

Shellability is NP-Complete

Xavier Goaoc¹

LIGM, Université Paris-Est

UMR 8049 CNRS, ENPC, ESIEE, UPEM, F-77454, Marne-la-Vallée, France


xavier.goaoc@u-pem.fr

Pavel Paták

Department of Mathematics and Statistics, Masaryk University

Brno, Czech Republic

patak@math.muni.cz


 <https://orcid.org/0000-0003-3016-0278>

Zuzana Patáková

IST Austria

Klosterneuburg, Austria

zuzka@kam.mff.cuni.cz

 <https://orcid.org/0000-0002-3975-1683>

Martin Tancer²

Department of Applied Mathematics, Charles University

Prague, Czech Republic


tancer@kam.mff.cuni.cz

Uli Wagner

IST Austria

Klosterneuburg, Austria

uli@ist.ac.at

 <https://orcid.org/0000-0002-1494-0568>

Abstract

We prove that for every $d \geq 2$, deciding if a pure, d -dimensional, simplicial complex is shellable is NP-hard, hence NP-complete. This resolves a question raised, e.g., by Danaraj and Klee in 1978. Our reduction also yields that for every $d \geq 2$ and $k \geq 0$, deciding if a pure, d -dimensional, simplicial complex is k -decomposable is NP-hard. For $d \geq 3$, both problems remain NP-hard when restricted to contractible pure d -dimensional complexes.

2012 ACM Subject Classification Mathematics of computing → Geometric topology, Theory of computation → Computational geometry

Keywords and phrases Shellability, simplicial complexes, NP-completeness, collapsibility

Digital Object Identifier 10.4230/LIPIcs.SoCG.2018.41

Related Version A full version of the paper is available at [14], <https://arxiv.org/abs/1711.08436>.

Funding Partially supported by the project EMBEDS II (CZ: 7AMB17FR029, FR: 38087RM) of Czech-French collaboration.

¹ Partially supported by IUF.

² Partially supported by the GAČR grant 16-01602Y and by Charles University project UNCE/SCI/004.



© Xavier Goaoc, Pavel Paták, Zuzana Patáková, Martin Tancer, and Uli Wagner; licensed under Creative Commons License CC-BY


34th International Symposium on Computational Geometry (SoCG 2018).

Editors: Bettina Speckmann and Csaba D. Tóth; Article No. 41; pp. 41:1–41:15

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



Acknowledgements We would like to thank Andrés D. Santamaría Galvis, Russ Woodroffe and anonymous referees for useful remarks on a preliminary version on this paper.

1 Introduction

A d -dimensional simplicial complex is called *pure* if all its facets (i.e., inclusion-maximal faces) have the same dimension d . A pure d -dimensional simplicial complex is *shellable* if there exists a linear ordering $\sigma_1, \sigma_2, \dots, \sigma_n$ of its facets such that, for every $i \geq 2$, $\sigma_i \cap (\cup_{j < i} \sigma_j)$ is a pure $(d - 1)$ -dimensional simplicial complex; such an ordering is called a *shelling* or *shelling order*. For example, the boundary of a simplex is shellable (any order works), but no triangulation of the torus is (the condition fails for the first triangle σ_i that creates a non-contractible 1-cycle).

The concept of shellings originated in the theory of convex polytopes (in a more general version for polytopal complexes), as an inductive procedure to construct the boundary of a polytope by adding the facets one by one in such a way that all intermediate complexes (except the last one) are contractible. The fact that this is always possible, i.e., that convex polytopes are shellable, was initially used as an unproven assumption in early papers (see the discussion in [15, pp. 141–142] for a more detailed account of the history), before being proved by Bruggesser and Mani [9].

The notion of shellability extends to more general objects (including non-pure simplicial complexes and posets [8]), and plays an important role in diverse areas including piecewise-linear topology [26, 3], polytope theory (e.g., McMullen’s proof of the *Upper Bound Theorem* [22]), topological combinatorics [5], algebraic combinatorics and commutative algebra [28, 24], poset theory, and group theory [4, 27]; for a more detailed introduction and further references see [31, §3].

One of the reasons for its importance is that shellability – a combinatorial property – has strong topological implications: For example, if a pure d -dimensional complex K is a *pseudomanifold*³ – which can be checked in linear time – and shellable, then K is homeomorphic to the sphere S^d (or the ball B^d , in case K has nonempty boundary) [10] – a property that is algorithmically undecidable for $d \geq 5$, by a celebrated result of Novikov [30, 23]. More generally, every pure d -dimensional shellable complex is homotopy equivalent to a wedge of d -spheres, in particular it is $(d - 1)$ -connected.

1.1 Results

From a computational viewpoint, it is natural to ask if one can decide efficiently (in polynomial time) whether a given complex is shellable. This question was raised at least as early as in the 1970’s [11, 12] (see also [18, Problem 34]) and is of both practical and theoretical importance (besides direct consequences for the experimental exploration of simplicial complexes, the answer is also closely related to the question there are simple conditions that would characterize shellability). Danaraj and Klee proved that shellability of 2-dimensional pseudomanifolds can be tested in linear time [11], whereas a number of related problems have been shown to be NP-complete [13, 19, 17, 20, 29, 2] (see Section 1.2), but the computational complexity of the shellability problem has remained open. Here we show:⁴

³ A pure, d -dimensional complex K is a pseudomanifold (with boundary) if every $(d - 1)$ -face of K is contained in exactly two (at most two) facets. (Sometimes, it is additionally required that the facet-adjacency graph of K is connected, but this does not matter in our setting, since shellable complexes always satisfy this connectivity property.)

⁴ For basic notions from computational complexity, such as NP-completeness or reductions, see, e.g., [1].

► **Theorem 1.** *Deciding if a pure 2-dimensional simplicial complex is shellable is NP-complete.*

Here, the input is given as a finite abstract simplicial complex (see Section 2).⁵

► **Remark.** The problem of testing shellability is easily seen to lie in the complexity class NP (given a linear ordering of the facets of a complex, it is straightforward to check whether it is a shelling). Thus, the nontrivial part of Theorem 1 is that deciding shellability of pure 2-dimensional complexes is NP-hard.

It is easy to check that a pure simplicial complex K is shellable if and only if the cone $\{v\} * K$ is shellable, where v is a vertex not in K (see Section 2). Thus, the hardness of deciding shellability easily propagates to higher-dimensional complexes, even to cones.

