FACULTY
OF MATHEMATICS
AND PHYSICS
Charles University

DOCTORAL THESIS

Michael Skotnica

Combinatorics, group theory, computational complexity & topology

Departement of Applied Mathematics

Supervisor of the doctoral thesis: doc. RNDr. Martin Tancer, Ph.D.
Study programme: Computer Science
Study branch: Theory of Computing, Discrete Models and Optimization

Prague 2023

I declare that I carried out this doctoral thesis independently, and only with the cited sources, literature and other professional sources. It has not been used to obtain another or the same degree.

I understand that my work relates to the rights and obligations under the Act No. 121/2000 Sb., the Copyright Act, as amended, in particular the fact that the Charles University has the right to conclude a license agreement on the use of this work as a school work pursuant to Section 60 subsection 1 of the Copyright Act.

In .... date ....

Author's signature

First of all, I would like to thank my supervisor Martin Tancer for his guidance, support, advice and cooperation.

Next, I would like to thank my coworkers Cornelius Brand, Denys Bulavka, Viktoriia Korchemna, Thomas Magnard, Josef Minařík, Shay Moran, Kirill Simonov and also Robert Ganian, Xavier Goaoc, Eran Nevo, and Uli Wagner who provided me internships abroad.

Last but not least, I thank my parents for supporting me during my studies.

The work which is this thesis based on was supported from the project “Grant Schemes at CU” (reg. no. CZ.02.2.69/0.0/0.0/19_073/0016935) and also from the grants no. 19-04113Y and no. 22-19073S of the Czech Science Foundation (GAČR).

Title: Combinatorics, group theory, computational complexity & topology

Author: Michael Skotnica

Department: Departement of Applied Mathematics

Supervisor: doc. RNDr. Martin Tancer, Ph.D., Departement of Applied Mathematics

Abstract: In this thesis, we present new results related to combinatorial properties of topological spaces given by abstract simplicial complexes, their relations and computational complexity.

First, we generalize a result of Hachimori on relations between shellability and collapsibility which are important combinatorial properties of simplicial complexes.

Next, we study the computational complexity of the PL geometric category of 2-dimensional polyhedra introduced by Borghini which is a combinatorial notion providing a natural upper bound for the Lusternik–Schnirelmann category. For 2-dimensional polyhedra it can be equal to 1, 2 or 3. While it is easy to decide whether the PL geometric category of a 2-dimensional polyhedron is equal 1, we show that it is NP-hard to decide whether this category is at most 2.

Finally, we show that computing the rank of higher homotopy groups of a simply connected topological space is #W[2]-hard using a problem called VEST, given by Anick, as an intermediate problem. We also establish results for the decision version of VEST and for its variants as self-contained problems. For most of them we show W[1]- or W[2]-hardness.

Keywords: simplicial complex, collapsibility, shellability, computational complexity, homotopy groups.

Contents

Introduction

An abstract simplicial complex is a structure for describing various topological spaces in an elegant way. From the set theoretical viewpoint, it is just a family of sets closed under taking subsets. Its geometric realization then forms a Hausdorff topological space and two geometric realizations of the same abstract simplicial complex are always homeomorphic.

Simplicial complexes were inspired by algebraic topology.[1] However, they are also used in another fields of mathematics, such as combinatorics. On the one hand, there are combinatorial problems which where solved using topology. (See, e.g., [Mat03].) On the other hand, establishing the notion of simplicial complexes, which can also be viewed as a generalization of graphs, led to new notions of several interesting combinatorial properties for simplicial complexes. Some of them can be viewed as "discrete", or "combinatorial", versions of important topological properties. For instance, collapsibility of simplicial complexes can be viewed as a discrete version of contractibility of topological spaces. Moreover, if a simplicial complex is collapsible, then its geometric representation is contractible.

Decomposition properties of simplicial complexes

Many combinatorial properties of simplicial complexes, especially decomposition properties and their relations, have been widely studied since 1980s. For the purpose of this thesis the following properties (to be defined later[2]) are most important:

- collapsibility,
- vertex-decomposability,
- shellability.

All these three properties are related: Vertex-decomposability implies shellability and complexes which are shellable and their geometric realizations are contractible are always collapsible. It is also interesting to ask when collapsibility implies shellability. This question was studied by Hachimori [Hac08] who showed that if a 2-dimensional simplicial complex becomes collapsible after removing a certain number of facets and the link of each its vertex is connected then such complex has a shellable subdivision.

In Chapter 2, we generalize his result in two ways: We show that the condition implies not only shellability but more general vertex-decomposability and we also prove that this holds in arbitrary dimension.

This chapter is essentially the content of the article [MST21] which is a joint work with Thomas Magnard and Martin Tancer.

[1]They are a special case of more general CW complexes.

[2]Shellability and collapsibility are defined in Chapter 1, vertex-decomposability in Chapter 2.

Computational problems

Note that an abstract simplicial complex is also a convenient way how to represent topological spaces on computers. Therefore, one can ask how hard is to decide, or whether it is possible to decide if a topological space represented by a simplicial complex satisfies a given (not only combinatorial) property. For instance, one can ask how hard it is to decide if a simplicial complex satisfies the properties mentioned in the previous paragraph: collapsibility, vertex-decomposability and shellability.

Their complexity was studied recently. Collapsibility was shown to be NP-complete by Tancer [Tan16] even for 3-dimensional complexes (it was previously known that the problem is polynomial for 2-dimensional complexes). Vertex-decomposability and shellability were also shown to be NP-complete even for 2-dimensional complexes by Goaoc, Paták, Patáková, Tancer and Wagner [GPP+19]. Let us point out that for this purpose, they used the result of Hachimori [Hac08] which was mentioned in the previous paragraph.

PL geometric category in dimension 2 In Chapter 3, we study computational complexity of another but related property, the PL geometric category of 2-dimensional polyhedra introduced by Borghini [Bor20]. It can be viewed as a discrete approximation and a natural upper bound of the famous Lusternik–Schnirelmann category, LS for short. We point out that a 2-dimensional topological space can have LS category equal to 1, 2 or 3. The same holds for polyhedra and their PL geometric category. While nothing is basically known about the actual complexity of determining LS category of 2-dimensional topological space, deciding whether a 2-dimensional polyhedron has PL geometric category 1 is in P. We prove that deciding whether a polyhedron represented by a simplicial complex has PL geometric category 2 is NP-hard. A useful step towards our proof is that we observe a relation between PL geometeric category 2 and shellability.

This chapter is essentially the content of the article [ST23] which is a joint work with Martin Tancer.

Computing higher homotopy groups and VEST The last chapter of the thesis, Chapter 4, is slightly motivated by the problem of computing higher homotopy groups and their rank, respectively.

An important result, implied by undecidability of the word problem for groups due to Novikov [Nov55] and independently to Boone [Boo59], was undecidability of computing the first homotopy group of a topological space (represented, e.g., by a simplicial complex).

However, for simply connected spaces and their higher homotopy groups the problem becomes decidable [Bro57]. It was later shown by Anick [Ani89] that the problem of computing the rank of higher homotopy groups for a simply connected space is #P-hard. As an intermediate problem whose hardness implies hardness of the original problem he used a problem called vector evaluated after a sequence of transformations, VEST for short, which is defined as follows. Let $\mathbf{v} \in \mathbb{Q}^d$ be a rational vector, $(T_1, T_2 \dots T_m)$ a list of $d \times d$ rational matrices, $S \in \mathbb{Q}^{h \times d}$ a rational matrix not necessarily square and a parameter k. The goal is to compute the number of ways one can choose k matrices $T_{i_1}, T_{i_2}, \dots, T_{i_k}$ from the list such that $ST_{i_k} \cdots T_{i_1}\mathbf{v} = \mathbf{0} \in \mathbb{Q}^h$.

It is also possible to look at the problem of computing the rank of the k-th homotopy group (for $k \geq 2$) from the viewpoint of parameterized complexity. The parameter is then the degree k. It was shown by Čadek, Krčál, Matoušek, Vokřínek and Wagner [ČKM$^+$14b] that this problem is in XP and Matoušek later showed $\#\mathsf{W}[1]$-hardness by showing $\#\mathsf{W}[1]$-hardness of VEST.

In Chapter 4, we show that the problem of VEST is $\#\mathsf{W}[2]$-hard which also implies $\#\mathsf{W}[2]$-hardness of the original problem.

Note that the parameterized complexity of VEST is a self-contained problem which can be also viewed as a generalization of the recently studied k-Sum problem [ALW14]. Therefore, we also further discuss a decision version of VEST and its several modifications for which we show $\mathsf{W}[1]-$ or $\mathsf{W}[2]$-hardness. In addition, we show that the decision version of VEST without the parameter k is an undecidable problem, and we give a fixed-parameter tractable algorithm for matrices of bounded size over finite fields, parameterized by the matrix dimensions and the order of the field.

This chapter is based on the articles [Sko22] and [BKSS23]. The results from the former were incorporated to the latter which is a joint work with Cornelius Brand, Viktoriia Korchemna and Kirill Simonov.

1. Preliminaries

In this chapter, we briefly recall basic terminology regarding simplicial complexes and computational complexity up to the level we need in this thesis. However, we generally assume that the reader is familiar with such notions. Thus, the main purpose of this chapter is to set up the notation.

1.1 Simplicial complexes and topology

Abstract simplicial complex. An *abstract simplicial complex* is a finite set system $\mathbf{K}$ such that if $\sigma \in \mathbf{K}$ and $\sigma' \subseteq \sigma$, then $\sigma' \in \mathbf{K}$. Elements of $\mathbf{K}$ are *faces*; a k-face is a face of *dimension* k, that is, a face of size $k+1$. Vertices correspond to 0-faces of $\mathbf{K}$ (specifically, a vertex v corresponds to a 0-face $\{v\}$); and the set of vertices is denoted $V(\mathbf{K})$. The dimension of $\mathbf{K}$ is the dimension of the largest face (or $-\infty$, if $\mathbf{K}$ is empty). The *k-skeleton* of a complex $\mathbf{K}$, is a subcomplex $\mathbf{K}^{(k)}$ consisting of faces of dimension at most k. The complex $\mathbf{K}$ is *pure* if all inclusion-wise maximal faces have the same dimension.

Such structure is essentially a combinatorial description of the following structure.

Geometric simplicial complex. A *geometric simplicial complex* is a collection of (geometric) simplices embedded in some $\mathbb{R}^m$ such that two simplices intersect in a face of both of them; and a face of any simplex in the complex belongs again to the complex. The *dimension* of a simplex is the number of its vertices minus one; the *dimension* of a simplicial complex is the maximum of the dimensions of simplices appearing in the complex.

Each geometric simplicial complex determines an abstract simplicial complex. Indeed, faces of an abstract simplicial complex can be viewed as sets of vertices of simplices of a geometric simplicial complex. The converse also holds. In other words, each abstract simplicial complex $\mathbf{K}$ has a geometric realization as a geometric simplicial complex:

- Let n be the number of vertices of the abstract simplicial complex.
- Pick an n-simplex. That is, a simplex on n vertices which lives in $\mathbb{R}^{n-1}$.
- Arbitrarily identify vertices of the simplicial complex with the vertices of the n-simplex.
- The corresponding geometric complex consists of $\{\mathrm{conv}(\sigma); \sigma \in \mathbf{K}\}$.

Since these two structures are basically two different descriptions of the same mathematical object we will usually say just simplicial complex. Let us only point out that for the purpose of Chapter 2 it is sufficient to speak only about abstract simplicial complexes. However, in Chapter 3 we also want to work with polyhedra. Therefore, we will be using geometric simplicial complexes (with a single exception that the input for any computational problem we consider is the corresponding abstract simplicial complex). For more details on simplicial complexes, we refer to textbooks such as [RS82, Mat03].

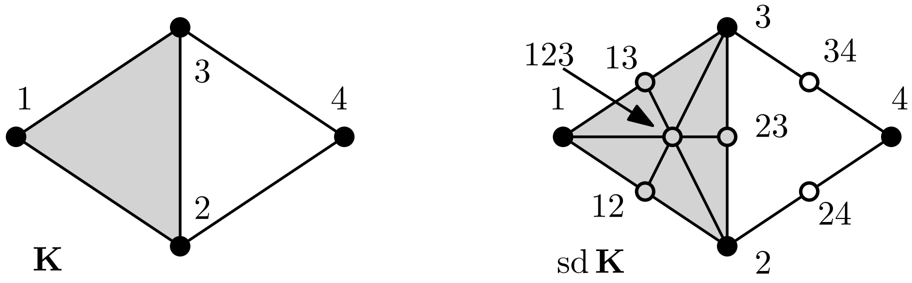

Figure 1.1: The barycentric subdivision $\mathrm{sd}\,\mathbf{K}$ of a complex $\mathbf{K}$. The notation on the right picture is simplified so that 12 stands for $\{1, 2\}$, etc.

Polyhedron. We work with *polyhedra* as defined in [RS82]. When we say 'polyhedron' we always mean a compact polyhedron. Because every compact polyhedron can be triangulated, an equivalent definition is that a polyhedron is the underlying space $|\mathbf{K}| := \bigcup_{\sigma\in\mathbf{K}} \sigma$ of some finite (geometric) simplicial complex $\mathbf{K}$ (a.k.a. the *polyhedron of* $\mathbf{K}$).

Simplicial mapping. A mapping $\Psi\colon V(\mathbf{K}_1) \to V(\mathbf{K}_2)$ between sets of vertices of simplicial complexes $\mathbf{K}_1, \mathbf{K}_2$ is called *simplicial* if $\Psi(\sigma) \in \mathbf{K}_2$ for each $\sigma \in \mathbf{K}_1$. Such a mapping also induces a continuous mapping between $|\mathbf{K}_1|$ and $|\mathbf{K}_2|$, i.e. the polyhedrons of $\mathbf{K}_1$ and $\mathbf{K}_2$, respectively.

Subdivision of simplicial complexes. A geometric simplicial complex $\mathbf{K}'$ is a *subdivision* of a complex $\mathbf{K}$ if $|\mathbf{K}'| = |\mathbf{K}|$ and every $\sigma' \in \mathbf{K}'$ is a subset of some $\sigma \in \mathbf{K}$. Given a subcomplex $\mathbf{L}$ of $\mathbf{K}$, then the *subcomplex* $\mathbf{L}'$ *of* $\mathbf{K}'$ *corresponding to* $\mathbf{L}$ *is the complex* $\mathbf{L}' := \{\sigma' \in K'\colon \sigma \subseteq |\mathbf{L}|\}$.

An important subdivision of a simplicial complex $\mathbf{K}$ is called *barycentric subdivision.* It is denoted by $\mathrm{sd}\,\mathbf{K}$ and on the level of abstract simplicial complexes it has the following nice combinatorial description.

$$\mathrm{sd}\,\mathbf{K} := \{\{\sigma_1, \dots, \sigma_n\}\colon \sigma_1, \dots, \sigma_n \in \mathbf{K},\ \emptyset \neq \sigma_1 \subsetneq \sigma_2 \subsetneq \cdots \subsetneq \sigma_n\}.$$

The geometric idea behind the definition of barycentric subdivision is the following: According to the definition, the vertices of $\mathrm{sd}\,\mathbf{K}$ are nonempty faces of $\mathbf{K}$. Place a vertex of $\mathrm{sd}\,\mathbf{K}$ into the barycenter of the face it represents in $\mathbf{K}$ (viewed as a geometric simplicial complex). Then $\mathrm{sd}\,\mathbf{K}$ represents a (geometric) subdivision of $\mathbf{K}$; see Figure 1.1.

Join, star, link of abstract simplicial complexes. A *join* of two abstract simplicial complexes $\mathbf{K}_1$ and $\mathbf{K}_2$ is the complex $\mathbf{K}_1 * \mathbf{K}_2 := \{\sigma_1 \sqcup \sigma_2\colon \sigma_1 \in \mathbf{K}_1, \sigma_2 \in \mathbf{K}_2\}$ where $\sqcup$ stands for disjoint union.[1] In our inductive arguments, we will carefully distinguish the empty complex $\emptyset$ and the complex $\{\emptyset\}$ containing a single face, which is $\emptyset$. Note that $\mathbf{K} * \emptyset = \emptyset$, whereas $\mathbf{K} * \{\emptyset\} = \mathbf{K}$.

[1] We can perform the disjoint union of two sets A and B even if A and B are not disjoint. The standard model in such case is to take $A \times \{1\} \cup B \times \{2\}$.

Given a face σ of an abstract simplicial complex $\mathbf{K}$, the *link* of σ in $\mathbf{K}$ is defined as $\mathrm{lk}(\sigma, \mathbf{K}) := \{\sigma' \setminus \sigma \colon \sigma' \in \mathbf{K}, \sigma \subseteq \sigma'\}$. The (closed) *star* of σ in $\mathbf{K}$ is defined as $\mathrm{st}(\sigma, \mathbf{K}) := \{\sigma' \in \mathbf{K} \colon \sigma' \cup \sigma \in \mathbf{K}\}$.

Note that it is possible to define join, star and link also on the level of geometric simplicial complexes. However, we only need these notions in Chapter 2 where we work only with abstract simplicial complexes.

Collapsibility and PL collapsibility. Let $\mathbf{K}$ be a simplicial complex (abstract or geometric) and $\sigma \in \mathbf{K}$ be a face which is contained in only one face $\tau \in K$ with $\sigma \subsetneq \tau$. (Necessarily $\dim \tau = \dim \sigma + 1$ and τ is a *facet* $\mathbf{K}$, that is, an inclusion-wise maximal face of $\mathbf{K}$). In this case, we say that σ is a *free* face of $\mathbf{K}$ and we also say that a complex $\mathbf{K}'$ *arises from* $\mathbf{K}$ *by an elementary collapse* if there are σ and τ as above such that $\mathbf{K}' = \mathbf{K} \setminus \{\sigma, \tau\}$, we denote this by $\mathbf{K} \searrow \mathbf{K}'$. A complex $\mathbf{K}$ is *collapsible*, if there is a sequence $(\mathbf{K}_1, \dots, \mathbf{K}_r)$ of complexes such that $\mathbf{K}_1 = \mathbf{K}$, $\mathbf{K}_r$ is a point, and $\mathbf{K}_1 \searrow \mathbf{K}_2 \searrow \cdots \searrow \mathbf{K}_r$. An important property of collapsibility is that the elementary collapses preserve the homotopy type, a fortiori, the homology groups.

A polyhedron P is *PL collapsible* if some triangulation of P is a collapsible simplicial complex. Similarly, a simplicial complex $\mathbf{K}$ is *PL collapsible* if $|\mathbf{K}|$ is a PL collapsible polyhedron. Here, we should point out a certain subtlety in the definition of PL collapsible simplicial complex: If $\mathbf{K}$ is PL collapsible, then there is some triangulation $\mathbf{K}'$ of $|\mathbf{K}|$ which is collapsible (in the simplicial sense). This triangulation $\mathbf{K}'$ need not be a priori a subdivision of $\mathbf{K}$. However, by [Hud69, Theorem 2.4] we may assume that $\mathbf{K}'$ actually is a subdivision of $\mathbf{K}$.

In general, collapsibility and PL collapsibility of a simplicial complex differ because PL collapsibility allows an arbitrarily fine subdivision before starting the collapses. In Chapter 3, we need both and we carefully distinguish these two notions.

Shellability. Let $\mathbf{K}$ be a pure k-dimensional simplicial complex (abstract or geometric). For its face ϑ, we denote $\mathbf{K}(\vartheta)$ the inclusion-wise minimal subcomplex of $\mathbf{K}$ containing ϑ.[2]

A total order $\vartheta_1, \dots, \vartheta_m$ of facets of $\mathbf{K}$ is called a *shelling* if for every $k \in \{2, \dots, m\}$ the complex $\mathbf{K}(\vartheta_k) \cap \left(\bigcup_{i=1}^{k-1} \mathbf{K}(\vartheta_i)\right)$ is a pure $(k-1)$-dimensional complex.

A simplicial complex $\mathbf{K}$ is then said to be *shellable* if it admits a shelling order. For comparison with collapsibility, we will also use the *reverse shelling order* $\vartheta_m, \dots, \vartheta_1$.

Homology. In our auxiliary computations in Chapters 2 and 3, we will often need homology groups, including the exact sequence for pairs, the Mayer-Vietoris exact sequence and the Lefschetz duality. In general, we refer to the literature such as [Hat02, Mun84] for details (in case of Lefschetz duality, we will recall its statement when used).

In all our computations, we work with homology with $\mathbb{Z}_2$-coefficients. When working with simplicial complexes, we use simplicial homology. In particular,

[2] In other words, $\mathbf{K}(\vartheta)$ consists of ϑ and all its faces. In the case of an abstract complex, $\mathbf{K}(\vartheta)$ is in fact the power set of ϑ.

when we speak of k-chains, then we can identify a k-chain with a collection of k-simplices (in its support). (Similarly, a k-cycle is such a collection with trivial boundary, i.e., each $(k-1)$-simplex is in an even number of k-simplices of the cycle.) In case of polyhedra, we use singular homology. However, we of course implicitly use that the simplicial and singular homology groups are (naturally) isomorphic (for a simplicial complex and its polyhedron).

Homotopy groups. The main problem of Chapter 4 is inspired by computation of homotopy groups of a topological space X, denoted by $\pi_k(X)$, for $k \geq 1$. However, we actually do not need their precise definition as we only deal with an intermediate problem related to the computation of them.

We briefly mention that the most intuitive of them is the group $\pi_1(X)$, which is often called the *fundamental group* of the space X. Intuitively, it is a group of continuous mappings $S^1 \to X$, where S^1 is the 1-dimensional sphere, up to homotopy. The higher homotopy groups ($k > 1$) are then groups of mappings from higher dimensional spheres to the topological space up to homotopy. In general, they carry more information than homology groups.

1.2 Computational complexity

Here, we briefly overview a few notions from computational complexity we need in this thesis. For more details see, e.g., [AB09, Chapter 2].

A decision problem belongs to the *class* NP if an affirmative answer to it can be verified in polynomial time using a certificate of polynomial size. A decision problem A is NP-*hard* if for each problem B from the class NP there is a polynomial time reduction from B to A. More precisely, given an instance q of the problem B one can construct in polynomial time in the size of q an instance p of the problem A such that the answer to q is yes if and only if the answer to p is yes.

An important NP-hard problem is the so called 3-satisfiability problem (it also belongs to NP). An input for the 3-satisfiability problem is a 3-CNF formula ϕ, that is, a boolean formula in conjunctive normal form where every clause contains exactly three literals.[3] The output is the answer whether the formula is satisfiable, that is, whether it is possible to assign the variables TRUE or FALSE so that the formula evaluates to TRUE in this assignment.

It is well known that 3-satisfiability is NP-hard. In order to show that another problem X is NP-hard, it is sufficient to construct a polynomial time reduction from 3-satisfiability to X.

Note that *if a problem is shown to be* NP*-hard one should not expect existence of a polynomial time algorithm solving this problem.* (This is equivalent to the standard conjecture $\mathsf{P} \neq \mathsf{NP}$ in theory of computation.)

[3] A *literal* is some variable x or its negation $\neg x$; a *clause* with 3 literals is a (sub)formula of form $(\ell_1 \vee \ell_2 \vee \ell_3)$ where ℓ_i are literals. A formula ϕ is in conjunctive normal form if it can be written as $\phi = c_1 \wedge c_2 \wedge \cdots \wedge c_m$ where c_j are clauses. An example of a 3-CNF formula is $(x \vee \neg y \vee z) \wedge (\neg x \vee \neg y \vee t)$.

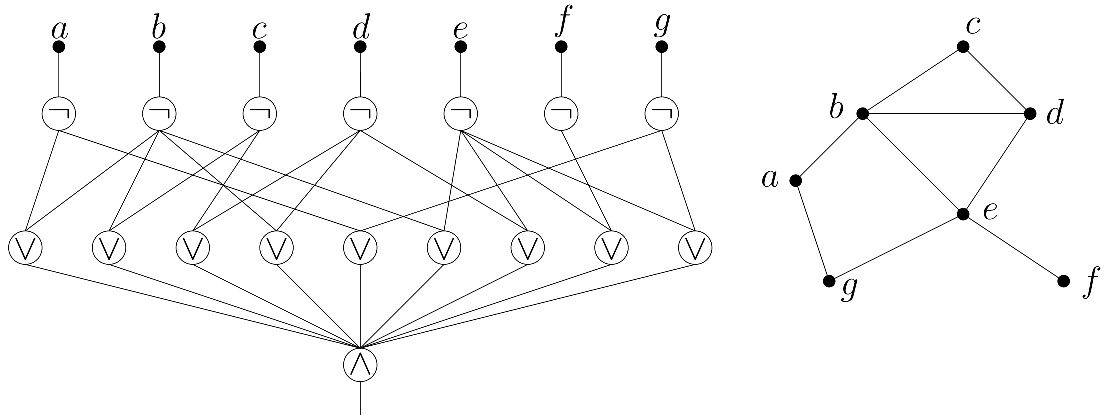

Figure 1.2: A boolean circuit solving the problem of existence of an independent set of size k in the graph on the left. There is an independent set of size k in the graph if and only if the boolean circuit outputs TRUE for an input consisting of exactly k true values.

1.3 Parameterized complexity

For the purpose of Chapter 4, we also need basic notions from parameterized complexity which classifies decision or counting computational problems with respect to a given parameter(s). For instance, one can ask if there exists an independent set of size k in a given graph or how many independent sets of size k (for counting version) are in a given graph, respectively, where k is the parameter. From this viewpoint, we can divide problems into several groups which form the W-*hierarchy*.

$$\mathsf{FPT} \subseteq \mathsf{W}[1] \subseteq \mathsf{W}[2] \subseteq \cdots \subseteq \mathsf{W}[\mathsf{P}] \subseteq \mathsf{XP}$$

The class FPT consists of decision problems solvable in time $f(k)n^c$, where $f(k)$ is a computable function of the parameter k, n is the size of input and c is a constant, while the class XP consists of decision problems solvable in time $cn^{f(k)}$. The class $\mathsf{W}[1]$ consists of all problems which admit a parameterized reduction to the satisfiability problem of a boolean circuit of constant depth with AND, OR and NOT gates such that there is at most 1 gate of higher input size than 2 on each path from the input gate to the final output gate (such number of larger gates in a circuit is called *weft*), where the parameter is the number of input gates set to TRUE. Here, a parameterized reduction from a parameterized problem A to a parameterized problem B is an algorithm that, given an instance (p, k) of A, in time $f(k)n^c$ produces an equivalent instance (q, k') of B such that $k' \leq g(k)$, for some computable functions $f(\cdot)$, $g(\cdot)$, and a constant c. See Figure 1.2 for an example of a reduction showing $\mathsf{W}[1]$-completeness of finding independent set of size k.

The class $\mathsf{W}[i]$ then consists of problems that admit a parameterized reduction to the satisfiability problem of a boolean circuit of a constant depth and weft at most i, parameterized by the number of input gates set to TRUE. Finally, the class $\mathsf{W}[\mathsf{P}]$ can be defined as a class of problems that can be solved by a non-deterministic Turing machine that can make at most $\mathcal{O}(g(k)\log n)$ non-deterministic choices and that works in time $f(k)n^c$ where $f(\cdot)$ is a computable function and c is a constant.

