteq \mathbb{R}^d$ separating $K(a)$ and $K(b)$ (we can assume that the sets in the representation are compact). Then the nerve of the family

$$\{K_1 \cap H, \dots, K_t \cap H\}$$

of convex sets in $H \simeq \mathbb{R}^{d-1}$ is isomorphic to K. □

Since 2-REPRESENTABILITY is NP-hard, we have the following corollary of Lemma A.1 (considering complexes that are obtained as $(d-2)$-tuple suspensions):

Theorem A.2. *d-REPRESENTABILITY is NP-hard for $d \geq 2$.*

□