► **Corollary 2.** *For $d \geq 3$, deciding if a pure d -dimensional complex is shellable is NP-complete even when the input is assumed to be a cone (hence contractible).*

Moreover, our hardness reduction (from 3-SAT) used in the proof of Theorem 1 (see Section 3) turns out to be sufficiently robust to also imply hardness results for a number of related problems.

Hardness of k -decomposability and CL-shellability. Let $d \geq 2$ and $k \geq 0$. A pure d -dimensional simplicial complex K is k -decomposable if it is a simplex or if there exists a face σ of K of dimension at most k such that (i) the link of σ in K is pure $(d - |\sigma|)$ -dimensional and k -decomposable, and (ii) deleting σ and faces of K containing σ produces a d -dimensional k -decomposable complex. This notion, introduced by Provan and Billera [25], provides a hierarchy of properties (k -decomposability implies $(k + 1)$ -decomposability) interpolating between *vertex-decomposable* complexes ($k = 0$) and shellable complexes (shellability is equivalent to d -decomposability [25]). The initial motivation for considering this hierarchy was to study the *Hirsch conjecture* on combinatorial diameters of convex polyhedra, or in the language of simplicial complex, the diameter of the facet-adjacency graphs of pure simplicial complexes: at one end, the boundary complex of every d -dimensional simplicial polytope is shellable [9], and at the other end, every vertex-decomposable simplicial complex has small diameter (it satisfies the *Hirsch bound* [25]).

► **Theorem 3.** *Let $d \geq 2$ and $k \geq 0$. Deciding if a pure d -dimensional simplicial complex is k -decomposable is NP-hard. For $d \geq 3$, the problem is already NP-hard for pure d -dimensional simplicial complexes that are cones (hence contractible).*

Another notion related to shellability is the *CL-shellability* of a poset, introduced in [6]. We do not reproduce the definition here, but note that a simplicial complex is shellable if and only if the dual of its face lattice is CL-shellable [7, Corollary 4.4]. For any fixed dimension d the face lattice can be computed in polynomial time, so we get:

► **Corollary 4.** *Deciding whether a given poset is CL-shellable is NP-hard.*

⁵ There are several different ways of encoding an abstract simplicial complex – e.g., we can list the facets, or we can list all of its simplices –, but since we work with complexes of fixed dimension, these encodings can be translated into one another in polynomial time, so the precise choice does not matter.

1.2 Related work on collapsibility and our approach

Our proof of Theorem 1 builds on earlier results concerning *collapsibility*, a combinatorial analogue, introduced by Whitehead [32], of the topological notion of contractibility.⁶ A face σ of a simplicial complex K is *free* if there is a unique inclusion-maximal face τ of K with $\sigma \subsetneq \tau$. An *elementary collapse* is the operation of deleting a free face and all faces containing it. A simplicial complex K *collapses* to a subcomplex $L \subseteq K$ if L can be obtained from K by a finite sequence of elementary collapses; K is called *collapsible* if it collapses to a single vertex.

The problem of deciding whether a given 3-dimensional complex is collapsible is NP-complete [29]; the proof builds on earlier work of Malgouyres and Francés [20], who showed that it is NP-complete to decide whether a given 3-dimensional complex collapses to some 1-dimensional subcomplex. By contrast, collapsibility of 2-dimensional complexes can be decided in polynomial time (by a greedy algorithm) [17, 20]. It follows that for any *fixed* integer k , it can be decided in polynomial time whether a given 2-dimensional simplicial complex can be made collapsible by deleting at most k faces of dimension 2; by contrast, the latter problem is NP-complete if k is part of the input [13].⁷

Our reduction uses the gadgets introduced by Malgouyres and Francés [20] and reworked in [29] to prove NP-hardness of deciding collapsibility for 3-dimensional complexes. However, these gadgets are not pure: they contain maximal simplices of two different dimensions, 2 and 3. Roughly speaking, we fix this by replacing the 3-dimensional subcomplexes by suitably triangulated 2-spheres and modifying the way in which they are glued. Interestingly, this also makes our reduction robust to subdivision and applicable to other types of decomposition.

Collapsibility and shellability. Furthermore, we will use the following connection between shellability and collapsibility, due to Hachimori [16] (throughout, $\tilde{\chi}$ denotes the reduced Euler characteristic).

► **Theorem 5** ([16, Theorem 8]). *Let K be a 2-dimensional simplicial complex. The second barycentric subdivision $\text{sd}^2 K$ is shellable if and only if the link of each vertex of K is connected and there exists $\tilde{\chi}(K)$ triangles in K whose removal makes K collapsible.*

At first glance, Hachimori’s theorem might suggest to prove Theorem 1 by a direct polynomial-time reduction of collapsibility to shellability. However, for 2-dimensional complexes this would not imply hardness, since, as mentioned above, collapsibility of 2-dimensional complexes is decidable in polynomial time [17, 20]. Instead, we will use the existential part of Hachimori’s theorem (“there exists $\tilde{\chi}(K)$ triangles”) to encode instances of the 3-SAT problem, a classical NP-complete problem.

2 Notation and terminology

We give here an overview of the basic terminology, including the notions used but not defined in the introduction. We assume that the reader is familiar with standard concepts regarding simplicial complexes, and mostly list the notions we use and set up the notation.

⁶ Collapsibility implies contractibility, but the latter property is undecidable for complexes of dimension at least 4 (this follows from Novikov’s result [30], see [29, Appendix A]), whereas the problem of deciding collapsibility lies in NP.

⁷ We remark that building on [13], a related problem, namely computing *optimal discrete Morse matchings* in simplicial complexes (which we will not define here), was also shown to be NP-complete [19, 17].

We recall that the input in Theorem 1 is assumed to be described as an abstract simplicial complex,⁸ i.e., a purely combinatorial object. For the purposes of the exposition, however, it will be more convenient to use a description via geometric simplicial complexes.⁹ In fact, in our construction, we will sometimes first describe a polyhedron¹⁰ and only then a geometric simplicial complex triangulating the polyhedron, with the understanding that this is simply a convenient way to specify the associated abstract simplicial complex.

A *subdivision* of a (geometric) complex K is a complex K' such that the polyhedra of K and of K' coincide and every simplex of K' is contained in some simplex of K . The *reduced Euler characteristic* of a complex K is defined as $\tilde{\chi}(K) = \sum_{i=-1}^{\dim K} (-1)^i f_i(K)$ where $f_i(K)$ is the number of i -dimensional faces of K and, by convention, $f_{-1}(K)$ is 0 if K is empty and 1 otherwise.

For the definitions of *links*, the *barycentric subdivision*, $\text{sd} K$, or the *join* $K * L$ of two complexes K and L we refer to the standard sources such as [21, Chapter 1] (or to the full version [14]). We denote by Δ_ℓ the simplex of dimension ℓ .