It is only known that $\mathsf{FPT} \subsetneq \mathsf{XP}$, while the other inclusions in the W-hierarchy are not known to be strict. However, it is strongly believed that $\mathsf{FPT} \subsetneq \mathsf{W}[1]$. *Therefore, one should not expect existence of an algorithm solving a* $\mathsf{W}[1]$*-hard problem in time $f(k)n^c$ where $f(k)$ is a computable function of k and c is a constant.* For the detailed presentation of W-hierarchy and parameterized complexity in general we refer the reader to [FG04a].

Analogously, one can define classes FPT and XP for counting problems. That is, a class of counting problems solvable in time $f(k)n^c$ or $cn^{f(k)}$, respectively. Problems for which there is a parameterized counting reduction to a problem of counting solutions for a boolean circuit of constant depth and weft at most i then form the class $\#\mathsf{W}[i]$. Note that there are decision problems from FPT whose counting versions are in $\#\mathsf{W}[1]$, e.g., counting paths or cycles of length k parameterized by k [FG04b]. Similarly to the decision case, *if a counting problem is shown to be* $\#\mathsf{W}[i]$*-hard for some i, one should not expect existence of an algorithm solving this problem in time $f(k)n^c$.* For more details on parameterized counting we refer the reader to [FG04b].

2. Shellings and sheddings induced by collapses

2.1 Introduction

Shellability and collapsibility are two widely used approaches for combinatorial decomposition of a simplicial complex. They are similar in spirit, yet there are important differences among those two notions. There are shellable complexes homotopy equivalent to a wedge of spheres, whereas no non-trivial wedge can be collapsible. On the other hand, two triangles sharing a vertex provide an example of a collapsible complex that is not shellable. Yet in some important cases, one can relate these two notions.

The easy direction is that shellability implies collapsibility whenever the complex is contractible (in fact, whenever the complex has trivial homology). We will focus here on a more interesting direction: when collapsibility implies shellability?

In this spirit, Hachimori [Hac08] proved the following theorem.

Theorem 2.1 ([Hac08]). *Let $\mathbf{K}$ be a 2-dimensional simplicial complex. Then the following statements are equivalent:*

(i) The complex $\mathbf{K}$ has a shellable subdivision.

(ii) The second barycentric subdivision $\operatorname{sd}^2 \mathbf{K}$ is shellable.

(iii) The link of each vertex of K is connected and K becomes collapsible after removing $\tilde{\chi}(\mathbf{K})$ triangles where $\tilde{\chi}$ denotes the reduced Euler characteristic.[1]

As Hachimori points out, one cannot expect that such an equivalence would be achievable in higher dimensions. Namely, the implication (i) $\Rightarrow$ (ii) cannot hold in higher dimensions due to the examples by Lickorish [Lic91]. However, we will show that it is possible to generalize the interesting implication (iii) $\Rightarrow$ (ii). The equivalence of (iii) and (ii) was one of the important steps in a recent proof of NP-hardness of recognition of shellable complexes [GPP$^+$19]. Though the hardness reduction requires the implication only in dimension 2, we find it interesting to provide a higher-dimensional generalization. For example, the computational complexity status of recognition of shellable/collapsible 3-spheres is unknown and the implication (iii) $\Rightarrow$ (ii) could provide a link between the two notions.

For explaining our generalization, we first introduce a removal-collapsibility condition.

Removal-collapsibility condition. We will say that a pure complex $\mathbf{K}$ satisfies the *removal-collapsibility* condition, abbreviated to (RC) condition, if $\mathbf{K}$ is either empty or $\mathbf{K}$ becomes collapsible after removing some number of facets. We remark that if $\dim \mathbf{K} = d$ the number of removed facets can be easily computed as

[1] The equivalence of (i) and (iii) in particular implies that $\tilde{\chi}(\mathbf{K})$ cannot be negative if $\mathbf{K}$ has a shellable subdivision. This can be deduced directly from the fact that those shelling steps that change the homotopy type decrease the Euler characteristic by 1 while the other shelling steps keep the Euler characteristic. For more details see [Hac08].

$\tilde{\beta}_d(\mathbf{K}; \mathbb{Z}_2)$ where $\tilde{\beta}_d(\mathbf{K}; \mathbb{Z}_2)$ denotes the reduced d-th Betti number, i.e., the rank of the reduced homology group $\tilde{H}_d(\mathbf{K}; \mathbb{Z}_2)$.[2] Indeed, by a routine application of the Mayer-Vietoris exact sequence, removing a facet either decreases $\tilde{\beta}_d(\mathbf{K}; \mathbb{Z}_2)$ by one or increases $\tilde{\beta}_{d-1}(\mathbf{K}; \mathbb{Z}_2)$ by one. But we cannot afford the latter case if the complex becomes collapsible after removing some number of facets. In addition, the lower dimensional homology remains unaffected when removing a facet (directly from the definition of simplicial homology or again by a Mayer-Vietoris exact sequence), therefore a complex satisfying (RC) condition also satisfies $\tilde{\beta}_i(\mathbf{K}; \mathbb{Z}_2) = 0$ for $0 \leq i \leq d-1$. In particular, $\tilde{\chi}(\mathbf{K}) = (-1)^d \tilde{\beta}_d(\mathbf{K}; \mathbb{Z}_2)$.

We also observe that if $d = 1$, that is, if $\mathbf{K}$ is a graph, then the (RC) condition is equivalent with stating that $\mathbf{K}$ is connected. Also, every 0-complex satisfies the (RC) condition.

Altogether, Hachimori's condition (iii) for 2-complexes is equivalent to saying that the link of the empty face (i. e., $\mathbf{K}$) and the link of every vertex satisfies the (RC) condition. This is furthermore equivalent with saying that the link of every face of $\mathbf{K}$ satisfies the (RC) condition as links of dimension at most 0 always satisfy the (RC) condition. We say that $\mathbf{K}$ satisfies the *hereditary removal-collapsibility* condition, abbreviated to (HRC) condition, if the link of every face of $\mathbf{K}$ satisfies the (RC) condition. In particular, (HRC) is equivalent to Hachimori's condition (iii) for 2-complexes. This condition is hereditary in the following sense: If $\mathbf{K}$ satisfies (HRC) and $\sigma \in \mathbf{K}$, then the link $\operatorname{lk}(\sigma, \mathbf{K})$ also satisfies (HRC). Indeed, the link of σ' in $\operatorname{lk}(\sigma, \mathbf{K})$ is just the link of $\sigma \cup \sigma'$ in $\mathbf{K}$.[3]

We establish the following generalization of Hachimori's implication (iii) $\Rightarrow$ (ii).

Theorem 2.2. *Let $\mathbf{K}$ be a pure simplicial d-complex satisfying the (HRC) condition, then the second barycentric subdivision $\operatorname{sd}^2 \mathbf{K}$ is shellable.*

We suspect that the reverse implication does not hold but we are not aware of a concrete complex violating the reverse implication. Possibly interesting examples could be the non-collapsible triangulations of the 3-ball $B_{15,66}$ and $B_{17,95}$ constructed by Benedetti and Lutz [BL13] but we do not know if their second barycentric subdivisions are shellable.

For the proof of Theorem 2.2, we will define two coarser notions than shellability called *star decomposability* and *star decomposability in vertices*, which may be of independent interest. Together with *vertex decomposability* of Provan and Billera [PB80] we will establish the following chain of implications, where the last implication is a result of Provan and Billera.

star decomposable in vertices $\Rightarrow$ star decomposable $\Rightarrow$ vertex decomposable $\Rightarrow$ shellable

Therefore, for a proof of Theorem 2.2 it is sufficient to prove the following generalization (together with the first two promised implications).

Theorem 2.3. *Let $\mathbf{K}$ be a pure simplicial d-complex satisfying the (HRC) condition, then the second barycentric subdivision $\operatorname{sd}^2 \mathbf{K}$ is star decomposable in vertices.*

[2] The choice of coefficients $\mathbb{Z}_2$ is not very important here. We could choose an arbitrary field.

[3] Note that we do not claim that (HRC) is hereditary with respect to subcomplexes or induced subcomplexes.

Additional motivation and background. Both notions, collapsibility and shellability, play an important role in PL topology because they may help to determine not only the homotopy type of a given collapsible/shellable space but sometimes even the (PL) homeomorphism type. For example, a collapsible PL manifold is a ball, and a shelling of a PL-manifold (if it does not change the homotopy type) preserves the homeomorphism type [RS82].

A relation between collapsibility or shellability of some subdivision of a complex and of some barycentric subdivision has been studied by Adiprasito and Benedetti [AB17]. Namely, they show that a simplicial complex is PL homeomorphic to a shellable complex if and only if it is shellable after finitely many barycentric subdivisions,[4] and they show an analogous result for collapsibility. If we were interested only in shellability of some barycentric subdivision of $\mathbf{K}$ in Theorem 2.2, it is possible that the proof could be easier, because it would be possible to use arbitrary suitable subdivisions in the intermediate steps.

Hachimori's implication (iii)$\Rightarrow$(ii), as well as its generalization, Theorem 2.2, can be understood as a tool for showing that a concrete complex is shellable. A lot of effort has been devoted to developing such tools in various contexts; see e.g. [BW83, Koz97]. The advantage of Theorem 2.2 could be that the (HRC) condition may naturally follow from the topological/combinatorial properties of a considered problem as it is in the case of the application of Hachimori's result in [GPP$^+$19]. A possible disadvantage could be that we have to allow some flexibility on the target complex (it has to be the second barycentric subdivision of another complex).

An additional piece of motivation may come from commutative algebra. For example, Herzog and Takayama [HT02] found out that if $\mathbf{K}$ is a complex (not necessarily pure) and $I_{\mathbf{K}}$ is the Stanley-Reisner ideal corresponding to $\mathbf{K}$, then $I_{\mathbf{K}}$ has linear quotients if and only the Alexander dual $\mathbf{K}^*$ is shellable (in the non-pure sense, but the pure case is a special case, of course). Thus, Theorem 2.2 may serve as a tool showing that certain Stanley-Reisner ideals have linear quotients.

Finally, the notions of star decomposability and star decomposability in vertices that we introduce along the way may be of independent interest as inductive tools similar to collapsibility, shellability, vertex-decomposability, etc. Although their definitions are slightly technical, they appear very naturally in our context, as we sketch in the proof strategy below. It would also be interesting to know whether these notions admit some counterpart in terms of commutative algebra (similarly to the Herzog-Takayama equivalence above).

Proof strategy. Here we first sketch Hachimori's proof (iii) $\Rightarrow$ (ii), in our words though. Then we sketch the necessary steps for upgrading the proof to higher dimensions.

Let $\mathbf{K}$ be a pure 2-complex satisfying the conditions of (iii). We want to sketch a strategy how to shell $\operatorname{sd}^2 \mathbf{K}$. For simplicity of pictures, we will assume that $\mathbf{K}$ is already collapsible (as we want to avoid the non-trivial second homology in the pictures).

The second barycentric subdivision $\operatorname{sd}^2 \mathbf{K}$ is covered by stars of vertices of $\operatorname{sd}^2 \mathbf{K}$ which correspond to original faces of $\mathbf{K}$; see Figure 2.1. The stars may

[4] The result is stated in terms of derived subdivisions but there is no difference on the combinatorial level.

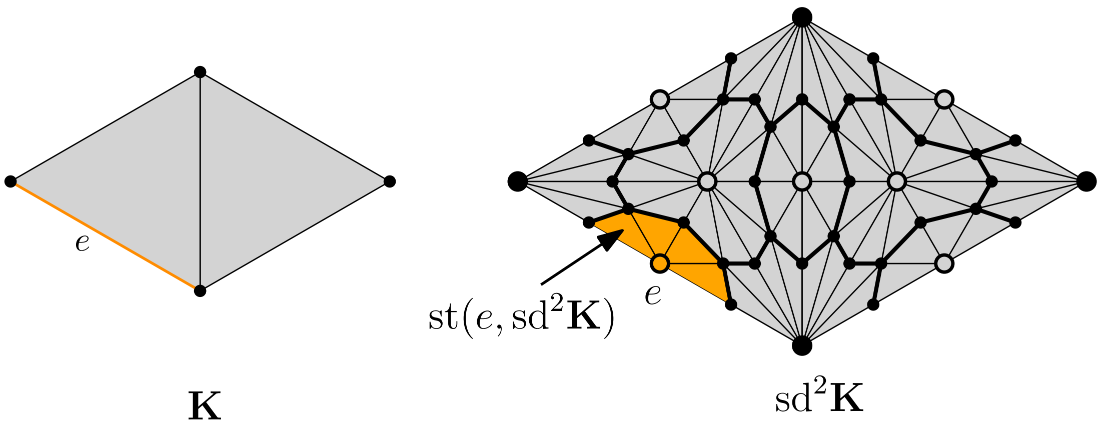

Figure 2.1: Decomposition of $\mathrm{sd}^2\,\mathbf{K}$ into stars. For example, an edge e of $\mathbf{K}$ becomes a vertex in $\mathrm{sd}^2\,\mathbf{K}$. Consequently, its star in $\mathrm{sd}^2\,\mathbf{K}$ is one of the stars in the decomposition.

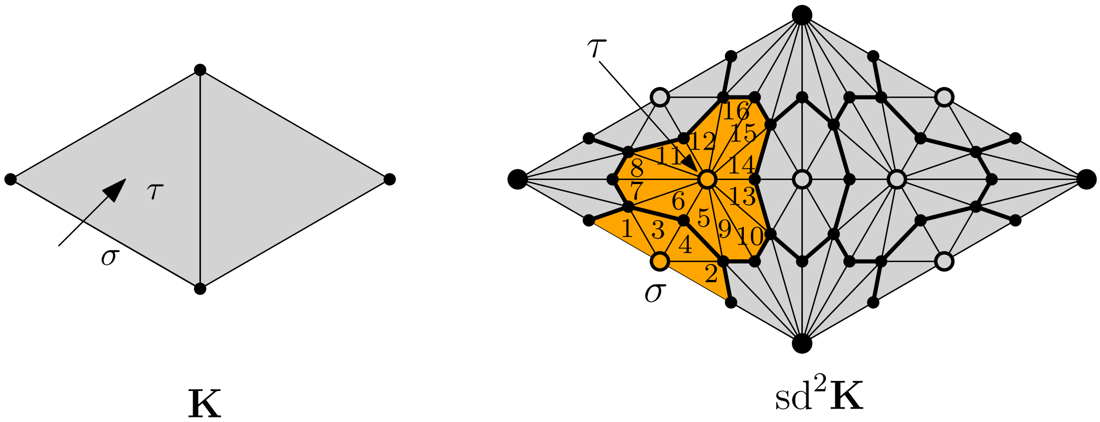

Figure 2.2: Reverse shelling of $\mathrm{sd}^2\,\mathbf{K}$ following an elementary collapse of $\mathbf{K}$. The numbers in triangles indicate a valid order of removing triangles.

overlap, but they overlap only in their boundaries (in links). Now, let us consider an elementary collapse $\mathbf{K} \searrow \mathbf{K}'$ while removing a free face σ and a maximal face τ containing σ. Naturally, in $\mathrm{sd}^2\,\mathbf{K}$ we want to emulate this by a reverse shelling removing the triangles first in $\mathrm{st}(\sigma, \mathrm{sd}^2\,\mathbf{K})$ and then in $\mathrm{st}(\tau, \mathrm{sd}^2\,\mathbf{K})$;[5] see Figure 2.2. This is indeed a good strategy as Hachimori [Hac08] showed. However, this quite heavily depends on the fact that the dimension of the complex is 2 as the structure of $\mathrm{sd}^2\,\mathbf{K}$ is so simple that all steps are obvious.

In general dimension we want to proceed similarly. However it seems out of reach to describe directly the order of removals of facets of $\mathrm{sd}^2\,\mathbf{K}$ and check that this is a shelling order due to a complicated structure of $\mathrm{sd}^2\,\mathbf{K}$. At least we initially tried this approach but we quickly got lost in addressing too many cases. Therefore, we instead use the aid of some coarser notions.

The first helpful notion is vertex decomposability introduced by Provan and

[5] Formally speaking, $\mathrm{st}(\sigma, \mathrm{sd}^2\,\mathbf{K})$ stands for $\mathrm{st}(\{\{\sigma\}\}, \mathrm{sd}^2\,\mathbf{K})$, etc.; see our convention in the preliminaries.

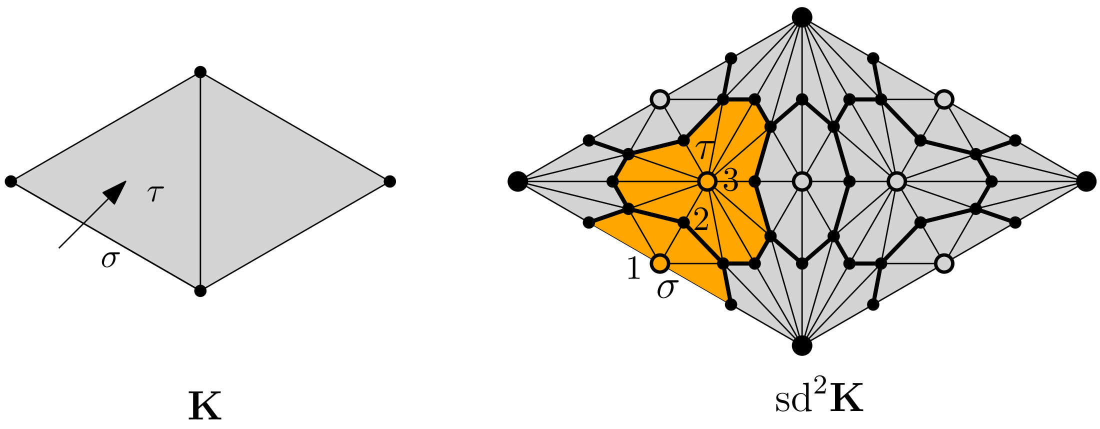

Figure 2.3: Vertex decomposition (shedding) of $\mathrm{sd}^2\,\mathbf{K}$ following an elementary collapse of $\mathbf{K}$. In this case, we first remove σ, then the vertex in between of σ and τ and finally τ.

Billera [PB80]. A simplicial d-complex $\mathbf{K}$ is *vertex decomposable* if it is pure and

- $\mathbf{K}$ is a d-simplex, or
- there is a vertex $v \in V(\mathbf{K})$ such that $\mathbf{K} - v$ is d-dimensional vertex decomposable (where $\mathbf{K} - v$ denotes the complex obtained by removing v and all faces containing v from $\mathbf{K}$) and $\mathrm{lk}(v, \mathbf{K})$ is $(d-1)$-dimensional vertex decomposable.

This recursive definition induces an order $v_1, \ldots, v_{n-(d+1)}$ of $n-(d+1)$ vertices of $\mathbf{K}$ according to the sequence of vertex removals in the second item (where n is the number of vertices of $\mathbf{K}$). This order is called a *shedding order* and we artificially extend any shedding order to all vertices of $\mathbf{K}$ so that the remaining vertices follow in arbitrary order. (Intuitively, as soon as we reach a d-simplex in the first item, we allow removing vertices in arbitrary order.)

Proving that $\mathrm{sd}^2\,\mathbf{K}$ is vertex decomposable is stronger than showing that $\mathrm{sd}^2\,\mathbf{K}$ is shellable, and it also seems easier to specify the shedding order as we deal with a smaller number of objects. For example, in case of the collapse from Figure 2.2, we specify the order only on three vertices; see Figure 2.3.

On the other hand, it is even easier to start removing the closed stars of vertices (and then taking a closure to get again a simplicial complex). In case of Figure 2.3, we would first remove the closed star of σ in $\mathrm{sd}^2\,\mathbf{K}$. Subsequently, when taking the closure, we reintroduce the full link of σ. Thus in this case, our first step coincides with removing σ (and therefore the open star of σ). The second step is, however, more interesting (see Figure 2.4): First we remove the closed star of τ. Then, when taking the closure, we do not reintroduce the vertex in between of σ and τ. Therefore, this second step removes simultaneously two vertices.

This will be our notion of star decomposability; however, one of the key steps in our approach is to identify an appropriate property of order of removals as above, which implies vertex decomposability of our complex. For sketching the idea, let us again consider the case of removing the closed star of τ in the second

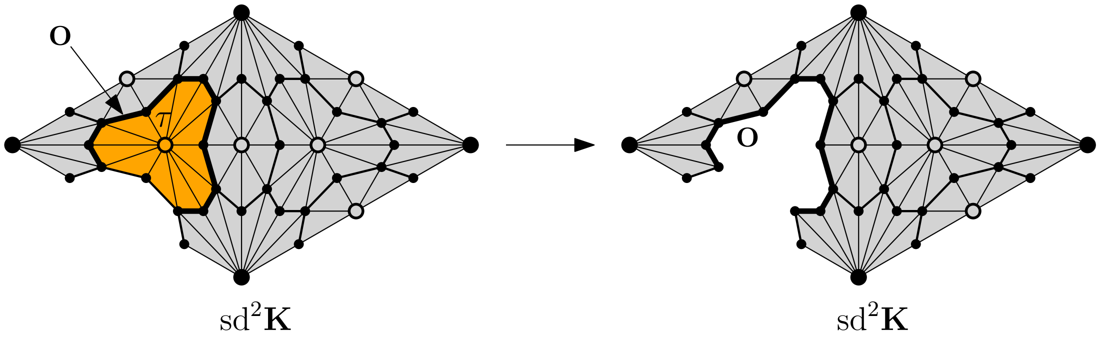

Figure 2.4: Overlap of the link of τ and the rest of the complex.

step above. Similarly as in the case of vertex decomposability, we will need that the link of the center of the removed star (in this case the link of τ) is star decomposable. However, this is not the only condition that we require. Let $\mathbf{O}$ be the overlap of the link of τ and the remainder of the complex after removing the star of τ (see Figure 2.4). We will actually need a star decomposition of the link of τ such that $\mathbf{O}$ is an intermediate step in this decomposition. Overall, this additional condition ensures a well working induction for deducing vertex decomposability. We postpone the precise definition of star decomposability to Section 2.2.

Finally, we will utilize the fact that we are interested in star decomposability of the complex $\mathrm{sd}^2\,\mathbf{K}$ which is a barycentric subdivision of another complex, namely $\mathrm{sd}^2\,\mathbf{K} = \mathrm{sd}\,\mathbf{L}$ where $\mathbf{L} = \mathrm{sd}\,\mathbf{K}$. We will introduce the notion of star decomposability in vertices which will mean that we are removing only stars centered in vertices of $\mathrm{sd}\,\mathbf{L}$ which are simultaneously vertices of $\mathbf{L}$ as in Figure 2.4. (Note that vertices of $\mathbf{L}$ are faces of $\mathbf{K}$.) This brings one more advantage. We will essentially need claims of the following spirit: If $\mathrm{sd}(\mathbf{X})$ and $\mathrm{sd}(\mathbf{Y})$ are star decomposable in vertices, then $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$ is star decomposable in vertices as well (here $\mathbf{X} * \mathbf{Y}$ denotes the join of $\mathbf{X}$ and $\mathbf{Y}$). In addition, we will also need to describe the order of the star decomposition in vertices of $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$. Though it is probable that analogous claims are valid also for star decomposability, vertex decomposability and/or shellability, the notion of star decomposability in vertices removes at least one layer of complications in the proof: It is just sufficient to describe the order of the decomposition of $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$ as some total order on $V(\mathbf{X} * \mathbf{Y}) = V(\mathbf{X}) \sqcup V(\mathbf{Y})$ via a suitable way of interlacing the total orders on $V(\mathbf{X})$ and $V(\mathbf{Y})$ (here $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ denotes the disjoint union of $V(\mathbf{X})$ and $V(\mathbf{Y})$).

Convention regarding notation. Let us recall that we use the following combinatorial description of barycentric subdivision for an abstract simplicial complex $\mathbf{K}$. (For more details see Chapter 1.)

$$\mathrm{sd}\,\mathbf{K} := \{\{\sigma_1, \dots, \sigma_n\} \colon \sigma_1, \dots, \sigma_n \in \mathbf{K},\ \emptyset \neq \sigma_1 \subsetneq \sigma_2 \subsetneq \cdots \subsetneq \sigma_n\}.$$

Therefore, if v is a vertex $\mathbf{K}$, then $\{v\}$ is a vertex of $\mathrm{sd}\,\mathbf{K}$. If there is no risk of confusion, we write v instead of $\{v\}$ in formulas such as $\mathrm{lk}(v, \mathrm{sd}\,\mathbf{K})$. We apply similar conventions to the second barycentric subdivision, so we write $\mathrm{lk}(v, \mathrm{sd}^2\,\mathbf{K})$ instead of the cumbersome $\operatorname{lk}(\{\{v\}\}, \operatorname{sd}^2 \mathbf{K})$, or $\operatorname{lk}(\sigma, \operatorname{sd}^2 \mathbf{K})$ instead of $\operatorname{lk}(\{\sigma\}, \operatorname{sd}^2 \mathbf{K})$ if σ is a face of $\mathbf{K}$.

2.2 Star decomposability

Given a simplicial complex $\mathbf{X}$ and a set $W \subseteq V(\mathbf{X})$, we say that W *induces a star partition* of $\mathbf{X}$ if

(i) $\mathbf{X} = \bigcup_{w \in W} \operatorname{st}(w, \mathbf{X})$, and

(ii) any two distinct vertices $w_1, w_2 \in W$ are not neighbors in $\mathbf{X}$.

An example of a set inducing a star partition is the set $\{w_1, w_2, w_3, w_4\}$ in Figure 2.5.

Now, let us assume that W induces a star partition. Given a total order $\prec$ on W, $W' \subseteq W$, and $w \in W$, we set $W'_{\succ w} := \{w' \in W' \colon w' \succ w\}$ and $W'_{\succeq w} := \{w' \in W' \colon w' \succeq w\}$. We will also use the notation

$$\operatorname{st}(W', \mathbf{X}) := \bigcup_{w' \in W'} \operatorname{st}(w', \mathbf{X})$$

for an arbitrary subset W' of $V(\mathbf{X})$. Furthermore, given $x \in W$ and a set $W' \subseteq W$, we define[6]

$$\begin{aligned} \mathbf{O}(x, W') &:= \operatorname{lk}(x, \mathbf{X}) \cap \operatorname{st}(W', \mathbf{X}) = \operatorname{lk}(x, \mathbf{X}) \cap \bigcup_{w' \in W'} \operatorname{st}(w', \mathbf{X}) \\ &= \operatorname{lk}(x, \mathbf{X}) \cap \bigcup_{w' \in W'} \operatorname{lk}(w', \mathbf{X}). \end{aligned} \tag{2.1}$$

See Figure 2.5. Note that this is the overlap mentioned in the introduction. Occasionally, if we need to emphasize dependency on $\mathbf{X}$, we write $\mathbf{O}_{\mathbf{X}}(x, W')$.

Now, we are ready to introduce star decomposability. Following the sketch in the introduction, we want to introduce star decomposability of a simplicial complex $\mathbf{X}$. However, in order to formulate all conditions correctly, we need to state this definition for pairs.

Definition 2.4 (Star decomposability)**.** Let $(\mathbf{X}, X)$ be a pair where $\mathbf{X}$ is a simplicial complex which is pure and k-dimensional, $k \geq -1$ (that is, $\mathbf{X} \neq \emptyset$), and $X \subseteq V(\mathbf{X})$. We inductively define *star decomposability* of the pair $(\mathbf{X}, X)$. We also say that $\mathbf{X}$ is *star decomposable* if there is $X \subseteq V(\mathbf{X})$ for which the pair $(\mathbf{X}, X)$ is star decomposable.

