


Charles University in Prague
Faculty of Mathematics and Physics
Department of Applied Mathematics

Doctoral Thesis



Martin Tancer

Topological and Geometrical Combinatorics

Supervisor

Prof. RNDr. Jiří Matoušek, DrSc.

Study program

Computer Science

Discrete Models and Algorithms

Prague 2011

Acknowledgement

First of all, I would like to thank my advisor, Prof. Jiří Matoušek, for stable support during my PhD studies.

I would also like to thank my collaborators (and friends) who influenced me in many respects including (but not limited to) Jonathan Barmak, Xavier Goaoc, Riste Škrekovski, Kathrin Vorwerk and Uli Wagner.

During my studies, I could spend few months at ETH Zürich and KTH Stockholm. For this options I would like to thank Prof. Emo Welzl and Prof. Anders Björner who cordially invited me to their groups. I also again thank my advisor who negotiated these stays. A significant part of the work on my thesis was done during my stay in Coimbra, Portugal. Thus I would like to thank Prof. Aleš Pultr and Prof. Jorge Picado for an agreement due to which I had a good background at Coimbra University.

I should also emphasize that all members and PhD students at our department provided me an option to work in a very friendly and pleasant atmosphere. Special thanks belong to Jirka Fink for a constant help me with TeX and Unix related questions. The LaTeX template for my thesis is also due to him.

Finally, I would like to thank my family, my friends and my girlfriend Zuzka. I really enjoy the time we spend together. I also gain support and new energy from *You* when I feel desperate.

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources.

I understand that Charles University retains the copyright. However, my wish would be if the contents of the thesis could be publicly available for a non-commercial use. In particular, I would agree with making xerographic/electronic copies of the thesis.¹

In Prague, 11th July 2011

¹When I was working on my thesis, I needed to have an access to the thesis of Gert Wegner. It was very difficult to get an access to it because of copyright issues.

Název práce: Topologická a geometrická kombinatorika
 Autor: Martin Tancer
 Katedra: Katedra aplikované matematiky
 Vedoucí práce: Prof. RNDr. Jiří Matoušek, DrSc.
 E-mail vedoucího: matousek@kam.mff.cuni.cz
 Klíčová slova: Simpliciální komplex, topologická metoda v kombinatorice,
 d -reprezentovatelnost, d -kolabovatelnost, vnořitelnost.

Abstrakt: Cílem práce je prezentovat několik nových výsledků v oblasti topologických metod v kombinatorice. Výsledky lze zařadit do dvou hlavních oblastí.

První oblast pokrývá průsečíkové struktury konvexních množin. V práci je ukázáno, že konečné projektivní roviny nemůžou být průsečíkovými strukturami konvexních množin pevné dimenze, což odpovídá na otázku Alona, Kalaie, Matouška a Meshulama. Dále je ukázáno, že d -kolabovatelnost (nutná podmínka na vlastnosti průsečíkových struktur konvexních množin v dimenzi d) je NP-těžká k rozpoznání pro $d \geq 4$. A také je ukázáno, že d -kolabovatelnost není nutnou podmínkou na vlastnosti průsečíkových vzorů dobrých pokrytí, což vyvrací domněnku G. Wegnera z roku 1975.

Do druhé oblasti spadá několik výsledků ohledně algoritmické obtížnosti rozpoznávání simplicialních komplexů vnořitelných do \mathbb{R}^d . Konkrétněji, je algoritmicky nerozhodnutelné, zda lze k -rozměrný simplicialní komplex po částech lineárně vnořit do \mathbb{R}^d , pokud $d \geq 5$ a $k \in \{d-1, d\}$. Dále je tento problém NP-těžký, pokud $d \geq 4$ a $d \geq k \geq \frac{2d-2}{3}$.

Title: Topological and Geometrical Combinatorics
 Author: Martin Tancer
 Department: Department of Applied Mathematics
 Supervisor: Prof. RNDr. Jiří Matoušek, DrSc.
 Supervisor's e-mail: matousek@kam.mff.cuni.cz
 Keywords: Simplicial complex, topological method in combinatorics
 d -representability, d -collapsibility, embeddability

Abstract: The task of the thesis is to present several new results on topological methods in combinatorics. The results can be split into two main streams.

The first stream regards intersection patterns of convex sets. It is shown in the thesis that finite projective planes cannot be intersection patterns of convex sets of fixed dimension which answers a question of Alon, Kalai, Matoušek and Meshulam. Another result shows that d -collapsibility (a necessary condition on properties of intersection patterns of convex sets in dimension d) is NP-complete for recognition if $d \geq 4$. In addition it is shown that d -collapsibility is not a necessary condition on properties of intersection patterns of good covers, which disproves a conjecture of G. Wegner from 1975.

The second stream considers algorithmic hardness of recognition of simplicial complexes embeddable into \mathbb{R}^d . The following results are proved: It is algorithmically undecidable whether a k -dimensional simplicial complex piecewise-linearly embeds into \mathbb{R}^d for $d \geq 5$ and $k \in \{d-1, d\}$; and this problem is NP-hard if $d \geq 4$ and $d \geq k \geq \frac{2d-2}{3}$.

Contents

| | | |
|----------|--|-----------|
| 1 | Introduction | 9 |
| 2 | Preliminaries on simplicial complexes | 11 |
| 3 | Intersection patterns of convex sets | 13 |
| 3.1 | d -representable complexes | 13 |
| 3.2 | d -collapsible and d -Leray complexes | 14 |
| 3.3 | d -representability of complexes of small dimension | 16 |
| 3.4 | Gaps among the notions | 17 |
| 3.4.1 | The gap between representability and collapsibility | 17 |
| 3.4.2 | The gap between collapsibility and Leray number | 18 |
| 3.5 | Algorithmic perspective | 18 |
| 3.6 | Good covers | 19 |
| 3.6.1 | Nerve theorems | 20 |
| 3.6.2 | Good covers versus collections of convex sets | 20 |
| 4 | Embedding simplicial complexes | 21 |
| 5 | Representability of simplicial complexes of a given dimension | 25 |
| 6 | Non-representability of projective planes | 29 |
| 6.1 | Proof of Theorem 3.7 | 30 |
| 6.2 | Proof of the gap between d -representability and d -collapsibility | 32 |
| 6.3 | A modest generalization | 33 |
| 6.4 | A random almost disjoint set system | 35 |
| 7 | Computational complexity of d-collapsibility | 37 |
| 7.1 | 2-collapsibility | 38 |
| 7.2 | d -collapsible complex with a bad d -collapse | 41 |
| 7.2.1 | Bad complex | 41 |
| 7.2.2 | The complex C | 43 |
| 7.2.3 | The complex $C(\rho)$ | 43 |
| 7.2.4 | The complex $C(\rho; \zeta_1, \dots, \zeta_t)$ | 44 |
| 7.3 | NP-completeness | 46 |
| 7.3.1 | Sketch of the reduction | 46 |
| 7.3.2 | Simplicial gadgets | 47 |
| 7.3.3 | The reduction | 49 |
| 7.4 | Technical properties of d -collapsing | 53 |
| 7.4.1 | d -collapsing faces of dimension strictly less than $d - 1$ | 53 |

| | | |
|----------|---|-----------|
| 7.4.2 | d -collapsing to a subcomplex | 54 |
| 7.4.3 | Gluing distant faces | 56 |
| 7.5 | The complexity of d -representability | 57 |
| 8 | A counterexample to Wegner's conjecture | 59 |
| 8.1 | Planar case | 59 |
| 8.1.1 | Topological representability | 59 |
| 8.1.2 | Non-collapsibility by case analysis | 59 |
| 8.1.3 | Sketch of non-collapsibility | 61 |
| 8.2 | Higher dimensions | 63 |
| 8.3 | Conclusion | 63 |
| 9 | Hardness of embeddability | 65 |
| 9.1 | Preliminaries on PL topology | 65 |
| 9.2 | Undecidability: Proof of Theorem 4.1 | 68 |
| 9.3 | Hardness of embedding 2-dimensional complexes in \mathbb{R}^4 | 69 |
| 9.3.1 | The clause gadget | 69 |
| 9.3.2 | The conflict gadget | 71 |
| 9.3.3 | The reduction | 72 |
| 9.4 | NP-hardness for higher dimensions | 75 |
| 9.4.1 | The clause gadget | 75 |
| 9.4.2 | The conflict gadget | 77 |
| 9.4.3 | The reduction | 82 |
| | Bibliography | 85 |

Chapter 1

Introduction

Combinatorics is a mathematical discipline attempting to enumerate the number of objects with a given property. Usually it deals with discrete objects rather than continuous. Although enumeration is the main question of combinatorics (at least it has given the name to combinatorics “number of combinations”) there are also other important branches which are interconnected with the enumeration due to similarity of appearing objects or similarity of mathematical ideas. Thus, e.g., the questions of whether there is an object with a given property, or whether there is an algorithm constructing/enumerating all the objects with a given property are also an integral part of combinatorics.

Combinatorics strongly interacts with graph theory where the objects in the question are graphs, vertices of graphs, edges of graphs, etc. (Actually, graph theory can be seen as a subbranch of combinatorics.) It also interacts with geometry when the objects in the question are, e.g., points/segments/lines/triangles in a plane/space or various other geometric objects.

Surprisingly, introducing continuous objects can often help to solve discrete combinatorial problems. One of the most famous results of this spirit is Lovász’s breakthrough solution of Kneser’s conjecture [Lov78]. He managed to provide a tight lower bound for chromatic number of so-called Kneser’s graph by showing that a certain topological space is highly connected.

The task of this thesis is to solve several open problems on the boundary of combinatorics, geometry and topology.

Important objects for linking combinatorics and topology are *simplicial complexes*. Roughly speaking, they are topological spaces formed by gluing simplices of various dimensions together (i.e., points, edges, triangles, tetrahedra, ...). Their advantage is that they have a simple combinatorial description. For a reader not familiar with simplicial complexes we point out that the simplicial complexes are discussed in more detail in the following chapter.

Results of the thesis. Results of the thesis focus on two main branches.

First branch regards intersection patterns of convex sets. The main question of this branch is to describe possible intersection patterns of collections of convex sets in the Euclidean space \mathbb{R}^d . There is perhaps no simple description of possible intersection patterns (e.g., because of some algorithmic hardness results). There are, however, some necessary and sufficient conditions that help understanding this area. We describe a formal viewpoint on this study and then also several new results (mostly already published—see the bibliographic remarks below) are presented here. They

include the fact that the gap between so called d -representability and d -collapsibility can be arbitrary large; clarifying the complexity status of several algorithmic questions from this area; and an example showing that Wegner's d -collapsibility does not extend to good covers. The precise statements of the results are mentioned in the introductory chapter to this branch—Chapter 3. The results are then separately proved in Chapters 5, 6, 7, and 8.

The second branch regards the question of whether a given simplicial complex embeds into an Euclidean space of given dimension. This area was intensively theoretically studied by topologists; however, not from an algorithmic perspective. Thus we focus on this question from an algorithmic point of view and we show some hardness results, or even undecidability results depending on the dimensions of the complex and the space in the question. Chapter 4 is an introductory chapter to this area. Main results are proved in Chapter 9.

Bibliographic remarks. Most of the contents of the thesis has already been published or is being considered for publication. (Some parts of the text were, of course, modified in order to obtain a coherent text.) Here we would like to explain the bibliographic status of the results in the thesis.

A significant part of Chapter 3 coincide with the survey article [Tan11b]. Chapters 4 and 9 are based on a common work with my advisor Jiří Matoušek and Uli Wagner from ETH Zürich [MTW11]. Chapter 5 is new.¹ Chapter 6 follows [Tan10b]; however, the contents of the Sections 6.3 and 6.4 is not published anywhere. The contents of Chapter 7 is mainly from [Tan10a] and the contents of Chapter 8 from [Tan11a].

¹However, an article based on the contents of the chapter was recently submitted to a journal.

Chapter 2

Preliminaries on simplicial complexes

Simplicial complexes are central objects for translating combinatorics into topology and vice versa. We guess that the reader is already familiar with simplicial complexes; however, we still prefer to introduce simplicial complexes in full detail. The reader familiar with simplicial complexes can skip this chapter and consult it only if a problem occurs. We present here only fundamental properties of simplicial complexes used throughout the thesis. More advanced properties are introduced exactly before they (first) usage.

We also refer to another sources covering simplicial complexes such as [Hat01, Mat03, Mun84].

We deal with finite *abstract simplicial complexes*, i.e., collections \mathbf{K} of subsets of a finite set X such that if $\alpha \in \mathbf{K}$ and $\beta \subset \alpha$, then $\beta \in \mathbf{K}$. Elements of \mathbf{K} are *faces* of \mathbf{K} . The *dimension* of a face $\alpha \in \mathbf{K}$ is defined as $\dim \alpha = |\alpha| - 1$. The *dimension* of a simplicial complex is the maximum of dimensions of its faces. The set of *vertices* of \mathbf{K} is defined as $V(\mathbf{K}) := \{v \in X : \{v\} \in \mathbf{K}\}$.

In some cases we simplify the notation (if there is no risk of confusing reader). For instance we may write v instead of $\{v\}$ for a face containing a single vertex or ab instead of $\{a, b\}$ for an *edge*.

A *geometric simplicial complex* can be obtained from an abstract simplicial complex \mathbf{K} in such a way that a face $\alpha \in \mathbf{K}$ is replaced by a (nondegenerated!) simplex $s(\alpha)$ of dimension $\dim \alpha$ in an Euclidean space \mathbb{R}^d in the following way (see Figure 2).

- A vertex $v \in V(\mathbf{K})$ is mapped to a point $p(v) \in \mathbb{R}^d$.
- A face $\alpha \in \mathbf{K}$ is mapped to $s(\alpha) = \text{conv}\{p(v); v \in V(\mathbf{K})\}$.
- For $\alpha, \beta \in \mathbf{K}$ we have $s(\alpha \cap \beta) = s(\alpha) \cap s(\beta)$.

If we need to distinguish geometric and abstract simplicial complexes we call $s(\alpha)$ a *geometric face* of \mathbf{K} . For many purposes it is not necessary to distinguish abstract and geometric simplicial complexes; and we will not distinguish them if there is no risk of confusing the reader. However, if we write a *simplicial complex* or even just a *complex*, we primarily always mean an *abstract simplicial complex*.

Graphs coincide with 1-dimensional simplicial complexes. If V' is a subset of vertices of \mathbf{K} then the *induced subcomplex* $\mathbf{K}[V']$ is a complex of faces $\alpha \in \mathbf{K}$ such that $\alpha \subseteq V'$. We use the notation $\mathbf{L} \leq \mathbf{K}$ for pointing out that \mathbf{L} is an induced




Figure 2.1: A geometric simplicial complex corresponding to an abstract simplicial complex with faces $\{1\}, \{2\}, \{3\}, \{4\}, \{5\}, \{1, 2\}, \{1, 4\}, \{2, 4\}, \{2, 5\}, \{4, 5\},$ and $\{1, 2, 4\}$.




Figure 2.2: Barycentric subdivision of a complex. For example, the vertex b_{13} denotes the barycenter of the face $13 = \{1, 3\}$ (in geometric setting).

subcomplex of K . An m -skeleton of K is a simplicial complex consisting of faces of K of dimension at most m . We denote it by $K^{(m)}$. The m -dimensional *full simplex*, Δ_m , is a simplicial complex with the vertex set $\{1, \dots, m+1\}$ and all possible faces. We also use the notation $\Delta(S)$ for a full simplex on a set S . The *geometric realization* of a complex K , denoted by $|K|$, is the topological space $\bigcup\{\text{conv } s(\alpha) : \alpha \in K\}$, where $s(\alpha)$ are geometric faces of K . The resulting topological space does not depend (up to a homeomorphism) on the choice of the geometric simplicial complex for a given abstract simplicial complex.

Barycentric subdivision. Given a simplicial complex K , sometimes we need to refine K while the geometric realization of K remains unaffected. A suitable tool for such a task is the barycentric subdivision of K .

From geometric point of view we put a new vertex into the barycenter of every geometric face of K . Then we form a new simplicial complex whose vertices are the barycenters and whose faces are simplices formed in between these barycenters.

It is perhaps more convenient to state the precise definition in abstract setting. Given a simplicial complex K the *barycentric subdivision* of K is a simplicial complex $\text{sd } K$ whose set of vertices is the set $K \setminus \emptyset$ and whose faces are collections $\{\alpha_1, \dots, \alpha_m\}$ of faces of K such that

$$\alpha_1 \supsetneq \alpha_2 \supsetneq \dots \supsetneq \alpha_m \neq \emptyset.$$

The vertices of $\text{sd } K$ play role of barycenters of faces of $K \setminus \emptyset$. The faces of $\text{sd } K$ are the simplices in between of these barycenters. See Figure 2.2.

The complexes K and $\text{sd } K$ have the same geometric realization, i.e., $|K| = |\text{sd } K|$.

Chapter 3

Intersection patterns of convex sets

An important branch of combinatorial geometry regards studying intersection patterns of convex sets. Research in this area was initiated by a theorem of Helly [Hel23] which can be formulated as follows: If C_1, \dots, C_n are convex sets in \mathbb{R}^d , $n \geq d + 1$, and any collection of $d + 1$ sets among C_1, \dots, C_n has a nonempty intersection, then all the sets have a common point. We will focus on results of similar spirit.

3.1 d -representable complexes

First we introduce d -representable simplicial complexes which provide a systematic way for studying intersection patterns of convex sets. Let \mathcal{C} be a collection of some subsets of a given set X . The *nerve* of \mathcal{C} , denoted by $\mathbf{N}(\mathcal{C})$, is a simplicial complex whose vertices are the sets in \mathcal{C} and whose faces are subcollections $\{C_1, \dots, C_k\} \subseteq \mathcal{C}$ such that the intersection $C_1 \cap \dots \cap C_k$ is nonempty. The notion of nerve is designed to record the ‘intersection pattern’ of the sets in \mathcal{C} .

A simplicial complex K is *d -representable* if it is isomorphic to the nerve of a finite collection of convex sets in \mathbb{R}^d . Such a collection of convex sets is called a *d -representation* for K . As it was mentioned above, d -representable complexes exactly record all possible intersection patterns of finite collections of convex sets in \mathbb{R}^d .

Using the notion of d -representability, the Helly theorem can be reformulated as follows: A d -representable simplicial complex does not contain a k -dimensional *simplicial hole* for $k \geq d$, i.e., a complex isomorphic to $\Delta_{k+1}^{(k)}$. (We remark that a geometric representation of k -dimensional simplicial hole is homeomorphic to the k -sphere S^d ; and it is the simplest way, how to obtain the k -sphere as a simplicial complex.)

Example 3.1. Figure 3.1 shows a collection $\mathcal{C} = \{C_1, C_2, C_3, C_4, C_5\}$ of convex sets (on left) and their nerve (on right). In another words the simplicial complex on right is 2-representable and \mathcal{C} is a 2-representation of it.

Let us also remark that it is possible to put various restrictions to the convex sets in the definition of d -representability without affecting the definition. These restrictions can be summarized in the following simple lemma.

Lemma 3.2. *Let K be a simplicial complex. The following conditions are equivalent.*

- (1) K is d -representable;
- (2) K is the nerve of a collection of convex polytopes;




Figure 3.1: A 2-representable complex and its nerve.

- (3) \mathbf{K} is the nerve of a collection of compact convex sets;
- (3') \mathbf{K} is the nerve of a collection of closed convex sets;
- (4) \mathbf{K} is the nerve of a collection of bounded open convex sets.
- (4') \mathbf{K} is the nerve of a collection of open convex sets;

Proof. First we show the implication (1) \Rightarrow (2). Let $\{C_v\}$ be a d -representation of \mathbf{K} (C_v corresponds to a vertex v). For a face α of \mathbf{K} we pick a point p_α belonging to all C_v where $v \in \alpha$. For $v \in V(\mathbf{K})$ let $C'_v = \text{conv}\{p_\alpha : v \in \alpha\}$. Then $\{C'_v\}$ is a representation of \mathbf{K} using convex polytopes. (Note that $C'_v \subseteq C_v$.)

The implications (2) \Rightarrow (3) \Rightarrow (3') are obvious.

It is not hard to see (3) \Rightarrow (4) if the compact convex sets representing \mathbf{K} are blown up a bit.

The implications (4) \Rightarrow (4') \Rightarrow (1) and (3') \Rightarrow (1) are again obvious. \square

3.2 d -collapsible and d -Leray complexes

There are two other important classes of simplicial complexes related to the d -representable ones. Informally, a simplicial complex is d -collapsible if it can be vanished by removing faces of dimension at most $d - 1$ which are contained in a single maximal face; a simplicial complex is d -Leray if its induced subcomplexes do not contain, homologically, d -dimensional holes.

Wegner [Weg75] proved that d -representable simplicial complexes are d -collapsible and also that d -collapsible complexes are d -Leray.

Now we precisely define d -collapsible complexes and then d -Leray complexes.

Let \mathbf{K} be a simplicial complex. Let T be the collection of inclusion-wise maximal faces of \mathbf{K} . A face σ is d -collapsible if there is only one face $\tau = \tau(\sigma) \in T$ containing σ (possibly $\sigma = \tau$), and moreover $\dim \sigma \leq d - 1$. The simplicial complex

$$\mathbf{K}' := \mathbf{K} \setminus \{\eta \in \mathbf{K} : \eta \supseteq \sigma\}$$

is an *elementary d -collapse* of \mathbf{K} . For such a situation, we use the notation $\mathbf{K} \rightarrow \mathbf{K}'$. A simplicial complex is d -collapsible if there is a sequence,

$$\mathbf{K} \rightarrow \mathbf{K}_1 \rightarrow \mathbf{K}_2 \rightarrow \cdots \rightarrow \emptyset,$$




Figure 3.2: A 2-collapsing of a simplicial complex.

of elementary *d*-collapses ending with an empty complex.

Example 3.3. A simplicial complex K consisting of a full tetrahedron, two full triangles and one hollow triangle in Figure 3.2 is 2-collapsible. For a proof there is a 2-collapsing of K drawn on the picture. In every step the faces σ and τ are indicated.

A simplicial K complex is *d-Leray* if the *i*th reduced homology group $\tilde{H}_i(L)$ (over \mathbb{Q}) vanishes for every induced subcomplex $L \leq K$ and every $i \geq d$.

We mention several remarks regarding *d*-collapsible and *d*-Leray complexes.

- The fact that *d*-collapsible complexes are *d*-Leray is simple (for a reader familiar with homology) since *d*-collapsing does not affect homology of dimension *d* or more.

It is a bit less trivial to show that a *d*-representable complex K is *d*-collapsible. The idea is to slide a generic hyperplane (from infinity to minus infinity) over a *d*-representation for K and gradually cut off whatever is on the positive side of the hyperplane. See Figure 3.3 and the text below the picture for a more detailed sketch. The reader is referred to [Weg75] for full details.

The inclusion of *d*-representable complexes in *d*-Leray complexes can be also deduced, without using Wegner’s results, from the nerve theorem (see Theorem 3.9).

- A *d*-dimensional simplicial complex is $(d + 1)$ -collapsible and hence also $(d + 1)$ -Leray. For a complex K the smallest possible ℓ such that K is *d*-Leray is traditionally called the *Leray number* of K .
- Neither *d*-representability, *d*-collapsibility nor the Leray number is an invariant under a homeomorphism: the full simplex Δ_m is 0-representable; however, its barycentric subdivision is not even $(m - 1)$ -Leray, since it contains an $(m - 1)$ -sphere as an induced subcomplex.
- It is not very difficult to see that every induced subcomplex of a *d*-collapsible complex is again *d*-collapsible. If $K[V'] \leq K$ and $K \rightarrow K_1 \rightarrow \dots \rightarrow \emptyset$ is a *d*-collapsing of K , then $K[V'] \rightarrow K_1[V'] \rightarrow \dots \rightarrow \emptyset = \emptyset[V']$ is a *d*-collapsing for $K[V']$, where some steps are possibly trivial, i.e., $K_i[V'] = K_{i+1}[V']$.




Figure 3.3: A schematic sketch of the proof of Wegner's theorem. A generic hyperplane h is slid from infinity to minus infinity until there is a nontrivial intersection of the convex sets on its positive side. In this case it slides to h' and it cuts off $A \cap B \cap C$ (it also cuts off $A \cap C$, but for the moment we consider a maximal collection). From genericity there is a single point $p \in A \cap B \cap C \cap h'$. It can be shown (using Helly's theorem) that there is only at most d sets of the starting collection necessary to obtain p . In this case $\{p\} = A \cap C \cap h'$. Thus we obtain a d -collapse with $\sigma = \{A, C\}$ and $\tau = \{A, B, C\}$. Finally, $A \cap (h')^-, \dots, D \cap (h')^-$ form a d -representation for the resulting collapsed complex thus the procedure can be repeated.

- The Helly theorem easily follows from the fact that d -representable complexes are contained in d -collapsible ones (or d -Leray ones). For we have that a d -dimensional simplicial hole is neither d -collapsible nor d -Leray.

On the other hand these two notions provide (much) stronger limitations to intersection patterns than the Helly theorem. For instance they also exclude (in dimension 2) the boundary of the octahedron (i.e., the simplicial complex with vertices $\{-3, -2, -1, 1, 2, 3\}$ and faces α such that there is no $i \in \{1, 2, 3\}$ with $-i, i \in \alpha$) or a triangulation of a torus.

3.3 d -representability of complexes of small dimension

Every finite simplicial complex is d -representable for d big enough. Let K be a simplicial complex on vertex set $\{1, \dots, n\}$. Let x_1, \dots, x_n be affinely independent points in \mathbb{R}^{n-1} (i.e., they form a simplex). For a nonempty face $\alpha = \{a_1, \dots, a_t\} \in K$ let b_α be the barycentre of the points x_{a_1}, \dots, x_{a_t} . Then for $i \in \{1, \dots, n\}$ we set $C_i := \text{conv}\{b_\alpha : i \in \alpha, \alpha \in K\}$. The reader is welcome to check that sets C_{i_1}, \dots, C_{i_k} intersect if and only if $\{i_1, \dots, i_k\} \in K$. Thus, the nerve of C_1, \dots, C_n is isomorphic to K . See Figure 3.4 for an illustration. (If we really would not care about the dimension, it would be even easier to check the situation where the points b_α are set to be the vertices of a simplex of dimension $|K| - 1$.)




Figure 3.4: Representing a complex.

There is, however, another way how to obtain a representation of a complex depending on the dimension of the complex.