3 The main proposition and its consequences

The cornerstone of our argument is the following construction:

► **Proposition 6.** *There is an algorithm that, given a 3-CNF formula¹¹ ϕ , produces, in time polynomial in the size of ϕ , a 2-dimensional simplicial complex K_ϕ with the following properties:*

- (i) *the link of every vertex of K_ϕ is connected,*
- (ii) *if ϕ is satisfiable, then K_ϕ becomes collapsible after removing some $\tilde{\chi}(K_\phi)$ triangles,*
- (iii) *if an arbitrary subdivision of K_ϕ becomes collapsible after removing some $\tilde{\chi}(K_\phi)$ triangles, then ϕ is satisfiable.*

The rest of this section derives our main result and its variants from Proposition 6. We then describe the construction of K_ϕ in Section 4 and prove Proposition 6 in Sections 5 and 6 (modulo a few claims, treated in detail in the full version [14]).

Hardness of shellability. Proposition 6 and Hachimori's theorem imply our main result:

Proof of Theorem 1. Let ϕ be a 3-CNF formula and let K_ϕ denote the 2-dimensional complex built according to Proposition 6. Since the link of every vertex of K_ϕ is connected, Theorem 5 guarantees that $\text{sd}^2 K_\phi$ is shellable if and only if there exist $\tilde{\chi}(K_\phi)$ triangles whose removal makes K_ϕ collapsible. Hence, by statements (ii) and (iii), the formula ϕ is satisfiable

⁸ A (finite) *abstract simplicial complex* is a collection K of subsets of a finite set V that is closed under taking subsets, i.e., if $\sigma \in K$ and $\tau \subseteq \sigma$, then $\tau \in K$. The elements $v \in V$ are called the *vertices* of K (and often identified with the singleton sets $\{v\} \in K$), and the elements of K are called *faces* or *simplices* of K . The *dimension* of a face is its cardinality minus 1, and the *dimension* of K is the maximum dimension of any face. This is a purely combinatorial description of a simplicial complex and a natural input model for computational questions.

⁹ A (finite) *geometric simplicial complex* is a finite collection K of geometric simplices (convex hulls of affinely independent points) in \mathbb{R}^d (for some d) such that (i) if $\sigma \in K$ and τ is a face of σ , then τ also belongs to K , and (ii) if $\sigma_1, \sigma_2 \in K$, then $\sigma_1 \cap \sigma_2$ is a face of both σ_1 and σ_2 . There is a straightforward translation between the two descriptions (see, e.g. [21, Chapter 1]), and this is the setting we will work in for the rest of the article.

¹⁰ The *polyhedron* of a geometric simplicial complex K is defined as the union of simplices contained in K , $\bigcup_{\sigma \in K} \sigma$. We also say that K *triangulates* $X \subseteq \mathbb{R}^d$ if X is the polyhedron of K . Note that a given polyhedron usually has many different triangulations.

¹¹ That is, a boolean formula in conjunctive normal form such that each clause consists of three literals.

41:6 Shellability is NP-Complete

if and only if $\text{sd}^2 K_\phi$ is shellable. Taking the barycentric subdivision of a two-dimensional complex multiplies its number of simplices by at most a constant factor. The complex $\text{sd}^2 K_\phi$ can thus be constructed from ϕ in polynomial time, and 3-SAT reduces in polynomial time to deciding the shellability of 2-dimensional pure complexes. ◀

Hardness of k -decomposability. Note that statement (iii) in Proposition 6 deals with arbitrary subdivision whereas we needed it without subdivisions in the proof above. This extra elbow room comes at no cost in the proof of Proposition 6 and yields the NP-hardness of k -decomposability.

Proof of Theorem 3. Assume without loss of generality that $k \leq d$. Let ϕ be a 3-CNF formula and K_ϕ the complex produced by Proposition 6. We have the following implications:¹²

$$\begin{aligned}
 \phi \text{ is satisfiable} &\Rightarrow K_\phi \text{ is collapsible after removal of some } \tilde{\chi}(K_\phi) \text{ triangles} \\
 &\Rightarrow \text{sd}^2 K_\phi \text{ is shellable} \\
 &\Rightarrow_{(b)} \text{sd}^3 K_\phi \text{ is vertex-decomposable} \\
 &\Rightarrow_{(c)} \Delta_{d-3} * \text{sd}^3 K_\phi \text{ is vertex-decomposable} \\
 &\Rightarrow_{(a)} \Delta_{d-3} * \text{sd}^3 K_\phi \text{ is } k\text{-decomposable} \\
 &\Rightarrow_{(a)} \Delta_{d-3} * \text{sd}^3 K_\phi \text{ is shellable (i.e., } d\text{-decomposable)} \\
 &\Rightarrow_{(d)} \text{sd}^3 K_\phi \text{ is shellable} \\
 &\Rightarrow \text{sd} K_\phi \text{ is collapsible after removal of some } \tilde{\chi}(K_\phi) \text{ triangles} \\
 &\Rightarrow \phi \text{ is satisfiable}
 \end{aligned}$$

The first and last implications are by construction of K_ϕ (Proposition 6). The second and second to last follow from Theorem 5, given that Proposition 6 ensures that links of vertices in K_ϕ are connected. The remaining implications follow from the following known facts:

- (a) if K is k -decomposable, then K is k' -decomposable for $k' \geq k$,
- (b) if K is shellable, then $\text{sd} K$ is vertex-decomposable [8],
- (c) K is vertex-decomposable if and only if $\Delta_\ell * K$ is vertex-decomposable [25, Prop. 2.4],
- (d) K is shellable if and only if $\Delta_\ell * K$ is shellable (see the full version [14] for details).

Since the first and last statement are identical, these are all equivalences. In particular, ϕ is satisfiable if and only if $\Delta_{d-3} * \text{sd}^3 K_\phi$ is k -decomposable. Since this complex can be computed in time polynomial in the size of K_ϕ , *i.e.*, polynomial in the size of ϕ , the first statement follows. Since $\Delta_{d-3} * \text{sd}^3 K_\phi$ is contractible for $d \geq 3$, the second statement follows. ◀

4 Construction

We now define the complex K_ϕ mentioned in Proposition 6. This complex consists of several building blocks, called *gadgets*. We first give a “functional” outline of the gadgets (in Section 4.1), insisting on the properties that guided their design, before moving on to the details of their construction and gluing (Sections 4.2 and 4.3).