For $k = -1$, the pair $(\{\emptyset\}, \emptyset)$ is star decomposable.

If $k \geq 0$, then $(\mathbf{X}, X)$ is star decomposable, if there is a set $W \neq \emptyset$ inducing a star partition and a total order $\prec$ on W with the following properties.

Order condition: $X = W_{\succeq w'}$ for some $w' \in W$.

Link condition: For any vertex $w \in W$ except for the last vertex in the order $\prec$, there is a nonempty set $U = U(w) \subseteq V(\operatorname{lk}(w, \mathbf{X}))$ such that

[6]The symbol $\mathbf{O}$ in the notation stands for the 'overlap' of $\operatorname{lk}(x, \mathbf{X})$ and $\operatorname{st}(W', \mathbf{X})$.

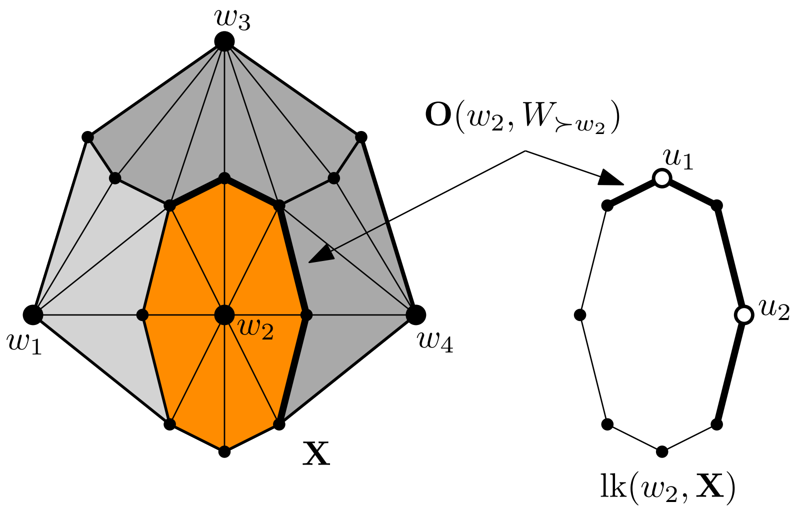

Figure 2.5: An example of the star decomposition induced by the set $W = \{w_1, w_2, w_3, w_4\}$ with the order $w_1 \prec w_2 \prec w_3 \prec w_4$ (left) and an example of the set $U(w_2) = \{u_1, u_2\}$ such that $\mathrm{st}(U(w_2), \mathrm{lk}(w, \mathbf{X})) = \mathbf{O}(w_2, W_{\succ w_2})$ and the pair $(\mathrm{lk}(w_2, \mathbf{X}), U(w_2))$ is star decomposable (right).

- $\mathrm{st}(U, \mathrm{lk}(w, \mathbf{X})) = \mathbf{O}(w, W_{\succ w})$ and
- the pair $(\mathrm{lk}(w, \mathbf{X}), U)$ is star decomposable.

Last vertex condition: For the last vertex $\hat{w} \in W$ in the order $\prec$, the link $\mathrm{lk}(\hat{w}, \mathbf{X})$ is star decomposable.

If the order $\prec$ on W satisfies the three conditions above, we say that $\prec$ induces a *star decomposition* of $(\mathbf{X}, X)$.

See Figure 2.5 for an example.

Remarks 2.5.

(i) Observe that the order condition implies $X \neq \emptyset$ if $k \geq 0$.

(ii) In the definition above, we remark that if $\mathbf{X}$ is k-dimensional and pure, for $k \geq 0$, then for any $w \in V(\mathbf{X})$, the link $\mathrm{lk}(w, \mathbf{X})$ is $(k-1)$-dimensional and pure. Therefore, in the last two conditions, we indeed refer to star decomposability of a pure complex of smaller dimension.

In addition, for any $W' \subseteq V(\mathbf{X})$, $W' \neq \emptyset$, $\mathrm{st}(W', \mathbf{X})$ is k-dimensional and pure. In particular, when replacing $\mathbf{X}$ with $\mathrm{lk}(w, \mathbf{X})$, we get that $\mathbf{O}(w, W_{\succ w}) = \mathrm{st}(U, \mathrm{lk}(w, \mathbf{X}))$ is $(k-1)$-dimensional and pure.

(iii) If $k = 0$, then every pair $(\mathbf{X}, X)$ is star decomposable if and only if $X \neq \emptyset$. Indeed, the only if part follows from (i). For the 'if' part, we observe that we can set $W = V(\mathbf{X})$ and we can use any order $\prec$ on W such that $X = W_{\succeq w'}$ for some w'. Both the link condition and the last vertex condition refer to star decomposability of $(\{\emptyset\}, \emptyset)$, which we assume.

(iv) If $k = 1$, then it is not difficult to show that $\mathbf{X}$ is star decomposable if and only if $\mathbf{X}$ is a connected bipartite graph. Note that requiring that $\mathbf{X}$ is connected is a must as we want to get that star decomposability implies vertex decomposability. Here is the place where the possibly slightly mysterious property '$X \neq \emptyset$ if $k \geq 0$' comes into the play. Indeed, this property and the link condition achieve that the overlap $\mathbf{O}(w, W_{\prec w})$ is nonempty, thus $\mathbf{X}$ must be a connected graph.

2.3 Star decomposability implies vertex decomposability.

In this section, we want to describe how star decomposability implies vertex decomposability. We start with a simple (folklore) lemma verifying that some order is a shedding order (with respect to our convention that we extend the shedding order also to the vertices of the last simplex). Given a simplicial complex $\mathbf{X}$, a total (or partial) order $\prec$ on $V(\mathbf{X})$, and $v \in V(\mathbf{X})$, by $\mathbf{X}_{\succ v}$ we denote the subcomplex of $\mathbf{X}$ induced by vertices that are greater than v. Similarly $\mathbf{X}_{\succeq v}$ is induced by v and the vertices that are greater than v.

Lemma 2.6. *Let $\mathbf{X}$ be a pure k-dimensional simplicial complex, $k \geq 0$. Let $\prec$ be a total order on $V(\mathbf{X})$. Then $\prec$ is a shedding order if and only if for every vertex v except for the last $k+1$ vertices, the link $\mathrm{lk}(v, \mathbf{X}_{\succeq v})$ is vertex decomposable and $(k-1)$-dimensional, and $\mathbf{X}_{\succ v}$ is pure k-dimensional.*

Proof. The 'only if' part of the statement follows immediately from the definition of vertex decomposability and the shedding order, thus we focus on the 'if' part.

If $\mathbf{X}$ has $k+1$ vertices, then $\mathbf{X}$ is a k-simplex and we are done. Otherwise, we proceed by induction on the number of vertices of $\mathbf{K}$.

Let v_1 be the first vertex in the order $\prec$. Then we need to check that $\mathrm{lk}(v_1, \mathbf{X}_{\succeq v_1})$ is vertex decomposable and $(k-1)$-dimensional, which is part of the assumptions. We also need to check that $\mathbf{X} - v_1 = \mathbf{X}_{\succ v_1}$ is vertex decomposable and k-dimensional. Again, k-dimensional is part of the assumptions, thus, it remains to check that $\mathbf{X} - v_1$ is vertex decomposable. However, this follows from the induction applied to $\mathbf{X}_{\succ v_1}$ and $\prec$ restricted to $V(\mathbf{X}) \setminus \{v_1\}$. □

Now, let $\mathbf{X}$ be a star decomposable simplicial complex, let W be a subset of $V(\mathbf{X})$ which induces a star partition of $\mathbf{X}$ and let $\prec$ be a total order which induces a star decomposition of $\mathbf{X}$. We will define a suitable partial order $\prec'$ on $V(\mathbf{X})$ extending $\prec$ such that the desired shedding order in the vertex decomposition of $\mathbf{X}$ will follow $\prec'$.

For arbitrary $v \in V(\mathbf{X})$, let $p(v)$ be the last vertex in the $\prec$ order among the vertices $w \in W$ such that $v \in \mathrm{st}(w, \mathbf{X})$. In particular $p(w) = w$ for any $w \in W$. If we want to emphasize $\prec$, we write $p(v, \prec)$ (which will be used in a single but important case of the proof of Theorem 2.8). Now, we define $\prec'$ in the following way. We set $v \prec' v'$ if $p(v) \prec p(v')$ for $v, v' \in V(\mathbf{X})$. In addition, we set $v \prec' w$ if $p(v) = w$ and $v \neq w$. Finally, if $p(v) = p(v')$ and $v, v' \notin W$, then v and v' are incomparable in $\prec'$. We say that $\prec'$ is *derived from* $\prec$. An example of this order is given in Figure 2.6 where $P(w) = \{v \in V(\mathbf{X}) \colon v \neq w,\ p(v) = w\}$ for $w \in W$; the elements in $P(w)$ are incomparable.

We will often need that $\mathrm{st}(W_{\succ w}, \mathbf{X})$ is an induced subcomplex of $\mathbf{X}$ for $w \in W \setminus \{\hat{w}\}$:

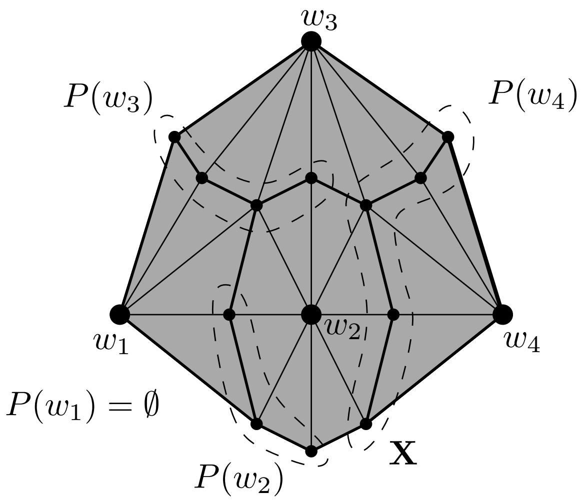

$$P(w_1) \prec' w_1 \prec' P(w_2) \prec' w_2 \prec' P(w_3) \prec' w_3 \prec' P(w_4) \prec' w_4$$

Figure 2.6: The set $P(w)$ and the auxiliary order $\prec'$ for the star decomposition in Figure 2.5.

Lemma 2.7. *Let $\mathbf{X}$ be a star decomposable complex, let W be a subset of $V(\mathbf{X})$ which induces a star partition of $\mathbf{X}$ and let $\prec$ be a total order on W which induces a star decomposition of $\mathbf{X}$. Let $\prec'$ be the partial order on $V(\mathbf{X})$ derived from $\prec$ and let $w \in W$ be different from the last vertex $\hat{w}$. Then $\mathrm{st}(W_{\succ w}, \mathbf{X})$ is the induced subcomplex $\mathbf{X}_{\succ' w}$ of $\mathbf{X}$.*

Proof. If $\dim \mathbf{X} = -1$, then the statement is void. If $\dim \mathbf{X} = 0$, the assertion easily follows from Remark 2.5(iii). Thus, we may assume $\dim \mathbf{X} \geq 1$, which we will implicitly when referring to the link condition.

Recall that $\mathrm{st}(W_{\succ w}, \mathbf{X}) = \bigcup_{w^+ \in W_{\succ w}} \mathrm{st}(w^+, \mathbf{X})$. It is easy to check the inclusion $\mathbf{X}_{\succ' w} \supseteq \bigcup_{w^+ \in W_{\succ w}} \mathrm{st}(w^+, \mathbf{X})$ because $\mathrm{st}(w^+, \mathbf{X}) \subseteq \mathbf{X}_{\succ' w}$ for every $w^+ \succ w$. (Note that if v is a neighbor of w^+ in $\mathbf{X}$, then $p(v) \succeq w^+ \succ w$. Thus, v belongs to $V(\mathbf{X}_{\succ' w})$.) Therefore, it remains to show $\mathbf{X}_{\succ' w} \subseteq \bigcup_{w^+ \in W_{\succ w}} \mathrm{st}(w^+, \mathbf{X}_{\succ' w})$.

Let $\sigma \in \mathbf{X}_{\succ' w}$. For contradiction, let us assume that $\sigma \notin \mathrm{st}(w^+, \mathbf{X})$ for all $w^+ \succ w$. (In particular, $\sigma \neq \emptyset$.) Let $w^- \preceq w$ be the largest vertex in W (according to the total order $\prec$) such that $\sigma \in \mathrm{st}(w^-, \mathbf{X})$. Such w^- must exist because W induces a star partition of $\mathbf{X}$. In addition, because $\sigma \in \mathbf{X}_{\succ' w}$ and $w^- \preceq w$, we get that $w^- \notin \sigma$. Thus, $\sigma \in \mathrm{lk}(w^-, \mathbf{X})$.

Now, we use that $\mathbf{X}$ is star decomposable. Namely, we use the link condition for w^-. There is $U \subseteq V(\mathrm{lk}(w, \mathbf{X}))$ such that $\mathrm{st}(U, \mathrm{lk}(w^-, \mathbf{X})) = \mathbf{O}(w^-, W_{\succ w^-})$ and the pair $(\mathrm{lk}(w^-, \mathbf{X}), U)$ is star decomposable. By the order condition for this pair, there is a set $Z \neq \emptyset$ inducing a star partition of $\mathrm{lk}(w^-, \mathbf{X})$ and a total order $\lhd$ on Z such that $U = Z_{\unrhd z'}$ for some $z' \in Z$. Because $\sigma \in \mathrm{lk}(w^-, \mathbf{X})$ and Z induces a star partition of $\mathrm{lk}(w^-, \mathbf{X})$ some vertex v of σ has to belong to Z. If $v \in U$, then $\sigma \in \mathrm{st}(U, \mathrm{lk}(w^-, \mathbf{X})) = \mathbf{O}(w^-, W_{\succ w^-})$ which contradicts the fact that w^- is the largest vertex such that $\sigma \in \mathrm{st}(w^-, \mathbf{X})$. If $v \in Z \setminus U$, then $p(v) = w^-$ which contradicts $\sigma \in \mathbf{X}_{\succ' w}$. $\square$

Now, we are ready to state and prove that star decomposability implies vertex decomposability. As the reader may expect, the order $\prec'$ appears in the statement to allow a well working induction.

Theorem 2.8. *Let* $\mathbf{X}$ *be a star decomposable simplicial complex; let* W *be a subset of* $V(\mathbf{X})$ *which induces a star partition of* $\mathbf{X}$*; and let* $\prec$ *be a total order which induces a star decomposition of* $\mathbf{X}$*. Let* $\prec'$ *be the partial order on* $V(\mathbf{X})$ *derived from* $\prec$*. Then* $\mathbf{X}$ *is vertex decomposable in a shedding order extending* $\prec'$*.*

Proof. We prove the statement by induction on k, the dimension of $\mathbf{X}$. If $k = -1$, the complex $\{\emptyset\}$ is vertex decomposable according to the definition of vertex decomposability (it is regarded as a -1-simplex). Although it could be covered by the second induction step, we can observe that the case $k = 0$ is also easy as any order of removing vertices from a 0-complex is a shedding order.

Now, let us prove the theorem for some $k \geq 1$ assuming that it is valid for lower values.

We first describe a total order $\prec''$ on $V(\mathbf{X})$ extending $\prec'$. Then we verify that $\prec''$ is a shedding order. Recall that for $w \in W$, $P(w)$ is the set of vertices $v \in V(\mathbf{X})$ such that $p(v) = w$ but $v \neq w$; see Figure 2.6. To describe $\prec''$ it remains to describe $\prec''$ on each $P(w)$ separately. We distinguish whether w is the last vertex in $\prec$.

If $w = \hat{w}$ is the last vertex, then $P(\hat{w}) = V(\mathrm{lk}(\hat{w}, \mathbf{X}))$. By the last vertex condition (for star decomposability) $\mathrm{lk}(\hat{w}, \mathbf{X})$ is star decomposable, therefore vertex decomposable by induction as well. We set $\prec''$ on $P(\hat{w})$ as an arbitrary shedding order of $\mathrm{lk}(\hat{w}, \mathbf{X})$.

If w is not the last vertex, then $P(w) = V(\mathrm{lk}(w, \mathbf{X})) \setminus V(\mathbf{O}(w, W_{\succ w}))$. By the link condition, the pair $(\mathrm{lk}(w, \mathbf{X}), U)$ is star decomposable where $U \subseteq V(\mathrm{lk}(w, \mathbf{X}))$ satisfies $\mathrm{st}(U, \mathrm{lk}(w, \mathbf{X})) = \mathbf{O}(w, W_{\succ w})$.

Claim 2.8.1. *Let* $w \in W$ *be different from the last vertex* $\hat{w}$*. Then the link* $\mathrm{lk}(w, \mathbf{X})$ *is vertex decomposable in some shedding order* $\lhd''$ *that starts on* $P(w) = V(\mathrm{lk}(w, \mathbf{X})) \setminus V(\mathbf{O}(w, W_{\succ w}))$ *and then continues on* $V(\mathbf{O}(w, W_{\succ w}))$*.*

Proof. Consider a set $Z \subseteq V(\mathrm{lk}(w, \mathbf{X}))$ inducing a star partition of $\mathrm{lk}(w, \mathbf{X})$ and a total order $\lhd$ on Z witnessing that the pair $(\mathrm{lk}(w, \mathbf{X}), U)$ is star decomposable. In particular, $U = Z_{\unrhd z'}$ for some $z' \in Z$ by the order condition. Let $\lhd'$ be the partial order on $V(\mathrm{lk}(w, \mathbf{X}))$ derived from $\lhd$. By induction, $\mathrm{lk}(w, \mathbf{X})$ is vertex decomposable in a shedding order $\lhd''$ extending $\lhd'$.

In addition, by the link condition (on star decomposable $\mathbf{X}$) we get

$$\mathbf{O}(w, W_{\succ w}) = \mathrm{st}(U, \mathrm{lk}(w, \mathbf{X})) = \mathrm{st}(Z_{\unrhd z'}, \mathrm{lk}(w, \mathbf{X})).$$

The vertices of $\mathrm{st}(Z_{\unrhd z'}, \mathrm{lk}(w, \mathbf{X}))$ are exactly the vertices of $\mathrm{lk}(w, \mathbf{X})$ with $p(v, \lhd) \in Z_{\unrhd z'}$. Therefore all vertices in $V(\mathrm{lk}(w, \mathbf{X})) \setminus V(\mathbf{O}(w, W_{\succ w}))$ precede the vertices in $V(\mathbf{O}(w, W_{\succ w})) = V(\mathrm{st}(Z_{\unrhd z'}, \mathrm{lk}(w, \mathbf{X})))$ in the order $\lhd'$, a fortiori, in the order $\lhd''$, as we need. □

Now, we set $\prec''$ on $P(w)$ as the shedding order $\lhd''$ on $\mathrm{lk}(w, \mathbf{X})$ from Claim 2.8.1, restricted to $P(w)$; see Figure 2.7.

It remains to check that $\prec''$ is the required shedding order which we do via Lemma 2.6. Namely, given a vertex $v \in V(\mathbf{X})$ which is not one of the last $k + 1$

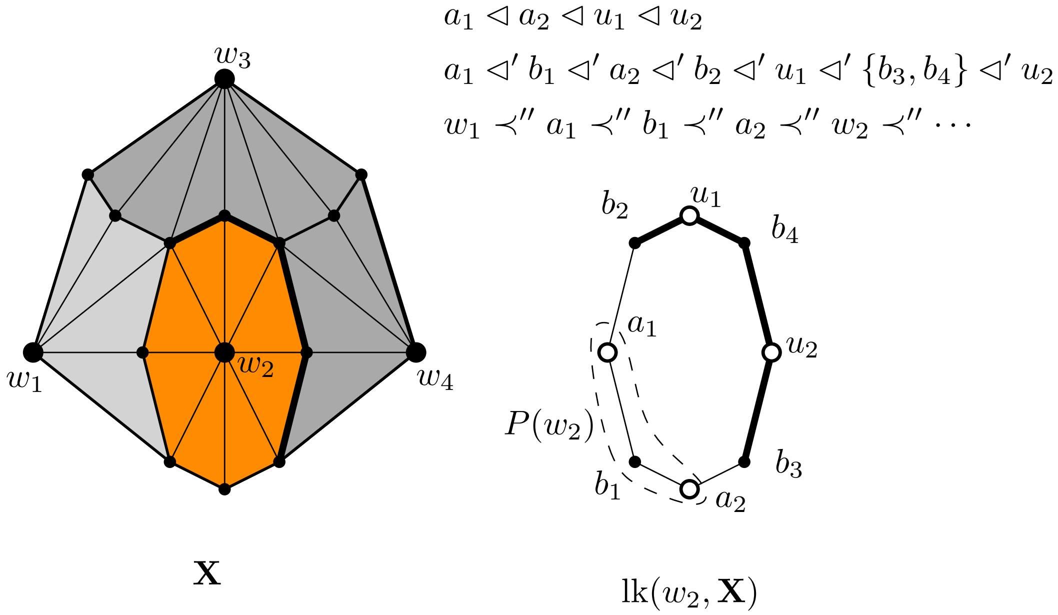

Figure 2.7: Setting up the order $\prec''$ on $P(w_2)$. The order $\lhd$ on $Z = \{a_1, a_2, u_1, u_2\}$ induces a star decomposition of $(\mathrm{lk}(w_2, \mathbf{X}), U)$ where $U = \{u_1, u_2\}$. Then $\lhd'$ is the corresponding partial order on $V(\mathrm{lk}(w_2, \mathbf{X}))$ (similarly as $\prec'$ corresponds to $\prec$). Finally, we take a shedding order $\lhd''$ on $\mathrm{lk}(w_2, \mathbf{X})$ extending $\lhd'$ (by induction) and restrict it to $P(w_2)$ obtaining $\prec''$.

vertices, we need to check that $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ is vertex decomposable and $(k-1)$-dimensional and that $\mathbf{X}_{\succ'' v}$ is pure k-dimensional. We distinguish whether $v \in W$.

Case 1, $v \in W$: We observe that v is not the last vertex $\hat{w}$ of $\prec$ as $\hat{w}$ is also the last vertex of $\prec''$. This allows to describe $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ as an overlap.

Claim 2.8.2. $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v}) = \mathbf{O}(v, W_{\succ v})$.

Proof. According to the definition of the overlap, we have $\mathbf{O}(v, W_{\succ v}) = \mathrm{lk}(v, \mathbf{X}) \cap \mathrm{st}(W_{\succ v}, \mathbf{X})$.

First, let us assume that $\sigma \in \mathrm{lk}(v, \mathbf{X}) \cap \mathrm{st}(W_{\succ v}, \mathbf{X})$. Each vertex v' of $\mathrm{st}(W_{\succ v}, \mathbf{X})$ satisfies $p(v') \succ v$ which implies $v' \succ'' v$. Therefore, each vertex of $\sigma \cup \{v\}$ belongs to $V(\mathbf{X}_{\succeq'' v})$. Because σ simultaneously belongs to $\mathrm{lk}(v, \mathbf{X})$, we get that it belongs to $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$.

Now, for the second inclusion, let us assume that $\sigma \in \mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$. Immediately, $\sigma \in \mathrm{lk}(v, \mathbf{X})$. Because $\sigma \in \mathbf{X}_{\succ' v} = \mathbf{X}_{\succ'' v}$, Lemma 2.7 gives $\sigma \in \mathrm{st}(W_{\succ v}, \mathbf{X})$. $\square$

By Claim 2.8.2, $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v}) = \mathbf{O}(v, W_{\succ v})$ which is $(k-1)$-dimensional by Remark 2.5(ii). In addition, $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ is vertex decomposable, as we checked that $\mathrm{lk}(v, \mathbf{X})$ is vertex decomposable in some shedding order starting with $P(v) = V(\mathrm{lk}(v, \mathbf{X})) \setminus V(\mathbf{O}(v, W_{\succ v}))$ and continuing with $V(\mathbf{O}(v, W_{\succ v}))$; see Claim 2.8.1. Also $\mathbf{X}_{\succ'' v} = \mathbf{X}_{\succ' v} = \mathrm{st}(W_{\succ v}, \mathbf{X})$ by Lemma 2.7. Therefore $\mathbf{X}_{\succ'' v}$ is pure k-dimensional by Remark 2.5(ii). This finishes Case 1.

Case 2, $v \notin W$: Let $w := p(v) \in W$. Note that, in particular, $w \succ' v$. We first check that $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ is vertex decomposable and $(k-1)$-dimensional. This will follow from the following two claims.

Claim 2.8.3. *The link* $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ *is the join of* w *and* $\mathrm{lk}(\{v, w\}, \mathbf{X}_{\succeq'' v})$.

Proof. The link $\mathrm{lk}(\{v, w\}, \mathbf{X}_{\succeq'' v})$ consists of simplices $\sigma \in \mathbf{X}_{\succeq'' v}$ satisfying $v, w \notin \sigma$, and $\sigma \cup \{v, w\} \in \mathbf{X}_{\succeq'' v}$. Therefore, the join of w and $\mathrm{lk}(\{v, w\}, \mathbf{X}_{\succeq'' v})$, considered as a subcomplex of $\mathbf{X}_{\succeq'' v}$, consists of simplices $\sigma \in \mathbf{X}_{\succeq'' v}$ satisfying

$$v \notin \sigma, \text{ and } \sigma \cup \{v, w\} \in \mathbf{X}_{\succeq'' v}. \tag{2.2}$$

On the other hand, $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ consists of simplices $\sigma \in \mathbf{X}_{\succeq'' v}$ satisfying

$$v \notin \sigma, \text{ and } \sigma \cup \{v\} \in \mathbf{X}_{\succeq'' v}. \tag{2.3}$$

A simplex $\sigma \in \mathbf{X}_{\succeq'' v}$ satisfying (2.2) immediately satisfies (2.3) as well. Thus, it remains to consider a simplex $\sigma \in \mathbf{X}_{\succeq'' v}$ satisfying (2.3); and to show that it satisfies (2.2).

First, we want to deduce that $\sigma \cup \{v\}$ belongs to $\mathrm{st}(w', \mathbf{X})$ for some $w' \succeq w$. If w is the first vertex of W in the order $\prec$, then this claim follows from the fact that W induces a star partition of $\mathbf{X}$. If w is not the first vertex of W, let w^- be the vertex that immediately precedes w in the order $\prec$. Note that $\sigma \in \mathbf{X}_{\succ' w^-}$. By Lemma 2.7, $\sigma \cup \{v\}$ belongs to $\mathrm{st}(w', \mathbf{X})$ for some $w' \succ w^-$, that is, $w' \succeq w$ as required.

Now, because $p(v) = w$, the only option is that $w' = w$. Therefore, $\sigma \cup \{v\} \in \mathrm{st}(w, \mathbf{X})$; that is, $\sigma \cup \{v, w\} \in \mathbf{X}$. Because all vertices of $\sigma \cup \{v, w\}$ belong to $\mathbf{X}_{\succeq'' v}$, σ satisfies (2.2). □

Claim 2.8.4. *The link* $\mathrm{lk}(\{v, w\}, \mathbf{X}_{\succeq'' v})$ *is vertex decomposable and* $(k-2)$-*dimensional.*

Proof. We will deduce the claim from the 'only if' part of Lemma 2.6 used with the pure $(k-1)$-dimensional complex $\mathrm{lk}(w, \mathbf{X})$ and the shedding order $\lhd''$, coming from Claim 2.8.1. Let us recall that $\prec''$ is defined so that it coincides with $\lhd''$ on $\mathrm{lk}(w, \mathbf{X})$ restricted to $P(w)$. Because $v \in P(w)$, we in particular get that $\mathrm{lk}(w, \mathbf{X})_{\succeq'' v} = \mathrm{lk}(w, \mathbf{X})_{\unrhd'' v}$.