Theorem 3.4 (Wegner [Weg67], Perel'man [Per85]). *Let K be a d -dimensional simplicial complex. Then K is $(2d + 1)$ -representable.*

The references for Theorem 3.4 are due to Eckhoff [Eck93]. (Perel'man rediscovered Wegner's result). Unfortunately, I have not been able to check the proof by Perel'man (in Russian). In Chapter 5 we supply an idea of the proof of the theorem (the proof is different from Wegner's one and it was communicated by Jiří Matoušek).

The value $2d + 1$ in Theorem 3.4 is the least possible; see Chapter 5 for a proof.

3.4 Gaps among the notions

In this section we overview how the notions of d -representable, d -collapsible and d -Leray complexes differ.

3.4.1 The gap between representability and collapsibility

For $d = 0$ all three notions 0-representable, 0-collapsible and 0-Leray coincide and they can be replaced with 'being a simplex'.

For $d = 1$: 1-representable complexes are clique complexes over *interval graphs*; 1-collapsible and 1-Leray complexes are clique complexes over *chordal graphs* (we remark that results in [LB63, Weg75] easily imply these statements).

For $d \geq 2$ there is perhaps no simple characterization of d -representable, d -collapsible and d -Leray complexes. Wegner [Weg75] gave an example of complex, which is 2-collapsible but not 2-representable. Alon, Kalai, Matoušek and Meshulam [AKMM02] asked how big can be the gap between representability and collapsibility. Matoušek and the author [MT09] found d -collapsible complexes that are not $(2d - 2)$ -representable. Later, the author [Tan10b] improved this result by finding 2-collapsible complexes that are not d -representable (for any fixed d). We sketch the first construction, and then we present the second construction in full detail. (Even the weaker construction contains some steps of their own interest.)

Let E be a $(d - 1)$ -dimensional simplicial complex which is not embeddable in \mathbb{R}^{2d-2} . Such a complex always exist, for example the *van Kampen complex* $\Delta_{2d}^{(d-1)}$; see [vK32], or the *Flores complex* [Flo34], which is the join of d copies of a set of three independent points. The first example is the nerve $N(E)$. It is d -collapsible but not $(2d - 2)$ -representable due to the following two propositions [MT09].

Proposition 3.5. *Let K be a simplicial complex such that the nerve $N(K)$ is n -representable. Then K embeds in \mathbb{R}^n , even linearly.*

Proposition 3.6. *Let \mathcal{F} be a family of sets, each of size at most n . Then the nerve $N(\mathcal{F})$ is n -collapsible.*

Another construction (coincidentally) attaining the weaker bound $2d - 2$ is discussed in Chapter 5 (the bound can be derived from Theorem 5.1).

The second (stronger) example regards finite projective planes seen as simplicial complexes. Let (P, \mathcal{L}) be a finite projective plane, where P is the set of its points and \mathcal{L} is the set of its lines. There is a natural simplicial complex \mathbf{P} associated to the projective plane. Its ground set is P and faces are the collections of points lying on a common line.

It is not hard to show that \mathbf{P} is 2-collapsible. Non-representability of \mathbf{P} is summarized in the following theorem [Tan10b].

Theorem 3.7. *For every $d \in \mathbb{N}$ there is a $q_0 = q_0(d)$ such that if a complex \mathbf{P} correspond to a projective plane of order $q \geq q_0$ then \mathbf{P} is not d -representable.*

We remark that the assumption $q \geq q_0$ cannot be left out since every projective plane (P, \mathcal{L}) of order q can be easily represented by convex sets in \mathbb{R}^{2q+1} ; see Theorem 3.4.

Theorem 3.7 is proved in Chapter 6. Some details on finite projective planes are also discussed there. In addition, Theorem 3.7 is also extended for some other *almost disjoint set systems*.

3.4.2 The gap between collapsibility and Leray number

Wegner showed an example of complex which is 2-Leray but not 2-collapsible, namely a triangulation D of the dunce hat. If we consider the multiple join $D \star \cdots \star D$ of d copies of D , we obtain a complex which is $2d$ -Leray but not $(3d - 1)$ -collapsible. See [MT09] for more details.

3.5 Algorithmic perspective

As we consider different criteria for d -representability, it is also natural to ask whether there is an algorithm for recognition d -representable complexes. We denote this algorithmic question as d -REPRESENTABILITY. More precisely, the input of this question is a simplicial complex. The size of the input is the number of faces of the complex. The value d is considered as a fixed integer. The output of the algorithm is the answer whether the complex is d -representable.

We can also ask similar questions for d -collapsible and d -Leray complexes as relaxations of the previous problem. Thus we have algorithmic problems d -COLLAPSIBILITY and d -LERAYNUMBER.

Representability. The first mentioned problem d -REPRESENTABILITY is perhaps the most difficult among the three algorithmic questions. It is NP-hard for $d \geq 2$. Reduction can be done in a very similar fashion as a reduction for hardness of recognition intersection graphs of segments [KM89, KM94]. Full details can be found in Section 7.5. It is not known to the author whether d -REPRESENTABILITY belongs

to NP. On the other hand it is not hard to see that there is a PSPACE algorithm for d -REPRESENTABILITY. It is based on solving systems of polynomial inequalities. See [KM94, Theorem 1.1(i)(a)] for a very similar reduction.

Collapsibility. The main result of [Tan10a] claims that d -COLLAPSIBILITY is NP-complete for $d \geq 4$ and it is polynomial time solvable for $d \leq 2$. For $d = 3$, the problem remains open. We prove the results in Chapter 7.

Leray number. The last question, d -LERAYNUMBER, is polynomial time solvable. An equivalent characterization of d -Leray complexes is when induced subcomplexes are replaced with *links* of faces (including an empty face). See [KM06, Proposition 3.1] for a proof. The tests on links can be done in polynomial time since it is sufficient to test the homology up to the dimension of the complex.

Greedy collapsibility. The algorithmic results above suggest that it is easier to test/compute the Leray number than collapsibility. However, if we are interested in them because of a hint for representability, computing collapsibility still can be more convenient, since d -collapsibility is closer to d -representability than the Leray number. An example from Section 3.4 is maybe a bit unconvincing; however, there is a more important example. As it is shown in Section 3.6 (and in Chapter 8 in more detail), d -collapsibility can distinguish collections of convex sets and good covers.

A useful tool for computation could be greedy d -collapsibility. We say that a simplicial complex K is *greedily d -collapsible* if it is d -collapsible and any sequence of d -collapses of K ends up in a complex which is still d -collapsible. In another words greedy collapsibility allows us to collapse the faces of K in whatever order without risk of being stuck. Thus, if a complex is greedily d -collapsible, then there is a simple (greedy) algorithm for showing that it is d -collapsible. Not all d -collapsible complexes are greedily d -collapsible. Complexes which are not greedily d -collapsible for $d \geq 3$ are constructed in Chapter 7. However, none of these complexes is d -representable. In summary there is a hope for obtaining a simple algorithm for showing that a complex is either d -collapsible or it is not d -representable if the answer to the following question is true.

Problem 3.8. *Is it true that every d -representable simplicial complex is greedily d -collapsible?*

3.6 Good covers

A *good cover* in \mathbb{R}^d is a collection of open sets in \mathbb{R}^d such that the intersection of any subcollection is either empty or a contractible (in particular, the sets in the collection are contractible).¹ We consider only finite good covers. A simplicial complex is *topologically d -representable* if it is isomorphic to the nerve of a (finite) good cover in \mathbb{R}^d . We should emphasize that (for our purposes) a good cover need not cover whole \mathbb{R}^d .

Topologically d -representable complexes generalize d -representable complexes since every collection of convex sets is a good cover.

¹The definition of a good cover is not fully standard in the literature. For example, it may be assumed that sets in the collection are closed instead of open, or that the intersections are homeomorphic to (open) balls instead of contractible. These differences are not essential for the most of the purposes mentioned here, because all these options satisfy the assumptions of a nerve theorem (see the text below).

3.6.1 Nerve theorems

Suppose we are given a collection \mathcal{F} of subsets of \mathbb{R}^d . If the sets are “sufficiently nice” and also all their intersections are sufficiently nice then the nerve of the collection, $\mathbf{N}(\mathcal{F})$, is homotopy equivalent to the union of the sets in the collection, $\bigcup \mathcal{F}$. For a weaker assumption on “sufficiently nice”, it is possible to derive not necessarily homotopy equivalence, but at least equivalence on homology (up to some level). Such results are known as *homotopic/homological* nerve theorems.

We mention here a one of possible versions (suitable for our purposes); see [Hat01, Corollary 4G.3].

Theorem 3.9 (A homotopy nerve theorem). *Let \mathcal{F} be a collection of open contractible sets in a paracompact space X such that $\bigcup \mathcal{F} = X$ and every nonempty intersection of finitely many sets in \mathcal{F} is contractible (or empty). Then the nerve $\mathbf{N}(\mathcal{F})$ and X are homotopy equivalent.*

Corollary 3.10. *The nerve of every good cover is d -Leray.*

3.6.2 Good covers versus collections of convex sets

Good covers have many similar properties as collections of convex sets. Many results on intersection patterns of convex sets can be generalized for good covers. See Section 6 of [Tan11b] for a collection of results of this spirit. An exceptional case is Theorem 3.7 which cannot be generalized for good covers.

On the other hand it is not hard to see that topologically d -representable complexes are strictly more general than d -representable complexes for $d \geq 2$. There is a less trivial example on Figure 8.1 showing that there is a complex which is topologically d -representable but not d -collapsible. Originally, Wegner conjectured that there is no such example. See Chapter 8 for more details.

There is another important difference among collections of convex sets and good covers from a computational point of view. It is algorithmically undecidable whether a simplicial complex is topologically d -representable for $d \geq 5$. (Let us recall that d -REPRESENTABILITY belongs to PSPACE.) This result is obtained in a common work with D. Tonkonog [TT11]; however, the current status is a manuscript in preparation.

Chapter 4

Embedding simplicial complexes

Does a given (finite) simplicial complex K of dimension at most k admit an embedding into \mathbb{R}^d ? We consider the computational complexity of this question, regarding k and d as fixed integers. Besides its intrinsic interest for the theory of computing, an algorithmic view of a classical subject such as embeddability may lead to new questions and also to a better understanding of known results. For example, computation complexity can be seen as a concrete “measuring rod” that allows one to compare the “relative strength” of various embeddability criteria, respectively of examples showing the necessity of dimension restrictions in the criteria. Moreover, hardness results provide concrete evidence that for a certain range of the parameters (outside the so-called metastable range), no simple structural characterization of embeddability (such as Kuratowski’s forbidden minor criterion for graph planarity) is to be expected.

For algorithmic embeddability problems, we consider *piecewise linear* (PL) embeddings. Let us remark that there are at least two other natural notions of embeddings of simplicial complexes in \mathbb{R}^d : *linear embeddings* (also called *geometric realizations*), which are more restricted than PL embeddings, and *topological embeddings*, which give us more freedom than PL embeddings. We will recall the definitions in Chapter 9; here we quickly illustrate the differences with a familiar example: embeddings of 1-dimensional simplicial complexes, a.k.a. simple graphs, into \mathbb{R}^2 . For a topological embedding, the image of each edge can be an arbitrary (curved) Jordan arc, for a PL embedding it has to be a polygonal arc (made of finitely many straight segments), and for a linear embedding, it must be a single straight segment. For this particular case ($k = 1, d = 2$), all three notions happen to give the same class of embeddable complexes, namely, all planar graphs (by Fáry’s theorem). For higher dimensions there are significant differences, though, which we also discuss in Section 9.1.

Here we are interested mainly in embeddability in the topological sense (as opposed to linear embeddability, which is a much more geometric problem and one with a very different flavor), but since it seems problematic to deal with arbitrary topological embeddings effectively, we stick to PL embeddings, which can easily be represented in a computer.

We thus introduce the decision problem $\text{EMBED}_{k \rightarrow d}$, whose input is a simplicial complex K of dimension at most k , and where the output should be YES or NO depending on whether K admits a PL embedding into \mathbb{R}^d .

We assume $k \leq d$, since a k -simplex cannot be embedded in \mathbb{R}^{k-1} . For $d \geq 2k + 1$ the problem becomes trivial, since it is well known that *every* finite k -dimensional simplicial complex embeds in \mathbb{R}^{2k+1} , even linearly (this result goes back to Menger).

In all other cases, i.e., $k \leq d \leq 2k$, there are both YES and NO instances; for the NO instances one can use, e.g., examples of k -dimensional complexes not embeddable in \mathbb{R}^{2k} due to Van Kampen [vK32] and Flores [Flo34]. These complexes were already defined in Subsection 3.4.1.

Let us also note that the complexity of this problem is monotone in k by definition, since an algorithm for $\text{EMBED}_{k \rightarrow d}$ also solves $\text{EMBED}_{k' \rightarrow d}$ for all $k' \leq k$.

Tractable cases. It is well known that $\text{EMBED}_{1 \rightarrow 2}$ (graph planarity) is linear-time solvable [HT74]. Based on planarity algorithms and on a characterization of complexes embeddable in \mathbb{R}^2 due to Halin and Jung [HJ64], it is not hard to come up with a polynomial-time decision algorithm for $\text{EMBED}_{2 \rightarrow 2}$. Outline of such an algorithm is given in Appendix A of [MTW11].

There are many problems in computational topology that are easy for low dimensions (say up to dimension 2 or 3) and become intractable from some dimension on (say 4 or 5); we will mention some of them later. For the embeddability problem, the situation is subtler, since there are tractable cases in arbitrarily high dimensions, namely, $\text{EMBED}_{k \rightarrow 2k}$ for every $k \geq 3$.

The algorithm is based on ideas of Van Kampen [vK32], which were made precise by Shapiro [Sha57] and independently by Wu [Wu65]. A treatment in an algorithmic context, and a self-contained elementary presentation of the algorithm (but not a proof of correctness) can be found in Appendix D of [MTW11].

Hardness. According to a celebrated result of Novikov ([VKF74]; also see, e.g., [Nab95] for an exposition), the following problem is algorithmically unsolvable: Given a d -dimensional simplicial complex, $d \geq 5$, decide whether it is homeomorphic to S^d , the d -dimensional sphere. By a simple reduction we obtain the following result:

Theorem 4.1. $\text{EMBED}_{(d-1) \rightarrow d}$ (and hence also $\text{EMBED}_{d \rightarrow d}$) is algorithmically undecidable for every $d \geq 5$.

This has an interesting consequence, which in some sense strengthens results of Brehm and Sarkaria [BS92]:

Corollary 4.2. For every computable (recursive) function $f: \mathbb{N} \rightarrow \mathbb{N}$ and for every $d \geq 5$ there exist n and a finite $(d-1)$ -dimensional simplicial complex K with n simplices that PL-embeds in \mathbb{R}^d but such that no subdivision of K with at most $f(n)$ simplices embeds linearly in \mathbb{R}^d .

Our main result is hardness for cases where $d \geq 4$ and k is larger than roughly $\frac{2}{3}d$.

Theorem 4.3. $\text{EMBED}_{k \rightarrow d}$ is NP-hard for every pair (k, d) with $d \geq 4$ and $d \geq k \geq \frac{2d-2}{3}$.

We prove a special case of this theorem, NP-hardness of $\text{EMBED}_{2 \rightarrow 4}$, in Section 9.3; the proof is somewhat more intuitive than for the general case and it contains most of the ideas. All the remaining cases are proved in Section 9.4.

Let us briefly mention where the dimension restriction $k \geq (2d-2)/3$ comes from. There is a certain necessary condition for embeddability of a simplicial complex into \mathbb{R}^d , called the *deleted product obstruction*. A celebrated theorem of Haefliger and Weber, which is a far-reaching generalization of the ideas of Van Kampen mentioned above, asserts that this condition is also *sufficient* provided that $k \leq \frac{2}{3}d - 1$ (these k are said to lie in the *metastable range*). The condition on k in Theorem 4.3 is exactly

| $k =$ | $d =$ | | | | | | | | | | | | |
|-------|-------|---|-----|-----|-----|-----|-----|-----|-----|-----|----|----|----|
| | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 | 13 | 14 |
| 1 | P | + | + | + | + | + | + | + | + | + | + | + | + |
| 2 | P | ? | NPh | + | + | + | + | + | + | + | + | + | + |
| 3 | | ? | NPh | NPh | P | + | + | + | + | + | + | + | + |
| 4 | | | NPh | UND | NPh | NPh | P | + | + | + | + | + | + |
| 5 | | | | UND | UND | NPh | NPh | ? | P | + | + | + | + |
| 6 | | | | | UND | UND | NPh | NPh | NPh | ? | P | + | + |
| 7 | | | | | | UND | UND | NPh | NPh | NPh | ? | ? | P |

Table 4.1: The complexity of $\text{EMBED}_{k \rightarrow d}$ (P = polynomial-time solvable, UND = algorithmically undecidable, NPh = NP-hard, + = always embeddable, ? = no result known).

that k must be outside of the metastable range (we refer to Appendix B of [MTW11] for a brief discussion and references).

There are examples showing that the restriction to the metastable range in the Haefliger–Weber theorem is indeed necessary, in the sense that whenever $d \geq 3$ and $d \geq k > (2d - 3)/3$, there are k -dimensional complexes that cannot be embedded into \mathbb{R}^d but the deleted product obstruction fails to detect this. We use constructions of this kind, namely, examples due to Segal and Spieź [SS92], Freedman, Krushkal, and Teichner [FKT94], and Segal, Skopenkov, and Spieź [SSS98], as the main ingredient in our proof of Theorem 4.3.

Discussion. The current complexity status of $\text{EMBED}_{k \rightarrow d}$ is summarized in Table 4.1. The most interesting currently open cases are perhaps $(k, d) = (2, 3)$ and $(3, 3)$. These are outside the metastable range, and it took the longest to find an example showing that they are not characterized by the deleted product obstruction; see [GS06]. That example does not seem to lend itself easily to a hardness reduction, though.

A variation on the proof of our undecidability result (Theorem 4.1) shows that both $\text{EMBED}_{2 \rightarrow 3}$ and $\text{EMBED}_{3 \rightarrow 3}$ are at least as hard as the problem of recognizing the 3-sphere (that is, given a simplicial complex, decide whether it is homeomorphic to S^3). The latter problem is in NP [Iva08, Sch04], but no hardness result seems to be known.

For the remaining questionmarks in the table (with $d \geq 9$), which all lie in the metastable range, it seems that existing tools of algebraic topology, such as Postnikov towers and/or suitable spectral sequences, could lead at least to decision algorithms, or even to polynomial-time algorithms in some cases. Here the methods of “constructive algebraic topology” mentioned below, which imply, e.g., the computability of higher homotopy groups, should be relevant. However, as is discussed, e.g., in [RRS06], computability issues in this area are often subtle, even for questions considered well understood in classical algebraic topology. Current work of Čadek, Krčál, Matoušek, Sergeraert, Vokřínek and Wagner [ČKM⁺11] focuses on this case and offer some new tools that could eventually lead to clarifying the complexity status of this case.

The NP-hardness results presented here are probably not the final word on the computational complexity of the corresponding embeddability problems; for example, some or all of these might turn out to be undecidable.

Related work. Among the most important computational problems in topology are the *homeomorphism problem* for manifolds, and the *equivalence problem* for knots. The first one asks if two given manifolds M_1 and M_2 (given as simplicial complexes, say) are homeomorphic. The second one asks if two given knots, i.e., PL embeddings $f, g: S^1 \rightarrow \mathbb{R}^3$, are equivalent, i.e., if there is a PL homeomorphism $h: \mathbb{R}^3 \rightarrow \mathbb{R}^3$ such that $f = h \circ g$. An important special case of the latter is the *knot triviality* problem: Is a given knot equivalent to the trivial knot (i.e., the standard geometric circle placed in \mathbb{R}^3)?

There is a vast amount of literature on computational problems for 3-manifolds and knots. For instance, it is algorithmically decidable whether a given 3-manifold is homeomorphic to S^3 [Rub95, Tho94], or whether a given polygonal knot in \mathbb{R}^3 is trivial [Hak61]. Indeed, both problems have recently shown to lie in NP [Iva08, Sch04], [HLP99]. The knot equivalence problem is also algorithmically decidable [Hak61, Hem79, Mat97], but nothing seems to be known about its complexity status. We refer the reader to the above-mentioned sources and to [AHT06] for further results, background and references.

In higher dimensions, all of these problems are undecidable. Markov [Mar58] showed that the homeomorphism problem for d -manifolds is algorithmically undecidable for every $d \geq 4$. For $d \geq 5$, this was strengthened by Novikov to the undecidability of recognizing S^d (or any other fixed d -manifold), as was mentioned above. Nabutovsky and Weinberger [NW96] showed that for $d \geq 5$, it is algorithmically undecidable whether a given PL embedding $f: S^{d-2} \rightarrow \mathbb{R}^d$ is equivalent to the standard embedding (placing S^{d-2} as the “equator” of the unit sphere S^{d-1} , say). For further undecidability results, see, e.g., [NW99] and the survey by Soare [Soa04].

Another direction of algorithmic research in topology is the computability of *homotopy groups*. While the fundamental group $\pi_1(X)$ is well-known to be uncomputable [Mar58], all higher homotopy groups of a given finite simply connected simplicial (or CW) complex are computable (Brown [Bro57]). There is also a $\#P$ -hardness result of Anick [Ani89] for the computation of higher homotopy, but it involves CW complexes presented in a highly compact manner, and thus it doesn’t seem to have any direct consequences for simplicial complexes. More recently, there appeared several works (Schön [Sch91], Smith [Smi98], and Rubio, Sergeraert, Dousson, and Romero, e.g. [RRS06]) aiming at making methods of algebraic topology, such as spectral sequences, “constructive”; the last of these has also resulted in an impressive software called KENZO.

A different line of research relevant to the embedding problem concerns linkless embeddings of graphs. Most notably, results of Robertson, Seymour, and Thomas [RST95] on linkless embeddings provide an interesting *sufficient* condition for embeddability of a 2-dimensional complex in \mathbb{R}^3 , and they can thus be regarded as one of the known few positive results concerning $\text{EMBED}_{2 \rightarrow 3}$. We refer to Section 6 of [MTW11] for further details.

Chapter 5

Representability of simplicial complexes of a given dimension

In this chapter we first show that every d -dimensional simplicial complex is $(2d + 1)$ -representable (c.f. Theorem 3.4). Then we show that the value $2d + 1$ is the least possible by constructing d -dimensional complexes which are not $2d$ -representable.

Sketch of a proof of Theorem 3.4. Let K be a d -representable complex with n vertices.

A k -neighborly polytope is a convex polytope such that every k vertices form a face of the polytope. It is well known that there are $2k$ -dimensional k -neighborly polytopes with arbitrary number of vertices for every $k \geq 1$. For instance *cyclic polytopes* satisfy this property. (See, e.g., [Mat02] for a background on convex polytopes including cyclic polytopes.)

Let Q be a $(2d + 2)$ -dimensional $(d + 1)$ -neighborly polytope with n vertices. Let Q^* be a polytope dual to Q . It has n facets and any $d + 1$ of its facets share a face of the polytope. Finally, we consider the *Schlegel diagram* of Q^* . The Schlegel diagram of an m -dimensional convex polytope is a projection of the polytope to $(m - 1)$ -space through a point beyond one of its facets (the point is very close to the facet). See Figure 5.1 for the Schlegel diagram of a cube. In particular the facets of Q^* project to convex sets C_1, \dots, C_n in \mathbb{R}^{2d+1} such that each $d + 1$ of them share the projection of a face of Q^* (on their boundary). Thus if we look at the nerve N of C_1, \dots, C_n , then it contains full d -skeleton of a simplex with n vertices. Therefore, without loss of generality, we can assume that K is a subcomplex of N . Let $\vartheta = \{C_{i_1}, \dots, C_{i_j}\}$ be a face of N which does not belong to K . The sets C_{i_1}, \dots, C_{i_j} intersect on their boundaries and it is possible to remove their intersection by removing a small neighborhood of $C_{i_1} \cap \dots \cap C_{i_j}$ in each of the sets while keeping the sets convex. Hence only ϑ and the superfaces of ϑ disappear from the nerve during this procedure. After repeating the procedure we obtain a collection of convex sets with the nerve K . \square

Theorem 5.1. *The barycentric subdivision $\text{sd } \Delta_{2d+2}^{(d)}$ of d -skeleton of the full $(2d + 2)$ -simplex is not $2d$ -representable.*

Theorem 5.1 can be easily extended to the barycentric subdivision of any d -dimensional simplicial complex K such that the Van Kampen obstruction of K is not zero.

For the proof we will need two auxiliary results.




Figure 5.1: Schlegel diagram of a cube. The reference point is rather more far away in order to make the picture more lucid.

Theorem 5.2 (Van Kampen - Flores theorem; see, e.g., [Mat03, Theorem 5.1.1]). *Let $\mathbf{K} = \Delta_{2d+2}^{(d)}$. Then for any continuous map $f: |\mathbf{K}| \rightarrow \mathbb{R}^{2d}$ there are two disjoint d -dimensional simplices γ and δ of \mathbf{K} such that their images $f(|\gamma|)$ and $f(|\delta|)$ intersect.*

We remark that the conclusion of the theorem remains true if \mathbf{K} is replaced with any d -dimensional complex with non-zero Van Kampen obstruction. The interested reader is referred to [MTW11, Appendix D] for an elementary exposition of Van Kampen obstruction or to [Mel09] for a survey on it.

Let α and β be faces of a simplicial complex \mathbf{K} . We say that α and β are *remote* if there is no edge $ab \in \mathbf{K}$ with $a \in \alpha, b \in \beta$.

Lemma 5.3. *Let \mathcal{K} be a collection of convex sets in \mathbb{R}^m and let $\mathbf{K} := \mathbf{N}(\mathcal{K})$ be the nerve of \mathcal{K} . Then there is a linear map $g: |\text{sd } \mathbf{K}| \rightarrow \mathbb{R}^m$ such that $g(|\text{sd } \alpha|) \cap g(|\text{sd } \beta|) = \emptyset$ for any remote $\alpha, \beta \in \mathbf{K}$.*

Proof. First we specify g on the vertices of $\text{sd } \mathbf{K}$ then we extend it linearly to the whole $\text{sd } \mathbf{K}$. See Figure 5.2

A vertex of $\text{sd } \mathbf{K}$ is a simplex of \mathbf{K} , i.e., a subcollection \mathcal{K}' of \mathcal{K} with a nonempty intersection. Let us pick a point $p(\mathcal{K}')$ inside $\cap \mathcal{K}'$. We set $g(\mathcal{K}') := p(\mathcal{K}')$ for $\mathcal{K}' \in \mathbf{K}$. As we already mentioned, we extend g linearly to $\text{sd } \mathbf{K}$.