We use the notational convention that complexes that depend on a variable u are denoted with round brackets, *e.g.* $f(u)$, whereas complexes that depend on a literal are denoted with square brackets, *e.g.* $f[u]$ or $f[\neg u]$.

¹²In the case $d = 2$, we use the convention that $\Delta_{-1} * L = L$ for any simplicial complex L .

4.1 Outline of the construction

The gadgets forming K_ϕ are designed with two ideas in mind. First, every gadget, when considered separately, can only be collapsed starting in a few *special edges*. Next, the special edges of each gadgets are intended to be glued to other gadgets, so as to create dependencies in the flow of collapses: if an edge f of a gadget \mathbf{G} is attached to a triangle of another gadget \mathbf{G}' , then \mathbf{G} cannot be collapsed starting by f before some part of \mathbf{G}' has been collapsed.

Variable gadgets. For every variable u we create a gadget $\mathbf{V}(u)$. This gadget has three special edges; two are associated, respectively, with TRUE and FALSE; we call the third one “unlocking”. Overall, the construction ensures that any removal of $\tilde{\chi}(K_\phi)$ triangles from K_ϕ either frees exactly one of the edges associated with TRUE or FALSE in every variable gadget, or makes K_ϕ obviously non-collapsible. This relates the removal of triangles in K_ϕ to the assignment of variables in ϕ . We also ensure that part of each variable gadget remains uncollapsible until the special unlocking edge is freed.

Clause gadgets. For every clause $c = \ell_1 \vee \ell_2 \vee \ell_3$ we create a gadget $\mathbf{C}(c)$. This gadget has three special edges, one per literal ℓ_i . Assume that $\ell_i \in \{u, \neg u\}$. Then the special edge associated with ℓ_i is attached to $\mathbf{V}(u)$ so that it can be freed if and only if the triangle removal phase freed the special edge of $\mathbf{V}(u)$ associated with TRUE (if $\ell_i = u$) or with FALSE (if $\ell_i = \neg u$). This ensures that the gadget $\mathbf{C}(c)$ can be collapsed if and only if one of its literals was “selected” at the triangle removal phase.

Conjunction gadget. We add a gadget \mathbf{A} with a single special edge, that is attached to every clause gadget. This gadget can be collapsed only after the collapse of every clause gadget has started (hence, if every clause contains a literal selected at the triangle removal phase). In turn, the collapse of \mathbf{A} will free the unlocking special edge of every variable gadget, allowing to complete the collapse.

Notations. For any variable u , we denote the special edges of $\mathbf{V}(u)$ associated with TRUE and FALSE by, respectively, $f[u]$ and $f[\neg u]$; we denote the unlocking edge by $f(u)$. For every clause $c = \ell_1 \vee \ell_2 \vee \ell_3$, we denote by $f[\ell_i, c]$ the special edge of $\mathbf{C}(c)$ associated with ℓ_i . We denote by f_{and} the special edge of the conjunction gadget \mathbf{A} . The attachment of these edges are summarized in the table below.

gadget	special edges	attached to	freed by
$\mathbf{V}(u)$	$f[u]$ $f[\neg u]$ $f(u)$	- - \mathbf{A}	triangle deletion triangle deletion freeing f_{and}
$\mathbf{C}(u_2 \vee \neg u_4 \vee u_9)$	$f[u_2, c]$ $f[\neg u_4, c]$ $f[u_9, c]$	$\mathbf{V}(u_2)$ $\mathbf{V}(u_4)$ $\mathbf{V}(u_9)$	freeing $f[u_2]$ freeing $f[\neg u_4]$ freeing $f[u_9]$
\mathbf{A}	f_{and}	every clause gadget	collapsing all clause gadgets

Flow of collapses. Let us summarize the mechanism sketched above. Assume that ϕ is satisfiable, and consider a satisfying assignment. Remove the triangles from each $\mathbf{V}(u)$ so that the edge that becomes free is $f[u]$ if u was assigned TRUE, and $f[\neg u]$ otherwise. This will allow to collapse each clause gadget in order to make f_{and} free. Consequently, we will be


able to collapse \mathbf{A} and make all unlocking edges $f(u)$ free. This allows finishing the collapses on all $\mathbf{V}(u)$.

On the other hand, to collapse K_ϕ we must collapse f_{and} at some point. Before this can happen, we have to collapse in each clause $c = \ell_1 \vee \ell_2 \vee \ell_3$ one of the edges $f[\ell_i, c]$. This, in turn, requires that $f[\ell_i]$ has been made free. If we can ensure that $f[\neg \ell_i]$ cannot also be free, then we can read off from the collapse an assignment of the variables that must satisfy every clause, and therefore ϕ . (If $\ell_i = u$, then we set u to TRUE, if $\ell_i = \neg u$, then we set u to FALSE. If there are unassigned variables after considering all clauses, we assign them arbitrarily.)

4.2 Preparation: modified Bing's houses

Our gadgets rely on two modifications of Bing's house, a classical example of a 2-dimensional simplicial complex that is contractible but not collapsible. Bing's house consists of a box split into two parts (or *rooms*); each room is connected to the outside by a tunnel through the other room; each tunnel is attached to the room that it traverses by a rectangle (or *wall*). The modifications that we use here make the complex collapsible, but restricts its set of free faces to exactly one or exactly three edges.

One free edge. We use here a modification due to Malgouyres and Francés [20]. In one of the rooms (say the top one), the wall has been thickened and hollowed out, see the figure below. We call the resulting polyhedron a Bing's house with a single free edge, or a *1-house* for short. Two special elements of a 1-house are its *free edge* (denoted f and in thick stroke in the figure below) and its *lower wall* rectangle (denoted L and colored in light blue in the figure below). We only consider triangulations of 1-houses that subdivide the edge f and the lower wall L . We use 1-houses for the following properties:




► **Lemma 7.** *Let B be a 1-house, f its free edge and L its lower wall. In any triangulation of B , the free faces are exactly the edges that subdivide f . Moreover, B collapses to any subtree of the 1-skeleton of B that is contained in L and shares with the boundary of L a single endpoint of f .*

The first statement follows from the fact that the edges that subdivide f are the only ones that are not part of two triangles; see [20, Remark 1]. The second statement was proven in [29, Lemma 7] for certain trees, but the argument holds for arbitrary trees; see the full version [14] for details.

When working with 1-houses, we will usually only describe the lower wall to clarify which subtree we intend to collapse to.