In order to apply Lemma 2.6, we also check that v is not among the last k vertices of the aforementioned shedding $\lhd''$ of $\mathrm{lk}(w, \mathbf{X})$. If $w = \hat{w}$, we get this because we assume that v is not among the last $k+1$ vertices in the $\prec''$ order on $V(\mathbf{X})$ (the last one is $\hat{w}$, and the vertices of $P(\hat{w})$ immediately precede). If $w \neq \hat{w}$ we get this because the overlap $\mathbf{O}(w, W_{\succ w})$ is $(k-1)$-dimensional (see Remark 2.5(ii)), and the vertices of this overlap belong to $V(\mathrm{lk}(w, \mathbf{X}))$ while they do not belong to $P(w)$.

Now, using Lemma 2.6 as explained above, we get that $\mathrm{lk}(v, \mathrm{lk}(w, \mathbf{X})_{\succeq'' v}) = \mathrm{lk}(v, \mathrm{lk}(w, \mathbf{X})_{\unrhd'' v})$ is vertex decomposable and $(k-2)$-dimensional. Finally, $\mathrm{lk}(v, \mathrm{lk}(w, \mathbf{X})_{\succeq'' v}) = \mathrm{lk}(\{v, w\}, \mathbf{X}_{\succeq'' v})$ because $\mathbf{X}_{\succeq'' v}$ is an induced subcomplex of $\mathbf{X}$. □

It follows immediately from Claims 2.8.3 and 2.8.4 that $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ is $(k-1)$-dimensional. In addition, because the join of two vertex decomposable complexes

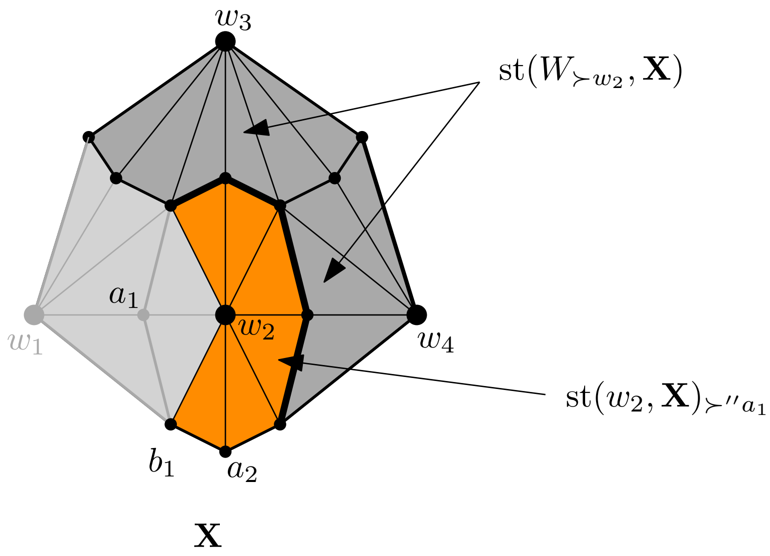

Figure 2.8: The complex $\mathbf{X}_{\succ'' a_1}$ as the union of $\mathrm{st}(W_{\succ w_2}, \mathbf{X})$ and $\mathrm{st}(w_2, \mathbf{X})_{\succ'' a_1}$. Here we use the order $\succ''$ from Figure 2.7.

is vertex decomposable [PB80, Proposition 2.4], we also get that $\mathrm{lk}(v, \mathbf{X}_{\succeq'' v})$ is vertex decomposable.

Finally, we need to check that $\mathbf{X}_{\succ'' v}$ is pure k-dimensional. We need one more claim; see also Figure 2.8.

Claim 2.8.5. *If $w \neq \hat{w}$, then $\mathbf{X}_{\succ'' v} = \mathrm{st}(W_{\succ w}, \mathbf{X}) \cup \mathrm{st}(w, \mathbf{X})_{\succ'' v}$. If, $w = \hat{w}$, then $\mathbf{X}_{\succ'' v} = \mathrm{st}(w, \mathbf{X})_{\succ'' v}$.*

Proof. If $w \neq \hat{w}$, then $\mathrm{st}(W_{\succ w}, \mathbf{X}) = \mathbf{X}_{\succ' w} = \mathbf{X}_{\succ'' w}$ by Lemma 2.7. Therefore it is sufficient to show that every $\sigma \in \mathbf{X}_{\succ'' v}$ which contains a vertex v' with $v' \preceq w$ belongs to $\mathrm{st}(w, \mathbf{X})_{\succ'' v}$. This will resolve both cases, $w = \hat{w}$ and $w \neq \hat{w}$, simultaneously. The ideas in the reminder of the proof are very similar to the ideas in the proof of Claim 2.8.3.

First, we check that $\sigma \in \mathrm{st}(w', \mathbf{X})$ for some $w' \succeq w$. If w is the first vertex of W, then this follows from the fact that W induces a star decomposition of $\mathbf{X}$. If w is not the first vertex of W, let w^- be the vertex of W that immediately precedes w. By Lemma 2.7, $\mathrm{st}(W_{\succ w^-}, \mathbf{X}) = \mathbf{X}_{\succ'' w^-}$. Because $\sigma \in \mathbf{X}_{\succ'' w^-}$, this implies that there is $w' \succ w^-$ with $\sigma \in \mathrm{st}(w', \mathbf{X})$.

On the other hand, σ cannot belong to $\mathrm{st}(w'', \mathbf{X})$ with $w'' \succ w$ as σ contains v' with $v' \preceq w$. Therefore, $w' = w$. Given that $\mathrm{st}(w, \mathbf{X})_{\succ'' v} = \mathbf{X}_{\succ'' v} \cap \mathrm{st}(w, \mathbf{X})$, we deduce that $\sigma \in \mathrm{st}(w, \mathbf{X})_{\succ'' v}$. $\square$

The union of two pure k-dimensional complexes is a pure k-dimensional complex. Therefore, due to Claim 2.8.5, it remains to check that $\mathrm{st}(W_{\succ w}, \mathbf{X})$ and $\mathrm{st}(w, \mathbf{X})_{\succ'' v}$ are pure k-dimensional (the former case applies only if $w \neq \hat{w}$).

Checking that $\mathrm{st}(W_{\succ w}, \mathbf{X})$ is pure k-dimensional is easy; see Remark 2.5(ii).

For checking that $\mathrm{st}(w, \mathbf{X})_{\succ'' v}$ is pure k-dimensional, we need that $\mathrm{lk}(w, \mathbf{X})_{\succ'' v}$ is pure $(k-1)$-dimensional. Because $v \in P(w)$, $\mathrm{lk}(w, \mathbf{X})_{\succ'' v} = \mathrm{lk}(w, \mathbf{X})_{\trianglerighteq'' v}$ where $\triangleleft''$ is the shedding of $\mathrm{lk}(w, \mathbf{X})$ as introduced below Claim 2.8.1. This means that $\mathrm{lk}(w, \mathbf{X})_{\succ'' v}$ is an intermediate step in the shedding $\trianglerighteq''$ of $\mathrm{lk}(w, \mathbf{X})$. If we realize that v is not among the last k vertices of the order $\lhd''$ on $\mathrm{lk}(w, \mathbf{X})$, then we can deduce that $\mathrm{lk}(w, \mathbf{X})_{\succ'' v}$ is pure and $(k-1)$-dimensional.

If $w \neq \hat{w}$, then $\mathrm{lk}(w, \mathbf{X})_{\succ'' v}$ still contains the overlap $\mathbf{O}(w, W_{\succ} w)$ which is $(k-1)$-dimensional by Remark 2.5(ii). If $w = \hat{w}$, then we assume that v is not among the last $k+1$ vertices of $\prec''$ while $\prec''$ and $\lhd''$ coincide on $P(\hat{w})$ and the vertices of $P(\hat{w})$ immediately precede $\hat{w}$ in $\prec''$. This finishes Case 2 and thereby the proof of the theorem. □

2.4 Star decomposability in vertices

Star decomposability of a barycentric subdivision. In our approach, we will need to consider the star decomposability of the barycentric subdivision $\mathrm{sd}(\mathbf{X})$ of a complex $\mathbf{X}$. In fact, we will consider only a special type of star decomposition of $\mathrm{sd}(\mathbf{X})$ using only stars of vertices of $\mathbf{X}$, that is, the faces of $\mathbf{X}$ which are actually vertices of X. For a well working induction, we will need that this property is kept also in the link condition and the last vertex condition of Definition 2.4. For stating this precisely, first, we need a more explicit description of $\mathrm{lk}(\vartheta, \mathrm{sd}(\mathbf{X}))$ if ϑ is a face (possibly a vertex) of $\mathbf{X}$.

Lemma 2.9. *Let ϑ be a face of a simplicial complex $\mathbf{X}$, then*

$$\mathrm{lk}(\vartheta, \mathrm{sd}\,\mathbf{X}) \cong \mathrm{sd}\,\partial\vartheta * \mathrm{sd}\,\mathrm{lk}(\vartheta, \mathbf{X}).$$

In particular, if x is a vertex of $\mathbf{X}$, then

$$\mathrm{lk}(x, \mathrm{sd}\,\mathbf{X}) \cong \mathrm{sd}\,\mathrm{lk}(x, \mathbf{X}).$$

Proof. We will construct a simplicial isomorphism

$$\Psi \colon V(\mathrm{lk}(\vartheta, \mathrm{sd}\,\mathbf{X})) \to V(\mathrm{sd}\,\partial\vartheta * \mathrm{sd}\,\mathrm{lk}(\vartheta, \mathbf{X})).$$

First, we observe that

$$V(\mathrm{sd}\,\partial\vartheta * \mathrm{sd}\,\mathrm{lk}(\vartheta, \mathbf{X})) = V(\mathrm{sd}\,\partial\vartheta) \sqcup V(\mathrm{sd}\,\mathrm{lk}(\vartheta, \mathbf{X})) = \partial\vartheta \sqcup \mathrm{lk}(\vartheta, \mathbf{X}).$$

Next, we realize that the vertices of $\mathrm{lk}(\vartheta, \mathrm{sd}\,\mathbf{X})$ are all the faces $\lambda \neq \emptyset, \vartheta$ of $\mathbf{X}$ such that $\{\lambda, \vartheta\}$ forms a simplex of $\mathrm{sd}\,\mathbf{X}$, that is, either $\emptyset \neq \lambda \subsetneq \vartheta$ or $\vartheta \subsetneq \lambda$. Thus, we can define Ψ in the following way

$$\Psi(\lambda) = \begin{cases} \lambda \in \partial\vartheta & \text{if } \emptyset \neq \lambda \subsetneq \vartheta, \\ \lambda \setminus \vartheta \in \mathrm{lk}(\vartheta, \mathbf{X}) & \text{if } \vartheta \subsetneq \lambda. \end{cases}$$

From the description above, it immediately follows that Ψ is a bijection. It is also routine to check that Ψ is a simplicial isomorphism. Indeed, a simplex of $\mathrm{lk}(\vartheta, \mathrm{sd}\,\mathbf{X})$ is a collection $\{\alpha_1, \ldots, \alpha_k, \beta_1, \ldots, \beta_\ell\}$ satisfying

$$\emptyset \neq \alpha_1 \subsetneq \cdots \subsetneq \alpha_k \subsetneq \vartheta \subsetneq \beta_1 \subsetneq \cdots \subsetneq \beta_\ell.$$

Such a simplex maps to a simplex $\{\alpha_1, \ldots, \alpha_k, \beta_1 \setminus \vartheta, \ldots, \beta_\ell \setminus \vartheta\}$ of $\mathrm{sd}\,\partial\vartheta * \mathrm{sd}\,\mathrm{lk}(\vartheta, \mathbf{X})$ and the inverse map works analogously (note that $\beta_i \setminus \vartheta$ is disjoint from ϑ whereas α_i are subsets of ϑ). □

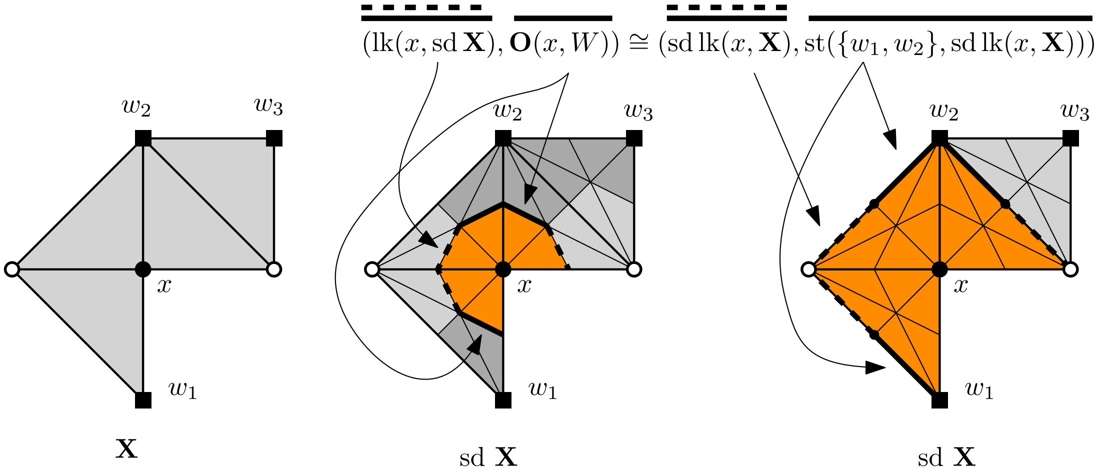

Figure 2.9: Isomorphism from Lemma 2.10 with $W = \{w_1, w_2, w_3\}$. The left hand side of the formula in Lemma 2.10 is depicted in the middle picture and the right hand side is in the right picture. Note that $W \cap V(\mathrm{lk}(x, \mathbf{X})) = \{w_1, w_2\}$ as w_3 does not belong to $\mathrm{lk}(x, \mathbf{X})$.

Now, we extend the isomorphism above to certain pairs; for the statement, recall that $\mathbf{O}(x, W)$ is defined via formula (2.1).

Lemma 2.10. *Let x be a vertex and W a subset of vertices of the simplicial complex $\mathbf{X}$ such that $x \notin W$. Then*

$$(\mathrm{lk}(x, \mathrm{sd}\,\mathbf{X}), \mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, W)) \cong (\mathrm{sd}\,\mathrm{lk}(x, \mathbf{X}), \mathrm{st}(W \cap V(\mathrm{lk}(x, \mathbf{X})), \mathrm{sd}\,\mathrm{lk}(x, \mathbf{X})))\,.$$

Though the formula in Lemma 2.10 may seem complicated at first sight, it has a nice geometric interpretation. All objects are subcomplexes of sd $\mathbf{X}$ and the isomorphism in the formula pushes the pair on the left hand side farther away from x; see Figure 2.9.

Proof. From Lemma 2.9 we have a simplicial isomorphism Ψ from $\mathrm{lk}(x, \mathrm{sd}\,\mathbf{X})$ to $\mathrm{sd}\,\mathrm{lk}(x, \mathbf{X})$. Hence, it remains to show that Ψ maps $\mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, w) := \mathrm{lk}(x, \mathrm{sd}\,\mathbf{X}) \cap \mathrm{lk}(w, \mathrm{sd}\,\mathbf{X})$ to $\mathrm{st}(w, \mathrm{sd}\,\mathrm{lk}(x, \mathbf{X}))$ for $w \in W \cap V(\mathrm{lk}(x, \mathbf{X}))$, where we use the explicit Ψ from the proof of Lemma 2.9, and that $\mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, w) = \emptyset$ for $w \in W \setminus V(\mathrm{lk}(x, \mathbf{X}))$. (Note that $\mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, W) = \bigcup_{w \in W} \mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, w)$.)

The faces of $\mathbf{O}_{\mathrm{sd}\,\mathbf{X}}(x, w)$ are collections $\{\beta_1, \dots, \beta_\ell\}$ of faces of $\mathbf{X}$ satisfying

$$\{x, w\} \subseteq \beta_1 \subsetneq \cdots \subsetneq \beta_\ell.$$

Let us emphasize that the first inclusion need not be strict. Therefore, $\mathbf{O}_{\mathrm{sd}\,X}(x, w)$ is non-empty if and only if $\{x, w\} \in \mathbf{X}$, that is, if and only if $w \in W \cap V(\mathrm{lk}(x, \mathbf{X}))$ as required. In sequel, we assume that $w \in W \cap V(\mathrm{lk}(x, \mathbf{X}))$.

The collections $\{\beta_1, \dots, \beta_\ell\}$ are mapped under Ψ to $\{\beta_1 \setminus \{x\}, \dots, \beta_\ell \setminus \{x\}\}$ satisfying the same condition due to the description of Ψ in the proof of Lemma 2.9. Setting $\gamma_j = \beta_j \setminus \{x\}$ we get

$$\{w\} \subseteq \gamma_1 \subsetneq \cdots \subsetneq \gamma_\ell$$

for γ_j not containing x but such that $\gamma_j \cup \{x\}$ is a face of $\mathbf{X}$, which is exactly a description of $\mathrm{st}(w, \mathrm{sd}(\mathrm{lk}(x, \mathbf{X})))$. $\square$

Now, we can define star decomposibility in vertices:

Definition 2.11 (Star decomposability in vertices)**.** Let $\mathbf{X}$ be a pure simplicial complex of dimension k where $k \geq -1$ and let $X \subseteq V(\mathbf{X})$. We inductively define *star decomposability in vertices* of the pair $(\operatorname{sd}\mathbf{X}, X)$. We also say that $\operatorname{sd}\mathbf{X}$ is *star decomposable in vertices* if the pair $(\operatorname{sd}\mathbf{X}, V(\mathbf{X}))$ is star decomposable in vertices.

If $k = -1$, then $(\operatorname{sd}\{\emptyset\}, \emptyset) = (\{\emptyset\}, \emptyset)$ is star decomposable in vertices. (This is the same as star decomposability in this case.)

If $k \geq 0$, then $(\operatorname{sd}\mathbf{X}, X)$ is star decomposable in vertices, if there is a total order $\prec$ on the set $V(\mathbf{X})$, inducing a star partition of $\operatorname{sd}\mathbf{X}$, with the following properties.[7]

Order condition: $X = V(\mathbf{X})_{\succeq w'}$ for some $w' \in V(\mathbf{X})$.

Link condition: For any vertex $w \in V(\mathbf{X})$ except for the last vertex in the order $\prec$, the pair $(\operatorname{sd}\operatorname{lk}(w, \mathbf{X}), V(\operatorname{lk}(w, \mathbf{X}))_{\succ w})$ is star decomposable in vertices.

Last vertex condition: For the last vertex $\hat{x} \in V(\mathbf{X})$ in the order $\prec$, the link $\operatorname{sd}\operatorname{lk}(\hat{x}, \mathbf{X})$ is star decomposable in vertices.

If the order $\prec$ on W satisfies the three conditions above, we say that $\prec$ induces a *star decomposition of* $(\operatorname{sd}\mathbf{X}, X)$ *in vertices.*

Lemma 2.10 implies the following proposition.

Proposition 2.12. *Let us assume that the pair* $(\operatorname{sd}\mathbf{X}, X)$ *is star decomposable in vertices, then it is star decomposable.*

Proof. We check that the order condition, the link condition and the last vertex condition in Definition 2.4 imply the corresponding conditions in Definition 2.11. The rest of the proof is a straightforward induction given that in dimensions -1 and 0 the notions coincide.

The order condition in Definitions 2.4 and 2.11 is actually identical.

For checking the link condition in Definition 2.4, for a given $w \in V(\mathbf{X})$ we need to find a set $U \subseteq V(\operatorname{lk}(w, \operatorname{sd}\mathbf{X}))$ such that (i) $\operatorname{st}(U, \operatorname{lk}(w, \operatorname{sd}\mathbf{X})) = \mathbf{O}_{\operatorname{sd}\mathbf{X}}(w, V(\mathbf{X})_{\succ w})$ and (ii) the pair $(\operatorname{lk}(w, \operatorname{sd}\mathbf{X}), U)$ is star decomposable in vertices (therefore star decomposable by induction). By Lemma 2.10 we have an isomorphism Ψ mapping the pair

$$(\operatorname{lk}(w, \operatorname{sd}\mathbf{X}), \mathbf{O}_{\operatorname{sd}\mathbf{X}}(w, V(\mathbf{X})_{\succ w}))$$

to the pair

$$(\operatorname{sd}\operatorname{lk}(w, \mathbf{X}), \operatorname{st}(V(\operatorname{lk}(w, \mathbf{X}))_{\succ w}, \operatorname{sd}\operatorname{lk}(w, \mathbf{X}))),$$

using that $V(\mathbf{X})_{\succ w} \cap V(\operatorname{lk}(w, \mathbf{X})) = V(\operatorname{lk}(w, \mathbf{X}))_{\succ w}$. We set

$$U := \Psi^{-1}(V(\operatorname{lk}(w, \mathbf{X}))_{\succ w}),$$

then (i) follows immediately from the isomorphism above. On the other hand, $(\operatorname{lk}(w, \operatorname{sd}\mathbf{X}), U)$ is isomorphic to $(\operatorname{sd}\operatorname{lk}(w, \mathbf{X}), V(\operatorname{lk}(w, \mathbf{X}))_{\succ w})$ by applying Ψ. Therefore, (ii) indeed follows from the link condition of Definition 2.11.

Finally the last vertex condition of Definition 2.11 implies the same condition of Definition 2.4 via Lemma 2.9 (and the induction). $\square$

[7] Note that $V(\mathbf{X})$ induces a star partition of $\operatorname{sd}\mathbf{X}$ for an arbitrary complex $\mathbf{X}$.

Merging orders inducing a star decomposition in vertices. Given simplicial complexes $\mathbf{X}$ and $\mathbf{Y}$ such that $\mathrm{sd}(\mathbf{X})$ and $\mathrm{sd}(\mathbf{Y})$ are star decomposable in vertices, we want to provide an order on $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ which induces a star decomposition in vertices of $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$. For the proof of our main result we need some flexibility how to merge the orders on $V(\mathbf{X})$ and $V(\mathbf{Y})$. First we provide a recipe that works in general but does not give all we need. This is the contents of the forthcoming Proposition 2.13. Then we also provide a more specific recipe which gives more under additional assumptions on $\mathbf{Y}$ (see Proposition 2.15).

Proposition 2.13. *Let $\mathbf{X}$ and $\mathbf{Y}$ be pure simplicial complexes such that $\mathrm{sd}(\mathbf{X})$ and $\mathrm{sd}(\mathbf{Y})$ are star decomposable in vertices. Let $\prec$ be an arbitrary total order on $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ satisfying that*

(i) the restriction of $\prec$ to $V(\mathbf{X})$ induces a star decomposition in vertices of $\mathrm{sd}(\mathbf{X})$,

(ii) the restriction of $\prec$ to $V(\mathbf{Y})$ induces a star decomposition in vertices of $\mathrm{sd}(\mathbf{Y})$,

(iii) if both $\mathbf{X}$ and $\mathbf{Y}$ are nonempty, then the last two elements in $\prec$ are the last element of $V(\mathbf{X})$ and the last element of $V(\mathbf{Y})$ (in arbitrary order).

Then $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$ is star decomposable in vertices in the order $\prec$ on $V(\mathbf{X} * \mathbf{Y}) = V(\mathbf{X}) \sqcup V(\mathbf{Y})$.

Corollary 2.14. *Let $\mathbf{X}$ and $\mathbf{Y}$ be simplicial complexes and $X \subseteq V(\mathbf{X})$, $Y \subseteq V(\mathbf{Y})$. Assume that the pairs $(\mathrm{sd}\,\mathbf{X}, X)$ and $(\mathrm{sd}\,\mathbf{Y}, Y)$ are star decomposable in vertices. Then the pair $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), X \sqcup Y)$ is star decomposable in vertices as well. In addition, if $|Y| = 1$, then the pair $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), Y)$ is star decomposable in vertices.*

Proof of Corollary 2.14. First, let us assume that $X = \emptyset$. Because $(\mathrm{sd}\,\mathbf{X}, X)$ is star decomposable, we deduce that $\mathbf{X} = \{\emptyset\}$. Consequently, $(\mathrm{sd}(\mathbf{X}*\mathbf{Y}), X \sqcup Y) = (\mathrm{sd}\,\mathbf{Y}, Y)$, which is star decomposable in vertices. Similarly, we resolve the case $Y = \emptyset$.

Now we can assume $X, Y \neq \emptyset$. Let $\prec_{\mathbf{X}}$ be a total order on $V(\mathbf{X})$ inducing a star decomposition of $(\mathrm{sd}\,\mathbf{X}, X)$ in vertices and let $\prec_{\mathbf{Y}}$ be a total order on $V(\mathbf{Y})$ inducing a star decomposition of $(\mathrm{sd}\,\mathbf{Y}, Y)$ in vertices. Let $\hat{x}$ be the last vertex of $V(\mathbf{X})$ in $\prec_{\mathbf{X}}$ and $\hat{y}$ be the last vertex of $V(\mathbf{Y})$ in $\prec_{\mathbf{Y}}$. Necessarily, $\hat{x} \in X$ and $\hat{y} \in Y$ as $X, Y \neq \emptyset$.

We define a total order $\prec$ on $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ so that we consider the vertices of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ in the order $[V(\mathbf{X}) \setminus X, V(\mathbf{Y}) \setminus Y, X \setminus \{\hat{x}\}, Y \setminus \{\hat{y}\}, \hat{x}, \hat{y}]$, where the individual sets $V(\mathbf{X}) \setminus X$, $V(\mathbf{Y}) \setminus Y$, $X \setminus \{\hat{x}\}$, and $Y \setminus \{\hat{y}\}$ are sorted according to $\prec_{\mathbf{X}}$ and $\prec_{\mathbf{Y}}$ respectively. Then $\prec$ satisfies the assumptions of Proposition 2.13. Therefore, $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$ is star decomposable in vertices in the order $\prec$.

Given that $\mathrm{st}(X \sqcup Y, \mathrm{sd}(\mathbf{X} * \mathbf{Y})) = \mathrm{st}((V(\mathbf{X}) \sqcup V(\mathbf{Y}))_{\succeq z}, \mathrm{sd}(\mathbf{X} * \mathbf{Y}))$ where z is the first vertex of $X \cup Y$ in $\prec$, we deduce that $\prec$ gives also a star decomposition of $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), X \sqcup Y)$ in vertices.

Finally, if $|Y| = 1$, then $Y = \{\hat{y}\}$. Thus $\mathrm{st}(Y, \mathrm{sd}(\mathbf{X} * \mathbf{Y})) = \mathrm{st}((V(\mathbf{X}) \sqcup V(\mathbf{Y}))_{\succeq \hat{y}}, \mathrm{sd}(\mathbf{X} * \mathbf{Y}))$ which means that $\prec$ gives a star decomposition of $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), Y)$ in vertices as well. $\square$

Proof of Proposition 2.13. First, similarly as in the previous proof, the statement is trivial if $\mathbf{X} = \{\emptyset\}$ or $\mathbf{Y} = \{\emptyset\}$ as a join with $\{\emptyset\}$ yields the same complex. Therefore, we can assume $\mathbf{X}, \mathbf{Y} \neq \{\emptyset\}$. In particular, the item (iii) of the statement is non-void.