If $\alpha = \mathcal{K}' \in \mathbf{K}$, then $g(|\text{sd } \alpha|) \subseteq \cup \mathcal{K}'$. Thus $g(|\text{sd } \alpha|) \cap g(|\text{sd } \beta|) = \emptyset$ for remote $\alpha, \beta \in \mathbf{K}$. □

Proof of Theorem 5.1. Let $\mathbf{K} = \text{sd } \Delta_{2d+2}^{(d)}$. For contradiction we assume that \mathbf{K} is $2d$ -representable. Let \mathcal{K} be the $2d$ -representation of it. (Without loss of generality $\mathcal{K} = \mathbf{N}(\mathbf{K})$.) According to Lemma 5.3 there is a map $g: |\text{sd } \mathbf{K}| \rightarrow \mathbb{R}^{2d}$ such that $g(|\text{sd } \alpha|) \cap g(|\text{sd } \beta|) = \emptyset$ for any remote $\alpha, \beta \in \mathbf{K}$.

Since $\text{sd } \mathbf{K} = \text{sd } \text{sd } \Delta_{2d+2}^{(d)}$, we have $|\Delta_{2d+2}^{(d)}| = |\mathbf{K}| = |\text{sd } \mathbf{K}|$, and thus we can also apply g to simplices of $\Delta_{2d+2}^{(d)}$.

Let γ and δ be disjoint simplices of $\Delta_{2d+2}^{(d)}$. Let α be a simplex of $\text{sd } \gamma$ and β a simplex of $\text{sd } \delta$. Then α and β are remote in \mathbf{K} . Thus $g(|\text{sd } \alpha|) \cap g(|\text{sd } \beta|) = \emptyset$.




Figure 5.2: Mapping $\text{sd } K$ into K . The notation is simplified. For instance 12 stands for $\{1, 2\}$, p_{123} stands for $p(\{1, 2, 3\})$, etc.

Consequently, $g(|\gamma|) \cap g(|\delta|) = \emptyset$ for any choice of γ and δ . However, this contradicts the Van Kampen-Flores theorem.

□

Chapter 6

Non-representability of projective planes

Let us recall that the main purpose of this chapter is to prove the following theorem.

Theorem 3.7. *For every $d \in \mathbb{N}$ there is a $q_0 = q_0(d)$ such that if a complex \mathbf{P} correspond to a projective plane of order $q \geq q_0$ then \mathbf{P} is not d -representable.*

First we recall some basis facts about finite projective planes.

Finite projective planes. A *finite projective plane of order $q \geq 2$* is a pair (P, \mathcal{L}) where P is a finite set of *points*, and $\mathcal{L} \subseteq 2^P$ is a set of subsets of P (called *lines*) such that (i) every two points are contained in a unique line, (ii) every two lines intersect in a unique point, and (iii) every line contains $q + 1$ points. It follows that every point is contained in $q + 1$ lines and $|P| = |\mathcal{L}| = q^2 + q + 1$, see e.g. [MN98] for more details.

It is well known that a projective plane of order q exists whenever q is a power of a prime. We remark that it is a well known open problem to decide whether there are projective planes of other orders.

It is also known that a finite projective plane cannot be represented in \mathbb{R}^d so that the points of the projective plane are points in \mathbb{R}^d and the lines of the projective plane are the inclusionwise maximal collections of points lying on a common Euclidean line. This fact follows for example from Sylvester-Gallai theorem: if p_1, \dots, p_n are points in the plane not all of them lying on a common line, then there is a line in the plane which intersects exactly two of these points (see [Gal44] for original solution and [Kel86] for an elegant proof). Our task is to obtain a similar result where the Euclidean lines are replaced by convex sets.

Representability of set systems. Let (X, \mathcal{B}) be a set system where X is a finite set and \mathcal{B} is a set of some subsets of X . The elements of \mathcal{B} will be called *blocks*. (The motivation for name ‘blocks’ comes from balanced incomplete block designs (BIBDs) that will be defined later.) We say that (X, \mathcal{B}) is *representable by convex sets* in \mathbb{R}^d if there are convex sets $C_B \subset \mathbb{R}^d$ for every block $B \in \mathcal{B}$ such that for any $B_1, \dots, B_k \in \mathcal{B}$ the convex sets C_{B_1}, \dots, C_{B_k} intersect if and only if the blocks B_1, \dots, B_k have a common point in X .

We **strongly distinguish** the terms d -representable simplicial complex and a set system representable by convex sets in \mathbb{R}^d ; they have a different meaning. In fact, they are dual in a certain sense. A set system (X, \mathcal{B}) is representable by convex sets in \mathbb{R}^d if and only if the nerve of \mathcal{B} is d -representable.¹

¹Ultimately we want to translate our results on representability of set systems (X, \mathcal{B}) by convex




Figure 6.1: Positive-fraction selection theorem: every triangle formed by the sets Z_i contains a .

Our main tools are the positive-fraction selection lemma and the fact that the projective planes (considered as bipartite graphs) are expanders.

We also have the following consequence of Theorem 3.7 which answers the question of Alon et al. [AKMM02] (as announced above) and which is proved in Section 6.2.

Corollary 6.1. *Let $d > 1$ be an integer and let $q_0 = q_0(d)$ be the integer from Theorem 3.7. Let (P, \mathcal{L}) be the projective plane of order $q \geq q_0$. Let K_q be a simplicial complex whose vertices are points in P and whose faces are subsets of lines in \mathcal{L} . Then K_q is 2-collapsible and is not d -representable.*

6.1 Proof of Theorem 3.7

In this section we prove Theorem 3.7. We need few preliminaries. Assume that (Z_1, \dots, Z_k) is a k -tuple of sets. By a *transversal* of this k -tuple we mean any set $T = \{t_1, \dots, t_k\}$ such that $t_i \in Z_i$ for every $i \in [k]$. We need the following result due to Pach [Pac98]; see also [Mat02, Theorem 9.5.1]. See Figure 6.1.

Theorem 6.2 (Positive-fraction selection theorem; a special case). *For every natural number d , there exists $c = c(d) > 0$ with the following property. Let $X \subset \mathbb{R}^d$ be a finite set of points in general position (i.e., there are no $d + 1$ points lying in a common hyperplane). Then there is a point $a \in \mathbb{R}^d$ and disjoint subsets Z_1, \dots, Z_{d+1} , with $|Z_i| \geq c|X|$ such that the convex hull of every transversal of (Z_1, \dots, Z_{d+1}) contains a .*

Later we refer to the constant $c(d)$ from the theorem as to *Pach's constant*.

We remark that the proof of Theorem 6.2 uses several involved tools such as weak hypergraph regularity lemma or same-type lemma (therefore we do not reproduce any details of the proof here). We should also remark that this is only a special case of Pach's theorem (but general enough); Pach moreover assumes that $Z_i \subseteq X_i$, where $X_1 \cup \dots \cup X_{d+1}$ is a partition of X , and in this setting $|Z_i| \geq c|X_i|$.

We also need the following expansion property of the projective plane [Alo85, Theorem 2.1], [Alo86].

sets into d -representability of the nerve of \mathcal{B} . We could work only with d -representable complexes from the beginning. However, this would lead to more complicated and less lucid statements of results. The disadvantage is that we have to introduce new (dual) terminology. See Table 6.1 for a dictionary.




Figure 6.2: Almost all lines intersect all of P_1, \dots, P_{d+1} .

Theorem 6.3. *Let (P, \mathcal{L}) be a projective plane of order q . Let $A \subseteq P$. Then $|\{\ell \in \mathcal{L} : \ell \cap A = \emptyset\}| \leq n^{3/2}/|A|$, where $n = q^2 + q + 1$.*

Alon, Haussler and Welzl [AHW87] used this expansion property in the context of range searching problems. They showed that the points of a projective plane (of high enough order) cannot be partitioned into a small number of sets P_1, \dots, P_m so that for every projective line ℓ the set $\bigcup_{\ell \cap P_i \neq \emptyset, P_i} P_i$ contains only a given fraction of all the points. Known results on range searching problems imply that a projective plane of a high order cannot be represented by halfspaces or simplices in \mathbb{R}^d . However, the author is not aware that this approach would imply the result for convex sets.

For completeness, we also reproduce a short proof of Theorem 6.3.

Proof of Theorem 6.3. Let $M = (m_{p\ell})$ be an $n \times n$ matrix with rows indexed by the points of P and columns indexed by the lines of \mathcal{L} . We set $m_{p\ell} := 1$ if $p \in \ell$ and $m_{p\ell} := 0$ otherwise. The matrix MM^T has real nonnegative eigenvalues $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$.

By a theorem of Tanner [Tan84]

$$|N(A)| \geq \frac{(q+1)^2|A|}{((q+1)^2 - \lambda_2)|A|/n + \lambda_2}$$

where $N(A)$ denotes $\{\ell \in \mathcal{L} : \ell \cap A \neq \emptyset\}$, the neighborhood of A .

It is not hard to compute that $\lambda_1 = (q+1)^2$ and $\lambda_2 = \dots = \lambda_n = q$. Consequently,

$$|N(A)| \geq \frac{(q+1)^2|A|}{|A| + q} = n - \frac{q(n - |A|)}{|A| + q} \geq n - \frac{n^{3/2}}{|A|}.$$

□

Proof of Theorem 3.7. For contradiction, we assume that (P, \mathcal{L}) is representable by convex sets in \mathbb{R}^d ; i.e., there are convex sets C_ℓ for $\ell \in \mathcal{L}$ such that $C_{\ell_1}, \dots, C_{\ell_k}$ intersect if and only if ℓ_1, \dots, ℓ_k contain a common point. By standard tricks, we can assume that these sets are open (see Lemma 3.2).

Let $p \in P$. We know that $\bigcap_{p \in \ell} C_\ell$ is nonempty (and open). Let x_p be a point of this intersection. We define $X := \{x_p : p \in P\}$. Because of the openness of the intersections we can assume that X is in general position.

Let $c = c(d) > 0$, $a \in \mathbb{R}^d$, and Z_1, \dots, Z_{d+1} be the (output) data from Theorem 6.2 (when applied to X). We know that $|Z_i| \geq c|X|$. Let us set $P_i := \{p \in P : x_p \in Z_i\}$ and $\mathcal{M}_i := \{\ell \in \mathcal{L} : \ell \cap P_i = \emptyset\}$. By Theorem 6.3 with $A = P_i$ we get

$$|\mathcal{M}_i| \leq \frac{n^{3/2}}{|P_i|} = \frac{n^{3/2}}{|Z_i|} \leq \frac{n^{3/2}}{c|X|} = \frac{n^{3/2}}{cn} \leq \frac{n}{2(d+1)} = \frac{|\mathcal{L}|}{2(d+1)}$$

provided that q (and hence n as well) is sufficiently large (depending on d and c).

Hence the set $\mathcal{L}' := \mathcal{L} \setminus (\mathcal{M}_1 \cup \dots \cup \mathcal{M}_{d+1})$ of lines that intersect each of P_1, \dots, P_{d+1} contains at least half of the lines of \mathcal{L} . Now let $\ell \in \mathcal{L}'$ and let $p_i \in \ell \cap P_i$. Then $(x_{p_1}, \dots, x_{p_{d+1}})$ is a transversal of (Z_1, \dots, Z_{d+1}) . Thus $a \in \text{conv}\{x_{p_1}, \dots, x_{p_{d+1}}\} \subseteq C_\ell$, and so a is contained in at least $\frac{|\mathcal{L}'|}{2}$ of the C_ℓ . This is a contradiction since at most $q+1$ sets among the C_ℓ can have a nonempty intersection. \square

6.2 Proof of the gap between d -representability and d -collapsibility

Proof of Corollary 6.1. The fact that the complex K_q is 2-collapsible is essentially mentioned in [AKMM02, discussion below Problem 15] (without a proof).

All the inclusionwise maximal faces of K_q are of the form $\sigma_\ell = \{p : p \in \ell\}$ for $\ell \in \mathcal{L}$. Two such faces intersect only in a vertex, thus it is possible to 2-collapse these faces gradually to the vertices; the details are given in Lemma 6.4 below. After these collapsings, it is sufficient to remove the vertices (which are already inclusionwise maximal).

It remains to show that K_q is not d -representable. We consider the dual projective plane (\mathcal{L}, \bar{P}) , where $\bar{P} := \{\{\ell \in \mathcal{L} : p \in \ell\} : p \in P\}$. In particular we can identify a point $p \in P$ with a dual line $\{\ell \in \mathcal{L} : p \in \ell\} \in \bar{P}$.

Theorem 3.7 applied for this dual plane (\mathcal{L}, \bar{P}) essentially states that K_q is not d -representable (convex sets in the statement now correspond to the lines in \bar{P} , which we have identified with P —the set of vertices of K_q). \square

Lemma 6.4. *Let Δ be a d -simplex, i.e., a simplicial complex with $[d+1]$ as the set of vertices and with all the possible faces. Then there is a sequence of elementary 2-collapses that starts with Δ and ends with the simplicial complex that contains all the vertices of Δ and no faces of higher dimension.*

Proof. In every elementary 2-collapse we only mention the smaller face σ (here we adopt the notation from the definition of an elementary d -collapse), since a 2-collapse is uniquely determined by σ .

The following sequence of choices of σ provides the required 2-collapsing (the faces are ordered in the lexicographical order, see Figure 6.3).

$$\{1, 2\}, \{1, 3\}, \dots, \{1, d+1\}, \{2, 3\}, \dots, \{2, d+1\}, \{3, 1\}, \dots, \{d, d+1\}.$$

\square




Figure 6.3: Collapsing a simplex to its vertices.

| set system | simplicial complex |
|--|--|
| (X, \mathcal{B}) | $\mathbf{N}(\mathcal{B})$ |
| representable by convex sets in \mathbb{R}^d | d -representable |
| size of X | number of maximal simplices of $\mathbf{N}(\mathcal{B})$ (*) |
| size of \mathcal{B} | number of vertices of $\mathbf{N}(\mathcal{B})$ |
| size of a block $B \in \mathcal{B}$ | number of maximal simplices containing B (*) |
| $\max\{\deg x : x \in X\}$ | $\dim \mathbf{N}(\mathcal{B})$ |

Table 6.1: A dictionary for translating properties of set systems into properties of simplicial complexes. The lines denoted with (*) assume that the system is almost disjoint and it does not contain blocks of size 1. Then the maximal simplices of $\mathbf{N}(\mathcal{B})$ are in one to one correspondence with points of X . The symbol $\deg x$ denotes the number of blocks containing a point $x \in X$.

6.3 A modest generalization

The approach in the proof of Theorem 6.3 can be generalized to other almost disjoint set systems then just projective planes. Let (X, \mathcal{B}) be a set system. We say that (X, \mathcal{B}) is *almost disjoint* if $|B \cap B'| \leq 1$ for any two blocks $B, B' \in \mathcal{B}$. (It is *k-almost disjoint* if $|B \cap B'| \leq k$.)

From a set system to a simplicial complex. Here we summary how to get a simplicial complex from a set system, which was briefly mentioned in previous sections. Let (X, \mathcal{B}) be a set system. We consider the simplicial complex $\mathbf{N}(\mathcal{B})$, i.e., the nerve of \mathcal{B} . We recall that in this case the blocks in \mathcal{B} are the vertices of $\mathbf{N}(\mathcal{B})$ and faces of $\mathbf{N}(\mathcal{B})$ are collections of blocks having a point $x \in X$ in common.

Representability of (X, \mathcal{B}) by convex sets in \mathbb{R}^d is equivalent to d -representability of $\mathbf{N}(\mathcal{B})$. If (X, \mathcal{B}) is almost disjoint then $\mathbf{N}(\mathcal{B})$ is 2-collapsible using the same approach as in Section 6.2. More generally, if (X, \mathcal{B}) is k -almost disjoint then $\mathbf{N}(\mathcal{B})$ is $(k + 1)$ -collapsible; however, we focus only on almost disjoint set systems. See Table 6.1.

A generalization. Let us now consider d as a fixed integer. We have shown that there is $q_0 = q_0(d)$ such that a projective plane (P, \mathcal{L}) of order q which is at least q_0 is not representable by convex sets in \mathbb{R}^d . However, the value $q_0(d)$ is huge when compared with d (in our approach it depends on Pach’s constant). In particular, it means the size of blocks in $(X, \mathcal{B}) := (P, \mathcal{L})$ is huge and also each point x of X belongs to a huge number of blocks (in another words degree of x is large). It is natural to ask whether the approach can apply for smaller values of these parameters.

For the size of blocks we offer the following affirmative answer.

Theorem 6.5. *There is an almost disjoint set system (X, \mathcal{B}) with blocks of size $d + 1$ which is not representable by convex sets in \mathbb{R}^d .*

The question on degree we pose as a problem.

Problem 6.6. *What is the smallest possible maximum degree of vertices of an almost disjoint set system which is not representable by convex sets in \mathbb{R}^d ?*

It is perhaps even more natural to ask this in the terms of simplicial complexes; see Table 6.1.

Problem 6.7. *What is the smallest dimension of a 2-collapsible simplicial complex which is not d -representable?*

Note that the above mentioned problems need not be equivalent; however, a small value in Problem 6.6 can be translated into a small value in 6.7.

Proof of Theorem 6.5 is probabilistic and is given in Section 6.4. The remainder of this section is devoted to a discussion on explicit constructions of almost disjoint set systems.

BIBDs. A *balanced incomplete block design* with parameters v, b, r, k and λ is a set system (X, \mathcal{B}) satisfying the following conditions.

- $|X| = v$;
- $|\mathcal{B}| = b$;
- every point $x \in X$ is contained in exactly r blocks $B \in \mathcal{B}$;
- $|B| = k$ for every $B \in \mathcal{B}$; and
- every two distinct points $x, y \in X$ are together contained in exactly λ blocks.

The parameters have to satisfy the relations $bk = vr$ and $\lambda(v-1) = r(k-1)$; thus b and r are uniquely determined by v, k and λ . We abbreviate such a block design as (v, k, λ) -BIBD. In addition $(v, k, 1)$ -BIBDs are almost disjoint.

It is not fully characterized for which values of parameters do (v, k, λ) exist. On the other hand if k and λ are fixed, then for every v_0 there is $v \geq v_0$ such that a (v, k, λ) -BIBD exist [Wil75] (it is even sufficient if v is large enough and satisfy certain number theoretic conditions). There are several known explicit constructions of BIBDs; see, e.g., [VM04] and the references therein (for case $\lambda = 1$).

Advantages and disadvantages of using BIBDs for non-representability. An advantage of BIBDs is that the eigenvalues of MM^T can be very easily computed where M is the matrix of incidence of points of X and blocks of \mathcal{B} . Then the theorem of Tanner [Tan84] can be used in a very similar way as in the proof of Theorem 6.3 obtaining a generalization of Theorem 6.3 for BIBDs. Then we can use this generalization similarly as in proof of Theorem 3.7.

On the downside, using our approach, we obtain a reasonable bound for non-representability of $(v, k, 1)$ -BIBD only if the number of blocks b is proportional to v (i.e., v is not much larger than k^2).

This lead us to a following problem.

Problem 6.8. *Is there an explicit construction of a set system in Theorem 6.5? Is there a construction coming from a BIBD?*

6.4 A random almost disjoint set system

Let d be a fixed integer. Depending on d we chose a large enough integer $v = v(d)$ and we set $b = \lceil C(d) \cdot v \ln v \rceil$ where $C = C(d)$ is another large enough integer. We specify values of v and C later.

We set X to be a v -element set. Next we set \mathcal{B}' to be a multiset $\{B_i\}_{i=1}^b$ where B_i is a uniformly chosen random $(d+1)$ -element subset of X (the choices for different i are independent).

The target is to exclude some of the blocks of \mathcal{B}' obtaining a set \mathcal{B} such that (X, \mathcal{B}) is almost disjoint but not d -representable with a positive probability. Then we are done since the size of blocks is $d+1$.

Let $c = c(d)$ be the Pach's constant. Let $\mathcal{A} = \{A_1, \dots, A_{d+1}\}$ be a collection of disjoint subsets of X of size at least cv . We call such \mathcal{A} a c -partition. Let B be a set of size $d+1$. We say that B hits \mathcal{A} if $B \cap A_i \neq \emptyset$ for every $i \in [d+1]$. The probability that a block $B \in \mathcal{B}'$ hits \mathcal{A} is at least $p = p(d) := (d+1)! \cdot c(d)$.

The proof of Theorem 6.5 can be deduced from the following three lemmas. The lemmas are proved at the end of this section.

Lemma 6.9. *Let us assume that $v = v(d)$ and $C = C(d)$ are large enough. Then the probability that for every c -partition \mathcal{A} there are at least $bp/2$ blocks of \mathcal{B}' hitting \mathcal{A} is at least $\frac{9}{10}$.*

Lemma 6.10. *Let us assume that $v = v(d)$ and $C = C(d)$ are large enough. Let M denotes $\max\{\deg(x) : x \in X\}$ where $\deg(x)$ is the number of blocks of \mathcal{B}' containing x . Then the probability of the event $\{M < bp/4\}$ is at least $\frac{9}{10}$.*

Lemma 6.11. *Let us assume that $v = v(d)$ and $C = C(d)$ are large enough. We say that a block $B_i \in \mathcal{B}'$ is conflict if it has at least two common elements with some other block $B_j \in \mathcal{B}'$. Then the probability that the number of conflict blocks is less than $bp/4$ is at least $\frac{9}{10}$.*

Proof of Theorem 6.5. The system \mathcal{B}' satisfy with probability at least $\frac{7}{10}$ the events of Lemmas 6.9, 6.10, and 6.11 simultaneously. Thus there has to exist a system \mathcal{B}'' for which all the events are valid. Let \mathcal{B} be the subsystem of \mathcal{B}'' of all blocks that are not conflict. In particular, \mathcal{B} is just a set (and no multiset).

It is obvious that (X, \mathcal{B}) is almost disjoint and the size of blocks equal $d+1$. We show that (X, \mathcal{B}) is not representable by convex sets in \mathbb{R}^d . We proceed analogically as in the proof of Theorem 3.7.

For contradiction, we assume that (X, \mathcal{B}) is representable by convex sets in \mathbb{R}^d ; i.e., there are convex sets C_B for $B \in \mathcal{B}$ such that C_{B_1}, \dots, C_{B_k} intersect if and only if B_1, \dots, B_k contain a common point. Again we can assume that these sets are open.

Let $x \in X$. There is again a point $y_x \in \bigcap_{x \in B} C_B$. We define $Y := \{y_x : x \in X\}$. We also define $Y_B := \{y_x : x \in B\}$ for $B \in \mathcal{B}$. Because of the openness of the intersections we can assume that Y is in general position.

Again, let $c(d)$ be the Pach's constant; and $a \in \mathbb{R}^d$, and Z_1, \dots, Z_{d+1} be the (output) data from Theorem 6.2 (when applied to Y). We know that $|Z_i| \geq c|Y|$. By Lemmas 6.9 and 6.11 there are at least $bp/4$ blocks B such that Y_B intersects all sets Z_i . In particular a belongs to $\text{conv } Y_B$ for at least $bp/4$ blocks $B \in \mathcal{B}$. This contradicts Lemma 6.10. \square

A common tool for proving Lemmas 6.9 and 6.10, is the well known Chernoff Bound; see, e.g., [MR02]. We will need it in the following form ($\mathbf{P}(S)$ denotes the probability of an event S):

Theorem 6.12 (Chernoff Bound). *Let $p \in [0, 1]$ and let $\{V_i\}_{i=1}^n$ be independent random variables such that $\mathbf{P}(V_i = 1) = p$ and $\mathbf{P}(V_i = 0) = 1 - p$. Let $V = V_1 + \dots + V_n$. Then*

$$\mathbf{P}(|V - np| > t) < 2e^{-t^2/3np}$$

for $0 \leq t \leq np$.

Proof of Lemma 6.9. First we fix a c -partition \mathcal{A} . Let V_i be the random variable attaining 1 if the block B_i hits \mathcal{A} and 0 otherwise. As we observed above $\mathbf{P}(V_i = 1) = p$ and $\mathbf{P}(V_i = 0) = 1 - p$. In addition variables V_i are independent due to our construction. We set $V = \sum V_i$. Then, using Chernoff's Bound, we have:

$$\mathbf{P}(V < bp/2) < \mathbf{P}(|V - bp| > bp/2) < 2e^{-bp/12}.$$

In another words, the probability that at least $bp/2$ blocks of \mathcal{B}' hit \mathcal{A} is at least $1 - 2e^{-bp/12}$.

It is easy to see that there are less then $v^{vc(d+1)}$ choices of \mathcal{A} . Thus the probability that for every c -partition \mathcal{A} there are at least $bp/2$ blocks of \mathcal{B}' hitting \mathcal{A} is at least

$$1 - v^{vc(d+1)}2e^{-bp/12} \geq 1 - e^{(v \cdot \ln v)c(d+1) - C(d)(v \cdot \ln v)p/12} \geq 1 - \frac{1}{10}$$

assuming that $C(d)$ is large enough. \square

Proof of Lemma 6.10. Let $x \in X$. Let V_i be the random variable attaining 1 if $x \in B_i$ and 0 otherwise. It is easy to see that $\mathbf{P}(V_i = 1) = \frac{d+1}{v}$ and $\mathbf{P}(V_i = 0) = 1 - \frac{d+1}{v}$. The variables V_i are again independent and we again set $V = \sum V_i$.

Similarly as in the proof of Lemma 6.9 we have:

$$\begin{aligned} \mathbf{P}\left(M < \frac{bp}{4}\right) &> 1 - v\mathbf{P}\left(V \geq \frac{bp}{4}\right) > 1 - v\mathbf{P}\left(V > \frac{b(d+1)}{2v}\right) > \\ &> 1 - 2ve^{-b(d+1)/(12v)} > 1 - \frac{1}{10}. \end{aligned}$$

We recall that $b \geq Cv \ln v$ for the very last inequality. \square

Proof of Lemma 6.11. Let $B_i, B_j \in \mathcal{B}'$. The probability that $|B_i \cap B_j| \geq 2$ is less than

$$\frac{\binom{d+1}{2} \binom{v-2}{d-1}}{\binom{v}{d+1}} \leq K(d)v^{-2}$$

where $K(d)$ is independent of v . Thus the probability, that B_i is a conflict block is less than $bK(d)v^{-2} \leq K'(d)v^{-1} \ln v$, where $K'(d)$ is again independent of v . In summary, the expected number of conflict blocks is less than $bK'(d)v^{-1} \ln v \leq K''(d) \ln^2 v$ for $K''(d)$ independent of v .

Let \mathbf{E} denotes the expectation. Using the inequality $\mathbf{P}(U \geq m\mathbf{E}(U)) \leq \frac{1}{m}$ for a nonnegative random variable U we derive that with probability at least $\frac{9}{10}$ there is less than $10K''(d) \ln^2 v$ conflict blocks which is less than $bp/4$ assuming that v is large enough. \square

Chapter 7

Computational complexity of d -collapsibility

The main purpose of this chapter is to prove the following theorem.