Three free edges. We also use the Bing's houses with three collapsed walls introduced in [29]; we call them *3-houses* for short. These are 2-dimensional complexes whose construction is more involved; we thus state its main properties, so that we can use it as a black box, and

refer the reader interested in its precise definition to [29, §4]. Refer to the figure on the right (which corresponds to Figure 9 in [29]). The 3-house has exactly three free edges f_1, f_2, f_3 , and has three distinguished paths p_1, p_2, p_3 sharing a common vertex v and such that each p_i shares exactly one vertex with f_i and no vertex with f_j for $j \neq i$. In addition, it contains an edge e incident to v so that the union of $p_1, p_2, p_3, f_1, f_2, f_3$ and e forms a subdivided star (in graph-theoretic sense) with four rays.



Let C denote the 3-house as described above. In [29], the polyhedron of C is described in detail but no triangulation is specified. We are happy with any concrete triangulation for which Lemma 8 below holds; we can in addition require that the paths p_1, p_2 and p_3 each consist of two edges.¹³


► **Lemma 8** ([29, Lemma 8]). *In any subdivision of C , the free faces are exactly the edges that subdivide f_1, f_2 and f_3 . Moreover, C collapses to the 1-complex spanned by e, p_1, p_2, p_3 and any two of $\{f_1, f_2, f_3\}$.*

4.3 Detailed construction

Section 4.1 gave a quick description of the intended functions of the various gadgets. We now flesh them out and describe how they are glued together.


Triangulations. For some parts of the complex, it will be convenient to first describe the polyhedron, then discuss its triangulation. Our description of the triangulation may vary in precision: it may be omitted (if any reasonable triangulation works), given indirectly by the properties it should satisfy, or given explicitly (for instance to make it clear that we can glue the gadgets as announced).

Conjunction gadget. The conjunction gadget **A** is a 1-house. We let f_{and} denote its free edge and v_{and} one of the endpoints of f_{and} . We further triangulate the lower wall so that v_{and} has sufficiently high degree, allowing to assign every variable u to an internal edge $f(u)$ of the lower wall incident to v_{and} . See the lower left wall on the right picture. Any triangulation satisfying these prescriptions and computable in time polynomial in the size of ϕ suits our purpose.






Variable gadget. The variable gadget **V**(u) associated with the variable u has four parts.

1. The first part is a triangulated 2-sphere $S(u)$ that consists of two disks $D[u]$ and $D[\neg u]$ sharing a common boundary circle $s(u)$. The circle $s(u)$ contains a distinguished vertex $v(u)$. The disk $D[u]$ (resp. $D[\neg u]$) has a distinguished edge $f[u]$ (resp. $f[\neg u]$) that joins $v(u)$ to its center.




¹³The value two is not important here; what matters is to fix some value that can be used throughout the construction.

2. The second part is a 2-complex $O(u)$ that consists of two “boundary” circles sharing a vertex. The vertex is identified with the vertex $v(u)$ of $S(u)$. One of the circles is identified with $s(u)$. The other circle is decomposed into two arcs: one is a single edge named $b(u)$, the other is a path with two edges which we call $p(u)$. The vertex common to $b(u)$ and $p(u)$, distinct from $v(u)$, is identified with the vertex v_{and} of the conjunction gadget.
3. The third part is a 1-house $B(u)$ intended to block the edge $b(u) \in O(u)$ from being free as long as the conjunction gadget has not been collapsed. The free edge of $B(u)$ is identified with the edge $f(u)$ in the conjunction gadget \mathbf{A} and the edge $b(u) \in O(u)$ is identified with an edge of the lower wall of $B(u)$ that shares the vertex v_{and} with $f(u)$.
4. The fourth part consists of two complexes, $X[u]$ and $X[\neg u]$. Let $\ell \in \{u, \neg u\}$ and refer to the figure on the right. The complex $X[\ell]$ is a 1-house whose free edge is identified with the edge $f[\ell]$ from $D[\ell]$, and whose lower wall contains a path identified with $p(u)$. Hence, $p(u)$ is common to $X[u]$, $X[\neg u]$, and $O(u)$. For every clause c_i containing the literal ℓ , we add in the lower wall a two-edge path $p[\ell, c_i]$ extended by an edge $f[\ell, c_i]$; the path $p[\ell, c_i]$ intersects $p(u)$ in exactly v_{and} (in particular, these paths and edges form a subdivided star, in graph theoretic sense, centered at v_{and}).



Clause gadget. The clause gadget $\mathbf{C}(c)$ associated with the clause $c = \ell_1 \vee \ell_2 \vee \ell_3$ is a 3-house C where:

- the edges f_i of C are identified with the edges $f[\ell_i, c]$ in $X[\ell_i]$;
- the paths p_i of C are identified with the paths $p[\ell_i, c]$ in $X[\ell_i]$;
- the vertex v of C is identified with the vertex v_{and} ; and
- the edge e of C is identified with the edge f_{and} .



Putting it all together. Let ϕ be a 3-CNF formula with variables u_1, u_2, \dots, u_n and clauses c_1, c_2, \dots, c_m . The complex K_ϕ is defined as

$$K_\phi = \mathbf{A} \cup \left(\underbrace{\bigcup_{i=1}^n S(u_i) \cup O(u_i) \cup B(u_i) \cup X[u_i] \cup X[\neg u_i]}_{\mathbf{V}(u_i)} \right) \cup \left(\bigcup_{j=1}^m \mathbf{C}(c_j) \right).$$

4.4 Properties of K_ϕ

It remains to prove that K_ϕ has the properties required by Proposition 6. Item (i) is quite straightforward (although tedious) and it is proved in the full version [14]. Here we provide a sketch only.

Sketch of Proposition 6(i). We first check that the link of every vertex is connected within a 1-house, within a 3-house and within $S(u) \cup O(u)$ for every variable u . Then K_ϕ is obtained by gluing such subcomplexes together along edges which preserves the connectedness of the links. ◀

The proof of items (ii) and (iii) is given in the forthcoming sections. However, first we need to determine the reduced Euler characteristic of K_ϕ (again, we provide a sketch only; see the full version [14] for full proof):

► **Proposition 9.** $\tilde{\chi}(K_\phi)$ equals the number of variables of ϕ .

Sketch. Many of the gadgets are contractible. For example, if we replace a 1-house with its unique free edge, we do not affect the Euler characteristic. After a series of similar reductions we find out that the only contribution to the (reduced) Euler characteristic arises from the sphere $S(u)$, one for each variable. ◀

5 Satisfiability implies collapsibility

In this section we prove Proposition 6(ii), *i.e.* that if ϕ is satisfiable, then there exists a choice of $\tilde{\chi}(K_\phi)$ triangles of K_ϕ whose removal makes the complex collapsible.