Now, we prove the proposition by induction on $\dim(\mathbf{X} * \mathbf{Y})$. The start of the induction, when $\dim(\mathbf{X} * \mathbf{Y}) \leq 0$, is covered by the observation above.

We are given the order $\prec$ on $V(\mathbf{X}*\mathbf{Y})$; therefore it remains to check the order condition, the link condition and the last vertex condition.

As we check star decomposability of $\mathrm{sd}(\mathbf{X} * \mathbf{Y})$, that is, the pair $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), V(\mathbf{X}) \sqcup V(\mathbf{Y}))$, the order condition is trivial. (It is sufficient to take the first vertex of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ for checking the order condition.)

For checking the link condition, we consider arbitrary $x \in V(\mathbf{X}) \sqcup V(\mathbf{Y})$ distinct from the last vertex. Without loss of generality, we can assume $x \in V(\mathbf{X})$ as the argument is symmetric for a vertex from $V(\mathbf{Y})$. We need to check star decomposability of the pair

$$(\mathrm{sd}(\mathrm{lk}(x, \mathbf{X} * \mathbf{Y})), V(\mathrm{lk}(x, \mathbf{X} * \mathbf{Y}))_{\succ x}).$$

Given that $x \in V(\mathbf{X})$, this equals

$$(\mathrm{sd}(\mathrm{lk}(x, \mathbf{X}) * \mathbf{Y}), (V(\mathrm{lk}(x, \mathbf{X})) \sqcup V(\mathbf{Y}))_{\succ x}). \tag{2.4}$$

From the assumption on star decomposability of $\mathrm{sd}\,\mathbf{Y}$ in the order $\prec$, we deduce that the pair

$$(\mathrm{sd}(\mathbf{Y}), V(\mathbf{Y})_{\succ x}) \tag{2.5}$$

is star decomposable in vertices as long as $V(\mathbf{Y})_{\succ x}$ is nonempty. However, $V(\mathbf{Y})_{\succ x}$ is indeed nonempty as x is not the last vertex of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ in $\prec$ whereas there is a vertex from $V(\mathbf{Y})$ among the last two vertices.

From the assumption on star decomposability of $\mathbf{X}$ in the order $\prec$, checking the link condition gives that the pair

$$(\mathrm{sd}\,\mathrm{lk}(x, \mathbf{X}), V(\mathrm{lk}(x, \mathbf{X}))_{\succ x}) \tag{2.6}$$

is star decomposable in vertices if x is not the last vertex of $V(\mathbf{X})$. Therefore, if x is not the last vertex of $V(\mathbf{X})$, we will use the induction. From Corollary 2.14 for pairs (2.6) and (2.5) we deduce that the pair in (2.4) is indeed star decomposable in vertices as required. (Note that this is a correct use of the induction as we deduced Corollary 2.14 from Proposition 2.13 in the same dimension.)

It remains to consider the case when x is a last vertex of $V(\mathbf{X})$. In this case, x is the second to last vertex of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$. Let $\hat{y}$ be the last vertex of $V(\mathbf{Y})$, that is, the last vertex of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ as well. Then the pair (2.4) simplifies to

$$(\mathrm{sd}(\mathrm{lk}(x, \mathbf{X}) * \mathbf{Y}), \{\hat{y}\}).$$

Now, we can use Corollary 2.14 again with pairs $(\mathrm{sd}\,\mathrm{lk}(x, \mathbf{X}), V(\mathrm{lk}(x, \mathbf{X})))$ and $(\mathrm{sd}(\mathbf{Y}), \{\hat{y}\})$, using the 'in addition' part.

Finally, it remains to check the last vertex condition. Let us therefore assume that $\hat{x}$ is the last vertex of $V(\mathbf{X}) \sqcup V(\mathbf{Y})$. Again, we can without loss of generality assume that $\hat{x} \in V(\mathbf{X})$. We need to check star decomposability in vertices of

$\operatorname{sd}\operatorname{lk}(\hat{x}, \mathbf{X} * \mathbf{Y}) = \operatorname{sd}(\operatorname{lk}(\hat{x}, \mathbf{X}) * \mathbf{Y})$. By the last vertex condition on $\operatorname{sd}(\mathbf{X})$ we get that $\operatorname{sd}\operatorname{lk}(\hat{x}, \mathbf{X})$ is star decomposable in vertices. Therefore, by the induction applied to $\operatorname{sd}(\operatorname{lk}(\hat{x}, \mathbf{X}))$ and $\operatorname{sd}\mathbf{Y}$, we get that $\operatorname{sd}(\operatorname{lk}(\hat{x}, \mathbf{X}) * \mathbf{Y})$ is star decomposable in vertices as required. $\square$

Now, we state a more specialized version of Proposition 2.13 with an additional condition on homology. Let us recall that given a simplicial complex $\mathbf{Y}$ and $Y \subseteq V(\operatorname{sd}\mathbf{Y})$, the star $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ is defined as $\bigcup_{v \in Y} \operatorname{st}(v, \operatorname{sd}\mathbf{Y})$. Following our convention of neglecting a difference between $v \in V(\mathbf{Y})$ and $\{v\} \in V(\operatorname{sd}\mathbf{Y})$, we also set $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y}) := \bigcup_{v \in Y} \operatorname{st}(v, \operatorname{sd}\mathbf{Y})$ for $Y \subseteq V(\mathbf{Y})$.

Proposition 2.15. *Let $\mathbf{X}$ and $\mathbf{Y}$ be pure simplicial complexes, $\dim \mathbf{X}, \dim \mathbf{Y} \geq 0$, and Y be a nonempty subset of $V(\mathbf{Y})$. Assume that $\operatorname{sd}\mathbf{X}$ and $(\operatorname{sd}\mathbf{Y}, Y)$ are star decomposable in vertices and $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ has trivial reduced homology groups. Let $\prec$ be an arbitrarily total order on $V(\mathbf{X}) \sqcup V(\mathbf{Y})$ satisfying:*

(i) The restriction of $\prec$ to $V(\mathbf{X})$ induces a star decomposition in vertices of $\operatorname{sd}(\mathbf{X})$;

(ii) The restriction of $\prec$ to $V(\mathbf{Y})$ induces a star decomposition in vertices of $\operatorname{sd}(\mathbf{Y}, Y)$; and

(iii) $Y = (V(\mathbf{X}) \sqcup V(\mathbf{Y}))_{\succ \hat{x}}$ where $\hat{x}$ is the last vertex of $V(\mathbf{X})$ in $\prec$.

Then $\operatorname{sd}(\mathbf{X} * \mathbf{Y}, Y)$ is star decomposable in vertices in the order $\prec$ on $V(\mathbf{X} * \mathbf{Y}) = V(\mathbf{X}) \sqcup V(\mathbf{Y})$.

For the proof, we need a following auxiliary lemma which will be useful in the induction.

Lemma 2.16. *Let $\mathbf{Y}$ be a pure simplicial complex and $Y \subseteq V(\mathbf{Y})$. Assume that the pair $(\operatorname{sd}\mathbf{Y}, Y)$ is star-decomposable in vertices in some total order $\prec$ on $V(\mathbf{Y})$ and also that $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ has trivial reduced homology groups. Then $\operatorname{st}(V(\operatorname{lk}(y, \mathbf{Y}))_{\succ y}, \operatorname{sd}(\operatorname{lk}(y, \mathbf{Y})))$ has trivial reduced homology groups as well for all $y \in Y$ except for the last vertex in Y.*

Proof. Let $y \in Y$ be different from the last vertex in the order $\prec$. First, we show that $\operatorname{st}(Y_{\succ y}, \mathbf{Y})$ has trivial reduced homology groups.

Since the pair $(\operatorname{sd}\mathbf{Y}, Y)$ is star decomposable in vertices, Theorem 2.8 implies that $\operatorname{sd}\mathbf{Y}$ is vertex decomposable. In addition, we get that $\operatorname{sd}\mathbf{Y}$ is vertex decomposable in a shedding order $\succ''$ extending $\succ'$ where is derived from $\succ$. (We recall that the definition of the derived order is given above the statement of Lemma 2.7.) In particular, $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ and later $\operatorname{st}(Y_{\succ y}, \operatorname{sd}\mathbf{Y})$ are intermediate steps in the sequence of complexes obtained by gradually removing vertices of $\mathbf{Y}$ in the given shedding order $\succ''$.

We also get that $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ and $\operatorname{st}(Y_{\succ y}, \operatorname{sd}\mathbf{Y})$ are shellable by [PB80] (see Theorem 2.8 and the note below Definition 2.1 in [PB80]). Therefore, each of them is homotopy equivalent to a wedge of d-spheres where $d = \dim \mathbf{Y}$; see [Koz08, Theorem 12.3]. Since $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ has trivial homology groups, this must be a trivial wedge. However, following the shedding order from $\operatorname{st}(Y, \operatorname{sd}\mathbf{Y})$ to $\operatorname{st}(Y_{\succ y}, \operatorname{sd}\mathbf{Y})$, we cannot introduce homology in dimension d when gradually removing vertices.

Therefore, $\mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y})$ has to be homotopy equivalent to a trivial wedge as well showing that $\mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y})$ has trivial reduced homology groups.

Note that $\mathrm{st}(Y_{\succeq y}, \mathrm{sd}\,\mathbf{Y})$ has trivial reduced homology groups as well by analogous reasoning.

Now, by Lemma 2.10,

$$\mathrm{st}(V(\mathrm{lk}(y, \mathbf{Y}))_{\succ y}, \mathrm{sd}(\mathrm{lk}(y, \mathbf{Y}))) \cong \mathbf{O}_{\mathrm{sd}\,\mathbf{Y}}(y, V(\mathbf{Y})_{\succ y}).$$

We use a Mayer-Vietoris sequence for $\mathrm{st}(Y_{\succeq y}, \mathrm{sd}\,\mathbf{Y})$ covered by $\mathrm{st}(y, \mathrm{sd}(\mathbf{Y}))$ and $\mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y})$. Then

$$\mathrm{st}(y, \mathrm{sd}(\mathbf{Y})) \cap \mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y}) = \mathbf{O}_{\mathrm{sd}\,\mathbf{Y}}(y, Y_{\succ y}) = \mathbf{O}_{\mathrm{sd}\,\mathbf{Y}}(y, V(\mathbf{Y})_{\succ y})$$

and we get the following long exact sequence

$$\begin{aligned} \cdots \longrightarrow \tilde{H}_{n+1}(\mathrm{st}(Y_{\succeq y}, \mathrm{sd}\,\mathbf{Y})) \longrightarrow \tilde{H}_n(\mathbf{O}_{\mathrm{sd}\,\mathbf{Y}}(y, V(\mathbf{Y})_{\succ y})) \longrightarrow \\ \longrightarrow \tilde{H}_n(\mathrm{st}(y, \mathrm{sd}\,\mathbf{Y})) \oplus \tilde{H}_n(\mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y})) \longrightarrow \tilde{H}_n(\mathrm{st}(Y_{\succeq y}, \mathrm{sd}\,\mathbf{Y})) \longrightarrow \cdots \end{aligned}$$

All $\mathrm{st}(Y_{\succeq y}, \mathrm{sd}\,\mathbf{Y})$, $\mathrm{st}(y, \mathrm{sd}\,\mathbf{Y})$ and $\mathrm{st}(Y_{\succ y}, \mathrm{sd}\,\mathbf{Y})$ have trivial reduced homology groups. Therefore, $\tilde{H}_n(\mathbf{O}_{\mathrm{sd}\,\mathbf{Y}}(y, V(\mathbf{Y})_{\succ y})) \cong \tilde{H}_n(\mathrm{st}(V(\mathrm{lk}(y, \mathbf{Y}))_{\succ y}, \mathrm{sd}(\mathrm{lk}(y, \mathbf{Y})))$ is trivial for all $n \in \mathbb{Z}$. □

Proof of Proposition 2.15. Similarly, as in the proof of Proposition 2.13, we proceed by induction on $\dim(\mathbf{X} * \mathbf{Y})$.

First, we observe that the case $\dim \mathbf{Y} = 0$ is covered by Proposition 2.13. Indeed, the only issue is to verify (iii) of Proposition 2.13. If $\dim \mathbf{Y} = 0$, then Y must contain a single vertex $\hat{y}$ (due to the condition on homology of $\mathrm{st}(Y, \mathrm{sd}(\mathbf{Y}))$). Consequently, (iii) (of this proposition) implies that the last two vertices of $\prec$ are $\hat{x}$ and $\hat{y}$ which verifies (iii) of Proposition 2.13.

Now, let us assume $\dim \mathbf{X} \geq 0$ and $\dim \mathbf{Y} \geq 1$. The order condition is satisfied since Y is non-empty and it is equal to $(V(\mathbf{X}) \sqcup V(\mathbf{Y}))_{\succ \hat{x}}$ by (iii).

For checking the link condition, we consider arbitrary $z \in V(\mathbf{X}) \sqcup V(\mathbf{Y})$ distinct from the last vertex. We need to check star decomposability of the pair

$$(\mathrm{sd}(\mathrm{lk}(z, \mathbf{X} * \mathbf{Y})), V(\mathrm{lk}(z, \mathbf{X} * \mathbf{Y}))_{\succ z}). \tag{2.7}$$

If $z \in V(\mathbf{X}) \setminus \{\hat{x}\} \sqcup V(\mathbf{Y}) \setminus Y$, then the analysis is the same as in the proof of Proposition 2.13.

If $z = \hat{x}$, the pair (2.7) becomes

$$(\mathrm{sd}(\mathrm{lk}(\hat{x}, \mathbf{X}) * \mathbf{Y}), Y).$$

If $\dim \mathbf{X} = 0$, then we further get $(\mathrm{sd}\,\mathbf{Y}, Y)$ which is star decomposable in vertices by the assumptions. If $\dim \mathbf{X} \geq 1$, then $\dim \mathrm{lk}(\hat{x}, \mathbf{X}) \geq 0$ and we can use the induction (note that $\mathrm{sd}\,\mathrm{lk}(\hat{x}, \mathbf{X})$ is star decomposable in vertices by the last vertex condition for decomposition of $\mathrm{sd}\,\mathbf{X}$).

Finally, by assuming $z \in Y \setminus \{\hat{y}\}$, where $\hat{y}$ is the last vertex of $\prec$, we get the pair

$$(\mathrm{sd}(\mathbf{X} * \mathrm{lk}(z, \mathbf{Y})), V(\mathrm{lk}(z, \mathbf{Y})_{\succ z}). \tag{2.8}$$

By Lemma 2.16 the pair $\mathrm{st}(V(\mathrm{lk}(z, \mathbf{Y})_{\succ z}, \mathrm{sd}(\mathrm{lk}(z, \mathbf{Y})))$ has trivial reduced homology groups. Therefore, (2.8) is star-decomposable in vertices by the induction hypothesis. (Here we use that $\dim \mathrm{lk}(z, \mathbf{Y}) \geq 0$ and that $(\mathrm{sd}\,\mathrm{lk}(z, \mathbf{Y}), V(\mathrm{lk}(z, \mathbf{Y})_{\succ z}))$ is star decomposable in vertices by the link condition for the decomposition of $(\mathrm{sd}\,\mathbf{Y}, Y)$.)

Finally, we check the last vertex condition. We need star decomposability in vertices of $\mathrm{sd}\,\mathrm{lk}(\hat{y}, \mathbf{X} * \mathbf{Y})$. Note that $\mathrm{lk}(\hat{y}, \mathbf{X} * \mathbf{Y}) = \mathbf{X} * \mathrm{lk}(\hat{y}, \mathbf{Y})$ as both sides contain simplices of the form $\xi \cup \eta$, where $\xi \in \mathbf{X}$, $\eta \cup \{\hat{y}\} \in \mathbf{Y}$, and $\hat{y} \notin \eta$. Thus, we need star decomposability in vertices of $\mathrm{sd}(\mathbf{X} * \mathrm{lk}(\hat{y}, \mathbf{Y}))$. This is star decomposable in vertices by Proposition 2.13. (Here, we again use that $\dim \mathrm{lk}(\hat{y}, \mathbf{Y}) \geq 0$ and also that $\mathrm{sd}\,\mathrm{lk}(\hat{y}, \mathbf{Y})$ is star decomposable in vertices by the last vertex condition in the decomposition of $\mathrm{sd}\,\mathbf{Y}$.) □

2.5 Proof of the main result

In this section, we prove Theorem 2.3 which also finishes the proof of Theorem 2.2.

We first need two auxiliary observations that we will use in the proof.

Observation 2.17. *The boundary of a simplex $\partial\sigma$ satisfies the (HRC) condition.*

Proof. We prove the observation by induction on $\dim \sigma$, starting with $\dim \sigma = 0$, in which case $\partial\sigma = \emptyset$. If $\dim \sigma > 0$, let $\sigma' \subsetneq \sigma$. We need to check that $\mathrm{lk}(\sigma', \partial\sigma)$ satisfies the (RC) condition. This link is again a boundary of a simplex. If $\sigma' \neq \emptyset$, we get a simplex of small dimension, therefore, we can use the induction. If $\sigma = \emptyset$, then $\mathrm{lk}(\sigma', \partial\sigma) = \partial\sigma$ which is collapsible after removing an arbitrary facet (it is a cone then). □

Observation 2.18. *Let $\mathbf{K}$ be a collapsible complex and w be an arbitrary vertex of $\mathbf{K}$. Then $\mathbf{K}$ collapses to w.*

Proof. First, we use the well known fact that the collapses of $\mathbf{K}$ can be rearranged so that they are ordered by non-increasing dimension [Whi39, Section 3]. In particular, this means that $\mathbf{K}$ collapses to a graph $\mathbf{G}$ with $V(\mathbf{G}) = V(\mathbf{K})$. This graph must be a tree as $\mathbf{K}$ is collapsible, and we can further rearrange the collapses of $\mathbf{G}$ so that w is the last vertex. □

Now we prove Theorem 2.3 by induction on the dimension of $\mathbf{K}$. We know that $\mathbf{K}$ satisfies the (RC) condition. Therefore, there are facets $\phi_1, \ldots, \phi_t$ of $\mathbf{K}$ such that $\mathbf{K}' := \mathbf{K} - \{\phi_1, \ldots, \phi_t\}$ is collapsible. We further consider a sequence $(\mathbf{K}_1, \ldots, \mathbf{K}_s)$ of elementary collapses of $\mathbf{K}'$ where $\mathbf{K}' = \mathbf{K}_1$, $\mathbf{K}_s$ is a vertex (denoted by z), and $\mathbf{K}_{i+1}$ arises from $\mathbf{K}_i$ by removing faces σ_i and τ_i where $\sigma_i \subset \tau_i$ and $\dim \sigma_i = \dim \tau_i - 1$, and τ_i is the unique maximal face containing σ_i. Then we consider the following total order $\prec$ on nonempty faces of $\mathbf{K}$, that is, vertices of $\mathrm{sd}\,\mathbf{K}$:

$$\phi_1 \prec \cdots \prec \phi_t \prec \sigma_1 \prec \tau_1 \prec \sigma_2 \prec \tau_2 \prec \cdots \prec \sigma_{s-1} \prec \tau_{s-1} \prec \{z\}.$$

Our aim is to show that $\prec$ induces a star decomposition in vertices of $\mathrm{sd}^2\,\mathbf{K}$. This we will also use in the induction; that is, for complexes $\mathbf{L}$ of lower dimension satisfying the (HRC) condition, we assume that a sequence of removals of facets and collapses induces a star decomposition in vertices of $\operatorname{sd}^2 \mathbf{L}$ as above. The proof is easy if $\dim \mathbf{K} = 0$ (here no collapses are used), thus we may assume that $\dim \mathbf{K} > 0$ and proceed with the second induction step.

There is essentially nothing to check for the order condition as we provide a total order on vertices of $\operatorname{sd} \mathbf{K}$. Thus the only issue is to check the link condition and the last vertex condition.

In order to access the vertices of $\operatorname{sd} \mathbf{K}$ more easily in the given order, we also give them alternate names $\omega_1, \ldots, \omega_k$ so that

$$(\phi_1, \ldots, \phi_t, \sigma_1, \tau_1, \ldots, \sigma_{s-1}, \tau_{s-1}, \{z\}) = (\omega_1, \ldots, \omega_k)$$

where $k = t + 2s - 1$. That is, $\phi_1 = \omega_1$, $\sigma_1 = \omega_{t+1}$, etc.

Checking the last vertex condition. Because it is easier, we check the last vertex condition first. We need to check that $\operatorname{sd} \operatorname{lk}(\omega_k, \operatorname{sd} \mathbf{K})$ is star decomposable in vertices. Because ω_k is a vertex of $\mathbf{K}$, this complex is isomorphic to $\operatorname{sd}^2 \operatorname{lk}(\omega_k, \mathbf{K})$ by Lemma 2.9. Therefore, this complex is star decomposable in vertices by induction because $\operatorname{lk}(\omega_k, \mathbf{K})$ satisfies the (HRC) condition as this condition is hereditary for links.

Checking the link condition: For checking the link condition, we need to check that the pair $(\operatorname{sd} \operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}), V(\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}))_{\succ \omega_i})$ is star decomposable in vertices for $i \in \{1, \ldots, k-1\}$. For checking this condition we again need to 'simplify' this pair so that we remove the subdivision from the link. The tool for this is again Lemma 2.9. For the first entry it gives

$$\operatorname{sd} \operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}) \cong \operatorname{sd}(\operatorname{sd} \partial\omega_i * \operatorname{sd} \operatorname{lk}(\omega_i, \mathbf{K})).$$

We use the specific isomorphism Ψ from the proof of Lemma 2.9 and our next task is to describe $V(\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}))_{\succ \omega_i})$ after applying this isomorphism.

First of all, we briefly describe the set $V(\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}))_{\succ \omega_i}$. The vertices of $\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K})$ are the nonempty faces η of $\mathbf{K}$ such that either $\eta \subsetneq \omega_i$ or $\omega_i \subsetneq \eta$. Therefore, the set $V(\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}))_{\succ \omega_i}$ consists of faces η as above, which in addition satisfy $\eta \succ \omega_i$. The isomorphism Ψ from the proof of Lemma 2.9 maps η again to η if $\eta \subsetneq \omega_i$ and it maps η to $\eta \setminus \omega_i$ if $\omega_i \subsetneq \eta$. Hence

$$\Psi(V(\operatorname{lk}(\omega_i, \operatorname{sd} \mathbf{K}))_{\succ \omega_i}) = V(\operatorname{sd} \partial\omega_i)_{\succ \omega_i} \sqcup \{\eta \setminus \omega_i \colon \eta \supsetneq \omega_i, \eta \succ \omega_i\},$$

which we denote by W. Thus, we need to check the star decomposability in vertices of the pair

$$(\operatorname{sd}(\operatorname{sd} \partial\omega_i * \operatorname{sd} \operatorname{lk}(\omega_i, \mathbf{K})), W). \tag{2.9}$$

We distinguish several cases according to the type of ω_i.

1. $\omega_i = \phi_i$, that is, $i \leq t$:

   In this case, ϕ_i is a facet. Therefore, $\operatorname{lk}(\phi_i, \mathbf{K}) = \emptyset$. Also $\eta \succ \phi_i$ for all proper subfaces η. Therefore, the pair (2.9) simplifies to $(\operatorname{sd}(\operatorname{sd} \partial\phi_i), V(\operatorname{sd} \partial\phi_i))$; see Figure 2.10. That is, we only need that $\operatorname{sd}(\operatorname{sd} \partial\phi_i)$ is star decomposable in vertices which follows by the induction and Observation 2.17.

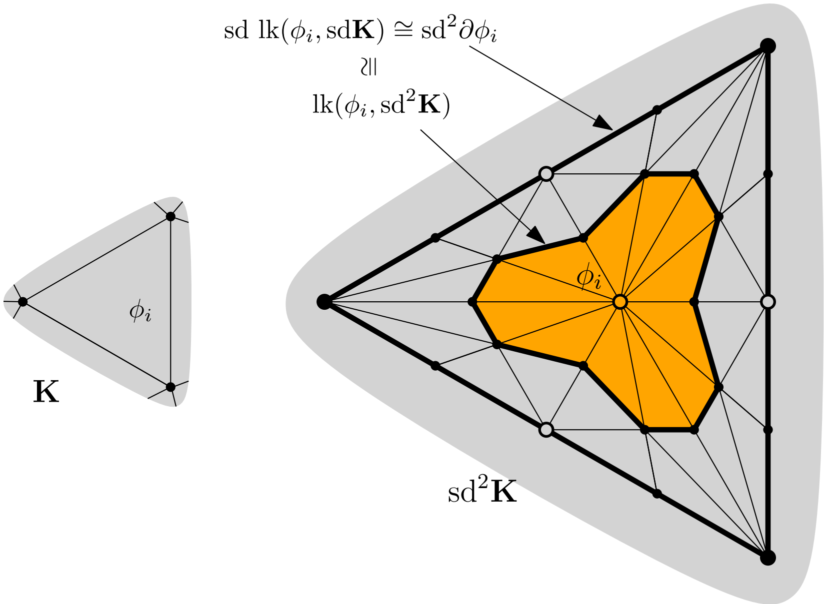

Figure 2.10: Isomorphisms for verifying the link condition in case 1. We consider the case of the removal of the facet ϕ_i. If we were checking star decomposability only, we would be interested in star decomposability of $\operatorname{lk}(\phi_i, \operatorname{sd}^2 \mathbf{K})$. For star decomposability in vertices, this translates to checking the link condition for $\operatorname{sd}\operatorname{lk}(\phi_i, \operatorname{sd}\mathbf{K})$ which is further isomorphic to $\operatorname{sd}^2 \partial\phi_i$ (in this case, the last isomorphism is even equality).

2. $\omega_i = \sigma_j$ for some j, that is, $i > t$ and $t - i$ is odd:

   We need to describe W, for which we need to describe the faces η such that $\eta \subsetneq \sigma_j$ or $\sigma_j \subsetneq \eta$ such that $\eta \succ \sigma_j$. As σ_j induces an elementary collapse in a sequence of collapses of $\mathbf{K}'$, we get $\tau_j \succ \sigma_j$ but $\eta \prec \sigma_j$ for any $\eta \supsetneq \sigma_j$ different from τ_j. On the other hand all proper subfaces of σ_j are removed only later on in collapsing of $\mathbf{K}'$. Altogether $W = V(\operatorname{sd}\partial\sigma_j) \sqcup \{\tau_j \setminus \sigma_j\}$. See Figure 2.11 for an example of the pair (2.9) in this case.

   Now, we aim to use Corollary 2.14 with
$$(\mathbf{X}, X) = (\operatorname{sd}\partial\sigma_j, V(\operatorname{sd}\partial\sigma_j))$$
   and
$$(\mathbf{Y}, Y) = (\operatorname{sd}\operatorname{lk}(\sigma_j, \mathbf{K}), \{\tau_j \setminus \sigma_j\}).$$
   The pair $(\operatorname{sd}\mathbf{X}, X)$ is star decomposable in vertices by Observation 2.17 and the induction. For checking star decomposability in vertices of $(\operatorname{sd}\mathbf{Y}, Y)$, we know that $\operatorname{lk}(\sigma_j, \mathbf{K})$ satisfies the (HRC) condition. In particular, $\operatorname{lk}(\sigma_j, \mathbf{K})$ is collapsible after removing some number of facets, and the subsequent collapses can be rearranged so that the vertex $\tau_j \setminus \sigma_j$ is the last vertex in the sequence of collapses. (If $\dim \operatorname{lk}(\sigma_j, \mathbf{K}) = 0$, then we instead rearrange the removals of the facets so that $\tau_j \setminus \sigma_j$ is the last.) Now, by induction, this sequence of removals of facets and collapses induces a star decomposition in vertices of $\operatorname{sd}\operatorname{sd}\operatorname{lk}(\sigma_j, \mathbf{K})$ such that $\{\tau_j \setminus \sigma_j\}$ is the last vertex in this decomposition. This exactly means that $(\operatorname{sd}\mathbf{Y}, Y)$ is star decomposable in vertices.