Theorem 7.1. (i) 2-COLLAPSIBILITY is polynomial time solvable.

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

In Section 7.5 we also focus on the complexity status of d -REPRESENTABILITY.

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 a d -collapsible face of K such that K_σ is not d -collapsible.

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 7.2. (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 7.1(i) is a straightforward consequence of Theorem 7.2(i). Indeed, if we want to test whether a given complex is 2-collapsible, it is sufficient to greedily collapse d -collapsible faces. Theorem 7.2(i) implies that we finish with an empty complex if and only if the original complex is 2-collapsible.

Our construction for Theorem 7.2(ii) is an intermediate step to proving Theorem 7.1(ii).

Related complexity results. Let us recall related complexity results (in a bit more detail than in Chapter 3).

By a modification of a 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 Section 7.5.

Finally, d -LERAYNUMBER 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*¹ 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—computing homology for this case is indeed only a linear algebra).

A particular example of computational interest. A collection of convex sets in \mathbb{R}^d has a (p, q) -property with $p \geq q \geq d + 1$ if among every p sets of the collection there is a subcollection of q sets with a nonempty intersection. The (p, q) -theorem of Alon and Kleitman states that for all integers p, q, d with $p \geq q \geq d + 1$ there is an integer c such that for every finite collection of convex sets in \mathbb{R}^d with (p, q) -property there are c points in \mathbb{R}^d such that every convex set of the collection contains at least one of the selected points. Let $c' = c'(p, q, d)$ be the minimum possible value of c for which the conclusion of the (p, q) -theorem holds. A significant effort was devoted to estimating c' . The first unsolved case regards estimating $c'(4, 3, 2)$. The best bounds² are due to Kleitman, Gyárfás and Tóth [KGT01]: $3 \leq c'(4, 3, 2) \leq 13$. It seems that the actual value of $c'(4, 3, 2)$ is rather closer to the lower bound in this case, and thus it would be interesting to improve the lower bound even by one.³

Here 2-collapsibility could come into the play. When looking for a small example one could try to generate all 2-collapsible complexes and check the other properties.

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 .

Let us mention that one of the important differences between d -collapsibility and collapsibility is that every finite simplicial complex is d -collapsible for d large enough; on the other hand not an every finite simplicial complex is collapsible.

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 to 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 section 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 7.4. The reader is encouraged to skip them for the first reading and look at them later for full details.

7.1 2-collapsibility

Here we prove Theorem 7.2(i).

The case $d = 1$ follows from the fact that d -collapsible complexes coincide with d -Leray ones ([LB63, Weg75]). Indeed, let K be a 1-collapsible complex and let σ be

¹A link of a face σ in a complex K is the complex $\{\eta \in K : \eta \cup \sigma \in K, \eta \cap \sigma = \emptyset\}$.

²Known to the author.

³Kleitman, Gyárfás and Tóth offer \$30 for such an improvement.

its 1-collapsible face. We have that \mathbf{K} is 1-Leray, which implies that \mathbf{K}_σ is 1-Leray (1-collapsing does not affect homology of dimensions 1 and more). This implies that \mathbf{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 $d = 2$.

Claim 7.3. *Let σ be a good face of \mathbf{K} and let σ' be a 2-collapsible face of \mathbf{K}_σ . Then σ' is a good face of \mathbf{K}_σ .*

Proof. The complex \mathbf{K}_σ is 2-collapsible since σ is a good face of \mathbf{K} . If σ' were a bad face of \mathbf{K}_σ , then \mathbf{K}_σ would be a smaller counterexample to Theorem 7.2(i) contradicting the choice of \mathbf{K} . \square

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 $\text{St}(\sigma, \mathbf{K})$ denotes the (open) *star* of a face σ in \mathbf{K} , which consists of all superfaces of σ in \mathbf{K} (including σ). We remark that $\text{St}(\sigma, \mathbf{K}) = [\sigma, \tau(\sigma)]$ in case that σ is collapsible.

Claim 7.4. *Let $\sigma, \sigma' \in \mathbf{K}$ be independent 2-collapsible faces. Then σ is a 2-collapsible face of $\mathbf{K}_{\sigma'}$, σ' is a 2-collapsible face of \mathbf{K}_σ , and $(\mathbf{K}_\sigma)_{\sigma'} = (\mathbf{K}_{\sigma'})_\sigma$.*

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

$$(\mathbf{K}_\sigma)_{\sigma'} = (\mathbf{K}_{\sigma'})_\sigma = \mathbf{K} \setminus \{\eta \in \mathbf{K} : \sigma \subseteq \eta \text{ or } \sigma' \subseteq \eta\}.$$

\square

Claim 7.5. *Any two 2-collapsible faces of \mathbf{K} are dependent.*

Proof. For contradiction, let σ, σ' be two independent 2-collapsible faces in \mathbf{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 \mathbf{K}_σ by Claim 7.4. Thus, $(\mathbf{K}_\sigma)_{\sigma'}$ is 2-collapsible by Claim 7.3. But $(\mathbf{K}_\sigma)_{\sigma'} = (\mathbf{K}_{\sigma'})_\sigma$ by Claim 7.4, 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 of 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 of σ_G . \square

Due to Claim 7.5 there exists a universal $\tau \in \mathbf{K}$ such that $\tau = \tau(\sigma)$ for every 2-collapsible $\sigma \in \mathbf{K}$. Let us remark that $\mathbf{K} \neq \Delta(\tau)$ 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 7.6. *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 \mathbf{K}_σ that contains σ' , so $(\mathbf{K}_\sigma)_{\sigma'}$ exists. Similarly, $(\mathbf{K}_{\sigma'})_\sigma$ exists and the same argument as in the proof of Claim 7.4 yields $(\mathbf{K}_\sigma)_{\sigma'} = (\mathbf{K}_{\sigma'})_\sigma$. Similarly as in the proof of Claim 7.5, $(\mathbf{K}_\sigma)_{\sigma'}$ is 2-collapsible (due to Claim 7.3), but it contradicts the fact that σ' is a bad face. \square


 Figure 7.1: The simplices τ , τ_k and η .

The complex K is 2-collapsible. Let $K = K_1 \rightarrow K_2 \rightarrow \dots \rightarrow K_m = \emptyset$ be a 2-collapsing of K , where $K_{i+1} = K_i \setminus [\sigma_i, \tau_i]$. Clearly, $\tau_1 = \tau$. Let k be the minimal integer such that $\tau_k \not\subseteq \tau$. Such k exists since $K \neq \Delta(\tau)$. (Let us recall that $\Delta(\tau)$ denotes the full simplex on set τ .) Moreover, we can assume that all the faces $\sigma_1, \dots, \sigma_k$ are edges. This assumption is possible since collapsing a vertex can be substituted by collapsing the edges connected to the vertex and then removing the isolated vertex at the very end of the process. See Lemma 7.15 for details.

Claim 7.7. *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 $\text{St}(\sigma_k, K) = \text{St}(\sigma_k, K_i)$ 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 $\tau_k \not\subseteq \tau$. \square

Claim 7.8. *The faces $\sigma_1, \sigma_2, \dots, \sigma_{k-1}$ are good faces of K .*

Proof. First we observe that each σ_i is 2-collapsible face of K for $0 \leq i \leq k-1$. If σ_i was not 2-collapsible then there is a face $\vartheta \in K$ containing σ_i such that $\vartheta \not\subseteq \tau$. Then $\vartheta \in K_i$ due to minimality of k . Consequently τ_i cannot be the unique maximal face of K_i containing σ_i since ϑ contains σ_i as well.

In order to show that the faces are good we proceed by induction. The face σ_1 is a good face of K since there is a d -collapsing of K starting with σ_1 .

Now we assume that $\sigma_1, \dots, \sigma_{i-1}$ are good faces of K for $i \leq k-1$. If there is an index $j < i$ such that $\sigma_j \cap \sigma_i \neq \emptyset$ then σ_i is good by Claim 7.6. If this is not the case then we set $\sigma_1 = \{x, y\}$. The faces $\sigma_i \cup \{x\}$ and $\sigma_i \cup \{y\}$ belong to K_i ; however, $\sigma_1 \cup \sigma_i$ does not belong to K_i since σ_1 was collapsed. Thus σ_i does not belong to a unique maximal face. \square

Let $\eta = \sigma_k \cup \sigma_B$. See Figure 7.1. Claim 7.7 implies that $\eta \subseteq \tau$. By Claim 7.6 (and the fact that σ_k is not a good face—a consequence of Claim 7.7) the face η does not contain a good face. Thus, $\eta \in K_k$ by Claim 7.8. 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 $K \supseteq K_k$ containing σ_B since σ_B is a 2-collapsible face, which implies $\tau_k \subseteq \tau$. It is a contradiction that $\tau_k \not\subseteq \tau$. \square

7.2 d -collapsible complex with a bad d -collapse

In this section we prove Theorem 7.2(ii).

We start with describing the intuition behind the construction. Given a full complex $K = \Delta(S)$ (the cardinality of S is $2d$), any $(d - 1)$ -face is d -collapsible. However, once we collapse one of them, say σ_B , the rest $(d - 1)$ -faces will be divided into two sets, those which are collapsible in K_{σ_B} (namely, Σ), and those which are not (namely, $\bar{\Sigma}$). For example, when $d = 2$, given a tetrahedron, after collapsing one edge, among the rest five edges, four are collapsible and one is not. The idea of the construction is to attach a suitable complex C to K in such a way that

- the faces of Σ are properly contained in faces of C (and thus they cannot be collapsed until C is collapsed);
- there is a sequence of d -collapses of some of the faces of $\bar{\Sigma}$ such that C can be subsequently d -collapsed.

In summary the resulting complex is d -collapsible because of the second requirement. However, if we start with σ_B , we get stuck because of the first requirement.

7.2.1 Bad complex

Now, for $d \geq 3$, we construct a *bad* complex B , which is d -collapsible but it contains a bad face. A certain important but technical step of the construction is still left out. This is to give the more detailed intuition to the reader. Details of that step are given in the subsequent subsections.

The complex C_{glued} .

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 $K = \Delta(\sigma) \cup \Delta(\gamma_1) \cup \dots \cup \Delta(\gamma_t)$. We attach a certain complex C to L' . The resulting complex is denoted by $C_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$. Here we leave out the details; however, the properties of C_{glued} are described in the forthcoming lemma (we postpone the proof of this lemma).


Lemma 7.9. *Let L, L' , and $C_{\text{glued}} = C_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$ be the complexes from the previous paragraph. Then we have:*

- (i) *If σ is a maximal face of L , then $L \cup C_{\text{glued}} \rightarrow L \setminus \{\sigma\}$.*
- (ii) *The only d -collapsible face of C_{glued} is the face σ .*
- (iii) *Suppose that d is a constant. Then the number of faces of C_{glued} is $O(t)$.*

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

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

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


 Figure 7.2: A schematic drawing of the complexes $\Delta(S)$ and B .

We define B by

$$B = \Delta(S) \cup C_{\text{glued}}(\iota; \alpha_1, \dots, \alpha_t).$$

See Figure 7.2 for a schematic drawing.

Proof of Theorem 7.2(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 7.9(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

$$R = \{\sigma \in \Delta(S) : \text{if } \{q_1, \dots, q_{d-1}\} \subseteq \sigma \text{ then } \sigma \subseteq \iota\}.$$

We observe that $R \setminus \{\iota\}$ is d -collapsible and also that $\Delta(S) \twoheadrightarrow R$ by collapsing all liberation faces (in any order). In fact, the first observation is a special case of Lemma 7.11(ii) used for the NP-reduction.

An auxiliary complex A is defined in a similar way to B :

$$A = R \cup C_{\text{glued}}(\iota; \alpha_1, \dots, \alpha_t).$$

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

$$B \twoheadrightarrow A \twoheadrightarrow R \setminus \{\iota\} \twoheadrightarrow \emptyset.$$

The fact that $B \twoheadrightarrow A$ is quite obvious—it is sufficient to d -collapse the liberation faces. More precisely, we use Lemma 7.16 with $K = B$, $K' = \Delta(S)$, and $L' = R$. The fact that $A \twoheadrightarrow R \setminus \{\iota\}$ follows from Lemma 7.9(i). We already observed that $R \setminus \{\iota\} \twoheadrightarrow \emptyset$ when defining R . \square

7.2.2 The complex \mathbf{C}

Our proof relies on constructing d -dimensional d -collapsible complex \mathbf{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 7.10.

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

Proposition 7.10. *Let $d \geq 2$ and $t \geq 0$ be integers. There is a d -dimensional complex $\mathbf{C} = \mathbf{C}(\rho; \zeta_1, \dots, \zeta_t)$ with the following properties:*

- (i) *It contains $(d - 1)$ -dimensional faces $\rho, \zeta_1, \dots, \zeta_t$ such that each two of them are distant faces.*
- (ii) *Let $\mathbf{C}' = \mathbf{C}'(\rho; \zeta_1, \dots, \zeta_t)$ be the subcomplex of \mathbf{C} given by $\mathbf{C}' = \Delta(\rho) \cup \Delta(\zeta_1) \cup \dots \cup \Delta(\zeta_t)$. Then $\mathbf{C} \twoheadrightarrow (\mathbf{C}' \setminus \{\rho\})$. In particular, \mathbf{C} is d -collapsible since $(\mathbf{C}' \setminus \{\rho\})$ is d -collapsible.*
- (iii) *The only d -collapsible face of \mathbf{C} is the face ρ .*
- (iv) *Suppose that d is a constant. Then the number of faces of \mathbf{C} is $O(t)$.*

7.2.3 The complex $\mathbf{C}(\rho)$

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

We remark that the construction of \mathbf{C} is in some respects similar to the construction of generalized dunce hats. We refer to [AMS93] for more background.

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

$$\text{conv} \{\mathbf{e}_1, -\mathbf{e}_1, \dots, \mathbf{e}_d, -\mathbf{e}_d\}$$

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

$$F_{\mathbf{s}} = \text{conv} \{s_1 \mathbf{e}_1, \dots, s_d \mathbf{e}_d\},$$

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 $\mathbf{u} = (1, \dots, 1)$ (see Figure 7.3).

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 $\mathbf{x} = \mathbf{x}_{\mathbf{a}, \mathbf{s}} = a_1 s_1 \mathbf{e}_1 + \dots + a_d s_d \mathbf{e}_d$ 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 $\{\mathbf{x}_{\mathbf{a}, \mathbf{s}} : \mathbf{s} \in \{-1, 1\}^d \setminus \{\mathbf{u}\}\}$; by X we denote the resulting space. We will construct \mathbf{C} in such a way that X is a geometric realization of \mathbf{C} .

Triangulations of the crosspolytope. We define two auxiliary triangulations of P —they are depicted in Figure 7.4. The simplicial complex \mathbf{J} is the simplicial complex with vertex set $\{\mathbf{0}, \mathbf{e}_1, -\mathbf{e}_1, \dots, \mathbf{e}_d, -\mathbf{e}_d\}$. The set of its faces is given by the maximal faces

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




Figure 7.3: The space X . The arrows denote, which facets are identified.

The complex J is a triangulation of P .

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

$$\{\{\sigma_1, \dots, \sigma_k\} \cup \tau\} \text{ where } \sigma_1 \supseteq \dots \supseteq \sigma_k \supseteq \tau; \sigma_1, \dots, \sigma_k \in J \setminus \Delta(\vartheta); \tau \subseteq \vartheta; k \in \mathbb{N}_0.$$

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

Formally, let \approx be an equivalence relation on $(J \setminus \Delta(\vartheta)) \cup \vartheta$ given by

$$\begin{aligned} \mathbf{e}_i &\approx \{-\mathbf{e}_i\} && \text{for } i \in [d], \\ \sigma_1 &\approx \sigma_2 && \text{for } \sigma_1, \sigma_2 \in J \setminus \Delta(\vartheta), \\ &&& \sigma_1 = \{s_1 \mathbf{e}_{k_1}, \dots, s_m \mathbf{e}_{k_m}\}, \sigma_2 = \{s'_1 \mathbf{e}_{k_1}, \dots, s'_m \mathbf{e}_{k_m}\} \\ &&& \text{where } s_i, s'_i \in \{-1, 1\} \text{ and } 1 \leq k_1 < \dots < k_m \leq d. \end{aligned}$$

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

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

The faces of C are the equivalence classes of \approx^+ . We use the notation $\langle \sigma \rangle$ for such an equivalence class given by $\sigma \in H$. By ρ we denote the face $\langle \{\mathbf{e}_1, \dots, \mathbf{e}_d\} \rangle$ of C .

7.2.4 The complex $C(\rho; \zeta_1, \dots, \zeta_t)$

Now we assume that $t \geq 1$ and we construct the complex $C(\rho; \zeta_1, \dots, \zeta_t)$, which is a refinement of $C(\rho)$. The idea of the construction is quite simple. We pick an interior simplex of $C(\rho)$; and we subdivide it in such a way that we obtain distant $(d - 1)$ -dimensional faces ζ_1, \dots, ζ_t (and also distant from ρ). For completeness we show a particular way how to get such a subdivision.

A suitable triangulation of a simplex. An example of the following construction is depicted in Figure 7.5. Let Δ be a d -dimensional (geometric) simplex with a set of vertices $V = \{\mathbf{v}_1, \dots, \mathbf{v}_{d+1}\}$, let \mathbf{b} be its barycentre, and let t be an integer. Next, we




Figure 7.4: The triangulations J (left) and H (right) of P with $d = 2$.

define

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

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

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

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

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]).

Let $D(\zeta_1, \dots, \zeta_t)$ denote 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, \dots, Q_{3t} and subsequently triangulating these polyhedra as described above.

The definition of $C(\rho; \zeta_1, \dots, \zeta_t)$. Let ξ be a d -face of H such that $\|\xi\| \subset \text{int}\|H\|$. Although there are multiple such d -faces only some of them are used as ξ . For example, in Figure 7.5, only one out of four such d -faces is chosen. Suppose that the set V (from above) is the set of vertices of ξ . We define

$$C(\rho; \zeta_1, \dots, \zeta_t) = (C(\rho) \setminus \{\langle \xi \rangle\}) \cup D(\zeta_1, \dots, \zeta_t)$$

while recalling that $\langle \xi \rangle$ denotes the equivalence class of \approx^+ from the definition of C . See Figure 7.5.

Proof of Proposition 7.10. 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 $C' \setminus \{\rho\}$. Formally, we use Lemma 7.17. \square

Gluing. Here we focus on gluing briefly discussed above Lemma 7.9. As the name of connecting gadget suggests, we want to use it (in Section 7.3) for connecting several other complexes (gadgets). In particular, we want to have some notation for gluing this gadget. We introduce this notation here.




Figure 7.5: The complex $D(\zeta_1, \zeta_2)$ (left) and $C(\rho; \zeta_1, \zeta_2)$ (right), here $d = 2$.

Again we suppose that $\sigma, \gamma_1, \dots, \gamma_t$ are already known $(d-1)$ -dimensional faces of a given complex L . They are assumed to be distinct, but not necessarily disjoint. There is a complex $K = \Delta(\sigma) \cup \Delta(\gamma_1) \cup \dots \cup \Delta(\gamma_t)$. We take a new copy of $C(\rho; \zeta_1, \dots, \zeta_t)$ and we perform identifications $\rho = \sigma, \zeta_1 = \gamma_1, \dots, \zeta_t = \gamma_t$. After these identifications, the complex $K \cup C(\rho; \zeta_1, \dots, \zeta_t)$ is denoted by $C_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_t)$. Note that C (before gluing) and C_{glued} are generally not isomorphic since the gluing procedure can identify some faces of C .

Proof of Lemma 7.9. The first claim follows from Lemma 7.19. The second claim follows from Proposition 7.10(i) and (iii). The last claim follows from Proposition 7.10(iv). \square

7.3 NP-completeness

Here we prove Theorem 7.1(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.

7.3.1 Sketch of the reduction

Let us recall the construction of the bad complex B . We have started with a simplex $\Delta(S)$ and we distinguished the initial face ι and the bad face σ_B . We were allowed to start the collapsing either with σ_B or with liberation faces and then with ι . As soon as one of the options was chosen the second one was unavailable. The idea is that these two options should represent an assignment of variables in the formula Φ .

A disadvantage is that we cannot continue after collapsing σ_B . Thus we rather need to distinguish two initial faces ι^+ and ι^- each of them having its own liberation faces. However, we need that these two collections of liberation faces do not interfere. That is why we have to assume $d \geq 4$.




Figure 7.6: A schematic pictures of simplicial gadgets; the liberation faces of the merge gadget are distinguished.

For every variable x_j of the formula Φ we thus construct a certain *variable gadget* V_j with two initial faces ι_j^+ and ι_j^- . For a clause C^i in the formula Φ there is a *clause gadget* G^i . Initially it is not possible to collapse clause gadgets. Assume, e.g., that C^i contains variables x_j and $x_{j'}$ in positive occurrence and $x_{j''}$ in negative occurrence. Then it is possible to collapse G^i as soon as ι_j^+ , $\iota_{j'}^+$, or $\iota_{j''}^-$ was collapsed. (This is caused by attaching a suitable copy of the connection gadget C from the previous section.) Thus the idea is that the complex F in the reduction is collapsible if and only if all clause gadgets can be simultaneously collapsed which happens if and only if Φ is satisfiable.

There are few more details to be supplied. Similarly as for the construction of B we have to attach a copy T of the connecting gadget C to the faces which are neither initial nor liberation (i.e., to attaching faces). This step is necessary for controlling which faces can be collapsed. This copy of connecting gadget is called a *tidy connection* and once it is activated (at least one of its faces is collapsed) then it is consequently possible to collapse the whole complex F . Finally, there are inserted certain gadgets called *merge gadgets*. Their purpose is to merge the information obtained by clause gadgets: they can be collapsed after collapsing all clause gadgets and then they activate the tidy connection. The precise definition of F will be described in following subsections. At the moment it can be helpful for the reader to skip few pages and look at Figure 7.7 (although there is a notation on the picture not introduced yet).

7.3.2 Simplicial gadgets

Now we start supplying the details. As sketched above we introduce several gadgets called *simplicial gadgets*. They consist of full simplices (on varying number of vertices) with several distinguished $(d-1)$ -faces. These gadgets generalize the complex $\Delta(S)$. 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* $V = V(\iota^+, \iota^-, \beta^+, \beta^-)$ is described by the following table:

$$\begin{aligned}
\text{vertices:} & \quad p^+, q_1^+, \dots, q_{d-1}^+, p^-, q_1^-, \dots, q_{d-1}^-; \\
\text{initial faces:} & \quad \iota^+ = \{p^+, q_1^+, \dots, q_{d-1}^+\}, \iota^- = \{p^-, q_1^-, \dots, q_{d-1}^-\}; \\
\text{bases:} & \quad \beta^+ = \{q_1^+, \dots, q_{d-1}^+\}, \beta^- = \{q_1^-, \dots, q_{d-1}^-\}.
\end{aligned}$$

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

$$\begin{aligned}
\text{vertices:} & \quad p_1, \dots, p_d, q; \\
\text{initial face:} & \quad \iota = \{p_1, p_2, \dots, p_d\}; \\
\text{bases:} & \quad \beta_1 = \iota \setminus \{p_1\}, \beta_2 = \iota \setminus \{p_2\}, \beta_3 = \iota \setminus \{p_3\}.
\end{aligned}$$

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

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

$$\begin{aligned}
\text{vertices:} & \quad p_1, \dots, p_d, q, r; \\
\text{initial face:} & \quad \iota_{\text{merge}} = \{p_1, p_2, \dots, p_d\}; \\
\text{base:} & \quad \iota_{\text{merge}} \setminus \{p_1\}.
\end{aligned}$$

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 7.11. *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) $R \setminus \{\iota\}$ is d -collapsible;
- (iii) $R \setminus \{\iota\} \rightarrow \Delta(\iota')$ where ι' is an initial face different from ι (if exists).

Proof. We prove each of the claims separately.

- (i) Let V be the 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 7.14.

- (ii) We observe that β is a maximal $(d-2)$ -face of $R \setminus \{\iota\}$ and $S_\beta = R \setminus \{\iota, \beta\}$, hence $R \setminus \{\iota\} \rightarrow S_\beta$. (We recall that K_σ denotes the resulting complex of an elementary d -collapse $K \rightarrow K_\sigma = K \setminus [\sigma, \tau(\sigma)]$.) Next, $S_\beta \rightarrow S_\emptyset = \emptyset$ by Lemma 7.14.
- (iii) Similarly as before we have $R \setminus \{\iota\} \rightarrow S_\beta$. Let v be a vertex of β , we have $S_\beta \rightarrow S_{\{v\}}$ by Lemma 7.14. The complex $S_{\{v\}}$ is a full simplex (S with removed v), this complex even 1-collapse to $\Delta(\iota')$ by collapsing vertices of $V \setminus (\iota' \cup \{v\})$ (in any order).

□

7.3.3 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 $V_j = V_j(\iota_j^+, \iota_j^-, \beta_j^+, \beta_j^-)$. 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 $G^i = G^i(\iota^i, \lambda_{j_1}^i, \lambda_{j_2}^i, \lambda_{j_3}^i)$. Moreover, for C^i with $i \geq 2$, we also take a new copy of the merge gadget and we denote it $M^i = M^i(\iota_{\text{merge}}^i, \lambda_{\text{merge},1}^i, \lambda_{\text{merge},2}^i)$.

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

$$O_j^+ = C_{\text{glued}}(\iota_j^+; \lambda_j^{i_1}, \dots, \lambda_j^{i_k}).$$

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

$$\begin{aligned} I_1^1 &= C_{\text{glued}}(\iota_1^1; \lambda_{\text{merge},2}^2); \\ I_1^i &= C_{\text{glued}}(\iota_1^i; \lambda_{\text{merge},1}^i) \quad \text{where } i \in \{2, \dots, n\}; \\ I_2^i &= C_{\text{glued}}(\iota_{\text{merge}}^i; \lambda_{\text{merge},2}^{i+1}) \quad \text{where } i \in \{2, \dots, n-1\}. \end{aligned}$$

For convenient notation we denote I_1^1 also by I_2^1 .

Finally, the *tidy connection* is defined by

$$T = C_{\text{glued}}(\iota_{\text{merge}}^n; \alpha_1, \dots, \alpha_t)$$

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

$$F = \bigcup_{j=1}^m V_j \cup \bigcup_{i=1}^n G^i \cup \bigcup_{i=2}^n M^i \cup \bigcup_{j=1}^m (O_j^+ \cup O_j^-) \cup \bigcup_{i=1}^n I_1^i \cup \bigcup_{i=2}^{n-1} I_2^i \cup T.$$

See Figure 7.7 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 7.9(iii).