Initial steps. Let us fix a satisfying assignment for ϕ . For every variable u , we set $\ell(u)$ to u if u is TRUE in our assignment, and to $\neg u$ otherwise. Next, for every variable u , we remove a triangle from the region $D[\ell(u)]$ of the sphere $S(u)$. Proposition 9 ensures that this removes precisely $\tilde{\chi}(K_\phi)$ triangles, as announced.

Constraint complex. It will be convenient to analyze collapses of K_ϕ locally within a subcomplex, typically a gadget. To do so formally, we use constraint complexes following [29]. Given a simplicial complex K and a subcomplex M of K , we define the *constraint complex* of (K, M) , denoted $\Gamma(K, M)$, as $\Gamma(K, M) := \{\vartheta \in M : \exists \eta \in K \setminus M \text{ s.t. } \vartheta \subset \eta\}$.

► **Lemma 10** ([29, Lemma 4]). *Let K be a simplicial complex and M a subcomplex of K . If M collapses to M' and $\Gamma(K, M) \subseteq M'$ then K collapses to $(K \setminus M) \cup M'$.*

Collapses. We now describe a sequence of collapses enabled by the removal of the triangles. Recall that we started from the complex

$$K_\phi = \mathbf{A} \cup \left(\bigcup_{i=1}^n O(u_i) \cup S(u_i) \cup B(u_i) \cup X[u_i] \cup X[\neg u_i] \right) \cup \left(\bigcup_{j=1}^m \mathbf{C}(c_j) \right)$$

where u_1, u_2, \dots, u_n and c_1, c_2, \dots, c_m are, respectively, the variables and the clauses of ϕ . We then removed a triangle from each $D[\ell(u)]$. The removal of a triangle of $D[\ell(u)]$ allows to collapse that subcomplex to $s(u) \cup f[\ell(u)]$. This frees $f[\ell(u)]$. The complex becomes:

$$K_a = \mathbf{A} \cup \left(\bigcup_{i=1}^n O(u_i) \cup D[\neg \ell(u_i)] \cup B(u_i) \cup X[u_i] \cup X[\neg u_i] \right) \cup \left(\bigcup_{j=1}^m \mathbf{C}(c_j) \right).$$

We can then start to collapse the subcomplexes $X[\ell(u)]$. We proceed one variable at a time. Assume that we are about to proceed with the collapse of $X[\ell(u)]$ and let K denote the current complex. Locally, $X[\ell(u)]$ is a 1-house with free edge $f[\ell(u)]$. Moreover, $\Gamma(K, X[\ell(u)])$ is the tree $T(u)$ formed by the path $p(u)$ and the union of the paths $p[\ell(u), c] \cup f[\ell(u), c]$ for every clause c using the literal $\ell(u)$. Lemma 7 ensures that $X[\ell(u)]$ can be locally collapsed to $T(u)$, and Lemma 10 ensures that K can be globally collapsed to $(K \setminus X[\ell(u)]) \cup T(u)$. We proceed in this way for every complex $X[\ell(u)]$. The complex becomes:

$$K_b = \mathbf{A} \cup \left(\bigcup_{i=1}^n O(u_i) \cup D[\neg \ell(u_i)] \cup B(u_i) \cup X[\neg \ell(u_i)] \right) \cup \left(\bigcup_{j=1}^m \mathbf{C}(c_j) \right).$$

The collapses so far have freed every edge of $f[\ell(u_i), c]$. We now consider every clause c_j in turn. Put $c_j = (\ell_1 \vee \ell_2 \vee \ell_3)$ and let K denote the current complex. The assignment that we

41:12 Shellability is NP-Complete

chose is satisfying, so at least one of ℓ_1, ℓ_2 or ℓ_3 coincides with $\ell(u_i)$ for some i ; let us assume without loss of generality that $\ell_1 = \ell(u_i)$. The edge $f[\ell_1, c]$ is therefore free and Lemma 8 yields that locally, $\mathbf{C}(c_j)$ collapses to the tree $T(c_j) = f_{\text{and}} \cup p[\ell_1, c] \cup p[\ell_2, c] \cup p[\ell_2, c] \cup f[\ell_2, c] \cup f[\ell_3, c]$. Moreover, $\Gamma(K, \mathbf{C}(c_j)) = T(c_j)$ so Lemma 10 ensure that K can be globally collapsed to $(K \setminus \mathbf{C}(c_j)) \cup T(c_j)$. After proceeding in this way for every complex $\mathbf{C}(c_j)$, we get:

$$K_c = \mathbf{A} \cup \left(\bigcup_{i=1}^n O(u_i) \cup D[\neg\ell(u_i)] \cup B(u_i) \cup X[\neg\ell(u_i)] \right) \cup \left(\bigcup_{j=1}^m T(c_j) \right).$$

The collapses so far have freed the edge f_{and} . We can then proceed to collapse \mathbf{A} . Locally, Lemma 7 allows to collapse \mathbf{A} to the tree $T = f(u_1) \cup f(u_2) \cup \dots \cup f(u_n)$. (From this point, we expect the reader to be able to check by her/himself that Lemma 10 allows to perform globally the collapse described locally.) The complex becomes:

$$K_d = \left(\bigcup_{i=1}^n O(u_i) \cup D[\neg\ell(u_i)] \cup B(u_i) \cup X[\neg\ell(u_i)] \right) \cup \left(\bigcup_{j=1}^m T(c_j) \right).$$

The collapses so far have freed every edge $f(u_i)$. Thus, Lemma 7 allows to collapse each complex $B(u_i)$ to its edge $b(u_i)$. This frees the edge $b(u_i)$, so the complex $O(u_i)$ can in turn be collapsed to $s(u_i) \cup p(u_i)$. At this point, the complex is:

$$K_e = \left(\bigcup_{i=1}^n s(u_i) \cup p(u_i) \cup D[\neg\ell(u_i)] \cup X[\neg\ell(u_i)] \right) \cup \left(\bigcup_{j=1}^m T(c_j) \right).$$

The collapses so far have freed every edge $s(u_i)$. We can thus collapse each $D[\neg\ell(u_i)]$ to $f[\neg\ell(u_i)]$. This frees every edge $f[\neg\ell(u_i)]$, allowing to collapse every subcomplex $X[\neg\ell(u_i)]$, again by Lemma 7, to the tree formed by the path $p(u_i)$ and the union of the paths $p[\neg\ell(u_i), c] \cup f[\neg\ell(u_i), c]$ for every clause c using the literal $\neg\ell(u_i)$.