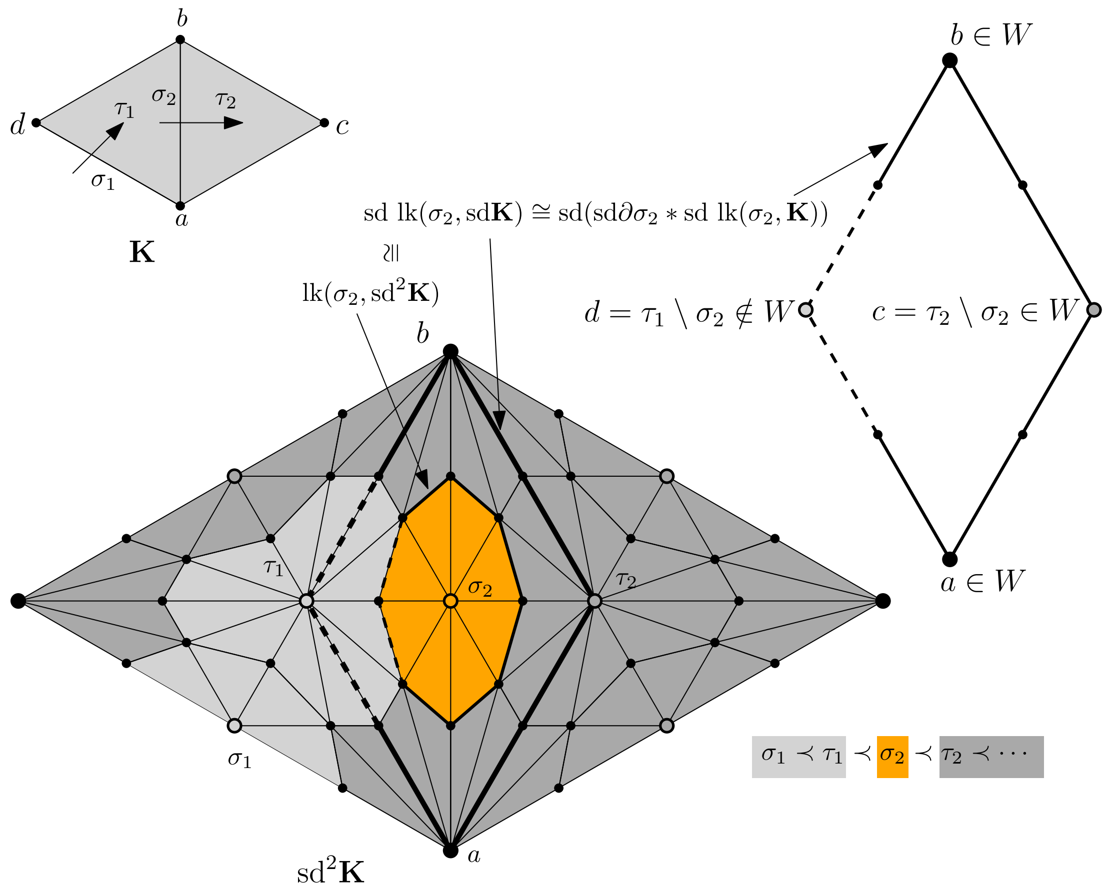

Figure 2.11: Isomorphisms for verifying the link condition in case 2. Here we consider the case $\sigma_j = \sigma_2$ coming from the collapses on the top left picture. The vertex decomposability of $(\mathrm{sd}\,\mathrm{lk}(\sigma_2, \mathrm{sd}\,\mathbf{K}), V(\mathrm{lk}(\sigma_2, \mathrm{sd}\,\mathbf{K}))_{\succ\sigma_2}) = (\mathrm{sd}\,\mathrm{lk}(\sigma_2, \mathrm{sd}\,\mathbf{K}), \{a, b, \tau_2\})$ in the middle picture translates to vertex decomposability of the pair $(\mathrm{sd}(\mathrm{sd}\,\partial\sigma_2 * \mathrm{sd}\,\mathrm{lk}(\sigma_2, \mathbf{K})), W)$ in the top right picture where $W = \{a, b, \tau_2 \setminus \sigma_2\}$, which coincides with $V(\mathrm{sd}\,\partial\sigma_j) \sqcup \{\tau_j \setminus \sigma_j\}$ as required.

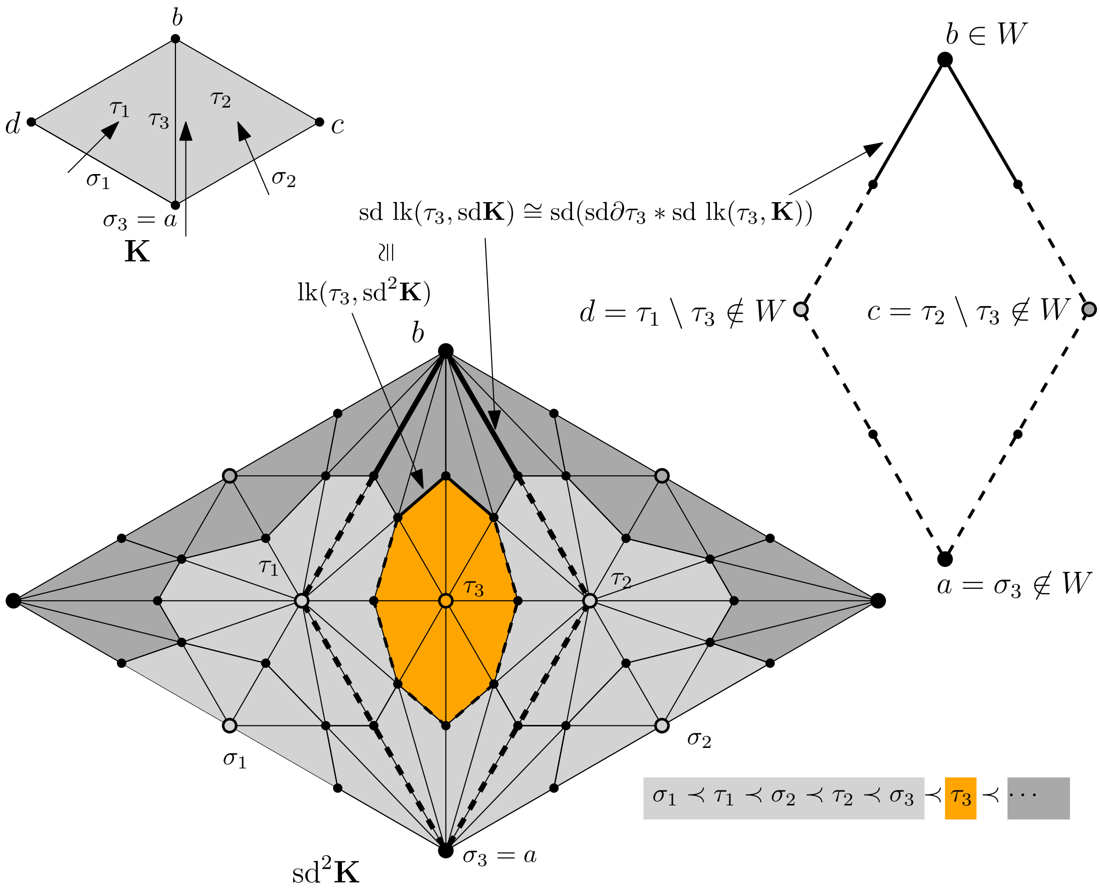

Figure 2.12: Isomorphisms for verifying the link condition in case 3. Here we consider the case $\tau_j = \tau_3$ coming from the collapses on the top left picture. The vertex decomposability of $(\mathrm{sd}\,\mathrm{lk}(\tau_3, \mathrm{sd}\,\mathbf{K}), V(\mathrm{lk}(\tau_3, \mathrm{sd}\,\mathbf{K}))_{\succ\tau_3}) = (\mathrm{sd}\,\mathrm{lk}(\tau_3, \mathrm{sd}\,\mathbf{K}), \{b\})$ in the middle picture translates to vertex decomposability of the pair $(\mathrm{sd}(\mathrm{sd}\,\partial\tau_3 * \mathrm{sd}\,\mathrm{lk}(\tau_3, \mathbf{K})), W)$ in the top right picture where $W = \{b\}$, which coincides with $V(\mathrm{sd}\,\partial\tau_j) \setminus \{\sigma_j\}$ as required.

Altogether, Corollary 2.14 implies that the pair $(\mathrm{sd}(\mathbf{X} * \mathbf{Y}), X \sqcup Y)$ is star decomposable in vertices which is exactly the required pair (2.9).

3. $\omega_i = \tau_j$ for some j, that is, $i > t$ and $t - i$ is even:

   We again first determine W. For each $\eta \supsetneq \tau_j$, we get $\eta \prec \tau_j$ as τ_j is a maximal face during the elementary collapse. On the other hand, for $\eta \subsetneq \tau_j$ we get $\eta \succ \tau_j$ unless $\eta = \sigma_j$ as all subfaces of τ_j have to be present at the moment of removing of σ_j, and τ_j immediately succeeds. Altogether, $W = (V(\partial\tau_j) \setminus \{\sigma_j\}) \sqcup \emptyset$. See Figure 2.12 for an example of the pair (2.9) in this case.

   We aim to use Proposition 2.15 with $\mathbf{X} = \mathrm{sd}\,\mathrm{lk}(\tau_j, \mathbf{K})$, $\mathbf{Y} = \mathrm{sd}\,\partial\tau_j$ and $Y = V(\mathrm{sd}\,\partial\tau_j) \setminus \{\sigma_j\}$. We get that $\mathbf{X}$ is star decomposable in vertices by induction as $\mathrm{lk}(\tau_j, \mathbf{K})$ satisfies the (HRC) condition. We also need that $(\mathrm{sd}\,\mathbf{Y}, Y)$ is star decomposable in vertices. For this we use Observation 2.17 and the induction while choosing σ_j to be the first face removed from $V(\mathrm{sd}\,\partial\tau_j)$. Then $Y = V(\mathrm{sd}\,\partial\tau_j)_{\succ'\{\sigma_j\}}$ where $\succ'$ is the corresponding order on $V(\operatorname{sd}\partial\tau_j)$. Altogether, for application of Proposition 2.15 we choose the order $\succ'$ on $V(\operatorname{sd}\operatorname{lk}(\tau_j,\mathbf{K}))\sqcup V(\operatorname{sd}\partial\tau_j)$ so that it starts with σ_j, it continues on $V(\operatorname{sd}\operatorname{lk}(\tau_j,\mathbf{K})$ in order of a star decomposition in vertices of $\operatorname{sd}\mathbf{X}$ and finally it continues on $Y = V(\operatorname{sd}\partial\tau_j)\setminus\{\sigma_j\}$ in the already prescribed order $\succ'$. Then we get the required conclusion that $(\operatorname{sd}(\mathbf{X}*\mathbf{Y}),Y)$, which is the pair (2.9), is star decomposable in vertices. This finishes the proof of Theorem 2.3. $\square$

3. NP-hardness of PL geometric category 2

3.1 Introduction

An important notion in homotopy theory is the Lusternik–Schnirelmann category (LS category) of a topological space. This notion is important not only as a purely mathematical object (see, e.g., the book [CLOT03]) but also in computer science as it is closely related to the topological complexity of motion planning; see, e.g, [Far03, Far04, FM20].

The *LS category*, $\mathrm{cat}(X)$, of a topological space X is the smallest n (if it exists) such that X can be covered by n open sets so that the inclusion of each of the open sets is nullhomotopic in X. One difficulty when working with the LS category is that it is often hard to determine. For example, determining whether $\mathrm{cat}(X) = 1$ is equivalent to contractibility of X. This is known to be undecidable if X is a simplicial complex of dimension at least 4; see [VKF74, §10] and [Tan16, Appendix] while it is an open problem whether this is decidable for simplicial complexes of dimension 2.[1] (We are mainly interested in complexes of dimension 2.)

In order to bound the LS category from above we can use some closely related notions. One of them is the geometric category, $\mathrm{gcat}(X)$, which requires that the open sets covering X are already contractible. (For more details see again [CLOT03].) If X is a polyhedron, this is equivalent to finding the minimum number of subpolyhedra covering X each of which is contractible. This may make estimating $\mathrm{gcat}(X)$ sometimes easier. However, determining whether $\mathrm{gcat}(X) = 1$ is still equivalent to contractibility of X.

Next step in this direction has been done by Borghini [Bor20] who introduced PL geometric category $\mathrm{plgcat}(P)$ of a compact (connected) polyhedron P. It is the minimum number of PL collapsible subpolyhedra of P that cover P. In this case determining whether $\mathrm{plgcat}(P) = 1$ is equivalent to asking whether P is PL collapsible. At least for 2-complexes this is a significant improvement as PL collapsibility of 2-complexes is a purely combinatorial notion which is easy to check. Indeed, it is not hard to derive from known results that this is a polynomially checkable criterion (by performing the collapses greedily on an arbitrary triangulation).

Proposition 3.1. *Given a 2-dimensional triangulated polyhedron P, it can be checked in polynomial time whether* $\mathrm{plgcat}(P) = 1$.

Borghini further proved [Bor20] that a connected d-dimensional polyhedron has PL geometric category at most $d+1$. For connected 2-polyhedra, this means that the only options are 1, 2, or 3. One of the main aims in [Bor20] is to provide a partial characterization of polyhedra P with $\mathrm{plgcat}(P) \leq 2$ (which we do not reproduce here). All these positive results suggest that determining $\mathrm{plgcat}(P)$

[1] Via tools in [HAMS93] (using the exercise on page 8) decidability of this problem is equivalent to determining whether a given balanced presentation of a group presents a trivial group. In this form, the problem is mentioned for example in [BMS02].

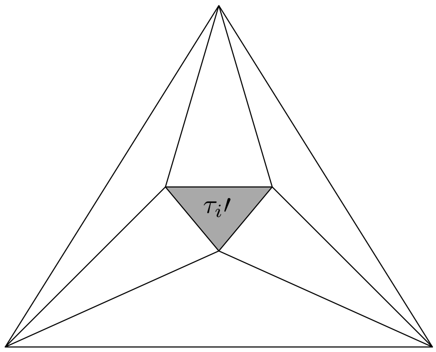

Figure 3.1: A subdivision of triangle τ_i into seven parts from the proof of Proposition 3.3.

could be easy for 2-polyhedra. In particular, one should be curious whether it is possible to extend Borghini's results to a *full characterization* that would distinguish 2-polyhedra with PL geometric category equal to 2 from those for which it equals 3.

We will show that this is essentially impossible, at least for an efficiently algorithmically checkable characterization. In technical terms, we show that determining whether $\mathrm{plgcat}(P) \leq 2$ is NP-hard. Let us recall that NP-hard problems are believed not to be solvable in polynomial time.

Theorem 3.2. *Given a 2-dimensional triangulated polyhedron P, it is NP-hard to decide whether* $\mathrm{plgcat}(P) \leq 2$.

We should also point out that we actually do not know whether recognition of triangulated polyhedra with $\mathrm{plgcat}(P) \leq 2$ belongs to the class NP(not even whether it is decidable). This could be certified by two subpolyhedra witnessing $\mathrm{plgcat}(P) \leq 2$ but we do not know whether we can bound their sizes.

A useful step towards our proof of Theorem 3.2 is that we observe a relation between $\mathrm{plgcat}(P) \leq 2$ and shellability (of some triangulation) of P.

Proposition 3.3. *If a 2-dimensional polyhedron P admits a (pure) shellable triangulation, then* $\mathrm{plgcat}(P) \leq 2$.

Proof. The proposition easily follows from the theorem of Hachimori (Theorem 2.1) which we have already mention in Chapter 2.

Let $\mathbf{K}$ be a pure shellable triangulation of P. By Theorem 2.1 there is a list of triangles $\tau_1, \ldots, \tau_\ell$ such that the resulting complex $\mathbf{K}'$ is collapsible after removing these triangles. Now we build an auxiliary complex $\mathbf{L}$ from $\mathbf{K}$ by subdividing each of the triangles $\tau_1, \ldots, \tau_\ell$ as in Figure 3.1. We also build a complex $\mathbf{L}'$ by removing the middle triangle τ_i' from each subdivided τ_i in $\mathbf{L}$. The complex $\mathbf{K}'$ is a subcomplex of $\mathbf{L}'$ and it is not hard to see that $\mathbf{L}'$ collapses to $\mathbf{K}'$. Hence L' is collapsible as well. Then $|\mathbf{L}'|$ is one of the two collapsible polyhedra covering P. The second polyhedron is obtained by taking the union of τ_i' and connecting them along the 1-skeleton of $\mathbf{L}$ so that the resulting complex is collapsible (the connection along the 1-skeleton of $\mathbf{L}$ can be, for example, obtained so that we pick two edges in each triangle and then we extend this forest to a spanning tree). $\square$

It has been shown by Goaoc, Paták, Patáková, Tancer and Wagner [GPP+19] that shellability is NP-hard already for 2-dimensional simplicial complexes. In addition, the reduction in [GPP$^+$19] is quite resistant with respect to subdivisions. Thus, we could hope to prove Theorem 3.2 in the following way: Consider a complex $\mathbf{K}$ that appears in the reduction in [GPP$^+$19]. If $\mathbf{K}$ is shellable, then $\mathrm{plgcat}(|\mathbf{K}|) \leq 2$ by Proposition 3.3 (let us recall that $|\mathbf{K}|$ stands for the polyhedron of $\mathbf{K}$). If we were able to show the other implication: 'if $\mathbf{K}$ is not shellable, then $\mathrm{plgcat}(|\mathbf{K}|) = 3$', we would immediately get a proof of Theorem 3.2. Unfortunately, the other implication, stated this way, is not true: with some more effort (which we do not do here), it could be shown that every complex $\mathbf{K}$ from the reduction in [GPP$^+$19] satisfies $\mathrm{plgcat}(|\mathbf{K}|) = 2$. However, this problem can be circumvented. We construct certain enriched complex $\mathbf{K}^+$ (by attaching a torus in a suitable way to every triangle of $\mathbf{K}$—it may be slightly surprising that this indeed helps). It turns out that $\mathrm{plgcat}(|\mathbf{K}^+|)$ stays 2 for shellable $\mathbf{K}$ but it grows to 3 for non-shellable $\mathbf{K}$ (coming from [GPP$^+$19]). This will prove Theorem 3.2.

We point out that Proposition 3.3 as stated is not really necessary in the proof of Theorem 3.2. But we state it here as it provides the motivation for our approach as well as it can be seen as a complementary result to the results of Borghini [Bor20] providing some sufficient (or necessary) conditions for $\mathrm{plgcat}(P) \leq 2$.

We also point out that instead of the reduction from [GPP$^+$19], it would be in principle possible to use also a modification of reduction by Santamaría-Galvis and Woodroofe [SGW21] where some of the gadgets are slightly simplified. However, some intermediate steps in [GPP$^+$19] are done via collapsibility thus, for our purposes, it is easier to adapt to the setting in [GPP$^+$19].

3.2 PL collapsibility of 2-complexes

It is a folklore result going back at least to Lickorish (according to [HAMS93]) that simplicial 2-complexes can be collapsed greedily:

Proposition 3.4 (see [HAMS93, page 20] or [MF08, Lemma 1 + Corollary 1])**.** *Let $\mathbf{K}$ be a collapsible 2-complex. Assume that $\mathbf{K}$ collapses to a subcomplex $\mathbf{L}$. Then $\mathbf{L}$ is collapsible as well. In particular, it can be checked in polynomial time whether a simplicial 2-complex is collapsible.*

For PL collapsibility we can essentially deduce the same conclusion as for collapsibility as soon as we observe that PL collapsibility of a 2-complex does not depend on the choice of the subdivision, which also might be a folklore result.

Lemma 3.5. *Let $\mathbf{K}$ be a simplicial complex of dimension at most 2 and $\mathbf{K}'$ be a subdivision of $\mathbf{K}$. Then $\mathbf{K}$ is collapsible if and only if $\mathbf{K}'$ is collapsible.*

In the proof of the lemma we use the following observation.

Observation 3.6. *Let τ be a triangle with vertices a, b, c. Let $\mathbf{K}'$ be an arbitrary subdivision of τ. Then $\mathbf{K}'$ collapses to the subcomplex $\mathbf{V}'$ formed by the subdivision of the edges ab and bc.*

Proof. We greedily perform collapses through free edges of $\mathbf{K}'$ which are not in $\mathbf{V}'$. Let $\mathbf{L}'$ be the resulting complex. We observe that $\mathbf{L}'$ contains no triangle. Indeed, every edge contained in some triangle of $\mathbf{L}'$ is either an edge of $\mathbf{V}'$ or it has to be contained in both neighboring triangles (otherwise we could continue with collapses). This means, because the dual graph of $\mathbf{K}'$ is connected, that once there is a single triangle of $\mathbf{K}'$ in $\mathbf{L}'$, then $\mathbf{L}'$ contains all triangles of $\mathbf{K}'$ which is a contradiction.

Thus, $\mathbf{L}'$ contains no triangles and it has the same homotopy type as $\mathbf{K}'$. That means that $\mathbf{L}'$ is a tree. Now we greedily perform collapses of edges not in $\mathbf{V}'$ through vertices of degree 1. By essentially the same argument as above, only the edges of $\mathbf{V}'$ remain (otherwise, we would find a cycle in $\mathbf{L}'$). □

Proof of Lemma 3.5. First, we show that if $\mathbf{K}$ is collapsible, then $\mathbf{K}'$ is collapsible by induction on the number of simplices of $\mathbf{K}$ (the case of one vertex is trivial). Assume that $\mathbf{K}_1$ arises from $\mathbf{K}$ by the first elementary collapse in some collapsing of $\mathbf{K}$. First, assume that it removes an edge ac and a triangle abc. Perform the collapses from Observation 3.6 on $\mathbf{K}'$ obtaining a complex $\mathbf{K}_1'$. Then $\mathbf{K}_1'$ is a subdivision of $\mathbf{K}_1$. Thus, it collapses by induction. The other option is that the first elementary collapse removes some vertex a and some edge ab. Then we obtain a subdivision $\mathbf{K}_1'$ of $\mathbf{K}_1$ by collapses on $\mathbf{K}'$ removing a and the subdivided edge ab in direction from a towards b.

Now, we show that if $\mathbf{K}'$ is collapsible, then $\mathbf{K}$ is collapsible again by induction on the number of simplices of $\mathbf{K}$. Assume that $\mathbf{K}'$ is collapsible. This implies that $\mathbf{K}'$ contains a free face σ' (a vertex or an edge) which subdivides a face σ of $\mathbf{K}$ which again has to be free. We perform a collapse on $\mathbf{K}$ through σ obtaining $\mathbf{K}_1$. As in the previous paragraph, we also collapse $\mathbf{K}'$ to a subdivision $\mathbf{K}_1'$ of $\mathbf{K}_1$. By Proposition 3.4 we get that $\mathbf{K}_1'$ is collapsible. Therefore, $\mathbf{K}_1$ is collapsible by induction which also implies that $\mathbf{K}$ is collapsible. □

Proof of Proposition 3.1. Let $\mathbf{K}$ be the input triangulation of P. By definition, $\mathrm{plgcat}(P) = 1$ if and only if P is PL collapsible which occurs if and only if some subdivision $\mathbf{K}'$ of $\mathbf{K}$ is collapsible. By Lemma 3.5, it is sufficient to check whether $\mathbf{K}$ is collapsible. This can be done in polynomial time due to Proposition 3.4. □

3.3 NP-hardness of PL geometric category 2

In this section, we prove Theorem 3.2. As we have sketched in the introduction, in our construction we need to attach a torus to every triangle of a certain intermediate complex. We start with the details regarding this attachment.

3.3.1 Attaching tori

First, let us us consider the standard torus $T = S^1 \times S^1$. An important curve in T is the *longitude* $\lambda = S^1 \times \{\cdot\}$ where '$\cdot$' stands for some fixed point in S^1.

Definition 3.7 (Enriched complex $\mathbf{K}^+$)**.** Given a simplicial complex $\mathbf{K}$, we define the *enriched* complex $\mathbf{K}^+$ as follows. For each triangle $\tau \in \mathbf{K}$ we consider a copy $\mathbf{T}_\tau$ of the standard torus with longitude λ_τ triangulated as in Figure 3.2. We get $\mathbf{K}^+$ as a result of gluing all tori $\mathbf{T}_\tau$ to $\mathbf{K}$ so that we identify λ_τ with $\partial\tau$. In the sequel, we consider $\mathbf{K}$ as well as all the tori $\mathbf{T}_\tau$ as subcomplexes of $\mathbf{K}^+$.

Note that the enriched complex $\mathbf{K}^+$ can be constructed in polynomial time in the size of $\mathbf{K}$.

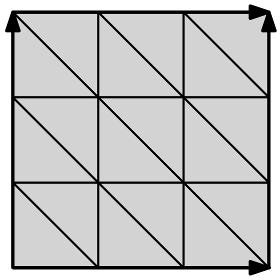

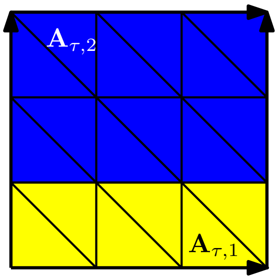

Figure 3.2: Left: The torus $\mathbf{T}_\tau$ with longitude λ_τ. Opposite edges are identified as usual. Right: Splitting $\mathbf{T}_\tau$ to two annuli.

Observation 3.8. *If $\mathbf{K}$ admits a covering by two collapsible subcomplexes $\mathbf{K}_1, \mathbf{K}_2$ such that both $\mathbf{K}_1$ and $\mathbf{K}_2$ contain the whole 1-skeleton of $\mathbf{K}$ then $\mathbf{K}^+$ can be also covered by two collapsible subcomplexes.*

Proof. Split each $\mathbf{T}_\tau$ to two annuli $\mathbf{A}_{\tau,1}$ and $\mathbf{A}_{\tau,2}$ as in Figure 3.2. (Both of them are subcomplexes of $\mathbf{T}_\tau$ and they share λ_τ on one of their boundaries.) Take $\mathbf{K}_i^+$ as the union of $\mathbf{K}_i$ and all annuli $\mathbf{A}_{\tau,i}$ for $i \in \{1, 2\}$. Then $\mathbf{K}_1^+$ and $\mathbf{K}_2^+$ cover $\mathbf{K}^+$. In addition, they are both collapsible because $\mathbf{K}_i^+$ collapses to $\mathbf{K}_i$ as each $\mathbf{A}_{\tau,i}$ collapses to λ_τ. □

We continue with the main technical lemma for our reduction.