Collapsibility for satisfiable formulae. We suppose that the formula is satisfiable and we describe a collapsing of F ; see Figure 7.8.

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 -collapse⁴ the liberation faces containing β_j^+ (see Lemma 7.11(i)), after that ι_j^+

⁴Note 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.

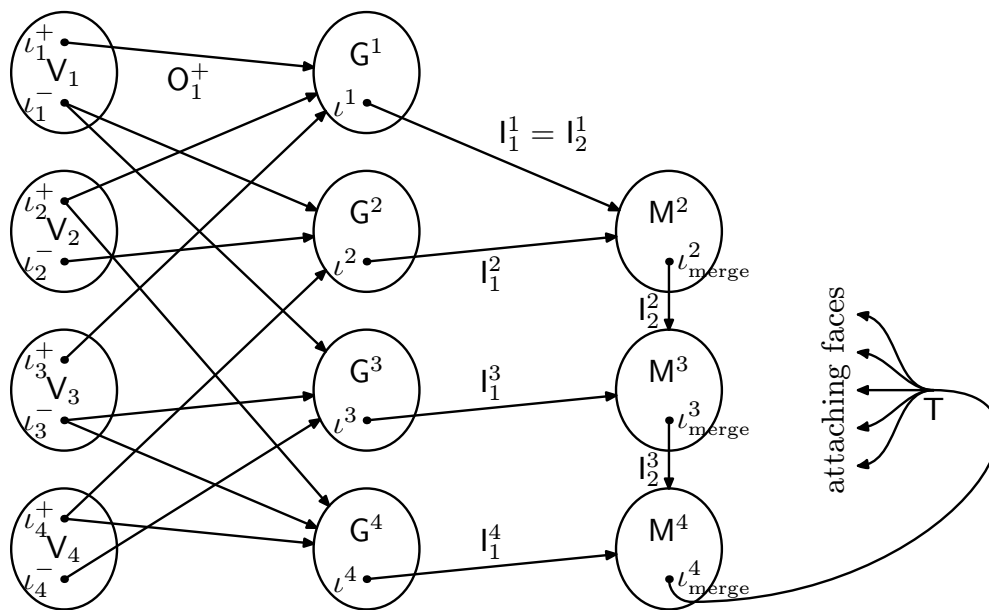


Figure 7.7: A schematic example of F for the formula $\Phi = (x_1 \vee x_2 \vee x_3) \wedge (\neg x_1 \vee \neg x_2 \vee x_4) \wedge (\neg x_1 \vee \neg x_3 \vee \neg x_4) \wedge (x_2 \vee \neg x_3 \vee x_4)$. Initial faces are drawn as points. (Multi)arrows denote connections. Each (multi)arrow points from the unique d -collapsible face of the connection to simplicial gadgets that are attached to the connection by some of its liberation faces.

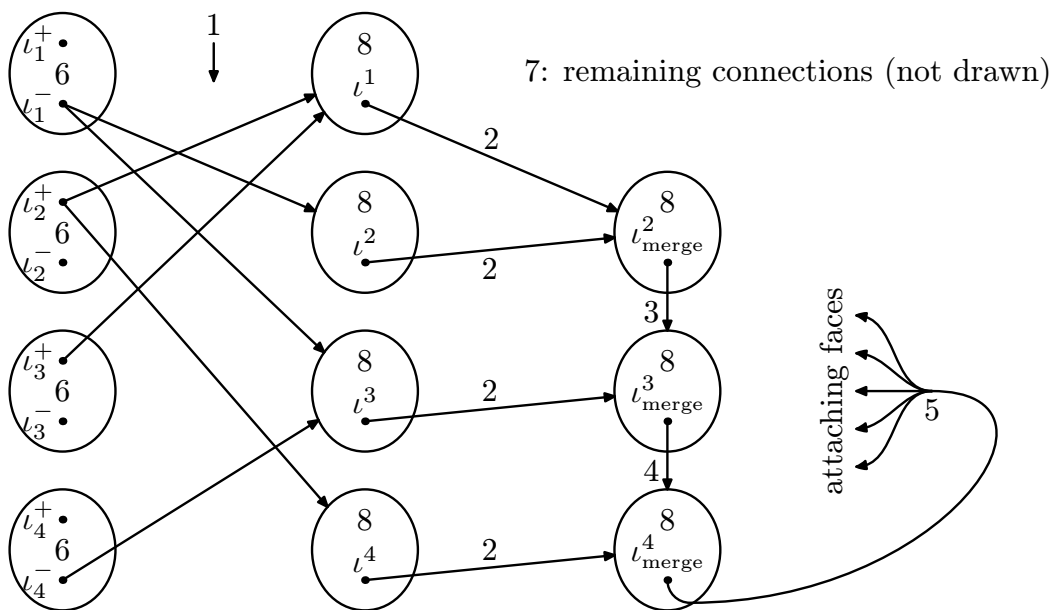


Figure 7.8: d -collapsing of F for the Φ from Figure 7.7 assigned (FALSE, TRUE, TRUE, FALSE). The numbers denote the order in which the parts of F vanish.

is d -collapsible, and we d -collapse O_j^+ (following Lemma 7.9(i) in the same way as in the proof of Theorem 7.2(ii)). Similarly, we d -collapse O_j^- if x_j is assigned FALSE.

We use several times Lemma 7.11(i) and Lemma 7.9(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 l_1^i for $i \in [n]$.

Next we gradually d -collapse the merge gadgets l_2^i for $i \in \{2, \dots, n-1\}$. For this, we have that both liberation faces of l_2^i are d -collapsible, we d -collapse them and we have that ι_{merge}^2 is d -collapsible. We d -collapse l_2^i and now we continue with the same procedure with l_2^3 , the l_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 7.11(iii)), then all remaining connections, and at the end all remaining simplicial gadgets due to Lemma 7.11(ii).

Non-collapsibility for unsatisfiable formulae. Now 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

$$F = F_1 \rightarrow F_2 \rightarrow \dots \rightarrow \emptyset$$

be a d -collapsing of F . We call it *our* d -collapsing. For a technical reason, according to Lemma 7.15, 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 $N_\ell = F_\ell \cap N$. We say that N is *activated* in F_ℓ if N_ℓ is a proper subcomplex of N .

The connection N is defined as $C_{\text{glued}}(\sigma; \gamma_1, \dots, \gamma_s)$ for some $(d-1)$ -faces $\sigma, \gamma_1, \dots, \gamma_s$ of simplicial gadgets in F . We remark that Lemma 7.9(ii) implies that if N is activated in F_ℓ then $\sigma \notin F_\ell$.

We also prove the following lemma about activated connections.

Lemma 7.12. *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 l_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 $i \in \{2, \dots, n-1\}$. If the merge connection l_2^i is activated in F_ℓ , then the merge connections l_1^i and l_2^{i-1} are activated in F_ℓ .*

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

$$F = F_1 \rightarrow F_2 \rightarrow \dots \rightarrow F_\ell,$$

where $F_{k+1} = F_k \setminus [\sigma_k, \tau_k]$ for $k \in [\ell-1]$. According to assumption on our d -collapsing, we have that $\sigma_1, \dots, \sigma_{\ell-1}$ 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 $\sigma_k \in V_j$. 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 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 $\sigma_{k^-} \cup \{s\} \in F_{k-1}$, because $\sigma_{k^-} \cup \{s\}$ does not contain a positive subspace (it does not contain β_j^+ since $|\sigma_{k^-} \cap \beta_j^+| \leq 1$, but $|\beta_j^+| \geq 3$). On the other hand $\sigma_{k^-} \cup S \notin K_{k-1}$ since it contains σ_{k^+} . I.e., σ_{k^-} is not in a unique maximal face.

- (ii) We again define a relevant index; this time $k \in [\ell - 1]$ is *relevant* if $\sigma_k \in G^i$. 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 7.9(ii) we would have that $I_1^i \subseteq F_{k'-1}$ and also $G^i \subseteq F_{k'-1}$ from minimality of k' , which would contradict that $\sigma_{k'}$ is a collapsible face of $F_{k'-1}$. Thus $\sigma_{k'}$ is a liberation face of G^i . This implies, again by Lemma 7.9(ii), that at least one of the occurrence gadgets attached to liberation faces is activated even in $F_{k'-1}$.
- (iii) By a similar discussion as in previous step we have that at least one of the liberation faces $\lambda_{\text{merge},1}^i$ and $\lambda_{\text{merge},2}^i$ of M^i have to be d -collapsed before d -collapsing ι_{merge}^i . However, we observe that d -collapsing only one of these faces is still insufficient for possibility of d -collapsing ι_{merge}^i . 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 also prove an analogy of Lemma 7.12 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 7.12(iii), therefore we omit it.

Lemma 7.13. *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 7.13. By this lemma and repeatedly using Lemma 7.12(iii) we have that all merge connections are activated in F_ℓ . By Lemma 7.12(ii), for every clause gadget G^i there is an occurrence gadget attached to G^i , which is activated in F_ℓ . Finally, Lemma 7.12(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.

□


 Figure 7.9: An example of 2-collapsing $K \rightarrow K_{\sigma'} \rightarrow K_{\sigma}$.

7.4 Technical properties of d -collapsing

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

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

Lemma 7.14. *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 $K_{\sigma'} \rightarrow K_{\sigma}$.*

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 $K_{\sigma'} \rightarrow K_{\sigma' \setminus \{v_1\}}$ 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

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

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 $(K_i)_{\eta_i} = K_{i+1}$. Thus, we have a d -collapsing

$$K = K_1 \rightarrow K_2 \rightarrow \dots \rightarrow K_{t+1}.$$


See Figure 7.9 for an example.

To finish the proof it remains to observe that $K_2 = K_{\sigma'}$ and K_{t+1} is a disjoint union of K_{σ} and $\{\sigma\}$, hence $K_{t+1} \rightarrow K_{\sigma}$. \square

As a corollary, we obtain the following lemma.

Lemma 7.15. *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 than $(d - 1)$ 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 $K' \rightarrow K'_\sigma$. Let $\sigma' \supseteq \sigma$ be such a face of K' that either $\dim \sigma' = d - 1$ or σ' is a maximal face. Then we replace this step by d -collapsing $K' \rightarrow K'_{\sigma'} \rightarrow K_{\sigma}$.


 Figure 7.10: Complexes K , K' , L and L' from the statement of Lemma 7.16.

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, while 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. \square

7.4.2 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 perform d -collapsing $K' \rightarrow L'$ in whole K ; see Figure 7.10 for an illustration. The precise statement is given in the following lemma.

Lemma 7.16 (*d -collapsing a subcomplex*). *Let K be a simplicial complex, K' a subcomplex of K , and L' a subcomplex of K' . Assume that if $\sigma \in K' \setminus L'$, $\eta \in K$, and $\eta \supseteq \sigma$, then $\eta \in K' \setminus L'$. Moreover assume that $K' \rightarrow L'$. Then $L = (K \setminus K') \cup L'$ is a simplicial complex and $K \rightarrow L$.*

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

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

In order to show $K \rightarrow L$, 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 $K'_{\sigma'} \supseteq L'$. Then we have

1. σ' is a d -collapsible face of K .
2. If $\sigma \in K'_{\sigma'} \setminus L'$, $\eta \in K_{\sigma'}$ and $\eta \supseteq \sigma$, then $\eta \in K'_{\sigma'} \setminus L'$.
3. $L = (K_{\sigma'} \setminus K'_{\sigma'}) \cup L'$.

We prove the claims separately:

1. We know that $\sigma' \notin L'$ since $K'_{\sigma'} \supseteq L'$. Thus, $\sigma' \in K' \setminus L'$. If $\eta' \in K$ and $\eta' \supseteq \sigma'$, then, by the assumption of the lemma, $\eta' \in K' \setminus L' \subseteq K'$. 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 .




Figure 7.11: In top right picture there are complexes K and L from Lemma 7.17; L is thick and dark. In top left picture there is the graph $G_2(K \setminus L)$. Collapsing $K \rightarrow L$ is in bottom pictures.

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

□

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:

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

Lemma 7.17 (*d -collapsing a d -dimensional complex*). *Suppose that K is a d -dimensional complex, L is its subcomplex and the following conditions are satisfied:*

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

Then $K \rightarrow L$.

Proof. See Figure 7.11 when following the proof. Let $\tau_0 = \tau(\sigma)$, τ_1, \dots, τ_j be an order of vertices of $G_d(\mathbf{K} \setminus \mathbf{L})$ 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

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

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

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

7.4.3 Gluing distant faces

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

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

we mean the resulting complex under the identification $u_1 = v_1, \dots, u_{k+1} = v_{k+1}$ (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

$$\mathbf{K}(\omega_1 = \eta_1, \dots, \omega_t = \eta_t)$$

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

Lemma 7.18 (Collapsing glued complex). *Suppose that ω and η are two distant faces in a simplicial complex \mathbf{K} . Let \mathbf{L} be a subcomplex of \mathbf{K} such that $\omega, \eta \in \mathbf{L}$. Suppose that \mathbf{K} d -collapses to \mathbf{L} . Then $\mathbf{K}(\omega = \eta)$ d -collapses to $\mathbf{L}(\omega = \eta)$.*

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

$$\mathbf{K}(\omega = \eta) \rightarrow \mathbf{K}_2(\omega = \eta) \rightarrow \mathbf{K}_3(\omega = \eta) \rightarrow \dots \rightarrow \mathbf{L}(\omega = \eta)$$

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

It is sufficient to show $\mathbf{K}(\omega = \eta) \rightarrow \mathbf{K}_2(\omega = \eta)$ 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 $\mathbf{K}_2 = \mathbf{K}_\sigma$ 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 $\sigma \cap \omega = \emptyset$ then $[\sigma] = \sigma$, and hence $\sigma' \subseteq \sigma$ since the vertices of σ are not glued to another vertices); 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 $[\sigma]$. \square

Lemma 7.19 (Collapsing of the connecting gadget). *Let t be an integer. Let $\bar{\mathbb{L}}$ be a complex with distinct d -dimensional faces $\sigma, \gamma_1, \dots, \gamma_t$ such that σ is a maximal face of $\bar{\mathbb{L}}$. Let $\mathbb{C} = \mathbb{C}(\rho, \zeta_1, \dots, \zeta_t)$ and $\mathbb{C}' = \mathbb{C}'(\rho, \zeta_1, \dots, \zeta_t)$ be complexes defined in Section 7.2.*

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

Proof. First, we observe that

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

This follows from Lemma 7.16 by setting $\mathbb{K} = (\bar{\mathbb{L}}\dot{\cup}\mathbb{C})(\sigma = \rho)$, $\mathbb{K}' = \mathbb{C}$, $\mathbb{L}' = \mathbb{C}' \setminus \{\sigma\}$, and then $\mathbb{L} = (\bar{\mathbb{L}}\dot{\cup}\mathbb{C}')(\sigma = \rho) \setminus \{\sigma\}$. Assumptions of the lemma are satisfied by Proposition 7.10(ii) and the inspection.

Now it is sufficient to iterate Lemma 7.18, assumptions are satisfied by Proposition 7.10(i). \square

7.5 The complexity of d -representability

In this section we prove that d -COLLAPSIBILITY 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 $\{\{F, F'\} : F, F' \in \mathcal{F}, F \neq F', F \cap F' \neq \emptyset\}$. 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 $\text{SEG}(\leq 2)$ 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 to another classes. We also establish a similar definition for simplicial complexes. Suppose that \mathbb{K} is a d -representable simplicial complex. A system \mathcal{C} of convex sets in \mathbb{R}^d such that \mathbb{K} is isomorphic to the nerve of \mathcal{C} is called a *d -representation* of \mathbb{K} .

We have $\text{STR} \supseteq \text{CON} \supseteq \text{SEG}$ (actually, it is known that the inclusions are strict). Furthermore, suppose that we are given a graph $G \in \text{SEG}$. 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 $\text{SEG} = \text{SEG}(\leq 2)$.

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 $G(\Phi) \in \text{SEG}$ if the formula is satisfiable, but $G(\Phi) \notin \text{STR}$ if the formula is unsatisfiable. Moreover, we already know that this yields $G(\Phi) \in \text{SEG}(\leq 2)$ 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 $\text{SEG}(\leq 2)$ -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-COLLAPSIBILITY is NP-hard.

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

$$\text{susp } \mathbf{K} = \mathbf{K} \cup \{\{a\} \cup \sigma : \sigma \in \mathbf{K}\} \cup \{\{b\} \cup \sigma : \sigma \in \mathbf{K}\}.$$

Lemma 7.20. *Let d be an integer. A simplicial complex \mathbf{K} is $(d - 1)$ -representable if and only if $\text{susp } \mathbf{K}$ is d -representable.*

Proof. First, we suppose that \mathbf{K} is $(d - 1)$ -representable and we show that $\text{susp } \mathbf{K}$ is d -representable. Let $K_1, \dots, K_t \subseteq \mathbb{R}^{d-1}$ be convex set from a $(d - 1)$ -representation of \mathbf{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

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

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

For the reverse implication, we suppose that $\text{susp } \mathbf{K}$ is d -representable and we show that \mathbf{K} is $(d - 1)$ -representable. Suppose that $K(a), K(b), K_1, \dots, K_t$ is a d -representation of $\text{susp } \mathbf{K}$ ($K(a)$ corresponds to a and $K(b)$ corresponds to b). We have that $\{a, b\} \notin \text{susp } \mathbf{K}$, 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 \mathbf{K} . □

Since 2-COLLAPSIBILITY is NP-hard, we have the following corollary of Lemma 7.20 (considering complexes that are obtained as $(d - 2)$ -tuple suspensions):

Theorem 7.21. *d -COLLAPSIBILITY is NP-hard for $d \geq 2$.*

□

Chapter 8

A counterexample to Wegner's conjecture

The purpose of this chapter is to prove the following theorem. The theorem is a counterexample to Wegner's conjecture [Weg75].

Theorem 8.1. *For every $d \geq 2$ there is a simplicial complex which is topologically d -representable but not d -collapsible.*

8.1 Planar case

We start this section with describing a complex L which will serve us for the planar case. Let $A_1, A_2, A_3, B_1, B_2, B_3, C_1, C_2, C_3, D, X_1, X_2, X_3, Y_1, Y_2, Y_3, Z_1, Z_2, Z_3$ be the (open) sets from the Figure 8.1. We also set $\mathcal{A} := \{A_1, A_2, A_3\}$, $\mathcal{B} := \{B_1, B_2, B_3\}$, $\mathcal{C} := \{C_1, C_2, C_3\}$, $\mathcal{D} := \{D\}$, $\mathcal{X} := \{X_1, X_2, X_3\}$, $\mathcal{Y} := \{Y_1, Y_2, Y_3\}$, and $\mathcal{Z} := \{Z_1, Z_2, Z_3\}$. Let \mathcal{L} be the collection of all these sets, i.e., $\mathcal{L} := \mathcal{A} \cup \mathcal{B} \cup \mathcal{C} \cup \mathcal{D} \cup \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$. Finally, L is the nerve of \mathcal{L} .

We will show that L is topologically 2-representable but not 2-collapsible.

8.1.1 Topological representability

It is sufficient to show that \mathcal{L} is a good cover. This property can be hand-checked; however, we offer an alternative approach.

First we realize that all sets of $\mathcal{L} \setminus \mathcal{Z}$ are convex. Thus $\mathcal{L} \setminus \mathcal{Z}$ is a good cover. It remains to check that adding sets of \mathcal{Z} does not violate this property.

Let $Z \in \mathcal{Z}$ and let $\mathcal{L}^Z := \{L \cap Z : L \in \mathcal{L}\}$. We are done as soon as we show that $\mathcal{L}^{Z_1}, \mathcal{L}^{Z_2}$, and \mathcal{L}^{Z_3} are good covers.

Because of the symmetry we show it only for \mathcal{L}^{Z_1} . The sets of \mathcal{L}^{Z_1} can be transformed into convex sets by a homeomorphism of \mathbb{R}^2 . See Figure 8.1.1. Thus they form a good cover.

8.1.2 Non-collapsibility by case analysis

Here we prove that L is not 2-collapsible by case analysis. We get a bit stronger results that will help us for higher dimensions. Disadvantage of this proof is that it does not give an explanation how is the complex constructed. Therefore we supply




Figure 8.1: The sets A_1, \dots, Z_3 . We rather supply more detailed description of the sets if the picture is print only in black and white: The sets A_* are the ovals on the boundary; B_* are the small discs close to the boundary; C_* are the trapezoid-shaped sets; D is the triangle in the center; X_* are the circles close to the center; Y_* are the deltoid-shaped sets; and Z_* are the boomerang-shaped sets

an additional heuristic explanation in the next subsection, although it would need a bit more effort to turn that explanation into a proof.

For a simplicial complex K we set

$$\gamma_0(K) := \min\{d : K \text{ has a } d\text{-collapsible face}\}.$$

The fact that L is not 2-collapsible is implied by the following proposition.

Proposition 8.2. $\gamma_0(L) = 3$.

In order to prove the proposition we need a simple lemma.

Lemma 8.3. *Let K be a simplicial complex and σ be a 1-face (edge) of it. Assume that u and v are vertices of K not belonging to σ such that $\sigma \cup \{u\} \in K$, $\sigma \cup \{v\} \in K$, but $\sigma \cup \{u, v\} \notin K$. Then σ is not a 2-collapsible face of K .*

Proof. If τ is a unique maximal face of K containing σ then $u, v \in \tau$ due to the conditions of the lemma. However, $\tau \supseteq \sigma \cup \{u, v\} \notin K$. \square

Proof of Proposition 8.2. In the spirit of Lemma 8.3 for every 1-face $\sigma \in L$ we find a couple of vertices $u, v \in L$ such that $\sigma \cup \{u\}, \sigma \cup \{v\} \in L$, but $\sigma \cup \{u, v\} \notin L$. It is




Figure 8.2: A transformation of \mathcal{L}^{Z_1} . Whatever is outside of Z_1 can be ignored.


sufficient to check 1-faces since if a 0-face (vertex) w is 1-collapsible then any 1-face containing w is 1-collapsible as well. Moreover, it is sufficient to check only some 1-faces because of the symmetries of the complex. The rest of the proof is given by the following table.

| σ | u, v | σ | u, v | σ | u, v |
|----------------|------------|----------------|------------|----------------|------------|
| $\{A_1, A_2\}$ | B_2, Z_2 | $\{A_1, B_1\}$ | C_1, A_3 | $\{A_1, C_1\}$ | B_1, B_2 |
| $\{A_1, Y_1\}$ | B_1, Z_1 | $\{A_1, Z_1\}$ | C_1, A_3 | $\{B_1, C_1\}$ | A_1, C_3 |
| $\{B_1, Y_1\}$ | C_1, A_3 | $\{C_1, C_2\}$ | B_2, D | $\{C_1, D\}$ | C_2, C_3 |
| $\{C_1, X_1\}$ | Y_1, Y_2 | $\{C_1, Y_1\}$ | B_1, Z_1 | $\{C_1, Z_1\}$ | Y_1, Z_2 |
| $\{D, X_1\}$ | Y_1, Y_2 | $\{D, Y_1\}$ | C_1, X_3 | $\{X_1, X_2\}$ | Y_2, X_3 |
| $\{X_1, Y_1\}$ | D, Z_1 | $\{X_1, Z_1\}$ | Y_1, Z_2 | $\{Y_1, Z_1\}$ | C_1, A_3 |
| $\{Z_1, Z_2\}$ | A_1, X_1 | | | | |

□

8.1.3 Sketch of non-collapsibility

The purpose of this subsection is to give a rough idea why the complex L should not be 2-collapsible. This description could be useful, for instance, for generalizations. However, the reader can easily skip this part. The author still prefer to include this discussion in order to explain how the complex is built up.

Figure 8.3: Decomposition of \mathcal{L} into two parts.

Let us split the collection \mathcal{L} into two parts $\mathcal{L}^+ := \mathcal{A} \cup \mathcal{B} \cup \mathcal{C} \cup \mathcal{D}$ and $\mathcal{L}^- := \mathcal{X} \cup \mathcal{Y} \cup \mathcal{Z}$. The nerve of \mathcal{L}^+ , resp. \mathcal{L}^- , is denoted by \mathcal{L}^+ , resp. \mathcal{L}^- . Both \mathcal{L}^+ and \mathcal{L}^- are triangulations of a disc with only three boundary edges $\{A_1, A_2\}$, $\{A_1, A_3\}$, and $\{A_2, A_3\}$; resp. $\{Z_1, Z_2\}$, $\{Z_1, Z_3\}$, and $\{Z_2, Z_3\}$; see Figure 8.1.3. Only these boundary faces are 2-collapsible faces of \mathcal{L}^+ , resp. \mathcal{L}^- .

By suitable overlapping of \mathcal{L}^+ and \mathcal{L}^- (i.e., obtaining \mathcal{L}) we get that also the above mentioned boundary faces are not 2-collapsible anymore (in whole \mathcal{L}). For instance $Z_1 \cap Z_2$ intersects A_1 (in addition to X_1 already in \mathcal{L}^-); however, A_1 and X_1 are disjoint. Thus $\{Z_1, Z_2\}$ is not a 2-collapsible face of \mathcal{L} .

It remains to check that merging \mathcal{L}^+ and \mathcal{L}^- does not introduce any new problems. It is, in fact, checked in a detail in the previous section. We just mention that there is no problem with 1-faces which already appear in \mathcal{L}^+ or \mathcal{L}^- . However; new 1-faces are introduced when one vertex comes from \mathcal{L}^+ and the second one from \mathcal{L}^- . For another triangulations these newly introduced faces can be 2-collapsible.¹

¹It would be perhaps possible to show that the complex is not 2-collapsible even if the newly introduced faces were 2-collapsible. Listing all 1-faces in the previous subsection seems, however, more convenient for the current purpose.

8.2 Higher dimensions

Joins of simplicial complexes will help us to generalize the counterexample to higher dimensions. Let K and K' be simplicial complexes with the vertex sets $V(K)$ and $V(K')$. Their *join* is a simplicial complex $K \star K'$ whose vertex set is the disjoint union $V(K) \sqcup V(K')$,² and whose set of faces is $\{\alpha \sqcup \beta : \alpha \in K, \beta \in K'\}$.

We need the following two lemmas.

Lemma 8.4 ([MT09, Lemma 4.2]). *For every two simplicial complexes K, K' we have $\gamma_0(K \star K') = \gamma_0(K) + \gamma_0(K')$.*

Lemma 8.5. *Let K be a convexly/topologically d -representable complex and K' be a convexly/topologically d' -representable complex. Then $K \star K'$ is a convexly/topologically $(d + d')$ -representable complex.*

Proof. Let \mathcal{F} be a collection of convex sets/good cover in \mathbb{R}^d such that K is isomorphic to the nerve of \mathcal{F} . Similarly \mathcal{F}' is a suitable collection in $\mathbb{R}^{d'}$ such that K' is isomorphic to the nerve of \mathcal{F}' .