At this point, we are left with a 1-dimensional complex. This complex is a tree (more precisely a subdivided star, in graph-theoretic sense, centered in v_{and} and consisting of the paths $p(u_i)$, the paths $p[\ell, c]$ and some of the edges $f[\ell, c]$). As any tree is collapsible, this completes the proof of Proposition 6(ii).

6 Collapsibility implies satisfiability

In this section we prove Proposition 6(iii), *i.e.* we consider some arbitrary subdivision K'_ϕ of K_ϕ , and prove that if K'_ϕ becomes collapsible after removing some $\tilde{\chi}(K_\phi)$ triangles, then ϕ is satisfiable. We thus consider a collapsible subcomplex \widehat{K} of K'_ϕ obtained by removing $\tilde{\chi}(K_\phi)$ triangles from K'_ϕ .

Notations. Throughout this section, we use the following conventions. In general, we use hats (for example \widehat{K}) to denote subcomplexes of K'_ϕ . Given a subcomplex M of K_ϕ , we also write M' for the subcomplex of K'_ϕ that subdivides M .

Variable assignment from triangle removal. We first read our candidate assignment from the triangle removal following the same idea as in Section 5. This relies on two observations:

- The set of triangles removed in \widehat{K} contains exactly one triangle from each sphere $S'(u)$. Indeed, since \widehat{K} is collapsible and 2-dimensional, it cannot contain a 2-dimensional sphere. Hence, every sphere $S'(u)$ had at least one of its triangles removed. By Proposition 9, $\chi(K_\phi) = \chi(K'_\phi)$ equals the number of variables of ϕ , so this accounts for all removed triangles.
- For any variable u , any removed triangle in $S'(u)$ is either in $D'[u]$ or in $D'[-u]$. We give u the TRUE assignment in the former case and the FALSE assignment in the latter case.

The remainder of this section is devoted to prove that this assignment satisfies ϕ . It will again be convenient to denote by $\ell(u)$ the literal corresponding to this assignment, that is, $\ell(u) = u$ if u was assigned TRUE and $\ell(u) = \neg u$ otherwise.

Analyzing the collapse. Let us fix some collapse of \widehat{K} . We argue that our assignment satisfies ϕ by showing that these collapses must essentially follow the logical order of the collapse constructed in Section 5. To analyze the dependencies in the collapse, it is convenient to consider the partial order that it induces on the simplices of \widehat{K} : $\sigma \prec \tau$ if and only if in our collapse, σ is deleted before τ . We also write $\sigma \prec \widehat{M}$ for a subcomplex \widehat{M} of \widehat{K} if σ was removed before removing any simplex of \widehat{M} .

The key observation is the following dependency:

► **Lemma 11.** *There exists an edge \widehat{e} of \mathbf{A}' such that $\widehat{e} \prec D'[\neg\ell(u)]$ for every variable u .*

Proof. We first argue that for every variable u , there exists an edge $\widehat{e}_1(u) \in b'(u) \cup p'(u)$ such that $\widehat{e}_1(u) \prec D'[\neg\ell(u)]$. To see this, remark that the complex $D'[\neg\ell(u)] \cup O'(u)$ is fully contained in \widehat{K} since the triangle removed from $S'(u)$ belongs to $D'[\ell(u)]$. It thus has to be collapsed. Since this complex is a disk, the first elementary collapse in $D'[\neg\ell(u)] \cup O'(u)$ has to involve some edge $\widehat{e}_1(u)$ of its boundary. This boundary is $b'(u) \cup p'(u)$, so it contains no edge of $D'[\neg\ell(u)]$. It follows that $\widehat{e}_1(u) \prec D'[\neg\ell(u)]$.

We next claim that $\widehat{e}_1(u) \in b'(u)$. Indeed, remark that every edge in $p'(u)$ belongs to two triangles of $X'[\neg\ell(u)]$. By Lemma 7, any collapse of $X'[\neg\ell(u)]$ must start by an elementary collapse using an edge of $f'[\neg\ell(u)]$ as a free face. Any edge of $f'[\neg\ell(u)]$ is, however, contained in two triangles of $D'[\neg\ell(u)]$ and thus cannot precede $D'[\neg\ell(u)]$ in \prec . It follows that $\widehat{e}_1(u) \in b'(u)$.

We can now identify \widehat{e} . Observe that $b'(u) \subset B'(u)$. As $B'(u)$ is a 1-house, Lemma 7 ensures that the first edge removed from $B'(u)$ must subdivide $f'(u)$. Hence, there is an edge $\widehat{e}_2(u) \subset f'(u)$ such that $\widehat{e}_2(u) \prec \widehat{e}_1(u)$. Since $f'(u) \subset \mathbf{A}'$, another 1-house, the same reasoning yields an edge $\widehat{e}_3(u)$ in f'_{and} such that $\widehat{e}_3(u) \prec \widehat{e}_2(u)$. Let \widehat{e} denote the first edge removed from \mathbf{A}' among all edges $\widehat{e}_3(u)$. At this point, we have for every variable u $\widehat{e} \prec \widehat{e}_2(u) \prec \widehat{e}_1(u) \prec D'[\neg\ell(u)]$, as announced. ◀

Let \widehat{e} denote the edge of \mathbf{A}' provided by Lemma 11, *i.e.* satisfying $\widehat{e} \prec D'[\neg\ell(u)]$ for every variable u . We can now check that the variable assignment does satisfy the formula:

- Consider a clause $c = (\ell_1 \vee \ell_2 \vee \ell_3)$ in ϕ . The complex $\mathbf{C}'(c)$ is a 3-house, so Lemma 8 restricts its set of free edges to the $f'[\ell_i, c]$. Hence, there is $i \in \{1, 2, 3\}$ and an edge $\widehat{e}_4(c)$ in $f'[\ell_i, c]$ such that $\widehat{e}_4(c) \preceq \mathbf{C}'(c)$. Note that, in particular, $\widehat{e}_4(c) \prec \widehat{e}$ as the edge f_{and} also belongs to $\mathbf{C}'(c)$ and must be freed before collapsing \mathbf{A}' (by Lemma 7).
- The subcomplex $f'[\ell_i, c]$ is contained not only in $\mathbf{C}'(c)$, but also in $X[\ell_i]$ which is a 1-house with free edge $f[\ell_i]$. By Lemma 7, the first elementary collapse of $X[\ell_i]$ uses as free face an edge $\widehat{e}_5(c)$ that subdivides $f'[\ell_i]$. In particular, $\widehat{e}_5(c) \prec f'[\ell_i, c]$ and $\widehat{e}_5(c) \prec \widehat{e}_4(c)$.
- Let u be the variable of the literal ℓ_i , that is, $\ell_i = u$, or $\ell_i = \neg u$; in particular $\ell_i \in \{\ell(u), \neg\ell(u)\}$. From $\widehat{e}_5(c) \prec \widehat{e}_4(c) \prec \widehat{e} \prec D'[\neg\ell(u)]$ it comes that $\widehat{e}_5(c)$ cannot belong to $D'[\neg\ell(u)]$. Yet, $\widehat{e}_5(c)$ belongs to $f'[\ell_i]$. It follows that $\ell_i \neq \neg\ell(u)$ and we must have $\ell_i = \ell(u)$. The definition of $\ell(u)$ thus ensures that our assignment satisfies the clause c .