Lemma 3.9. *Let P be a polyhedron which is a union of two subpolyhedra R and T. Assume that $T = S^1 \times S^1$ is the torus and assume that R and T intersect exactly in the longitude $\lambda = S^1 \times \{\cdot\}$ of T. Assume that P can be covered by two contractible subpolyhedra Q_1, Q_2. Then $\lambda \subseteq Q_1, Q_2$ and λ is nullhomologous in $R \cap Q_1$ as well as in $R \cap Q_2$.*

Proof. Let $A_i := T \cap Q_i$ for $i = 1, 2$. The lemma is implied by the following two claims where all the homology is considered with $\mathbb{Z}_2$ coefficients.

Claim 3.9.1.

(i) If $H_1(A_1) = 0$, then $\dim H_1(A_2) \geq 2$.

(ii) If $H_1(A_2) = 0$, then $\dim H_1(A_1) \geq 2$.

Claim 3.9.2.

(i) If $H_1(A_1) \neq 0$, then $\dim H_1(A_1) = 1$, λ belongs to Q_1 and λ is nullhomologous in $R \cap Q_1$.

(ii) If $H_1(A_2) \neq 0$, then $\dim H_1(A_2) = 1$, λ belongs to Q_2 and λ is nullhomologous in $R \cap Q_2$.

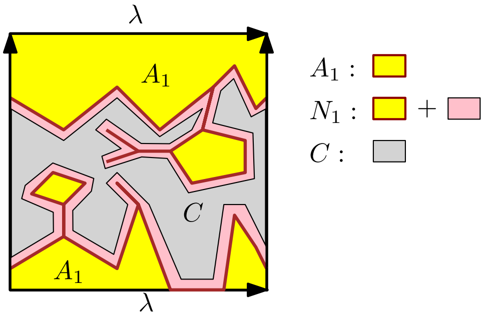

Figure 3.3: A_1, N_1 and C inside T.

Indeed, the conjunction of Claims 3.9.1 and 3.9.2 implies that only option is that $\dim H_1(A_1) = \dim H_1(A_2) = 1$ and thus we can use the conclusions of Claim 3.9.2. Therefore, it remains to prove the claims. In each of the claims, we only prove the first item as the other one is symmetric.

Proof of Claim 3.9.1(i). Let N_1 be the regular neighborhood[2] of A_1 inside T, which is homotopy equivalent to A_1; see Figure 3.3. Then N_1 is a surface with boundary. Thus we may apply the Lefschetz duality[3] obtaining

$$H_1(N_1, \partial N_1) \cong H^1(N_1) \cong H_1(N_1) \cong H_1(A_1) = 0 \tag{3.1}$$

where the second isomorphism follows from the fact that the homology and the cohomology groups are isomorphic over a field.

Now let C be the closure of the complement of N_1 in T, that is, $C := \overline{T \setminus N_1}$. By the excision property of homology, and then by (3.1)

$$H_1(T, C) \cong H_1(N_1, \partial N_1) = 0. \tag{3.2}$$

Finally, we consider the long exact sequence of the pair:

$$\cdots \to H_1(C) \xrightarrow{i_*} H_1(T) \to H_1(T, C) \to \cdots$$

The map i_* is induced by the inclusion $i\colon C \to T$. Because of (3.2), the map i_* is surjective. The inclusion i can be decomposed into inclusions $j\colon C \to A_2$ and $k\colon A_2 \to T$. (Note that $C \subseteq A_2$ as A_1 and A_2 cover T.) By functoriality of homology, $k_*\colon H_1(A_2) \to H_1(T)$ must be surjective as well. Therefore, $\dim H_1(A_2) \geq \dim H_1(T) = 2$. $\square$

Proof of Claim 3.9.2(i). Let $R_1 := R \cap Q_1$. Consider the Mayer-Vietoris exact sequence:

$$\cdots \to H_1(A_1 \cap R_1) \xrightarrow{f} H_1(A_1) \oplus H_1(R_1) \xrightarrow{g} H_1(Q_1) \to \cdots$$

As we assume that Q_1 is contractible, we get $H_1(Q_1) = 0$. Therefore, f is surjective (from exactness). As we also assume that $H_1(A_1) \neq 0$, there is

[2] In this case N_1 is a 2-manifold with boundary inside T which collapses to A_1. For a general definition of regular neighborhood see [RS82, Chapter 3].

[3] Lefschetz duality (see e.g. Theorem 3.43 in [Hat02]) over $\mathbb{Z}_2$: Let M be an n-dimensional compact manifold with boundary N. Then $H_i(M, N; \mathbb{Z}_2) \cong H^{n-i}(M; \mathbb{Z}_2)$ for every i.

a nonzero vector $v = (z, 0) \in H_1(A_1) \oplus H_1(R_1)$. We know $v \in \operatorname{im} f$ as f is surjective. In particular, $H_1(A_1 \cap R_1) \neq 0$. On the other hand, $A_1 \cap R_1 \subseteq T \cap R = \lambda$. Therefore, $A_1 \cap R_1 = \lambda$. This gives $\lambda \subseteq Q_1$ as we need. Using that f is surjective again, we get $\dim H_1(A_1) + \dim H_1(R_1) \leq \dim H_1(A_1 \cap R_1) = 1$. Because $H_1(A_1) \neq 0$ we actually get $\dim H_1(A_1) = 1$ and $\dim H_1(R_1) = 0$. This gives that λ is nullhomologous in $R_1 = R \cap Q_1$. □

□

3.3.2 Construction from [GPP+19]

As we sketched in the introduction, we use the construction from [GPP+19] as an intermediate step. Given that this construction is somewhat elaborated, we prefer to state it as a blackbox only mentioning the properties that we need in our reduction.

The NP-hardness in [GPP+19] is proved by a reduction from the classical 3-satisfiability problem.

Proposition 3.10 ([GPP+19])**.** *There is a polynomial time algorithm that produces from a given 3-CNF formula ϕ (with n variables) a pure 2-dimensional complex $\mathbf{K}_\phi$ with the following properties.*

(i) *$\mathbf{K}_\phi$ contains pairwise disjoint triangulated 2-spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$, one for each variable.*

(ii) *The second homology group, $H_2(\mathbf{K}_\phi)$, is generated by the spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$. In particular, $H_2(\mathbf{K}_\phi) \cong \mathbb{Z}_2^n$ and no triangle outside the spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$ is contained in a 2-cycle.*

(iii) *If ϕ is satisfiable, then there are triangles τ_i in $\mathbf{S}_i$ for every $i \in [n]$ such that $\mathbf{K}_\phi$ becomes collapsible after removing these triangles. In addition, for every $i \in [n]$, there are at least two options how to pick τ_i in $\mathbf{S}_i$. (Such a choice can be done independently in each $\mathbf{S}_i$ yielding at least 2^n collapsible subcomplexes.)*

(iv) *If an arbitrary subdivision of $\mathbf{K}_\phi$ becomes collapsible after removing some n triangles, then ϕ is satisfiable.*

Proof. The proof of the proposition consists mostly of references to [GPP+19]. However, a few items are not as explicitly stated in [GPP+19] as we need them here, thus we explain in detail how all the items of the proposition can be deduced from the text in [GPP+19].

The construction of $\mathbf{K}_\phi$ is given in Section 4 of [GPP+19]. The spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$ of item (i) are the spheres $S(u)$ introduced in §4.3 of [GPP+19]. For checking the other items, we first point out that [GPP+19, Proposition 12] states that the number of variables, n, is equal to the reduced Euler characteristic $\tilde{\chi}(\mathbf{K}_\phi)$.

It is stated in Remark 13 in [GPP+19] that $\mathbf{K}_\phi$ is homotopy equivalent to the wedge of n 2-spheres; in particular, $\dim H_2(\mathbf{K}_\phi) = n$. Then item (ii) immediately follows as the disjoint spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$ generate a subspace of dimension n in

$H_2(\mathbf{K}_\phi)$. Unfortunately, Remark 13 is only a side remark in [GPP+19] and it is not proved there. Therefore, we explain in the last section of this chapter (Section 3.3.4), Proposition 3.11, how Remark 13 of [GPP+19] follows from their tools.

Item (iii), using $n = \tilde{\chi}(\mathbf{K}_\phi)$, is the content of Proposition 8(ii) in [GPP+19] with the addendum that it is also necessary to check the proof: In the beginning of Section 7 of [GPP+19], it is specified that the triangles are removed in certain regions $D[\ell(u)]$. By checking the construction of $D[\ell(u)]$ in §4.3 of [GPP+19], these regions are in the correct spheres ($\mathbf{S}_i$ in our notation; $\mathbf{S}(u)$ in the notation of [GPP+19]) and in addition there are at least two choices of the removed triangle for every i (actually exactly three choices).

Item (iv), using $n = \tilde{\chi}(\mathbf{K}_\phi)$, is exactly the content of Proposition 8(iii). □

3.3.3 The final reduction

Proof of Theorem 3.2. Given a 3-CNF formula ϕ and its corresponding complex $\mathbf{K}_\phi$ we construct its enriched complex $\mathbf{K}_\phi^+$. (See Definition 3.7.) Theorem 3.2 is proved by showing that ϕ is satisfiable if and only if $\operatorname{plgcat}(\mathbf{K}_\phi^+) \leq 2$ as 3-satisfiability is an NP-hard problem.

(a) *ϕ is satisfiable $\implies$ $\mathbf{K}_\phi^+$ can be covered by two collapsible subcomplexes.*

Suppose that the formula ϕ is satisfiable. Then by Proposition 3.10(iii) $\mathbf{K}_\phi$ is collapsible after removal of n triangles, one from each sphere $\mathbf{S}_i$, and for each $\mathbf{S}_i$ there are at least two options, say $\tau_i^{(1)}, \tau_i^{(2)}$, how to pick such a triangle. Therefore, the subcomplexes

$$\mathbf{K}_1 := \mathbf{K}_\phi \setminus \{\tau_1^{(1)}, \ldots, \tau_n^{(1)}\},\ \mathbf{K}_2 := \mathbf{K}_\phi \setminus \{\tau_1^{(2)}, \ldots, \tau_n^{(2)}\}$$

are collapsible subcomplexes of $\mathbf{K}_\phi$ and they cover it.

Moreover, each of $\mathbf{K}_1$ and $\mathbf{K}_2$ contains the whole 1-skeleton of $\mathbf{K}_\phi$. Indeed, the complex $\mathbf{K}_\phi$ is pure thus every edge of $\mathbf{K}_\phi$ is contained in at least one triangle and in addition in at least two triangles if it is an edge in some of the spheres $\mathbf{S}_i$. In order to get $\mathbf{K}_1$ or $\mathbf{K}_2$, at most one triangle is removed from each $\mathbf{S}_i$. Therefore, each edge of $\mathbf{K}_\phi$ is still contained in at least one triangle of $\mathbf{K}_1$ and in at least one triangle of $\mathbf{K}_2$. Then Observation 3.8 implies that $\mathbf{K}_\phi^+$ can be covered by two collapsible subcomplexes.

(b) *A subdivision $\left(\mathbf{K}_\phi^+\right)'$ of $\mathbf{K}_\phi^+$ can be covered by two collapsible subcomplexes $\implies$ ϕ is satisfiable.*

First, we sketch the idea: Let $(\mathbf{K}_1^+)'$ and $(\mathbf{K}_2^+)'$ be the two collapsible subcomplexes of $(\mathbf{K}_\phi^+)'$ covering it. (We point out that $(\mathbf{K}_i^+)'$ is just a notation not implying that $(K_i^+)'$ is a subdivision of some complex $\mathbf{K}_i^+$.) We want to verify the assumption in Proposition 3.10(iv) in order to deduce that ϕ is satisfiable. For this, we need a subdivision of $\mathbf{K}_\phi$ such that removing n triangles from this subdivsion yields a collapsible complex. In fact, our subdivision will be trivial, thus we need to find n triangles in $\mathbf{K}_\phi$ such that their removal yields a collapsible complex. We will take $(\mathbf{K}_1^+)'$, say, and we will (essentially) deduce that in each $\mathbf{S}_i$ there must be τ_i such that $(\mathbf{K}_1^+)'$ must miss at least one triangle in the subdivided τ_i. These triangles τ_i are the triangles we want to remove from $\mathbf{K}_\phi$. However, we need several intermediate claims to deduce that the resulting complex is indeed collapsible. (We will use the second complex $(\mathbf{K}_2^+)'$ only very sparingly in order to verify the assumptions of Lemma 3.9.)

Let $\mathbf{K}'_\phi$ be the subcomplex of $(\mathbf{K}_\phi^+)'$ corresponding to $\mathbf{K}_\phi$ in this subdivision. (Let us recall that this means that $\mathbf{K}'_\phi$ is formed by simplices $\sigma \in (\mathbf{K}_\phi^+)'$ such that $\sigma \subseteq |\mathbf{K}_\phi|$.) Let $\mathbf{K}'_1 := \mathbf{K}'_\phi \cap (\mathbf{K}_1^+)'$.

Claim 3.10.1. *The complex $\mathbf{K}'_1$ is a collapsible subcomplex of $\mathbf{K}'_\phi$.*

Proof. Our aim is to show that $(\mathbf{K}_1^+)'$ collapses to $\mathbf{K}'_1$. Then it follows from Proposition 3.4 that $\mathbf{K}'_1$ is collapsible.

We pick an arbitrary triangle τ of $\mathbf{K}_\phi$. Recall that $\mathbf{T}_\tau$ is the torus attached to τ. (See Definition 3.7.) Let $\mathbf{T}'_\tau$ be the subcomplex of $(\mathbf{K}_\phi^+)'$ corresponding to $\mathbf{T}_\tau$. Note that (the subdivsion of) $\partial\tau$ belongs to $(\mathbf{K}_1^+)'$ by Lemma 3.9. We also observe that $\mathbf{T}'_\tau$ is not a subcomplex of $(\mathbf{K}_1^+)'$ otherwise $(\mathbf{K}_1^+)'$ would contain a nontrivial 2-cycle which is not possible if it is collapsible.

Now we proceed similarly as in the proof of Observation 3.6. We greedily perform collapses in $(\mathbf{K}_1^+)'$ on simplices of $\mathbf{T}'_\tau$ with the exception that we are not allowed to remove the simplices belonging to (the subdivision of) $\partial\tau$. (See Figure 3.4 for a realistic example of the intersection of $(\mathbf{K}_1^+)'$ and $\mathbf{T}'_\tau$.) Let $\mathbf{L}'$ be the resulting complex. We first observe that $\mathbf{L}'$ contains no triangles of $\mathbf{T}'_\tau$ as at least one triangle is missing and the dual graph to our triangulation of $\mathbf{T}'_\tau$ is connected even after removing the dual edges crossing $\partial\tau$. Therefore, $\mathbf{L}' \cap \mathbf{T}'_\tau$ is a graph. Due to our restriction on collapses, subdivided $\partial\tau$ is inside this graph. We observe that no other (graph theoretic) cycle may belong to this graph. Indeed, another cycle would contain an edge which is not in $\partial\tau$, thus not contained in any triangle of $\mathbf{L}'$. Therefore, such a cycle could not be filled with a 2-chain, and thus it would be necessarily homologically nontrivial in $\mathbf{L}'$ which is a contradiction with the fact that $\mathbf{L}'$ is contractible (obtained by collapses from a collapsible complex). Thus, we may conclude that $\mathbf{L}' \cap \mathbf{T}'_\tau$ is the subdivided $\partial\tau$ with a collection of pendant trees. However, these pendant trees have to be actually trivial as they get collapsed during the greedy collapses.

Altogether we have collapsed $(\mathbf{K}_1^+)'$ to a complex $\mathbf{L}'$ which agrees with $\mathbf{K}'_1$ on $\mathbf{K}'_\phi$ while we have removed all simplices of $\mathbf{T}'_\tau$ except those that belong to $\mathbf{K}'_\phi$. Now we pick another triangle σ of $\mathbf{K}_\phi$ and we remove (via collapses) the simplices of $\mathbf{T}'_\sigma$ except those belonging to $\mathbf{K}'_\phi$ by an analogous approach. After passing through every triangle of $\mathbf{K}_\phi$, we get exactly $\mathbf{K}'_1$ as required. □

Claim 3.10.2. *For every triangle $\tau \in \mathbf{K}_\phi$, $\partial\tau$ is contained in $|\mathbf{K}'_1|$ and it is nullhomologous in $|\mathbf{K}'_1|$.*

Proof. Let $P := |\mathbf{K}_\phi^+| = |(\mathbf{K}_\phi^+)'|$. Let R be the polyhedron of $\mathbf{K}_\phi$ and all tori of $\mathbf{K}_\phi^+$ except $\mathbf{T}_\tau$. Let $Q_1 := |(\mathbf{K}_1^+)'|$ and $Q_2 := |(\mathbf{K}_2^+)'|$. Then R, $|\mathbf{T}_\tau|$, Q_1 and Q_2 satisfy the assumptions of Lemma 3.9. Then we deduce that $\partial\tau$ is nullhomologous in $R \cap Q_1$. Assume that τ is such that $\mathbf{T}'_\tau$ is the first torus

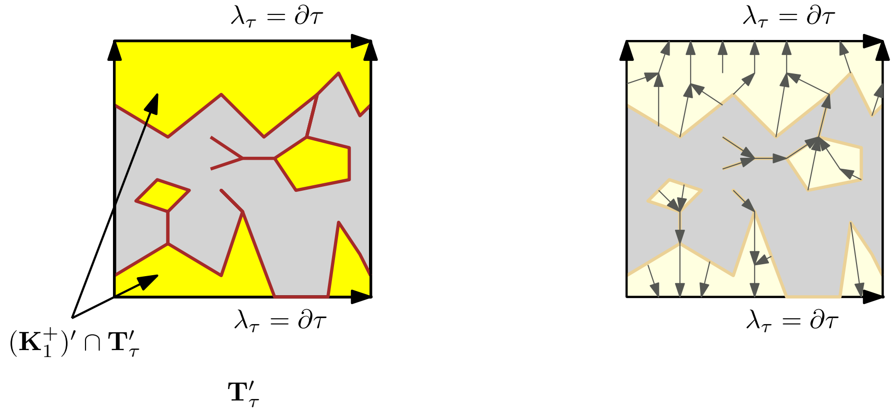

Figure 3.4: Left: A realistic example how may $(\mathbf{K}_1^+)'$ intersect T_τ. Right: Greedy collapses (only very schematically without emphasizing the triangulation).

to be removed in the proof of Claim 3.10.1, we can choose so. Then $R \cap Q_1$ is exactly the polyhedron of $\mathbf{L}'$ in the proof of Claim 3.10.1. In particular, $\mathbf{L}'$ collapses to $\mathbf{K}'_1$. As collapses provide a homotopy equivalence, we deduce that $\partial\tau$ is nullhomologous in $|\mathbf{K}'_1|$ as well. □

Now, for any triangle $\tau \in \mathbf{K}_\phi$ let τ' be the subcomplex of $\mathbf{K}'_\phi$ corresponding to this triangle.

Claim 3.10.3.

(i) If $\tau \in \mathbf{K}_\phi$ is a triangle which does not belong to any of the spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$, then τ' is a subcomplex of $\mathbf{K}'_1$.

(ii) For every $i \in [n]$, all triangles τ in $\mathbf{S}_i$ except exactly one satisfy that τ' is a subcomplex of $\mathbf{K}'_1$.

Proof. Let $\tau \in \mathbf{K}_\phi$. Due to Claim 3.10.2, it has to be possible to fill the subdivision of $\partial\tau$ by some 2-chain $c = c(\tau)$ in $\mathbf{K}'_1$.

If τ does not belong to any of the spheres $\mathbf{S}_1, \dots, \mathbf{S}_n$, then the only option for c is to contain all simplices of τ'. Indeed, if there is another such c', then considering τ' as a 2-chain, we get a nontrivial 2-cycle $\tau' + c'$ with support at least partially outside the spheres $S_1, \dots, S_n$ which contradicts Proposition 3.10(ii). Therefore, τ' must be a subcomplex of $\mathbf{K}'_1$ which concludes (i).

Now for (ii), take $i \in [n]$. Then $\mathbf{K}'_1$ has to miss at least one triangle in $|\mathbf{S}_i|$ otherwise subdivided $\mathbf{S}_i$ forms a non-trivial 2-cycle in $\mathbf{K}'_1$ which is a contraction with Claim 3.10.1. Assume that τ in $\mathbf{S}_i$ was chosen so that this missing triangle belongs to τ'. Then $\partial\tau$ splits (subdivided) $\mathbf{S}_i$ to two hemispheres; one of them is formed by τ' and another is formed by the union of subcomplexes σ' taken over all triangles σ in $\mathbf{S}_i$ different from τ. By using Proposition 3.10(ii) again, the only options are that $c = c(\tau)$ contains all the simplices of one or the other (subdivided) hemispheres. But the hemisphere of τ' is ruled out as τ' misses a triangle of $\mathbf{K}'_1$. Thus c has to be filled by

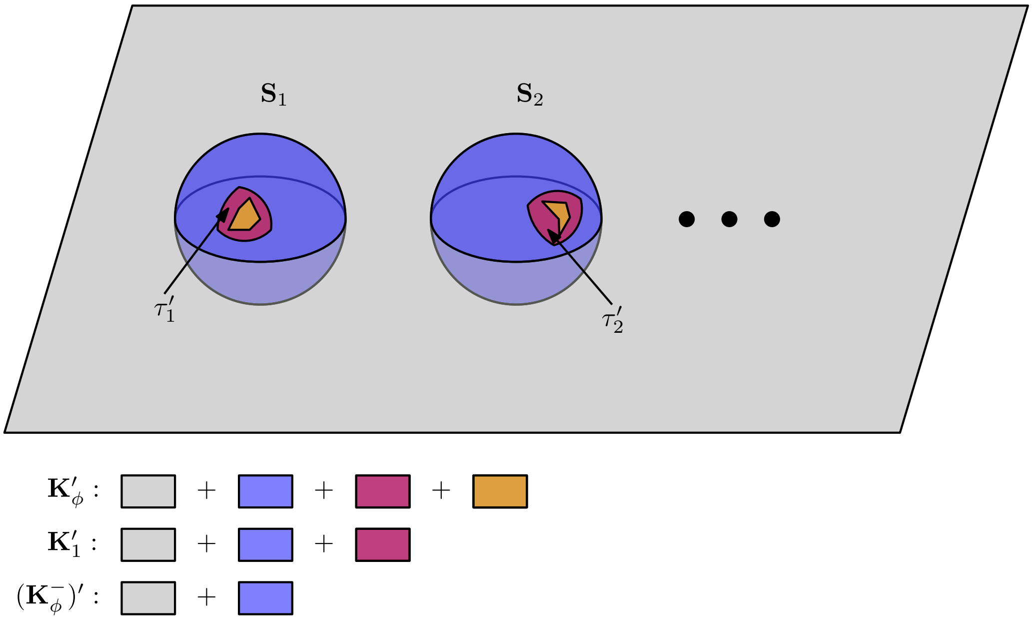

Figure 3.5: A schematic drawing of $\mathbf{K}_\phi'$, $\mathbf{K}_1'$ and $\mathbf{K}_\phi^-$. We emphasize that this not really a realistic drawing of $\mathbf{K}_\phi'$ (with the same polyhedron as $\mathbf{K}_\phi$) as constructed in [GPP+19]. We only attempt to draw as simple complex as possible satisfying conclusions (i) and (ii) of Proposition 3.10 and so that $\mathbf{K}_1'$ is collapsible. (The space inside the spheres is completely hollow.)

the other hemisphere. Then we conclude (ii) for all simplices σ in $\mathbf{S}_i$ except exactly τ as required. $\square$

In the light of Claim 3.10.3(ii), let τ_i be the unique triangle of $\mathbf{S}_i$ such that τ_i' is not a subcomplex of $\mathbf{K}_1'$. Let $\mathbf{K}_\phi^-$ be the subcomplex of $\mathbf{K}_\phi$ obtained by removing all triangles $\tau_1, \dots, \tau_n$ and let $(\mathbf{K}_\phi^-)'$ be the subcomplex of $\mathbf{K}_\phi'$ corresponding to $\mathbf{K}_\phi^-$. Note that Claim 3.10.3 implies that $(\mathbf{K}_\phi^-)'$ is a subcomplex of $\mathbf{K}_1'$. See Figure 3.5 for comparison of $\mathbf{K}_\phi'$, $\mathbf{K}_1'$ and $\mathbf{K}_\phi^-$ after using Claim 3.10.3.

Claim 3.10.4. $\mathbf{K}_1'$ *collapses to* $(\mathbf{K}_\phi^-)'$.

Proof. The complexes $\mathbf{K}_1'$ and $(\mathbf{K}_\phi^-)'$ differ only so that $\mathbf{K}_1'$ may contain some simplices of τ_i' for some i (except those that subdivide $\partial\tau_i$) which are not in $(\mathbf{K}_\phi^-)'$.

Now, we continue analogously as we did in the proof of Observation 3.6 or Claim 3.10.1. We greedily collapse all simplices of $\mathbf{K}_1'$ in τ_i' except those that subdivide $\partial\tau_i$. We first deduce that the resulting complex contains no triangles of τ_i' as at least one triangle was missing in the beginning. Then we deduce that there is no graph-theoretic cycle among simplices of τ_i' except the one corresponding to $\partial\tau_i$ by the same argument as in the proof of Claim 3.10.1 (using that $\mathbf{K}_1'$ is collapsible). Then, we deduce that among the simplices of τ_i' only the simplices subdividing $\partial\tau_i$ remain in the complex. After repeating this approach for every $i \in [n]$ we obtain $(\mathbf{K}_\phi^-)'$. $\square$

Now, we have acquired enough tools to conclude the case (b) and therefore to conclude the proof of the theorem. From Claims 3.10.1 and 3.10.4 and Proposition 3.4 we deduce that $(\mathbf{K}_\phi^-)'$ is collapsible. By Lemma 3.5 we deduce that $\mathbf{K}_\phi^-$ is collapsible. Finally, by Proposition 3.10(iv) we deduce that ϕ is satisfiable.

$\square$

3.3.4 Verification of Remark 13 from [GPP+19]

Our aim in this subsection is to verify Remark 13 in [GPP+19] which is stated but not proved in [GPP+19]. The exact statement we need is given by the following proposition. We will provide all the necessary detail in order to verify correctness of Remark 13 of [GPP+19]. On the other hand, we warn the reader that our proof is not self-contained but it relies on the construction of $\mathbf{K}_\phi$ and partially the notation in [GPP+19]; thus it is necessary to consult the contents of [GPP+19].

Proposition 3.11. *The complex $\mathbf{K}_\phi$ from [GPP+19] is homotopy equivalent to the wedge of n 2-spheres (where n is the number of variables).*

In the proof, we need the following simple lemma.

Lemma 3.12. *Let $\mathbf{K}_1, \mathbf{K}_2$ be simplicial complexes. Assume that $\mathbf{K}_1 \cap \mathbf{K}_2$ and $\mathbf{K}_2$ are contractible, then $\mathbf{K}_1$ and $\mathbf{K}_1 \cup \mathbf{K}_2$ are homotopy equivalent.*

Proof. It is well known that contracting a contractible subcomplex is a homotopy equivalence [Mat03, Proposition 4.1.5]. Therefore, we get

$$|\mathbf{K}_1 \cup \mathbf{K}_2| \simeq |\mathbf{K}_1 \cup \mathbf{K}_2|/|\mathbf{K}_2| = |\mathbf{K}_1|/|\mathbf{K}_1 \cap \mathbf{K}_2| \simeq |\mathbf{K}_1|$$

as required. $\square$

Proof of Proposition 3.11. We follow essentially in verbatim the proof of Proposition 12 in [GPP+19]. The only difference is that we use Lemma 3.12 instead of the weaker statement in [GPP+19]: If $\mathbf{K}_1 \cap \mathbf{K}_2$ and $\mathbf{K}_2$ are contractible, then $\tilde{\chi}(\mathbf{K}_1 \cup \mathbf{K}_2) = \tilde{\chi}(\mathbf{K}_1)$ where $\tilde{\chi}$ stands for the reduced Euler characteristic.