Let us set

$$\mathcal{F} \star \mathcal{F}' := \{F \times \mathbb{R}^{d'} : F \in \mathcal{F}\} \cup \{\mathbb{R}^d \times F' : F' \in \mathcal{F}'\}.$$

Then it is easy to check that $K \star K'$ is isomorphic to the nerve of $\mathcal{F} \star \mathcal{F}'$. Moreover $\mathcal{F} \star \mathcal{F}'$ is a collection of convex sets/good cover in $\mathbb{R}^{d+d'}$. \square

Now we can finish the proof of our main result.

Proof of Theorem 8.1. Let T be the simplicial complex consisting of two isolated points. The complex T is topologically 1-representable and $\gamma_0(T) = 1$. Let us set

$$J = L \star \underbrace{T \star \cdots \star T}_{d-2}.$$

In topology, the complex J would be called $(d - 2)$ -tuple *suspension* of L . Then $\gamma_0(J) = d + 1$ due to Proposition 8.2 and Lemma 8.4. On the other hand, J is topologically d -representable due to Lemma 8.5. \square

8.3 Conclusion

In the spirit of Helly-type theorems we could ask whether there is at least some weaker bound for collapsibility of topologically d -representable complexes.

Question 8.6. *For which $d \geq 2$ there is a $d' \in \mathbb{N}$ (as least as possible) such that every topologically d -representable complex is d' -collapsible?*

Using joins of multiple copies of L (instead of suspensions of L) we obtain the following bound.

Proposition 8.7. *For every $d \geq 2$ there is a simplicial complex which is topologically $2d$ -representable but not $(3d - 1)$ -collapsible.*

²If A and B are sets with $A \cap B \neq \emptyset$ then their disjoint union can be defined as $A \sqcup B := A \times \{1\} \cup B \times \{2\}$.

Proof. Consider the complex $\underbrace{\mathbb{L} \star \cdots \star \mathbb{L}}_d$. □

If there is a wider gap among these notions it will also reflect at the gap between d -representable and d -Leray complexes obtained (with a similar method) by Matoušek and the author [MT09].

Chapter 9

Hardness of embeddability

9.1 Preliminaries on PL topology

Here we review definitions and facts related to piecewise linear (PL) embeddings. We begin with very standard things but later on we discuss notions and results which we found quite subtle (although they might be standard for specialists), in an area where it is sometimes tempting to consider as “obvious” something that is unknown or even false. Some more examples and open problems, which are not strictly necessary for the purposes of the reduction are mentioned in Appendix C of [MTW11]. For more information on PL topology, and for facts mentioned below without proofs, we refer to Rourke and Sanderson [RS82], Bryant [Bry02], or Buoncrisiano [Buo03]. We also refer to a survey paper by Repovš and Skopenkov [RS99] on embeddability problems.

Simplicial complexes. For purposes of this chapter we work with geometric simplicial complexes; however, we assume that the input to the embeddability problem is given as an abstract simplicial complex.

We recall that $V(K)$ is the set of vertices of a complex K , and $|K|$ denotes the geometric realization of K , i.e., the union of all simplices in K . Often we do not strictly distinguish between a simplicial complex and its polyhedron; for example, by an embedding of K in \mathbb{R}^d we really mean an embedding of $|K|$ into \mathbb{R}^d .

A simplicial complex K' is a *subdivision* of K if $|K'| = |K|$ and each simplex of K' is contained in some simplex of K .

Linear and PL mappings of simplicial complexes. A *linear* mapping of a simplicial complex K into \mathbb{R}^d is a mapping $f: |K| \rightarrow \mathbb{R}^d$ that is linear on each simplex. More explicitly, each point $x \in |K|$ is a convex combination $t_0v_0 + t_1v_1 + \dots + t_s v_s$, where $\{v_0, v_1, \dots, v_s\}$ is the vertex set of some simplex $\sigma \in K$ and t_0, \dots, t_s are nonnegative reals adding up to 1. Then we have $f(x) = t_0f(v_0) + t_1f(v_1) + \dots + t_s f(v_s)$.

A *PL mapping* of K into \mathbb{R}^d is a linear mapping of some subdivision K' of K into \mathbb{R}^d .

Embeddings. A general *topological embedding* of K into \mathbb{R}^d is any continuous mapping $f: |K| \rightarrow \mathbb{R}^d$ that is a homeomorphism of $|K|$ with $f(|K|)$. Since we only consider finite simplicial complexes, this is equivalent to requiring that f be injective.

By contrast, for a *PL embedding* we require additionally that f be PL, and for a *linear embedding* we are even more restrictive and insist that f be (simplexwise) linear.

PL embeddings versus linear embeddings. In contrast to planarity of graphs, lin-

ear and PL embeddability do not always coincide in higher dimensions Brehm [Bre83] constructed a triangulation of the Möbius strip that does not admit a linear embedding into \mathbb{R}^3 . Using methods from the theory of oriented matroids, Bokowski and Guedes de Oliveira [BGdO00] showed that for any $g \geq 6$, there is a triangulation of the orientable surface of genus g that does not admit a linear embedding into \mathbb{R}^3 . In higher dimensions, Brehm and Sarkaria [BS92] showed that for every $k \geq 2$, and every d , $k + 1 \leq d \leq 2k$, there is a k -dimensional simplicial complex K that PL embeds into \mathbb{R}^d but does not admit a linear embedding. Moreover, for any given $r \geq 0$, there is such a K such that even the r -fold barycentric subdivision $K^{(r)}$ is not linearly embeddable into \mathbb{R}^d . Corollary 4.2 is another result of this kind.

On the algorithmic side, the problem of *linear* embeddability of a given finite simplicial complex into \mathbb{R}^d is at least algorithmically decidable, and for k and d fixed, it even belongs to PSPACE (since the problem can easily be formulated as the solvability over the reals of a system of polynomial inequalities with integer coefficients, which lies in PSPACE [Ren92]).

PL structures. Two simplicial complexes K and L are *PL homeomorphic* if there are a subdivision K' of K and a subdivision L' of L such that K' and L' are isomorphic.

Let us recall that Δ_d denote the simplicial complex consisting of all faces of a d -dimensional simplex (including the simplex itself), and let $\partial\Delta_d$ consist of all faces of Δ_d of dimension at most $d-1$. Thus, $|\Delta_d|$ is topologically B^d , the d -dimensional ball, and $|\partial\Delta_d|$ is topologically S^{d-1} .

A d -dimensional *PL ball* is a simplicial complex PL homeomorphic to Δ_d , and a d -dimensional *PL sphere* is a simplicial complex PL homeomorphic to $\partial\Delta_{d+1}$. Let us mention that a (finite) simplicial complex K is PL embeddable in \mathbb{R}^d iff it is PL homeomorphic to a subcomplex of a d -dimensional PL ball (and similarly, K is PL embeddable in $|\partial\Delta_{d+1}|$ iff it is PL homeomorphic to a subcomplex of a d -dimensional PL sphere).

One of the great surprises in higher-dimensional topology was the discovery that simplicial complexes with homeomorphic polyhedra need not be PL homeomorphic (the failure of the “Hauptvermutung”). In particular, there exist *non-PL spheres*, i.e., simplicial complexes homeomorphic to a sphere that fail to be PL spheres. More precisely, every simplicial complex homeomorphic to S^1 , S^2 , S^3 , and S^4 is a PL sphere,¹ but there are examples of non-PL spheres of dimensions 5 and higher (e.g., the double suspension of the *Poincaré homology 3-sphere*).

A weak PL Schoenflies theorem. The well-known Jordan curve theorem states that if S^1 is embedded (topologically) in \mathbb{R}^2 , the complement of the image has exactly two components. Equivalently, but slightly more conveniently, if S^1 is embedded in S^2 , the complement has two components. The *Schoenflies theorem* asserts that in the latter setting, the closure of each of the components is homeomorphic to the disk B^2 .

While the Jordan curve theorem generalizes to an arbitrary dimension (if S^{d-1} is topologically embedded in S^d , the complement has exactly two components), the Schoenflies theorem does not. There are embeddings $h: S^2 \rightarrow S^3$ such that the closure of one of the components of $S^3 \setminus h(S^2)$ is not a ball; a well known example is the *Alexander horned sphere*.

The Alexander horned sphere is an infinitary construction; one needs to grow infinitely many “horns” from the embedded S^2 to make the example work. In higher

¹The proof for S^4 relies on the recent solution of the Poincaré conjecture by Perelman.

dimensions, there are strictly finite examples, e.g., a 5-dimensional subcomplex \mathbf{K} of a 6-dimensional PL sphere S such that $|\mathbf{K}|$ is topologically an S^5 (and \mathbf{K} is a non-PL sphere), but the closure of a component of $|S| \setminus |\mathbf{K}|$ is not a topological ball (see Curtis and Zeeman [CZ61]).

Thus, one needs to put some additional conditions on the embedding to make a “higher-dimensional Schoenflies theorem” work. We will need the following version, in which we assume a $(d - 1)$ -dimensional PL sphere sitting in a d -dimensional PL sphere.

Theorem 9.1 (Weak PL Schoenflies Theorem). *Let f be a PL embedding of $\partial\Delta^d$ into $\partial\Delta^{d+1}$. Then the complement $|\partial\Delta^{d+1}| \setminus f(|\partial\Delta^d|)$ has two components, whose closures are topological d -balls.*

For a proof of this theorem, see, e.g., [New60] or [Gla71]. A simple, inductive proof is to appear in the upcoming revised edition of the book [Buo03] by Buoncrisiano and Rourke.

Let us remark that a “strong” PL Schoenflies theorem would claim that under the conditions of Theorem 9.1, the closure of each of the components is a PL ball, but the validity of this stronger statement is known only for $d \leq 3$, while for each $d \geq 4$ it is (to our knowledge) an open problem.

Genericity. First let us consider a linear mapping f of a simplicial complex \mathbf{K} into \mathbb{R}^d . We say that f is *generic* if $f(V(\mathbf{K}))$ is a set of distinct points in \mathbb{R}^d in general position. If $\sigma, \tau \in \mathbf{K}$ are disjoint simplices, then the intersection $f(\sigma) \cap f(\tau)$ is empty for $\dim \sigma + \dim \tau < d$ and it has at most one point for $\dim \sigma + \dim \tau = d$.

A PL mapping of \mathbf{K} into \mathbb{R}^d is generic if the corresponding linear mapping of the subdivision \mathbf{K}' of \mathbf{K} is generic.

A PL embedding can always be made generic (by an arbitrarily small perturbation).

Linking and linking numbers. Let k, ℓ be integers, and let $f: S^k \rightarrow \mathbb{R}^{k+\ell+1}$ and $g: S^\ell \rightarrow \mathbb{R}^{k+\ell+1}$ be PL embeddings with $f(S^k) \cap g(S^\ell) = \emptyset$ (so here we regard S^k and S^ℓ as PL spheres). We will need two notions capturing how the images of f and g are “linked” (the basic example is $k = \ell = 1$, where we deal with two disjoint simple closed curves in \mathbb{R}^3). For our purposes, we may assume that f and g are mutually generic (i.e. $f \sqcup g$, regarded as a PL embedding of the disjoint union $S^k \sqcup S^\ell$ into $\mathbb{R}^{k+\ell+1}$, is generic).

The images $f(S^k)$ and $g(S^\ell)$ are *unlinked* if f can be extended to a PL mapping $\bar{f}: B^{k+1} \rightarrow \mathbb{R}^{k+\ell+1}$ of the $(k + 1)$ -dimensional ball such that $\bar{f}(B^{k+1}) \cap g(S^\ell) = \emptyset$.

To define the *modulo 2 linking number* of $f(S^k)$ and $g(S^\ell)$, we again extend f to a PL mapping $\bar{f}: B^{k+1} \rightarrow \mathbb{R}^{k+\ell+1}$ so that \bar{f} and g are still mutually generic (but otherwise arbitrarily). Then the modulo 2 linking number is the number of intersections between $\bar{f}(B^{k+1})$ and $g(S^\ell)$ modulo 2 (it turns out that it does not depend on the choice of \bar{f}). In the sequel, we will use the phrase “odd linking number” instead of the more cumbersome “nonzero linking number modulo 2” (although “linking number” in itself has not been properly defined).

These geometric definitions are quite intuitive. However, alternative (equivalent in our setting but more generally applicable) definitions are often used, phrased in terms of homology or mapping degree, which are in some respects easier to work with (e.g., they show that linking is symmetric, i.e., $f(S^k)$ and $g(S^\ell)$ are unlinked iff $g(S^\ell)$ and $f(S^k)$ are unlinked).

9.2 Undecidability: Proof of Theorem 4.1

We begin with a statement of Novikov's result mentioned in the introduction (undecidability of S^d recognition for $d \geq 5$) in a form convenient for our purposes.

Theorem 9.2 (Novikov). *Fix $d \geq 5$. There is an effectively constructible sequence of simplicial complexes Σ_i , $i \in \mathbb{N}$, with the following properties:*

- (1) *Each $|\Sigma_i|$ is a homology d -sphere.*
- (2) *For each i , either Σ_i is a PL d -sphere, or the fundamental group of Σ_i is non-trivial (in particular, Σ_i is not homeomorphic to the d -sphere).*
- (3) *There is no algorithm that decides for every given Σ_i which of the two cases holds.*

We refer to the appendix in [Nab95] for a detailed proof. We begin the proof of Theorem 4.1 with the following simple lemma.

Lemma 9.3. *Let Σ be a simplicial complex whose polyhedron is a homology d -sphere, $d \geq 2$. (The same proof works for any homology d -manifold.) Let \mathbf{K} be the $(d-1)$ -skeleton of Σ . For every d -simplex $\sigma \in \Sigma$, the set $|\mathbf{K}| \setminus \partial\sigma$ is path connected (here $\partial\sigma$ is the relative boundary of σ).*

Proof. By Lefschetz duality (see, e.g., [Mun84, Theorem 70.2]), $|\Sigma| \setminus \sigma$ is path connected. Indeed, Lefschetz duality yields $H^0(|\Sigma| \setminus \sigma) \cong H_d(\Sigma, \sigma)$ (homology with \mathbb{Z}_2 coefficients, say). The exact homology sequence of the pair (Σ, σ) , together with the fact that σ is contractible, yields $H_d(\Sigma) \cong H_d(\Sigma, \sigma) \cong \mathbb{Z}_2$.

Next, we claim that if γ is a path in $|\Sigma| \setminus \sigma$ connecting two points $x, y \in |\mathbf{K}|$, then x and y can also be connected by a path in $|\mathbf{K}| \setminus \partial\sigma$. Indeed, given a d -dimensional simplex $\tau \in \Sigma \setminus \sigma$, we have $\partial\tau \setminus \sigma$ path-connected. Hence we can modify γ as follows: Letting $a := \min\{t : \gamma(t) \in \tau\}$ and $b := \max\{t : \gamma(t) \in \tau\}$, we replace the segment of γ between $\gamma(a)$ and $\gamma(b)$ by a path η in $\partial\tau \setminus \sigma$. Having performed this modification for every $\tau \in \Sigma \setminus \sigma$ (in some arbitrary order), we end up with a path connecting x and y that lies entirely within $|\mathbf{K}| \setminus \partial\sigma$. \square

Lemma 9.4. *Let $d \geq 2$. Suppose that Σ is a homology d -sphere, and let \mathbf{K} be its $(d-1)$ -skeleton.*

- (i) *If Σ is a PL sphere, then \mathbf{K} PL embeds into \mathbb{R}^d .*
- (ii) *If \mathbf{K} PL embeds into \mathbb{R}^d , then Σ is homeomorphic to S^d .*

Proof. Part (i) is clear.

For part (ii), let us suppose that f is a PL embedding of \mathbf{K} into \mathbb{R}^d . Since \mathbf{K} is compact, the image of f is contained in some big d -dimensional simplex, and by taking this simplex as one facet of Δ^{d+1} , we can consider f as a PL embedding of \mathbf{K} into $\partial\Delta^{d+1}$. Consider a d -simplex σ of Σ . By the weak PL Schoenflies theorem (Theorem 9.1), $|\partial\Delta^{d+1}| \setminus f(\partial\sigma)$ has two components, whose closures are topological d -balls. Moreover, since $|\mathbf{K}| \setminus \partial\sigma$ is path-connected, its image under f must be entirely contained in one of these components.

Therefore, we can use the closure of the other component to extend f to a topological embedding of σ . By applying this reasoning to each d -face, we obtain a topological embedding g of Σ into $\partial\Delta^{d+1}$. It follows for instance from Alexander duality (see, e.g., [Mun84, Theorem 74.1]) that g must be surjective, i.e., a homeomorphism. \square

Proof of Theorem 4.1. The undecidability of $\text{EMBED}_{(d-1) \rightarrow d}$ for $d \geq 5$ is an immediate consequence of Theorem 9.2 and Lemma 9.4. \square

Proof of Corollary 4.2. Let us suppose that there is a recursive function f contradicting the statement. That is, every $(d-1)$ -dimensional K with n simplices that PL-embeds in \mathbb{R}^d at all has a subdivision with at most $f(n)$ simplices that embeds linearly. Then, given a $(d-1)$ -dimensional complex K with n simplices, we could generate all subdivisions K' of K with at most $f(n)$ simplices (see Acquistapace et al. [ABB90], Proposition 2.15) and, using the PSPACE algorithms mentioned in Section 9.1, test the linear embeddability of each K' in \mathbb{R}^d . This would yield a decision algorithm for $\text{EMBED}_{(d-1) \rightarrow d}$, contradicting Theorem 4.1. \square

9.3 Hardness of embedding 2-dimensional complexes in \mathbb{R}^4

We will reduce the problem 3-SAT to $\text{EMBED}_{2 \rightarrow 4}$. Given a 3-CNF formula φ , we construct a 2-dimensional simplicial complex K that is PL embeddable in \mathbb{R}^4 exactly if φ is satisfiable.

First we define two particular 2-dimensional simplicial complexes \mathbf{G} (the *clause gadget*) and \mathbf{X} (the *conflict gadget*). They are closely related to the main example of Freedman et al. [FKT94]: \mathbf{X} is taken over exactly, and \mathbf{G} is a variation on a construction in [FKT94] (which, in turn, is similar in some respects to an example of Segal and Spieź [SS92], with some of the ideas going back to Van Kampen [vK32]).

9.3.1 The clause gadget

To construct \mathbf{G} , we begin with a 6-dimensional simplex on the vertex set $\{v_0, v_1, \dots, v_6\}$, and we let \mathbf{F} be the 2-skeleton of this simplex (\mathbf{F} for “full” skeleton). Then we make a hole in the interior of the three triangles (2-simplices) $v_0v_1v_2$, $v_0v_1v_3$, and $v_0v_2v_3$. That is, we subdivide each of the triangles and from each of these subdivisions we remove a small triangle in the middle, as is indicated in Fig. 9.1.²

This yields the simplicial complex \mathbf{G} .

Let $\omega_1, \omega_2, \omega_3$ be the three small triangles we have removed (where ω_1 comes from the triangle $v_0v_2v_3$ etc.). We call them the *openings* of \mathbf{G} and we let $O_{\mathbf{G}} := \{\omega_1, \omega_2, \omega_3\}$ be the set of openings. Thus, $\mathbf{G} \cup O_{\mathbf{G}}$ is a subdivision of the full 2-skeleton \mathbf{F} .

If we remove from \mathbf{F} the vertices v_0, v_1, v_2 and all simplices containing them, we obtain the boundary of the 3-simplex $\{v_3, v_4, v_5, v_6\}$. Topologically it is an S^2 , we call it the *complementary sphere* of the opening ω_3 , and we denote it by S_{ω_3} . The complementary spheres of the openings ω_1 and ω_2 are defined analogously. The following lemma is a variation on results in Van Kampen [vK32]:

Lemma 9.5.

- (i) *For every generic PL embedding f of \mathbf{G} into \mathbb{R}^4 there is at least one opening $\omega \in O_{\mathbf{G}}$ such that the images of the boundary $\partial\omega$ and of the complementary sphere S_{ω} have odd linking number.*

²Alternatively, we could also make the clause gadget by simply removing the the triangles $v_0v_1v_2$, $v_0v_1v_3$, and $v_0v_2v_3$ from \mathbf{F} . However, the embedding of the resulting complex \mathbf{K} for satisfiable formulas φ would become somewhat more complicated.




Figure 9.1: The clause gadget G , its openings, and one of the complementary spheres.

- (ii) For every opening $\omega \in O_G$ there exists an embedding of G into \mathbb{R}^4 in which only $\partial\omega$ is linked with its complementary sphere. More precisely, there exists a generic linear mapping of the full 2-skeleton F into \mathbb{R}^4 whose restriction to $|G \cup O_G \setminus \{\omega\}|$ is an embedding.

Proof of (i). This is very similar to Lemma 6 in [FKT94]. Let f_0 be a generic PL map (not necessarily an embedding) of \mathcal{F} into \mathbb{R}^4 . Van Kampen proved that

$$\sum_{\{\sigma, \tau\}} |f_0(\sigma) \cdot f_0(\tau)|$$

is always odd, where $|f_0(\sigma) \cdot f_0(\tau)|$ denotes the number of intersections between the image of σ and the image of τ , and the sum is over all unordered pairs of *disjoint* 2-dimensional simplices $\sigma, \tau \in F$ (the genericity of f_0 guarantees that the intersection $f_0(\sigma) \cap f_0(\tau)$ consists of finitely many points). (See Appendix D of [MTW11] for a wider context of this result.)

Now let us consider a generic PL embedding f of G into \mathbb{R}^4 , and let us extend it piecewise linearly and generically (and otherwise arbitrarily) to the openings of G . The resulting map can also be regarded as a generic PL map f_0 of F into \mathbb{R}^4 . For such an f_0 , $|f_0(\sigma) \cdot f_0(\tau)|$ can be nonzero only if σ contains an opening ω of G and τ belongs to its complementary sphere S_ω (or the same situation with σ and τ interchanged). Thus, for at least one $\omega \in O_G$, $f_0(\omega)$ intersects $f(S_\omega)$ in an odd number of points, and this means exactly that $f(\partial\omega)$ and $f(S_\omega)$ have odd linking number. \square

Proof of (ii). It suffices to exhibit a generic linear map f_0 of F into \mathbb{R}^4 such that the images of two disjoint 2-simplices intersect (at a single point), and this intersection is the only multiple point of f_0 . Such a mapping was constructed by Van Kampen [vK32]: 5 of the vertices are placed as vertices of a 4-dimensional simplex in \mathbb{R}^4 , and the remaining two are mapped in the interior of that simplex. \square




Figure 9.2: Attaching a disk to the polygonal line E. MODIFY D!




Figure 9.3: A 3-dimensional embedding of the conflict gadget.

9.3.2 The conflict gadget

To construct X , we start with the 1-dimensional simplicial complex E shown in Fig. 9.2 left, consisting of two triangular loops Σ_a and Σ_b and an edge c connecting them. We also fix an orientation of Σ_a , Σ_b , and c (marked by arrows). Then we take a disk D and we attach its boundary to E as indicated in Fig. 9.2 right; the disk is triangulated sufficiently finely so that the result of the attachment is still a simplicial complex. This is the complex X .

We observe that topologically, X is a “squeezed torus” (the reader may want to recall the usual construction of a torus by gluing the opposite sides of a square; this well-known construction would be obtained from the attachment as above if the edge c were contracted to a point). Fig. 9.3 shows such a squeezed torus embedded in \mathbb{R}^3 (with the loops Σ_a and Σ_b drawn circular rather than triangular).

Lemma 9.6.

- (i) [FKT94, Lemma 7] *Let S_a and S_b be PL 2-spheres. Then there is no PL embedding f of $S_a \sqcup S_b \sqcup X$ (disjoint union) into \mathbb{R}^4 such that*
 - *the 1-sphere $f(\Sigma_a)$ and the 2-sphere $f(S_a)$ have odd linking number, and so do $f(\Sigma_b)$ and $f(S_b)$;*
 - *$f(\Sigma_a)$ and $f(S_b)$ are unlinked, and so are $f(\Sigma_b)$ and $f(S_a)$.*
- (ii) *Let f be a generic linear embedding of E in \mathbb{R}^3 (not \mathbb{R}^4 this time) such that $f(\Sigma_a)$ and $f(\Sigma_b)$ are unlinked, and let $\delta > 0$. Then there is a PL embedding*

\bar{f} of \mathbf{X} in \mathbb{R}^3 extending f whose image is contained in the set $N = N(f, \delta) := N(T_a, \delta) \cup N(f(\Sigma_b), \delta) \cup N(f(c), \delta)$, where T_a is the triangle bounded by the loop $f(\Sigma_a)$ and $N(A, \delta)$ denotes the δ -neighborhood of a set A (in \mathbb{R}^3 in our case).³ (Symmetrically, and this is the main point of the construction, we can also embed \mathbf{X} into $N(f(\Sigma_a), \delta) \cup N(T_b, \delta) \cup N(f(c), \delta)$, thus leaving a hole on the other side.)

For a proof of part (i) we refer to (a few words about the basic approach of the proof will be said in the proof of Lemma 9.9 below), and for part (ii) to Fig. 9.3.

9.3.3 The reduction

Let the given 3-CNF formula be $\varphi = C_1 \wedge C_2 \wedge \cdots \wedge C_m$, where each C_i is a clause with three literals (each literal is either a variable or its negation). For each C_i , we take a copy of the clause gadget \mathbf{G} and we denote it by \mathbf{G}_i (the \mathbf{G}_i have pairwise disjoint vertex sets). We fix a one-to-one correspondence between the literals of C_i and the openings of \mathbf{G}_i , letting $\omega(\lambda)$ be the opening corresponding to a literal λ .

Let us say that a literal λ in a clause C_i is in conflict with a literal μ in a clause C_j if both λ and μ involve the same variable x but one of them is x and the other the negation \bar{x} . For convenience we assume, without loss of generality, that two literals from the same clause are never in conflict.

Let Ξ consist of all (unordered) pairs $\{\omega(\lambda), \omega(\mu)\}$ of openings corresponding to pairs $\{\lambda, \mu\}$ of conflicting literals in φ . For every pair $\{\omega, \psi\} \in \Xi$ we take a fresh copy $\mathbf{X}_{\omega\psi}$ of the conflict gadget \mathbf{X} . We identify the loop Σ_a in $\mathbf{X}_{\omega\psi}$ with the boundary $\partial\omega$ and the loop Σ_b with $\partial\psi$ (the rest of $\mathbf{X}_{\omega\psi}$ is disjoint from the clause gadgets and the other conflict gadgets).