Since our assignment satisfies every clause, it satisfies ϕ .

References

- 1 S. Arora and B. Barak. *Complexity Theory: A Modern Approach*. Cambridge University Press, Cambridge, 2009. URL: <http://www.cs.princeton.edu/theory/complexity/>.

- 2 D. Attali, O. Devillers, M. Glisse, and S. Lazard. Recognizing shrinkable complexes is NP-complete. *Journal of Computational Geometry*, 7(1):430–443, 2016.
- 3 R. H. Bing. *The geometric topology of 3-manifolds*, volume 40 of *American Mathematical Society Colloquium Publications*. American Mathematical Society, Providence, RI, 1983. doi:10.1090/coll/040.
- 4 A. Björner. Shellable and Cohen-Macaulay partially ordered sets. *Trans. Amer. Math. Soc.*, 260(1):159–183, 1980. doi:10.2307/1999881.
- 5 A. Björner. Topological methods. In *Handbook of combinatorics, Vol. 1, 2*, pages 1819–1872. Elsevier Sci. B. V., Amsterdam, 1995.
- 6 A. Björner and M. Wachs. Bruhat order of Coxeter groups and shellability. *Advances in Mathematics*, 43(1):87–100, 1982. doi:10.1016/0001-8708(82)90029-9.
- 7 A. Björner and M. Wachs. On lexicographically shellable posets. *Trans. Amer. Math. Soc.*, 277(1):323–341, 1983. doi:10.2307/1999359.
- 8 A. Björner and M. L. Wachs. Shellable nonpure complexes and posets. II. *Trans. Amer. Math. Soc.*, 349(10):3945–3975, 1997. doi:10.1090/S0002-9947-97-01838-2.
- 9 H. Bruggesser and P. Mani. Shellable decompositions of cells and spheres. *Mathematica Scandinavica*, 29(2):197–205, 1972.
- 10 G. Danaraj and V. Klee. Shellings of spheres and polytopes. *Duke Mathematical Journal*, 41(2):443–451, 1974.
- 11 G. Danaraj and V. Klee. A representation of 2-dimensional pseudomanifolds and its use in the design of a linear-time shelling algorithm. *Ann. Discrete Math.*, 2:53–63, 1978. Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).
- 12 G. Danaraj and V. Klee. Which spheres are shellable? *Ann. Discrete Math.*, 2:33–52, 1978. Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).
- 13 Ö. Eğecioğlu and T. F. Gonzalez. A computationally intractable problem on simplicial complexes. *Comput. Geom.*, 6(2):85–98, 1996. doi:10.1016/0925-7721(95)00015-1.
- 14 X. Goaoc, P. Paták, Z. Patáková, M. Tancer, and U. Wagner. Shellability is NP-complete, 2018. Preprint, <https://arxiv.org/abs/1711.08436>.
- 15 B. Grünbaum. *Convex polytopes*, volume 221 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2003. Prepared and with a preface by Volker Kaibel, Victor Klee and Günter M. Ziegler. doi:10.1007/978-1-4613-0019-9.
- 16 M. Hachimori. Decompositions of two-dimensional simplicial complexes. *Discrete Math.*, 308(11):2307–2312, 2008.
- 17 M. Joswig and M. E. Pfetsch. Computing optimal Morse matchings. *SIAM Journal on Discrete Mathematics*, 20(1):11–25, 2006.
- 18 V. Kaibel and M. E. Pfetsch. Some algorithmic problems in polytope theory. In *Algebra, geometry, and software systems*, pages 23–47. Springer, Berlin, 2003.
- 19 T. Lewiner, H. Lopes, and G. Tavares. Optimal discrete Morse functions for 2-manifolds. *Comput. Geom.*, 26(3):221–233, 2003. doi:10.1016/S0925-7721(03)00014-2.
- 20 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.
- 21 J. Matoušek. *Using the Borsuk-Ulam theorem*. Universitext. Springer-Verlag, Berlin, 2007.
- 22 P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17(2):179–184, 1970.
- 23 A. Nabutovsky. Einstein structures: Existence versus uniqueness. *Geom. Funct. Anal.*, 5(1):76–91, 1995.
- 24 I. Peeva, V. Reiner, and B. Sturmfels. How to shell a monoid. *Math. Ann.*, 310(2):379–393, 1998.
- 25 J. S. Provan and L. J. Billera. Decompositions of simplicial complexes related to diameters of convex polyhedra. *Math. Oper. Res.*, 5(4):576–594, 1980. doi:10.1287/moor.5.4.576.

- 26 C. P. Rourke and B. J. Sanderson. *Introduction to piecewise-linear topology*. Springer Study Edition. Springer-Verlag, Berlin-New York, 1982. Reprint.
- 27 J. Shareshian. On the shellability of the order complex of the subgroup lattice of a finite group. *Trans. Amer. Math. Soc.*, 353(7):2689–2703, 2001. doi:10.1090/S0002-9947-01-02730-1.
- 28 R. P. Stanley. *Combinatorics and commutative algebra*, volume 41 of *Progress in Mathematics*. Birkhäuser Boston, Inc., Boston, MA, second edition, 1996.
- 29 M. Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016.
- 30 I.A. Volodin, V.E. Kuznetsov, and A.T. Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Usp. Mat. Nauk*, 29(5):71–168, 1974. In Russian. English translation: *Russ. Math. Surv.* 29,5:71–172 (1974).
- 31 M. L. Wachs. Poset topology: Tools and applications. In *Geometric combinatorics*, volume 13 of *IAS/Park City Math. Ser.*, pages 497–615. American Mathematical Soc., 2007.
- 32 J. H. C. Whitehead. Simplicial spaces, nuclei and m-groups. *Proc. London Math. Soc. (2)*, 45(1):243–327, 1939. doi:10.1112/plms/s2-45.1.243.