The simplicial complex \mathbf{K} assigned to the formula φ is

$$\mathbf{K} := \left(\bigcup_{i=1}^m \mathbf{G}_i \right) \cup \left(\bigcup_{\{\omega, \psi\} \in \Xi} \mathbf{X}_{\omega\psi} \right).$$

It remains to show that \mathbf{K} is PL embeddable in \mathbb{R}^4 exactly if φ is satisfiable.

Nonembeddability for unsatisfiable formulas. This is a straightforward consequence of Lemma 9.5(i) and Lemma 9.6(i).

Indeed, if f is a PL embedding of \mathbf{K} into \mathbb{R}^4 , which we may assume to be generic, there is an opening in each clause gadget \mathbf{G}_i such that $f(\partial\omega_i)$ has odd linking number with the complementary sphere $f(S_{\omega_i})$; let us call it a *occupied opening* of \mathbf{G}_i . Since φ is not satisfiable, whenever we choose one literal from each clause, there are two of the chosen literals in conflict. Thus, there are two occupied openings $\omega \in O_{\mathbf{G}_i}$ and $\psi \in O_{\mathbf{G}_j}$ that are connected by a conflict gadget $\mathbf{X}_{\omega\psi}$.

Then the supposed PL embedding f provides us an embedding as in Lemma 9.6(i) with $S_a = S_\omega$, $S_b = S_\psi$, and $\mathbf{X} = \mathbf{X}_{\omega\psi}$. Concerning the assumptions in the lemma, we already know that $f(S_\omega)$ and $f(\partial\omega)$ have odd linking number, and so do $f(S_\psi)$ and $f(\partial\psi)$. It remains to observe that $f(\partial\omega)$ cannot be linked with $f(S_\psi)$ (and vice versa), since \mathbf{G}_i contains a disk bounded by $\partial\omega$: For example (refer to Fig. 9.1), $\partial\omega_3$ is the boundary of the disk consisting of the triangles $v_0v_1v_4$, $v_0v_2v_4$, $v_1v_2v_4$ and the

³Formally $N(A, \delta) = \{x \in \mathbb{R}^3 : \text{dist}(x, A) \leq \delta\}$, where $\text{dist}(x, A)$ is the Euclidean distance of x from the set A .

triangles in the subdivision of $v_0v_1v_2$ different from ω_3 . So the lemma applies and \mathbf{K} is not embeddable.

Embedding for satisfiable formulas. Given a satisfying assignment for φ , we choose a *witness literal* λ_i for each clause C_i that is true under the given assignment (and we will refer to the remaining two literals of C_i as *non-witness* ones). No two witness literals can be in conflict.

We describe an embedding of \mathbf{K} into \mathbb{R}^4 corresponding to this choice of witness literals.

Let us choose distinct points $p_1, \dots, p_m \in \mathbb{R}^4$. For each $i = 1, 2, \dots, m$, we let f_i be a generic linear embedding of the clause gadget \mathbf{G}_i into a small neighborhood of p_i (and far from the other p_j) as in Lemma 9.5(ii), where the role of ω in the lemma is played by the witness opening of \mathbf{G}_i (i.e., the one corresponding to the witness literal of C_i). In particular, the interiors of the triangles bounded by $f_i(\partial\omega')$ and by $f_i(\partial\omega'')$ are disjoint from $f_i(\mathbf{G}_i)$, where ω' and ω'' are the non-witness openings of \mathbf{G}_i .

Taking all the f_i together defines an embedding f of the union of the clause gadgets, and it remains to embed the conflict gadgets.

To this end, we will assign to each conflict gadget $\mathbf{X}_{\omega\psi}$ a “private” set $P_{\omega\psi} \subset \mathbb{R}^4$ homeomorphic to the 3-dimensional set N from Lemma 9.6(ii), and we will embed $\mathbf{X}_{\omega\psi}$ into $P_{\omega\psi}$. Each $P_{\omega\psi}$ will be disjoint from all other $P_{\omega'\psi'}$ and also from all the images $f(\mathbf{G}_i)$, *except* that $P_{\omega\psi}$ has to contain the loops $f(\partial\omega)$ and $f(\partial\psi)$ where the conflict gadget $\mathbf{X}_{\omega\psi}$ should be attached. In order to fit enough almost-disjoint homeomorphic copies of N into the space, we will “fold” them suitably.

We know that for every pair $\{\omega, \psi\}$ of openings connected by a conflict gadget, at least one of ω and ψ is non-witness. Let us choose the notation so that ω is non-witness and thus unoccupied in the embedding f .

We will build $P_{\omega\psi}$ from three pieces: a set $Q_{\omega\psi}^+$ that plays the role of $N(T_a, \delta)$ in Lemma 9.6(ii), a set $Q_{\psi\omega}$ that plays the role of $N(f(\Sigma_b), \delta)$, and a “connecting ribbon” in the role of $N(f(c), \delta)$.

Now let ω be an opening of some \mathbf{G}_i , witness or non-witness. Let t be the number of openings ψ that are connected to ω by a conflict gadget. The sets $Q_{\omega\psi}$ and $Q_{\omega\psi}^+$ we want to construct are indexed by these ψ , but with some abuse of notation, we will now regard them as indexed by an index j running from 1 to t , i.e., as $Q_{\omega 1}$ through $Q_{\omega t}$ (and similarly for $Q_{\omega\psi}^+$).

For concise notation let us write $\Sigma = f(\partial\omega)$ and let T be the triangle in \mathbb{R}^4 having Σ as the boundary. Let $\varepsilon > 0$ be a parameter and let $T^\varepsilon := \{x \in T : \text{dist}(x, \partial T) \leq \varepsilon\}$ be the part of T at most ε away from the boundary of T . Since the subdivided triangle in \mathbf{G}_i containing ω in its interior is embedded linearly by f , there is an $\varepsilon > 0$ such that if we start at a point $x \in T^\varepsilon$ and go distance at most ε in a direction orthogonal to T , we do not hit $f(\mathbf{G}_i)$. Moreover, if ω is non-witness and thus all of T is free of $f(\mathbf{G}_i)$, we can take any $x \in T$ with the same result. Fig. 9.4 tries to illustrate this in dimension one lower, where we have a segment T in \mathbb{R}^3 instead of a triangle T in \mathbb{R}^4 . Thus, there are a set $Q_\omega \subset \mathbb{R}^4$ with $Q_\omega \cap f(\mathbf{G}_i) = \Sigma$ and a homeomorphism (actually, a linear isomorphism) $h: Q_\omega \rightarrow T^\varepsilon \times B^2$ with $h(T^\varepsilon) = T^\varepsilon \times \{0\}$, where 0 is the center of the disk B^2 . Similarly, if ω is witness, there are Q_ω^+ and $h^+: Q_\omega^+ \rightarrow T \times B^2$ with $h^+(T) = T \times \{0\}$.

Let $W_1, \dots, W_t \subset B^2$ be disjoint wedges as in Fig. 9.5, and let w_j consist of the




Figure 9.4: A free region around the triangle T ; illustration in \mathbb{R}^3 instead of \mathbb{R}^4 .




Figure 9.5: The wedges.

two radii bounding W_j . We set

$$Q_{\omega_j} := h^{-1}((\Sigma \times W_j) \cup (T^\epsilon \times w_j)), \quad Q_{\omega_j}^+ := (h^+)^{-1}((\Sigma \times W_j) \cup (T \times w_j)).$$

As Fig. 9.6 tries to illustrate, $Q_{\omega_j}^+$ is homeomorphic to a 3-dimensional neighborhood of T (by a homeomorphism sending T to T), and Q_{ω_j} is similarly homeomorphic to a 3-dimensional neighborhood of Σ . Thus, the sets Q_{ω_j} and $Q_{\omega_j}^+$ can indeed play the roles of $N(f(\Sigma_b), \delta)$ and $N(T_a, \delta)$, respectively, in Lemma 9.5(ii).

It remains to construct the “connecting ribbons”: For every conflict gadget $X_{\omega\psi}$, we want to connect a vertex of $f(\partial\omega)$ to a vertex of $f(\partial\psi)$ by a narrow 3-dimensional “ribbon” (it need not be straight since we are looking only for PL homeomorphic




Figure 9.6: Folding a 3-dimensional neighborhood in \mathbb{R}^4 .




Figure 9.7: A schematic illustration of $F(3, 1)$.

copies of N).

We observe that each of the sets Q_{ω_j} and $Q_{\omega_j}^+$ can be deformation-retracted to the corresponding loop $f(\partial\omega)$ or to the corresponding triangle, respectively. It follows that the complement of the union U of all the Q_{ω_j} , $Q_{\omega_j}^+$, and $f(G_i)$ is path-connected (formally, this follows from Alexander duality, since this union is homotopy equivalent to a 2-dimensional space). Since all the considered embeddings are piecewise linear, any two points on the boundary of U can be connected by a PL path within $\mathbb{R}^4 \setminus U$.

Thus, the 3-dimensional “ribbon” connecting $f(\partial\omega)$ to $f(\partial\psi)$ can first go within the appropriate Q_{ω_j} to a point on the boundary, then continue along a path connecting this boundary point to a boundary point of $Q_{\psi_{j'}}$, and then reach $f(\partial\psi)$ within $Q_{\psi_{j'}}$.

In this way, we have allocated the desired “private” sets $P_{\omega\psi}$ for all conflict gadgets $X_{\omega\psi}$, and hence K can be PL embedded in \mathbb{R}^4 as claimed. This finishes the proof of the special case $k = 2, d = 4$ of Theorem 4.3. \square

9.4 NP-hardness for higher dimensions

In this section we prove all the remaining cases of Theorem 4.3. The proof is generally very similar to the case $k = 2, d = 4$ treated above: We will again reduce 3-SAT using clause gadgets and conflict gadgets, but the construction of the gadgets and of their embeddings require additional work.

By the monotonicity of $\text{EMBED}_{k \rightarrow d}$ in k mentioned in Chapter 4, it suffices to consider $d \geq 5$ and $k = \lceil (2d - 2)/3 \rceil$. In the construction we will often use the integer $\ell := d - k - 1$.

9.4.1 The clause gadget

The clause gadget $G = G(k, \ell)$ is very similar to a construction of Segal and Spiež [SS92]. We use the parameters k, ℓ, d as above. For the purposes of the present section we need that $1 \leq \ell < k$ and $d - \ell = k + 1 \geq 3$ (which are easy to verify using the definitions of k and ℓ and the assumption $d \geq 5$).

For the parameters k, ℓ, d as above, we first define a simplicial complex $F = F(k, \ell)$ on the vertex set $V := \{v_0, v_1, \dots, v_{d+1}, p\}$ as the union $F := F_0 \cup C_p$ of the following two sets of simplices:

- F_0 is the k -skeleton of the $(d + 1)$ -simplex with vertex set $\{v_0, \dots, v_{d+1}\}$;
- C_p consists of all the $(\ell + 1)$ -dimensional simplices on V that contain p .

See Fig. 9.7 for a schematic illustration; let us also note that for $d = 4, k = 2, \ell = 1$ we would get exactly the F as in Section 9.3.1.

Let us consider some $\sigma \in C_p$. By removing from \mathcal{F} all simplices intersecting σ (including σ), we obtain the k -skeleton of a $(k+1)$ -simplex, i.e., an S^k , which we call the *complementary sphere* S_σ .

Next, we fix three $(\ell+1)$ -dimensional simplices $\sigma_1, \sigma_2, \sigma_3 \in C_p$, say $\sigma_1 := pv_0v_2v_3 \cdots v_{\ell+1}$, $\sigma_2 := pv_0v_1v_3 \cdots v_{\ell+1}$, and $\sigma_3 := pv_0v_1v_2v_4 \cdots v_{\ell+1}$. As in Section 9.3.1, we make a hole in the interior of each σ_i , i.e., we subdivide each σ_i , $i = 1, 2, 3$, and we remove a small $(\ell+1)$ -simplex ω_i in the middle. This yields the simplicial complex $\mathbf{G} = \mathbf{G}(k, \ell)$.

The ω_i are again called the *openings* of \mathbf{G} , and we set $O_{\mathbf{G}} := \{\omega_1, \omega_2, \omega_3\}$. The *complementary sphere* S_{ω_i} is defined, with some abuse of notation, as the complementary sphere of the simplex $\sigma_i \in C_p$ that contains ω_i .

Lemma 9.7 (Higher-dimensional version of Lemma 9.5).

- (i) For every generic PL embedding f of \mathbf{G} into \mathbb{R}^d there is at least one opening $\omega \in O_{\mathbf{G}}$ such that the images of the boundary $\partial\omega$ and of the complementary sphere S_ω have odd linking number.
- (ii) For every opening $\omega \in O_{\mathbf{G}}$ there exists a generic linear embedding of \mathbf{G} into \mathbb{R}^d in which the boundaries of the two openings different from ω are unlinked with their complementary spheres.

Proof. Part (ii) is established in the proof of Lemma 1.1 in Segal and Spieź [SS92] (generalizing Van Kampen's embedding mapping mentioned in the proof of Lemma 9.5(ii)). They construct a PL embedding of $\mathbf{F}(k, \ell)$ (which they call $P(k, \ell)$, while their n is our $d-1$), but inspecting the first two paragraphs of their proof reveals that their embedding is actually linear (in the subsequent paragraphs, they modify the embedding on the interior of one of the $(\ell+1)$ -simplices from C_p , but this serves only to show the claim about linking number).

For part (i), it clearly suffices to prove the following:

Claim. For any generic PL mapping g of \mathbf{F} in \mathbb{R}^d whose restriction to \mathbf{F}_0 is an embedding, there is an $(\ell+1)$ -dimensional simplex $\sigma \in C_p$ such that $|g(\sigma) \cap g(S_\sigma)|$ is odd.


This claim follows easily from the *proof* of Lemma 1.4 in Segal and Spieź [SS92]. Indeed, they give a procedure that, given a generic PL map g_1 of \mathbf{F} into \mathbb{R}^d such that $|g_1(\sigma) \cap g_1(S_\sigma)|$ is even for some σ , constructs a new generic PL map g_2 with $g_2(\sigma) \cap g_2(S_\sigma) = \emptyset$ and such that there are no new intersections between images of *disjoint* simplices (compared to g_1).⁴

Assuming that there is a g contradicting the claim, after finitely many applications of the procedure we arrive at a generic PL mapping \tilde{g} such that $\tilde{g}(\sigma) \cap \tilde{g}(S_\sigma) = \emptyset$ for every $\sigma \in C_p$. We claim that then

$$\tilde{g}(\tau) \cap \tilde{g}(\tau') = \emptyset \text{ for every } \tau, \tau' \in \mathbf{F} \text{ with } \tau \cap \tau' = \emptyset. \quad (9.1)$$

Indeed, if (9.1) fails for some τ, τ' , one of τ, τ' (say τ') must belong to C_p , since g restricted to \mathbf{F}_0 is an embedding. But then we get $\tau \in S_{\tau'}$ —a contradiction. Hence (9.1) holds. But no generic PL mapping \tilde{g} satisfying (9.1) exists according to [SS92]

⁴The procedure requires $d - \ell \geq 3$, which is satisfied in our case. In [SS92] this inequality is reversed by mistake.


 Figure 9.8: A higher-dimensional version of \mathbf{E} .

(end of the proof of Lemma 1.4). This proves the claim and thus also part (i) of Lemma 9.7.

Let us remark that a perhaps more conceptual proof of part (i) can be obtained using the results of Shapiro [Sha57] on the “generalized Van Kampen obstruction”, but we would need many preliminaries for presenting it. \square

9.4.2 The conflict gadget

Here we construct the conflict gadget $\mathbf{X} = \mathbf{X}(\ell)$, which depends only on the parameter ℓ , and whose dimension is 2ℓ . The conflict gadget \mathbf{X} in Section 9.3.2 is essentially the same as the following construction for $\ell = 1$, up to minor formal differences. In addition to the inequalities among the parameters mentioned earlier, here we also need $2\ell \leq k$ (which again holds in our setting).

In the $\ell = 1$ case we attached a 2-dimensional disk by its boundary to the 1-dimensional complex \mathbf{E} . The ℓ -dimensional version of \mathbf{E} consists of two disjoint copies Σ_a^ℓ and Σ_b^ℓ of the boundary of the $(\ell + 1)$ -simplex connected by an edge c (see Fig. 9.8). To this \mathbf{E} we are going to attach the (2ℓ) -dimensional ball $B^{2\ell}$ by its boundary. For $\ell = 1$ the result was topologically a “squeezed” version of the 2-dimensional torus $S^1 \times S^1$; for larger ℓ it is going to be the higher-dimensional “torus” $S^\ell \times S^\ell$, again suitably squeezed.

Attaching a ball to $S^\ell \vee S^\ell$. Before defining \mathbf{X} itself, we define a certain mapping $g: S^{2\ell-1} \rightarrow S^\ell \vee S^\ell$, where $S^\ell \vee S^\ell$ is a wedge of two spheres, to be defined below. This construction is based on the *Whitehead product* in homotopy theory. As we will see, attaching the boundary of $B^{2\ell}$ to $S^\ell \vee S^\ell$ via g results topologically in $S^\ell \times S^\ell$ (without any squeezing).

The wedge $S^\ell \vee S^\ell$ consists of two copies of the sphere S^ℓ glued together at one point. For our purposes, we represent $S^\ell \vee S^\ell$ concretely as follows. We consider S^ℓ geometrically as the unit sphere in $\mathbb{R}^{\ell+1}$, we choose a distinguished point $s_0 = (1, 0, 0, \dots, 0) \in S^\ell$, and we let $S^\ell \vee S^\ell$ be the subspace $(S^\ell \times \{s_0\}) \cup (\{s_0\} \times S^\ell)$ of $\mathbb{R}^{\ell+1} \times \mathbb{R}^{\ell+1} = \mathbb{R}^{2\ell+2}$. For $\ell = 1$, we thus get two unit circles lying in perpendicular 2-flats in \mathbb{R}^4 and meeting at the point (s_0, s_0) .

For defining the map g , we need to represent the ball $B^{2\ell}$ not as the standard Euclidean unit ball, but rather as the product $B^\ell \times B^\ell$ (which is clearly homeomorphic to $B^{2\ell}$). Then we have

$$S^{2\ell-1} \cong \partial(B^\ell \times B^\ell) = (B^\ell \times S^{\ell-1}) \cup (S^{\ell-1} \times B^\ell); \quad (9.2)$$

see the left part of Fig. 9.9 for the (rather trivial) case $\ell = 1$. (Indeed, for arbitrary sets $A \subseteq \mathbb{R}^m$ and $B \subseteq \mathbb{R}^n$ we have $\partial(A \times B) = (A \times \partial B) \cup (\partial A \times B)$, as is easy to check.)




Figure 9.9: Representing S^1 as $(B^1 \times S^0) \cup (S^0 \times B^1)$ (left); mapping it to $S^1 \vee S^1$ (middle); squeezing the S^1 's to “lollipops” (right). We note that $S^1 \vee S^1$ and $L^1 \vee L^1$ actually live in \mathbb{R}^4 .




Figure 9.10: The map λ squeezing S^ℓ to the lollipop L^ℓ .

As is well known, if we shrink the boundary of an n -ball to a single point, the result is an n -sphere. Let us fix a mapping $\gamma: B^\ell \rightarrow S^\ell$ that sends all of ∂B^ℓ to the distinguished point s_0 and is a homeomorphism on the interior of B^ℓ . Now we are ready to define the map g . Namely, we define $\bar{g}: B^{2\ell} \rightarrow S^\ell \times S^\ell$ by

$$\bar{g}(x, y) = (\gamma(x), \gamma(y)),$$

where we still consider $B^{2\ell}$ as $B^\ell \times B^\ell$ and x comes from the first B^ℓ and y from the second. Then g is the restriction of \bar{g} to $S^{2\ell-1} = \partial B^{2\ell}$.

For the image of g we have, using (9.2),

$$g(S^{2\ell-1}) = g(B^\ell \times S^{\ell-1}) \cup g(S^{\ell-1} \times B^\ell) = (S^\ell \times \{s_0\}) \cup (\{s_0\} \times S^\ell) = S^\ell \vee S^\ell.$$

It remains to observe that \bar{g} restricted to $\text{int } B^{2\ell}$ is a homeomorphism onto $(S^\ell \times S^\ell) \setminus (S^\ell \vee S^\ell)$. Hence, the result of attaching the boundary of $B^{2\ell}$ to $S^\ell \vee S^\ell$ via g is indeed homeomorphic to $S^\ell \times S^\ell$ as claimed.

Squeezing. Now we define a “squeezing map” from $S^\ell \vee S^\ell$ to \mathbf{E} . We let the ℓ -lollipop L^ℓ be an ℓ -dimensional sphere of radius $\frac{1}{2}$ with attached segment (“stick”) of length 1; see Fig. 9.10. Formally, $L^\ell := \partial B(-\frac{1}{2}s_0, \frac{1}{2}) \cup [0, s_0]$, where $B(x, r)$ stands for the ball of radius r centered at x . We let $\lambda: S^\ell \rightarrow L^\ell$ be the projection that moves each point of S^ℓ in direction perpendicular to the axis $[-s_0, s_0]$.

Now, with $L^\ell \vee L^\ell := (L^\ell \times \{s_0\}) \cup (\{s_0\} \times L^\ell)$, we have the map $\lambda \vee \lambda: S^\ell \vee S^\ell \rightarrow L^\ell \vee L^\ell$ (given by $(x, y) \mapsto (\lambda(x), \lambda(y))$). Finally, $L^\ell \vee L^\ell$ can be identified with the




Figure 9.11: Contracting the wedge of two hemispheres.

complex \mathbf{E} as above by a suitable homeomorphism, and we arrive at the map

$$r = (\lambda \vee \lambda) \circ g: S^{2\ell-1} \rightarrow \mathbf{E}$$

(where the homeomorphism of $L^\ell \vee L^\ell$ with \mathbf{E} is not explicitly shown).

The clause gadget \mathbf{X} is obtained by attaching the boundary of $B^{2\ell}$ to \mathbf{E} via the map r . Of course, we want \mathbf{X} to be a simplicial complex, and so in reality we use a suitable PL version of the attaching map r (we have not presented it this way since the description above seems more accessible).

For the forthcoming proof of an analogue of Lemma 9.6, we need the following observation.

Observation 9.8. *Let $\kappa: L^\ell \vee L^\ell \rightarrow S^\ell \vee S^\ell$ be the quotient map corresponding to contracting the “stick” c of the double-lollipop to a single point. Then the composition $\kappa \circ (\lambda \vee \lambda): S^\ell \vee S^\ell \rightarrow S^\ell \vee S^\ell$ is homotopic to the identity on $S^\ell \vee S^\ell$.*

Proof. Let $H^\ell := \{x \in S^\ell : \langle s_0, x \rangle \geq 0\}$ be the closed hemisphere centered at s_0 . The assertion follows by observing that $\kappa \circ (\lambda \vee \lambda)$ is the quotient map corresponding to contracting the subset $H^\ell \vee H^\ell$ of $S^\ell \vee S^\ell$ to a single point; see Fig. 9.11. \square

Lemma 9.9 (Higher dimensional version of Lemma 9.6).

- (i) (Based on [SSS98, Lemma 2.2]). *Let Σ_a^ℓ and Σ_b^ℓ denote the two ℓ -spheres (boundaries of $(\ell+1)$ -simplices) contained in $\mathbf{E} \subset \mathbf{X}$. Let S_a^k and S_b^k be PL k -spheres. Then there is no PL embedding f of the disjoint union $S_a^k \sqcup S_b^k \sqcup \mathbf{X}$ into S^d such that*
- *the ℓ -sphere $f(\Sigma_a^\ell)$ and the k -sphere $f(S_a^k)$ have odd linking number, and so do $f(\Sigma_b^\ell)$ and $f(S_b^k)$;*
 - *$f(\Sigma_a^\ell)$ and $f(S_b^k)$ are unlinked, and so are $f(\Sigma_b^\ell)$ and $f(S_a^k)$.*
- (ii) *Let f be a generic linear embedding of \mathbf{E} in⁵ $\mathbb{R}^{2\ell+2}$, and let $\delta > 0$. Then there is a PL embedding \bar{f} of \mathbf{X} in $\mathbb{R}^{2\ell+2}$ extending f whose image is contained in the neighborhood $N = N(f, \delta) := N(T_a, \delta) \cup N(f(\Sigma_b^\ell), \delta) \cup N(f(c), \delta)$, where T_a is the $(\ell+1)$ -dimensional simplex bounded by $f(\Sigma_a^\ell)$.*

⁵It follows from our assumptions on d and k that $d \geq 2\ell+3$. Therefore, when, in the course of the reduction, we construct an embedding of a complex associated with a satisfiable formula, we can afford to embed each conflict gadget in its own “private” $(2\ell+2)$ -dimensional set. Since two ℓ -spheres in dimension $2\ell+2$ are never linked, we do not need to make an explicit unlinking assumption as in Lemma 9.6.

Proof of (i). Part (i) follows from the proof of [SSS98, Proof of Lemma 2.2] with only minor modifications. First, before giving a formal proof, we describe the basic approach of [SSS98], which also applies to the proof of Lemma 9.6(i).

Suppose that a PL embedding f as in (i) above exists. Let C denote the complement $\mathbb{R}^d \setminus f(S_a^k \sqcup S_b^k)$, let $r: S^{2\ell-1} \rightarrow \mathbf{E}$ be the attaching map used in the construction of \mathbf{X} , and let $\bar{r}: B^{2\ell} \rightarrow \mathbf{X}$ be the extension of r to $B^{2\ell}$ (formally, \bar{r} is the quotient map).

The basic strategy is as follows: On the one hand, using the assumptions about linking numbers, one shows that $f \circ r$ defines a nontrivial element of the homotopy group $\pi_{2\ell-1}(C)$. On the other hand, $f \circ \bar{r}$ witnesses that $f \circ r$ is homotopically trivial—a contradiction.

As in [SSS98], one distinguishes two cases: $\ell = 1$ and $\ell > 1$. In the case $\ell = 1$, we are dealing with the fundamental group $\pi_1(C)$, and the proof is essentially identical to that of [FKT94, Lemma 7], i.e., our Lemma 9.6, which we briefly summarize for the reader's convenience.

For showing that $f \circ r: S^1 \rightarrow C$ is homotopically nontrivial, one first observes that $\pi_1(\mathbf{E})$ is the free group on two generators a and b , and the attaching map $r: S^1 \rightarrow \mathbf{E}$ corresponds to the commutator $aba^{-1}b^{-1}$, which is a nontrivial element of $\pi_1(\mathbf{E})$. So it suffices to show that the map $f_*: \pi_1(\mathbf{E}) \rightarrow \pi_1(C)$ induced by the restriction $f|_{\mathbf{E}}$ is injective. To this end, one first considers the homomorphisms f_{*1} and f_{*2} induced by $f|_{\mathbf{E}}$ in the first and second *homology*.

By Alexander duality, the complement C has the same homology (with \mathbb{Z}_2 -coefficients, say) as $S^1 \vee S^1$, and thus $H_1(C; \mathbb{Z}_2) = \mathbb{Z}_2 \oplus \mathbb{Z}_2$ and $H_2(C; \mathbb{Z}_2) = 0$. For \mathbf{E} we have $H_1(\mathbf{E}; \mathbb{Z}_2) \cong \mathbb{Z}_2 \oplus \mathbb{Z}_2$ with a basis represented by the two circles Σ_a^1 and Σ_b^1 . The assumption on the linking numbers imply that f_{*1} is an isomorphism, and f_{*2} is trivially surjective. Then the injectivity of the homomorphism f_* of the fundamental groups follows from a theorem of Stallings [Sta65], which finishes the case $\ell = 1$.

In the case $\ell > 1$, the proof that $f \circ r$ defines a nontrivial element of $\pi_{2\ell-1}(C)$ requires somewhat more advanced machinery. Segal et al. [SSS98] prove essentially the same assertion as in part (i) of the lemma, with the following differences:

1. \mathbf{X} is replaced by X' , which is obtained by attaching $B^{2\ell}$ to $\Sigma_a^\ell \vee \Sigma_b^\ell$ via the map g (as described above) and hence homeomorphic to $S^\ell \times S^\ell$.
2. The disjoint union $S_a^k \sqcup S_b^k$ is replaced by the wedge⁶ $S_a^k \vee S_b^k$.

They show that if there were an embedding f of $(\Sigma_a^\ell \vee \Sigma_b^\ell) \sqcup (S_a^k \vee S_b^k)$ with the linking properties as in part (i) of the lemma, $f \circ g$ would be a nontrivial element of $\pi_{2\ell-1}(C)$.

Now we begin with a formal proof of Lemma 9.9. Instead of modifying the proof of [SSS98], we show how to reduce our assertion to theirs. Suppose there were a bad embedding f of $S_a^k \sqcup S_b^k \sqcup \mathbf{X}$ as in the lemma. Since the codimension of the image $f(S_a^k \sqcup S_b^k \sqcup \mathbf{X})$ is at least 2, we can grow a k -dimensional finger from $f(S_a^k)$ towards $f(S_b^k)$ avoiding $f(\mathbf{X})$ until the finger touches $f(S_b^k)$ in a single point. This results in an embedding of $(S_a^k \vee S_b^k) \sqcup \mathbf{X}$. For simplicity, we denote this modified embedding by f as well.

We observe that when pulling the finger, we can pull along a $(k+1)$ -dimensional image of B^{k+1} filling $f(S_a^k)$, and so the images are still linked or unlinked as in the assumption of the lemma.

⁶Wedges are used for technical reasons: By a theorem of Lickorish [Lic65], any embedding (PL or even topological) of a wedge of spheres of codimension at least 3 is unknotted, i.e., ambient isotopic to a standard embedding.

Next, consider the image $f(\mathbf{E})$ of the double lollipop in C . We modify f as follows. We deformation retract the arc $f(c)$ to its midpoint m , pulling along ℓ -dimensional fingers from the two ℓ -spheres $f(\Sigma_a^\ell)$ and $f(\Sigma_b^\ell)$, so that at the end of the deformation, the fingers touch in the single point m . This describes a continuous deformation of $f|_{\mathbf{E}}$ that only changes $f|_{\mathbf{E}}$ on the segment c and in two small neighborhoods U_a and U_b of the endpoints of c in the ℓ -spheres (these neighborhoods provide the ‘‘material’’ for the fingers). We have to take care to pull along the parts of $B^{2\ell}$ attached to U_a and U_b , respectively, i.e., we extend the deformation to a continuous deformation of f on all of X that changes f only on a small neighborhood V in X of $c \cup U_a \cup U_b$. The whole deformation can be carried out so that the image of V remains in a small ε -neighborhood of the original image $f(\mathbf{E})$ throughout the deformation. Let f' be the final modified map from $S_a^k \sqcup S_b^k \sqcup X$ into S^d (note that we made no changes on the two k -spheres). The map f' maps the ‘‘bent stick’’ c of the double lollipop constantly to m (in particular, it is not an embedding), and it induces a unique embedding $f'' : X' \rightarrow C$ such that f'' agrees with f' on the interior of $B^{2\ell}$ and $f'' \circ \kappa = f'$ on \mathbf{E} , where κ is the map from Observation 9.8. Moreover, the map $f \circ r = f \circ (\lambda \vee \lambda) \circ g : S^{2\ell-1} \rightarrow C$ is deformed into the map $f'' \circ \kappa \circ (\lambda \vee \lambda) \circ g : S^{2\ell-1} \rightarrow C$. Thus, $f \circ r$ and $f'' \circ \kappa \circ (\lambda \vee \lambda) \circ g$ define the same element of $\pi_{2\ell-1}(C)$. However, by the observation, the latter map is homotopic to $f'' \circ g$. Thus, $f \circ r$ and $f'' \circ g$ define the same element of $\pi_{2\ell-1}(C)$. But the former is trivial, as witnessed by $f \circ \bar{r}$, while the latter is not according to [SSS98]—a contradiction. This completes the proof of (i).

Proof of (ii). For easier presentation, we describe an embedding \bar{f} that is not a priori PL; it is routine to replace it by a PL embedding.

Applying a suitable homeomorphism $\mathbb{R}^{2\ell+2} \rightarrow \mathbb{R}^{2\ell+2}$, we may assume that $f(\mathbf{E})$ is actually $L^\ell \vee L^\ell$. Let \bar{L}^ℓ denote the ℓ -lollipop with its ℓ -sphere filled (i.e., $\bar{L}^\ell := B(-\frac{1}{2}s_0, \frac{1}{2}) \cup [0, s_0]$). It suffices to embed X in the δ -neighborhood of $L^\ell \vee \bar{L}^\ell$ for $\delta > 0$ arbitrarily small; actually, for notational convenience, we will eventually get 4δ instead of δ .

Instead of specifying the embedding $\bar{f} : X \rightarrow \mathbb{R}^{2\ell+2}$ directly, we define a mapping $\tilde{f} : S^\ell \times S^\ell \rightarrow \mathbb{R}^{2\ell+2}$ that coincides with $\lambda \vee \lambda$ on $S^\ell \vee S^\ell$ and maps the rest of $S^\ell \times S^\ell$ homeomorphically. Then \bar{f} can be given as (considering \mathbf{E} identified with $L^\ell \vee L^\ell$)

$$\bar{f}(z) = \begin{cases} z & \text{for } z \in \mathbf{E}, \\ \tilde{f}(z) & \text{for } z \notin \mathbf{E}. \end{cases}$$


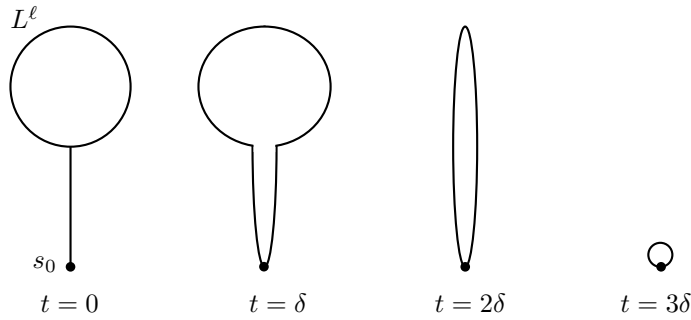
Writing a point of $S^\ell \times S^\ell$ as (x, y) , we define \tilde{f} using two auxiliary maps $u, v : S^\ell \times [0, \infty) \rightarrow \mathbb{R}^{\ell+1}$:

$$\tilde{f}(x, y) := \left(u(x, \text{dist}(y, s_0)), v(y, \text{dist}(x, s_0)) \right).$$

For defining $u(x, t)$, we think of t as time. For $t = 0$, the image $u(S^\ell, t)$ is the lollipop L^ℓ , while for all $t > 0$ it is topologically a sphere, which looks almost like the lollipop; see Fig. 9.12. Concretely, we set

$$u(x, t) := \begin{cases} tx + (1-t)\lambda(x) & \text{for } 0 \leq t \leq \delta, \\ \delta x + (1-\delta)\lambda(x) & \text{for } t \geq \delta. \end{cases}$$

As for v , we let it coincide with u for $t \leq \delta$ (see Fig. 9.13). For $t = 2\delta$, we set $v(x, 2\delta) := (x_1, \delta x_2, \delta x_3, \dots, \delta x_{\ell+1})$, and for all $t \geq 3\delta$ we set $v(x, t) := \delta(x - s_0) + s_0$

Figure 9.12: The (images of the) mappings $u(*, t)$.Figure 9.13: The (images of the) mappings $v(*, t)$.

(i.e., the sphere is shrunk by the factor of δ so that it still touches s_0). On the intervals $[\delta, 2\delta]$ and $[2\delta, 3\delta]$ we interpolate $v(x, t)$ linearly in t .

The \bar{f} defined in this way is clearly continuous and coincides with $\lambda \vee \lambda$ on $S^\ell \vee S^\ell$. Next, we want to show $\bar{f}(x, y) \neq \bar{f}(x', y')$ whenever $(x, y) \neq (x', y')$ and none of x, x', y, y' equals s_0 . First we note that $u(x, t) \neq u(x', t')$ whenever $x \neq x'$ and $t, t' > 0$, and thus we may assume $x = x', y \neq y'$. Then we just use injectivity of $v(*, t)$ for every $t > 0$.

It remains to check that the image of \bar{f} lies close to $\bar{L}^\ell \vee L^\ell$. The image $u(S^\ell, t)$ is δ -close to L^ℓ for all t , and the image $v(S^\ell, t)$ is 2δ -close to s_0 whenever $t \geq 3\delta$. Thus, whenever $\text{dist}(x, s_0) \geq 3\delta$, we have $\bar{f}(x, y)$ lying 3δ -close to $L^\ell \times \{s_0\}$.

Next, let us assume $\text{dist}(x, s_0) \leq 3\delta$. Then $u(x, t)$ is 3δ -close to s_0 for all t , and observing that $v(y, t)$ always lies δ -close to the filled lollipop \bar{L}^ℓ , we conclude that $\bar{f}(x, y)$ is 4δ -close to $\{s_0\} \times \bar{L}^\ell$. \square

9.4.3 The reduction

Having introduced the clause gadget and the conflict gadget, the rest of the reduction is almost the same as in Section 9.3.3, and so we mainly point out the (minor) differences.

Given a 3-CNF formula φ , the simplicial complex is pasted together from the gadgets exactly as in Section 9.3.3; we have $\dim K = \max(k, 2\ell) = k$. For φ unsatisfiable, nonembeddability of K is shown using Lemmas 9.7(i) and 9.9(i) instead of Lemmas 9.5(i) and 9.6(i), but otherwise in the same way as in Section 9.3.3.

Given a satisfiable formula φ , we again begin with embedding the clause gadgets, this time using Lemma 9.7(ii). For an opening ω of a clause gadget G_i , we can again

obtain a set $Q_\omega \subset \mathbb{R}^d$ with $Q_\omega \cap G_i = \Sigma$, where $\Sigma = f(\partial\omega)$, this time homeomorphic to $T^\varepsilon \times B^k$ (where T is the $(\ell + 1)$ -dimensional simplex bounded by Σ and T^ε is the part of it ε -close to Σ). Similarly we can build, for a non-witness opening ω , the set Q_ω^+ homeomorphic to $T \times B^k$.

Now we need to define the “private pieces” Q_{ω_j} and $Q_{\omega_j}^+$, $j = 1, 2, \dots, t$, within each Q_ω and Q_ω^+ , respectively. This time first we choose pairwise disjoint sets $B_1, \dots, B_t \subset \partial B^k$, each homeomorphic to $B^{\ell+1}$ (for this we need $k \geq \ell + 2$, which holds in our setting), we let W_j be the cone with base B_j and apex in the center of B^k , and we let w_j be the boundary of W_j (not including the interior of the base B_j). We have W_j homeomorphic to $B^{\ell+2}$ and w_j to $B^{\ell+1}$, and this allows us to construct Q_{ω_j} homeomorphic to a $(2\ell + 2)$ -dimensional neighborhood of Σ , and $Q_{\omega_j}^+$ homeomorphic to a $(2\ell + 2)$ -dimensional neighborhood of T .

The rest of the embedding construction can be copied from Section 9.3.3 almost verbatim. This concludes the proof of Theorem 4.3. \square

Bibliography

- [ABB90] F. Acquistapace, R. Benedetti, and F. Broglia. Effectiveness - non effectiveness in semialgebraic and PL geometry. *Invent. Math.*, 102(1):141–156, 1990.
- [AHT06] I. Agol, J. Hass, and W. Thurston. The computational complexity of knot genus and spanning area. *Trans. Am. Math. Soc.*, 358(9):3821–3850, 2006.
- [AHW87] N. Alon, D. Haussler, and E. Welzl. Partitioning and geometric embedding of range spaces of finite Vapnik-Chervonenkis dimension. In *Proc. 3rd ACM Sympos. on Computational Geometry*, pages 331–340, 1987.
- [AKMM02] N. Alon, G. Kalai, J. Matoušek, and R. Meshulam. Transversal numbers for hypergraphs arising in geometry. *Adv. in Appl. Math.*, 130:2509–2514, 2002.
- [Alo85] N. Alon. Expanders, sorting in rounds and superconcentrators of limited depth. In *Proc. 17th ACM Sympos. on Theory of Comput.*, pages 98–102, 1985.
- [Alo86] N. Alon. Eigenvalues, geometric expanders, sorting in rounds, and Ramsey theory. *Combinatorica*, 6(3):207–219, 1986.
- [AMS93] R. N. Andersen, M. M. Marjanović, and R. M. Schori. Symmetric products and higher-dimensional dunce hats. *Topology Proc.*, 18:7–17, 1993.
- [Ani89] D. J. Anick. The computation of rational homotopy groups is $\#\mathcal{P}$ -hard. *Computers in geometry and topology, Proc. Conf., Chicago/Ill. 1986*, Lect. Notes Pure Appl. Math. 114, 1–56, 1989.
- [BGdO00] J. Bokowski and A. Guedes de Oliveira. On the generation of oriented matroids. *Discrete Comput. Geom.*, 24(2-3):197–208, 2000. The Branko Grünbaum birthday issue.
- [Bre83] U. Brehm. A nonpolyhedral triangulated Möbius strip. *Proc. Amer. Math. Soc.*, 89(3):519–522, 1983.
- [Bro57] E. H. Brown (jun.). Finite computability of Postnikov complexes. *Ann. Math. (2)*, 65:1–20, 1957.
- [Bry02] J.L. Bryant. Piecewise linear topology. In *Daverman, R. J. (ed.) et al., Handbook of geometric topology.*, pages 219–259. Elsevier, Amsterdam, 2002.

- [BS92] U. Brehm and K. S. Sarkaria. Linear vs. piecewise linear embeddability of simplicial complexes. Tech. Report 92/52, Max-Planck-Institut f. Mathematik, Bonn, Germany, 1992.
- [Buo03] S. Buoncrisiano. *Fragments of Geometric Topology from the Sixties*. Geometry & Topology Monographs, Vol. 6, 2003. New edition in preparation, in collaboration with C. Rourke.
- [ČKM⁺11] M. Čadek, M. Krčál, J. Matoušek, F. Sergeraert, L. Vokřínek, and U. Wagner. Computing all maps into a sphere. Preprint; <http://arxiv.org/abs/1105.6257>, 2011.
- [Coo71] S. A. Cook. The complexity of theorem proving procedures. In *Proc. 3rd Ann. ACM Symp. on Theory of Computing*, pages 151–158, 1971.
- [CZ61] M.L. Curtis and E.C. Zeeman. On the polyhedral Schoenflies theorem. *Proc. Am. Math. Soc.*, 11:888–889, 1961.
- [Eck93] J. Eckhoff. Helly, Radon and Carathéodory type theorems. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*. North-Holland, Amsterdam, 1993.
- [FKT94] M. H. Freedman, V. S. Krushkal, and P. Teichner. Van Kampen’s embedding obstruction is incomplete for 2-complexes in \mathbb{R}^4 . *Math. Res. Lett.*, 1(2):167–176, 1994.
- [Flo34] A. Flores. Über n -dimensionale Komplexe die im R_{2n+1} absolut selbstverschlungen sind. *Ergeb. Math. Kolloq.*, 4:6–7, 1932/1934.
- [Gal44] T. Gallai. Solution of problem 4065. *Amer. Math. Monthly*, 51:169–171, 1944.
- [Gla71] L. C. Glaser. A proof of the most general polyhedral Schoenflies conjecture possible. *Pac. J. Math.*, 38:401–417, 1971.
- [GS06] D. Gonçalves and A. Skopenkov. Embeddings of homology equivalent manifolds with boundary. *Topology Appl.*, 153(12):2026–2034, 2006.
- [Hak61] W. Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961.
- [Hat01] A. Hatcher. *Algebraic Topology*. Cambridge University Press, Cambridge, 2001.
- [Hel23] E. Helly. Über mengen konvexer Körper mit gemeinschaftlichen Punkten. *Jahresber. Deutsch. Math.-Verein.*, 32:175–176, 1923.
- [Hem79] G. Hemion. On the classification of homeomorphisms of 2-manifolds and the classification of 3-manifolds. *Acta Math.*, 142(1-2):123–155, 1979.
- [HJ64] R. Halin and H. A. Jung. Charakterisierung der Komplexe der Ebene und der 2-Sphäre. *Arch. Math.*, 15:466–469, 1964.
- [HLP99] J. Hass, J. C. Lagarias, and N. Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999.

- [HT74] J. Hopcroft and R. Tarjan. Efficient planarity testing. *J. Assoc. Comput. Mach.*, 21:549–568, 1974.
- [Iva08] S. V. Ivanov. The computational complexity of basic decision problems in 3-dimensional topology. *Geom. Dedicata*, 131:1–26, 2008.
- [Kel86] L.M. Kelly. A resolution of the Sylvester-Gallai problem of J.-P. Serre. *Discrete Comput. Geom.*, 1(1):101–104, 1986.
- [KGT01] D. J. Kleitman, A. Gyárfás, and G. Tóth. Convex sets in the plane with three of every four meeting. *Combinatorica*, 21(2):221–232, 2001. Paul Erdős and his mathematics (Budapest, 1999).
- [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.
- [KM06] G. Kalai and R. Meshulam. Intersections of Leray complexes and regularity of monomial ideals. *J. Combin. Theory Ser. A*, 113(7):1586–1592, 2006.
- [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.
- [LB63] C. G. Lekkerkerker and J. Ch. Boland. Representation of a finite graph by a set of intervals on the real line. *Fund. Math.*, 51:45–64, 1962/1963.
- [Lic65] W. B. R. Lickorish. The piecewise linear unknotting of cones. *Topology*, 4:67–91, 1965.
- [Lov78] L. Lovász. Kneser’s conjecture, chromatic number, and homotopy. *J. Combin. Theory Ser. A*, 25(3):319–324, 1978.
- [Mar58] A. A. Markov. Insolubility of the problem of homeomorphy. In *Proc. Internat. Congress Math. 1958*, pages 300–306. Cambridge Univ. Press, New York, 1958.
- [Mat97] S. V. Matveev. Classification of sufficiently large 3-manifolds. *Uspekhi Mat. Nauk*, 52(5(317)):147–174, 1997. Translation in Russian Math. Surveys 52 (1997), no. 5, 1029–1055.
- [Mat02] J. Matoušek. *Lectures on Discrete Geometry*. Springer-Verlag New York, Inc., 2002.
- [Mat03] J. Matoušek. *Using the Borsuk-Ulam Theorem*. Springer, Berlin etc., 2003.

- [Mel09] S. A. Melikhov. The van Kampen obstruction and its relatives. *Tr. Mat. Inst. Steklova*, 266(Geometriya, Topologiya i Matematicheskaya Fizika. II):149–183, 2009.
- [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.
- [MN98] J. Matoušek and J. Nešetřil. *Invitation to Discrete Mathematics*. Oxford University Press, Oxford, 1998.
- [MR02] M. Molloy and B. Reed. *Graph colouring and the probabilistic method*, volume 23 of *Algorithms and Combinatorics*. Springer-Verlag, Berlin, 2002.
- [MT09] J. Matoušek and M. Tancer. Dimension gaps between representability and collapsibility. *Discrete Comput. Geom.*, 42(4):631–639, 2009.
- [MTW11] J. Matoušek, M. Tancer, and U. Wagner. Hardness of embedding simplicial complexes in \mathbb{R}^d . *J. Eur. Math. Soc. (JEMS)*, 13(2):259–295, 2011.
- [Mun84] J. R. Munkres. *Elements of Algebraic Topology*. Addison - Wesley, 1984.
- [Nab95] A. Nabutovsky. Einstein structures: Existence versus uniqueness. *Geom. Funct. Anal.*, 5(1):76–91, 1995.
- [New60] M. H. A. Newman. On the division of Euclidean n -space by topological $(n - 1)$ -spheres. *Proc. Roy. Soc. London Ser. A*, 257:1–12, 1960.
- [NW96] A. Nabutovsky and S. Weinberger. Algorithmic unsolvability of the triviality problem for multidimensional knots. *Comment. Math. Helv.*, 71(3):426–434, 1996.
- [NW99] A. Nabutovsky and S. Weinberger. Algorithmic aspects of homeomorphism problems. In *Tel Aviv Topology Conference: Rothenberg Festschrift (1998)*, volume 231 of *Contemp. Math.*, pages 245–250. Amer. Math. Soc., Providence, RI, 1999.
- [Pac98] J. Pach. A Tverberg-type result on multicolored simplices. *Comput. Geom.*, 10:71–76, 1998.
- [Per85] G. Ya. Perel'man. Realization of abstract k -skeletons of intersections of convex polyhedra in R^{2k-1} . *Geometric Questions in the Theory of Functions and Sets (Kalinin. Gos. Univ., Kalinin)*, pages 129–131, 1985. In Russian.
- [Ren92] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. III. Quantifier elimination. *J. Symbolic Comput.*, 13(3):329–352, 1992.
- [RRS06] A. Romero, J. Rubio, and F. Sergeraert. Computing spectral sequences. *J. Symb. Comput.*, 41(10):1059–1079, 2006. arXiv:cs/0602064.

- [RS82] C.P. Rourke and B.J. Sanderson. *Introduction to piecewise-linear topology*. Springer, Berlin, 1982. Rev. printing of “Ergebnisse der Mathematik und ihrer Grenzgebiete”, Vol. 69, 1972.
- [RS99] D. Repovš and A. Skopenkov. New results on embeddings of polyhedra and manifolds into Euclidean spaces. *Uspekhi Mat. Nauk*, 54(6(330)):61–108, 1999.
- [RST95] N. Robertson, P. D. Seymour, and R. Thomas. Sachs’ linkless embedding conjecture. *J. Combin. Theory Ser. B*, 64(2):185–227, 1995.
- [Rub95] J. H. Rubinstein. An algorithm to recognize the 3-sphere. In *Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994)*, pages 601–611, Basel, 1995. Birkhäuser.
- [Sch91] R. Schön. Effective algebraic topology. *Mem. Am. Math. Soc.*, 451:63 p., 1991.
- [Sch04] S. Schleimer. Sphere recognition lies in NP. Manuscript, available from <http://www.warwick.ac.uk/~masgar>, 2004.
- [Sha57] A. Shapiro. Obstructions to the imbedding of a complex in a euclidean space. I: The first obstruction. *Ann. of Math., II. Ser.*, 66:256–269, 1957.
- [Smi98] J. Smith. m -structures determine integral homotopy type. Preprint, arXiv:math/9809151v1, 1998.
- [Soa04] R. I. Soare. Computability theory and differential geometry. *Bull. Symbolic Logic*, 10(4):457–486, 2004.
- [SS92] J. Segal and S. Spież. Quasi embeddings and embeddings of polyhedra in \mathbb{R}^m . *Topology Appl.*, 45(3):275–282, 1992.
- [SSS98] J. Segal, A. Skopenkov, and S. Spież. Embeddings of polyhedra in \mathbb{R}^m and the deleted product obstruction. *Topology Appl.*, 85(1-3):335–344, 1998.
- [Sta65] J. Stallings. Homology and central series of groups. *J. Algebra*, 2:170–181, 1965.
- [Tan84] R. M. Tanner. Explicit concentrators from generalized n -gons. *SIAM J. Algebraic Discrete Methods*, 5(3):287–293, 1984.
- [Tan10a] M. Tancer. d -collapsibility is NP-complete for d greater or equal to 4. *Chicago Journal of Theoretical Computer Science*, 2010(3), June 2010.
- [Tan10b] M. Tancer. Non-representability of finite projective planes by convex sets. *Proc. Amer. Math. Soc.*, 138(9):3285–3291, 2010.
- [Tan11a] M. Tancer. A counterexample to Wegner’s conjecture on good covers. *Discrete & Computational Geometry*, pages 1–9, 2011. Online-first.
- [Tan11b] M. Tancer. Intersection patterns of convex sets via simplicial complexes, a survey. Preprint; <http://arxiv.org/abs/1102.0417>, 2011.

- [Tho94] A. Thompson. Thin position and the recognition problem for S^3 . *Math. Res. Lett.*, 1(5):613–630, 1994.
- [TT11] M. Tancer and D. Tonkonog. Good covers are algorithmically unrecognizable. Manuscript in preparation, 2011.
- [vK32] R. E. van Kampen. Komplexe in euklidischen Räumen. *Abh. Math. Sem. Hamburg*, 9:72–78, 1932. Berichtigung dazu, *ibid.* (1932) 152–153.
- [VKF74] 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).
- [VM04] Bane Vasic and Olgica Milenkovic. Combinatorial constructions of low-density parity-check codes for iterative decoding. *IEEE Trans. Inform. Theory*, 50(6):1156–1176, 2004.
- [Weg67] G. Wegner. *Eigenschaften der Nerven homologisch-einfacher Familien im R^n* . PhD thesis, Universität Göttingen, 1967. In German.
- [Weg75] G. Wegner. d -collapsing and nerves of families of convex sets. *Arch. Math.*, 26:317–321, 1975.
- [Wil75] Richard M. Wilson. An existence theory for pairwise balanced designs. III. Proof of the existence conjectures. *J. Combinatorial Theory Ser. A*, 18:71–79, 1975.
- [Wu65] W.-T. Wu. *A Theory of Imbedding, Immersion, and Isotopy of Polytopes in a Euclidean Space*. Science Press, Peking, 1965.