As described in the proof of Proposition 12 in [GPP+19], the complex $\mathbf{K}_\phi$ can be transformed into certain complex $\mathbf{K}'$ by a series of steps when we decompose some intermediate complex as $\mathbf{K}_1 \cup \mathbf{K}_2$ where $\mathbf{K}_2$ and $\mathbf{K}_1 \cap \mathbf{K}_2$ are contractible, and then we replace the intermediate complex with $\mathbf{K}_1$. Therefore, using Lemma 3.12 we get that the resulting complex $\mathbf{K}'$, after performing all these steps is homotopy equivalent to $\mathbf{K}_\phi$.

By a further homotopy equivalence Goaoc et al., [GPP+19], obtain another complex $\mathbf{K}''$ which is already (obviously) homotopy equivalent to the wedge of n 2-spheres. Therefore, $\mathbf{K}_\phi$ is homotopy equivalent to the wedge of n 2-spheres. $\square$

4. VEST and related problems from the viewpoint of parameterized complexity

4.1 Introduction

The homotopy groups π_k, for $k = 1, 2, \ldots$, are important invariants of topological spaces. The most intuitive of them is the group π_1 which is often called the fundamental group.

Many topological spaces can be described by finite structures, e.g., by abstract simplicial complexes. Such structures can be used as an input for a computer and therefore, it is natural to ask how hard it is to compute these homotopy groups of a given topological space represented by an abstract simplicial complex.

Novikov in 1955 [Nov55] and independently Boone in 1959 [Boo59] showed undecidability of the word problem for groups. Their result also implies undecidability of computing the fundamental group. In fact, even determining whether the fundamental group of a given topological space is trivial is undecidable.

On the other hand, for simply connected spaces (for those, whose π_1 is trivial) it is know that their π_k for $k \geq 2$ are finitely generated abelian groups which are always isomorphic to groups of the form

$$\mathbb{Z}^n \oplus \mathbb{Z}_{p_1} \oplus \mathbb{Z}_{p_2} \oplus \cdots \oplus \mathbb{Z}_{p_m},$$

where $p_1, \ldots, p_m$ are powers of prime numbers. An algorithm for computing π_k of a simply connected topological space, where $k \geq 2$, was first introduced by Brown in 1957 [Bro57].

In 1989, Anick [Ani89] proved that even computing the rank of π_k, that is, the number of direct summands isomorphic to $\mathbb{Z}$ (represented by n in the expression above) is #P-hard for 4-dimensional simply connected spaces.[1] Another computational problem called VEST, which we define below, was used in Anick's proof as an intermediate step. Briefly said, #P-hardness of the problem of VEST implies #P-hardness of computing the rank of π_k.

Vector evaluated after a sequence of transformations (VEST). The input of this problem defined by Anick [Ani89] is a vector $\mathbf{v} \in \mathbb{Q}^d$, a list $(T_1, \ldots, T_m)$ of rational $d \times d$ matrices and a rational matrix $S \in \mathbb{Q}^{h \times d}$ where $d, m, h \in \mathbb{N}$.

Now, let *M-sequence* be a sequence of integers $M_1, M_2, M_3, \ldots$, where

$$M_k := |\{(i_1, \ldots, i_k) \in \{1, \ldots, m\}^k; ST_{i_k} \cdots T_{i_1}\mathbf{v} = \mathbf{0}\}|.$$

Given an instance of VEST and $k \in \mathbb{N}$, the goal is to compute M_k.

From an instance of VEST, it is possible to construct a corresponding algebraic structure called $123H$*-algebra* in polynomial time whose *Tor-sequence* is equal to the M-sequence of the original instance of a VEST. This is stated

[1] When k is a part of the input and represented in unary.

in [Ani89, Theorem 3.4] and it follows from [Ani85, Theorem 1.3] and [Ani87, Theorem 7.6].

Given a presentation of a $123H$-algebra, one can construct a corresponding 4-dimensional simplicial complex in polynomial time whose sequence of ranks (rk π_2, rk $\pi_3, \ldots$) is related to the Tor-sequence of the $123H$-algebra. In particular, it is possible to compute that Tor-sequence from the sequence of ranks using an FPT algorithm. This follows from [Roo79] and [ČKM+14a].

Parameterized complexity and W-hierarchy It is also possible to look at the problem of computing π_k from the viewpoint of parameterized complexity. In our case, the dimension k of the homotopy group π_k plays the role of the parameter.

In 2014, Čadek et al. [ČKM+14b] proved that computing π_k (and thus, also computing the rank of π_k) of a simply connected space is in XP when parameterized by k.

A lower bound for the complexity from the parameterized viewpoint was obtained by Matoušek in 2013 [Mat13]. He proved that computing M_k of a VEST instance is #W[1]-hard. This also implies #W[1]-hardness for the original problem of computing the rank of higher homotopy groups π_k (for 4-dimensional simply connected spaces) when parameterized by k. Matoušek's proof also works as a proof for #P-hardness and it is shorter and considerably easier than the original proof of Anick in [Ani89].

In this chapter, we strengthen the result of Matoušek and show that computing M_k of a VEST instance is #W[2]-hard.

Theorem 4.1. *The problem of computing M_k of a VEST instance is #W[2]-hard when parameterized by k.*

Theorem 4.1 together with the result of Anick [Ani89] implies the following.

Corollary 4.2. *The problem of computing the rank of the k-th homotopy group for a 4-dimensional simply connected space is #W[2]-hard when parameterized by k.*

Remark 4.3. Note that computing M_k of a VEST instance is an interesting self-contained problem even without the topological motivation. We point out that our reduction showing #W[2]-hardness of this problem uses only 0, 1 values in the matrices and in the initial vector $\mathbf{v}$. Moreover, each matrix will have at most one 1 in each row. Therefore, such construction also shows #W[2]-hardness of computing M_k of a VEST instance in the $\mathbb{Z}_2$ setting. That is, for the case when $T_1, T_2, \ldots T_m \in \mathbb{Z}_2^{d\times d}, S \in \mathbb{Z}_2^{h\times d}$ and $\mathbf{v} \in \mathbb{Z}_2^d$.

The decision version of VEST We also provide a comprehensive overview of the parameterized complexity of VEST as a decision problem, where given an instance of VEST one needs to determine whether $M_k > 0$. In addition to the standard variant of the problem, we consider several simplified modifications of VEST: When the matrices are of constant size, when the special matrix S is the identity matrix, when we omit the initial vector and the target is the identity/zero matrix etc.

Unfortunately, even considering the simplifications above, we show that nearly all versions in our consideration are W[1]- or W[2]-hard. Table 4.1 is an overview of our results.

Size of T_i	a) no constraint	b) $S = I$		c) only S	d) no $\mathbf{v}$, no S
1. 1×1	in P	in P	$\mathbf{0}$	in P	in P
			I	W[1]-hard	W[1]-hard
2. 2×2	W[1]-hard	W[1]-hard	$\mathbf{0}$	W[1]-hard	W[1]-hard
			I		
3. input size	W[2]-hard	W[2]-hard	$\mathbf{0}$	W[2]-hard	W[2]-hard
			I	W[1]-hard	W[1]-hard

Table 4.1: The first column stands for the standard VEST while the second stands for the VEST without the special matrix S or alternatively, for the case when S is the identity matrix. Therefore, the hardness results for the first column follow from the second. The third and the fourth columns stands for the variants without the initial vector $\mathbf{v}$. In this case, it is natural to assume the following two targets for the result of the sought matrix product: the zero matrix (the rows labeled by $\mathbf{0}$) and the identity matrix (the rows labeled by I). Again, the hardness results for the third column follow from the fourth.

Regarding the 1×1 case, the only nontrivial case is when the target is $I = 1$. The W[1]-hardness results for the 1×1 case also implies W[1]-hardness for the 2×2 case and the input size case when the target is the identity matrix.

In Section 4.3, we show

- W[1]-hardness for the 1×1 case without S and $\mathbf{v}$ and with target $I = 1$, that is, "1 d) I" in Table 4.1 (Theorem 4.9).
- W[1]-hardness for the 2×2 case with $S = I$, that is, "2 b)" in Table 4.1 (Theorem 4.11).
- W[1]-hardness for the 2×2 case with without S and $\mathbf{v}$ and with target $\mathbf{0}$, that is, "2 d) $\mathbf{0}$" in Table 4.1 (Theorem 4.10).

The W[2]-hardness for the standard decision version of VEST, that is, for the case of input size without any constraint, "3 a)", follows from the proof of Theorem 4.1 (see Remark 4.6) and we show that the variant with $S = I$, "3 b)", and the variant without S and $\mathbf{v}$ and with target $\mathbf{0}$, "3 d) $\mathbf{0}$", are equivalent to the standard version, "3 a)", under parameterized reduction (Theorems 4.12, 4.13).

Note that it is also easy to observe that all discussed decision variants of VEST from Table 4.1 are in W[P].

Observation 4.4. *Each variant of the decision version of VEST from Table 4.1 is in* W[P] *for parameter k.*

Proof. Let n be the size of the input and m the number of matrices in the collection. In particular, $m \leq n$.

We can guess which k matrices we choose from the collection. Each matrix can be represented by an integer $\leq m$ which can be described by $\lceil \log m \rceil$ bits. Therefore, we need at most $k\lceil \log m \rceil \leq k(\log n + 1)$ non-deterministic choices.

Then, we need to multiply k (or $(k+1)$) matrices (together with 1 vector). This can be easily done in time $(k+2)n^3$. $\square$

Fixed-parameter tractability over finite fields Reductions from the previous section show that VEST remains hard even on highly restricted instances, such as binary matrices with all the ones located along the main diagonal (see Remark 4.6), or matrices of a constant size. However, it turns out that combination of this two restrictions – on the field size and the matrix sizes – makes even the counting version of VEST tractable.

We proceed by lifting tractability to the matrices of unbounded size but with all non-zero entries occurring in at most the p first rows.

Theorem 4.5. *Given an instance of VEST and $k \in \mathbb{N}$, computing M_k is* FPT *when parameterized by $|\mathbb{F}|$ and p, if all but the first p rows of the input matrices are zeros.*

The problem remains FPT with respect to $|\mathbb{F}|$ and p even if the task is to find the minimal k for which the vanishing sequence of length k exists, or to report that there is no such k.

Undecidability of VEST without parameter In contrast, we show in the last section (Section 4.5) that for $\mathbb{F} = \mathbb{Q}$ the problem of determining whether there exists k such that $M_k > 0$ for an instance of VEST is an undecidable problem (even for the case where $T_1, \dots, T_m$ are of size 4×4).

4.2 The proof of #W[2]-hardness of VEST

In this section, we prove that computing M_k of a VEST is #W[2]-hard (Theorem 4.1).

Our reduction is from the problem of counting dominating sets of size k which is known to be #W[2]-complete (see [FG04b]) and which we recall in the paragraph below.

For a graph $G(V,E)$ and its vertex $v \in V$ let $N[v]$ denote the closed neighborhood of a vertex v. That is, $N[v] := \{u \in V; \{u,v\} \in E\} \cup \{v\}$. A *dominating set* of a graph $G(V,E)$ is a set $U \subseteq V$ such that for each v there is $u \in U$ such that $v \in N[u]$.

Number of dominating sets of size k	
Input:	A graph $G(V,E)$ and a parameter k.
Question:	How many dominating sets of size k are in G?

Proof of Theorem 4.1. As we said, we show an FPT counting reduction from the problem of counting dominating sets of size k to VEST.

Let $G = (V,E)$ be the input graph and let $n = |V|$. The corresponding instance of VEST will consist of n matrices $\{T_u; u \in V\}$ of size $4n \times 4n$, one for each vertex, and matrix S of the same size. Whence, the initial vector $\mathbf{v}$ must be of size $4n$. For each vertex $u \in V$, we introduce four new coordinates $u_1, \dots, u_4$ and set $\mathbf{v}_{u_1} = 1, \mathbf{v}_{u_2} = \mathbf{v}_{u_3} = 0$ and $\mathbf{v}_{u_4} = 1$.

$$\begin{pmatrix} 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 \end{pmatrix}$$

Figure 4.1: The submatrix of T_u consisting of rows and columns $u_1, \ldots, u_4$. The rest of the non-diagonal entries of T_u are zero. The diagonal entries $T_u^{w_1,w_1}$ for $w \in N[u]$ are equal to zero, the rest of the diagonal entries are one.

We define the matrices $\{T_u; u \in V\}$ and S by describing their behavior. Let $\mathbf{x}$ be a vector which is going to be multiplied with a matrix T_u (that is, some intermediate vector obtained from $\mathbf{v}$ after potential multiplications). The matrix T_u sets $\mathbf{x}_{w_1}$ to zero for each $w \in N[u]$, which corresponds to domination of vertices in $N[u]$ by the vertex u, and also sets $\mathbf{x}_{u_2}$ to $\mathbf{x}_{u_3}$ and $\mathbf{x}_{u_3}$ to $\mathbf{x}_{u_4}$. The rest of the entries of $\mathbf{x}$ including $\mathbf{x}_{u_4}$ are kept, see Figure 4.1.

The matrix S then nullifies coordinates u_3, u_4 and keeps the coordinates u_1 and u_2 for each $u \in V$. In other words, S is diagonal such that $S^{u_1,u_1} = S^{u_2,u_2} = 1$ and $S^{u_3,u_3} = S^{u_4,u_4} = 0$.

The parameter remains equal to k.

For correctness, let $u^1, \ldots, u^k$ be any vertices from V, and let $\mathbf{r}$ be the vector obtained from $\mathbf{v}$ after multiplying by the matrices $T_{u^1}, \ldots, T_{u^k}$ (observe that the order of multiplication does not matter since all T_u, $u \in V$, pairwise commute). By construction, for every vertex $u \in V$, the entry $\mathbf{r}_{u_1} = 0$ if and only u is dominated by some u^i, $i \in \{1, \ldots k\}$, and $\mathbf{r}_{u_2} = 0$ if and only if T_u appears among $T_{u^1}, \ldots, T_{u^k}$ at most once. Indeed, if T_u is selected once then $\mathbf{r}_{u_2} = \mathbf{v}_{u_3} = 0$ while if it is selected more than once then $\mathbf{r}_{u_2} = \mathbf{v}_{u_4} = 1$. If T_u is not among $T_{u^1}, \ldots, T_{u^k}$ then $\mathbf{r}_{u_2} = \mathbf{v}_{u_2} = 0$.

Therefore, $\mathbf{r} = T_{u^1} \ldots T_{u^k}\mathbf{v}$ is a zero vector if and only if $u^1, \ldots, u^k$ are pairwise distinct and form the dominating set in G. This provides a one-to-one correspondence between subsets of matrices yielding the solution of VEST and dominating sets of size k in G. It remains to note that every such subset of matrices gives rise to $k!$ sequences that have to be counted in M_k. Hence, $M_k = k!D_k$ where D_k is the number of dominating sets of size k in G.

The reduction is clearly FPT since the construction does not use the parameter k and it is polynomial in the size of the input. □

Remark 4.6. Note that the decision version of the problem of Dominating Sets of Size k is W[2]-hard. For showing W[2]-hardness of the decision version of VEST we need not deal with the repetition of matrices. In particular, we do not need the special coordinates u_2, u_3, u_4 and therefore, the corresponding instance of VEST can consist only of 0, 1 diagonal matrices of size $n \times n$.

4.3 Modifications of VEST and their relationships

In this section, we prove hardness for the variants of the decision version of VEST we have discussed in the introduction. First of all, we recall a well-known W[1]-hard k-Sum problem. See also [ALW14].

k-Sum	
Input:	A set A of integers and parameter k.
Question:	Is it possible to choose k distinct integers from A such that their sum is equal to 0?

A similar problem appears if we allow to choose not exactly, but at most k integers.

At-Most-k-Sum	
Input:	A set A of integers and a parameter k.
Question:	Is it possible to choose *at most* k distinct integers from A such that their sum is equal to 0?

We note that in the versions of k-Sum studied in the literature, the goal is to pick *distinct* elements of the input set in order to achieve 0 or eventually another number. However, the motivation for VEST, to the contrary, does not suggest that the matrices chosen for the product have to be distinct. Thus, in order to model VEST by k-Sum, it is more natural to also allow repetition of numbers. For our particular proofs, we will use the following version with target number 1.

At-Most-k-Sum with Repetitions and Target 1	
Input:	A set A of integers and parameter k.
Question:	Is it possible to choose *at most* k integers from A (possibly with repetition) such that their sum is equal to 1?

We are not aware of any previous studies on parameterized complexity of At-Most-k-Sum with Repetitions and Target 1, nor does it seem that there exists a simple parameterized reduction from the original variant of the problem to the one with repetitions. Therefore, in the next theorem we prove W[1]-hardness of this problem directly. The starting point of our reduction is the problem of k-Exact Cover, which is known to be W[1]-hard (see [DF95]).

k-Exact Cover	
Input:	A universe U, a collection $\mathcal{C}$ of subsets of U and a parameter k.
Question:	Can U be partitioned into k sets from $\mathcal{C}$?

Theorem 4.7. *At-Most-k-Sum with Repetitions and Target 1 is* W[1]*-hard when parameterized by k.*

Proof. Consider an instance $(U, \mathcal{C}, k)$ of Unique Hitting Set. Intuitively, we would like to model the sets in $\mathcal{C}$ as their characteristic vectors over $|U|$ dimensions, where each dimension corresponds to an element from U, and the vector representing a set $C \in \mathcal{C}$ is set to one exactly in those dimensions which correspond to the elements contained in the set which is represented by the vector. To model this in an instance of At-Most-k-Sum with Repetitions and Target 1, we will represent said characteristic vectors as numbers in base $(k+2)$.

Formally, let $m = |U|$, $U = \{u_1, \ldots, u_m\}$, and $x = k+2$. For each $C \in \mathcal{C}$, we add an element $a_C := -\left(x^{m+1} + \sum_{j; u_j \in C} x^j\right)$ to the set A of numbers. Then we also add to A the number $y := kx^{m+1} + \sum_{j=0}^{m} x^j$ and we set the new parameter to $k+1$. Note that the numbers in A are bounded by x^{m+2}, thus can be represented by $\mathcal{O}(m \log k)$ bits, and $|A| = |\mathcal{C}| + 1$, meaning that the reduction can be done in polynomial time. It remains to verify that the produced instance of At-Most-k-Sum with Repetitions and Target 1 is equivalent to the original instance of k-Exact Cover.

First of all, let $C_1, \dots, C_k \in U$ be a solution to k-Exact Cover. Then $\{y, a_{C_1}, \dots, a_{C_k}\} \subset A$ is a solution to the instance $(A, k+1)$ of At-Most-k-Sum with Repetitions and Target 1. Indeed, by construction and since each element of U is covered exactly once, we have $a_{C_1} + \cdots + a_{C_k} = -\left(kx^{m+1} + \sum_{j=1}^{m} x^j\right) = -y + x^0 = -y + 1$. Therefore, $y + a_{C_1} + \cdots + a_{C_k} = 1$.

In the other direction, consider a solution $q_1, \dots, q_t \in A$ to At-Most-k-Sum with Repetitions and Target 1 where $t \leq k+1$. First of all, we observe that y must be chosen precisely once. The sum $\sum_{j=1}^{t} q_t = 1 = x^0$ and y is the only number with a coefficient $(= 1)$ of x^0. Therefore, y can be chosen $(\ell x + 1)$ times where $\ell \in \mathbb{N}_0$. However $t < x = k + 2$. Whence, $\ell = 0$. In other words, y is chosen precisely once and without loss of generality, we suppose that $q_1 = y$

Next, we show that $t = k + 1$. The number y which is chosen precisely once has k as the coefficient of x^{m+1} which has to be nullified. The only option how to do that is to choose k numbers other than y. (Such numbers have -1 as the coefficient of x^{m+1}.)

Finally, from the equality $\sum_{j=2}^{k+1} q_j = -y + 1 = -kx^{m+1} - \sum_{j=1}^{m} x^j$ we conclude that no $-x^i$ for $i \leq m$ is contained in more than one q_j as a summand since $k < k + 2 = x$. By the same argument we observe that each $-x^i$ is contained in some q_j as a summand. Indeed, addition of at most k terms $-x^i$ cannot affect coefficient of x^{i+1}. Therefore, each $-x^i$ for $i \leq m$ is contained in precisely one q_j and thus, $\{C; a_C \in \{q_2, \dots, q_{k+1}\}\}$ is a desired k-exact cover.

□

If we assume multiplication instead of addition the following problem arises.

k-Product with Repetitions	
Input:	A set A of rational numbers and a parameter k.
Question:	Is it possible to choose k numbers from A (possibly with repetitions) such that their product is equal to 1?

W[1]-hardness for this problem might be a folklore result but we present a complete proof using a reduction from k-Exact Cover. The idea will be in principle the same as in the previous proof. We start with a lemma which is essentially an easy consequence of the prime number theorem.

Lemma 4.8. *Let p_n denote the n-th prime. Then $p_n \leq n^2$ for $n \geq 2$.*

Proof. Let $\pi(x)$ denote the number of primes less than or equal to x. The lemma follows, e.g., from the following claims:

- $p_n < n(\ln n + \ln\ln n)$ for $6 \leq n \leq e^{95}$ (see [Bar41, Theorem 28]),
- $\frac{x}{\ln x + 2} \leq \pi(x)$ for $x \geq 55$ (see [Bar41, Theorem 29.A]),
- $p_2 = 3$, $p_3 = 5$, $p_4 = 7$, $p_5 = 11$.

□

Theorem 4.9. *k-*Product with Repetitions *is* W[1]*-hard when parameterized by k.*

Proof. We show a parameterized reduction from k-Exact Cover. For each element $u \in U$ we associate one prime p_u, then for each $C \in \mathcal{C}$ we set $a_C := p \prod_{u \in C} p_u$ where p is a prime which is not used for any element from U and $y := \frac{1}{p^k \prod_{u \in U} p_u}$.

The integers a_C, for each $C \in \mathcal{C}$, and y then form the input for $(k+1)$-Product with Repetitions

If $C_1, C_2, \ldots, C_k \in \mathcal{C}$ is a solution of k-Exact Cover then $y \prod_{i=1}^{k} a_{C_i} = 1$.

Conversely, let $q_1, q_2, \ldots, q_{k+1}$ be a solution of the constructed instance of $(k+1)$-Product with Repetitions. First of all, note that y must be chosen precisely once. Indeed, all numbers except for y are greater than 1 and thus, y must be chosen at least once. If it were chosen more than once it would not be possible to cancel a power of p^k in the denominator since the numerator would contain at most p^{k-1}. Therefore, the product of $q_1, q_2, \ldots, q_{k+1}$ is of the form $y a_{C_{j_k}} a_{C_{j_{k-1}}} \cdots a_{C_{j_1}} = 1$ which means that each prime representing an element of U in the denominator is canceled. In other words, each element of U is covered. Note also that since y is chosen precisely once, there cannot be any repetition within $a_{C_{j_k}} a_{C_{j_{k-1}}} \ldots a_{C_{j_1}}$.

The reduction is parameterized since we only need the parameter k for k multiplications of $\frac{1}{p}$ and the first $n+1$ primes, where $n = |U|$, can be generated in time $\leq \mathcal{O}(n^3)$ using, e.g., the Sieve of Eratosthenes for $(n+1)^2$. This follows from Lemma 4.8. □

Let us now call the variant of VEST without S and $\mathbf{v}$ Matrix k-Product with Repetitions. As we have mentioned in the introduction we consider two cases regarding the target matrix. Namely, the Identity matrix and the Zero matrix:

Matrix k-Product with Repetitions resulting to Zero Matrix	
Input:	A list of $d \times d$ rational matrices and a parameter k.
Question:	Is it possible to choose k matrices from the list (possibly with repetitions) such that their product is the $d \times d$ zero matrix?

Matrix k-Product with Repetitions resulting to Identity Matrix	
Input:	A list of $d \times d$ rational matrices and a parameter k.
Question:	Is it possible to choose k matrices from the list (possibly with repetitions) such that their product is the $d \times d$ identity matrix?

Note that Matrix k-Product with Repetitions resulting to Identity Matrix for 1×1 matrices is exactly k-Product with Repetitions. Therefore W[1]-hardness for Matrix k-Product with Repetitions resulting to Identity Matrix for all matrix sizes follows from Theorem 4.9.

Regarding Matrix k-Product with Repetitions resulting to Zero Matrix, we can easily see that it is solvable in linear time for 1×1 matrices. Indeed, it is sufficient to check whether $T_i = 0$ for some i. However, already for 2×2 matrices the problem becomes hard.

Theorem 4.10. *Matrix k-Product with Repetitions resulting to Zero Matrix is* W[1]*-hard for parameter k even for 2×2 integer matrices.*

Proof. We reduce from At-Most-k-Sum with Repetitions and Target 1. For every integer x let us define

$$U_x := \begin{pmatrix} 1 & x \\ 0 & 1 \end{pmatrix}.$$

It is easy to see that $U_x U_y = U_{x+y}$. Let $\mathcal{I}$ be an instance of At-Most-k-Sum with Repetitions and Target 1 with the set of integers A and parameter k. We create an equivalent instance $\mathcal{I}'$ of Matrix $(k+2)$-product with Repetitions with the set of matrices $\{U_a; a \in A\} \cup \{X\}$, where

$$X = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix}.$$

For correctness, assume that $\mathcal{I}$ is a YES-instance and $a_1, \ldots, a_\ell \in A$ are such that $\ell \le k$ and $\sum_{i=1}^{\ell} a_i = 1$. Consider the following product of $\ell + 2$ matrices:

$$X \cdot \prod_{i=1}^{\ell} U_{a_i} \cdot X = X \cdot U_{\sum_{i=1}^{\ell} a_i} \cdot X = X U_1 X = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}.$$

For the other direction, assume that $\mathcal{I}'$ is a YES-instance. Let ℓ, $1 \le \ell \le k+2$, be the minimal integer such that there are matrices $T_1, T_2, \ldots, T_\ell$ from $\{U_a; a \in A\} \cup \{X\}$ with $T_\ell T_{\ell-1} \cdots T_1 = \mathbf{0} \in \mathbb{Q}^{2\times 2}$. Since the matrix X is idempotent (i.e. $X^2 = X$), it does not appear two times in a row, otherwise we could reduce the length of the product. Notice that X should appear at least once, since the determinants of all U_a are non-zero. Assume that there is precisely one occurrence of X, then the product has form:

$$\begin{aligned} U_r X U_s &= \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} -r & r \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} -r & -rs + r \\ -1 & 1 - s \end{pmatrix} \neq \begin{pmatrix} 0 & 0 \\ 0 & 0 \end{pmatrix}. \end{aligned}$$

Hence, X appears at least twice. Let us fix any two consequent occurrences and consider the partial product between them:

$$X U_r X = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & r \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} = \begin{pmatrix} 0 & 0 \\ r-1 & 1-r \end{pmatrix} = (1-r) \cdot X.$$

If $r \neq 1$, we would get a shorter product resulting in zero, which contradicts to minimality of ℓ. Hence $r = 1$, so the product of U_a that appear between two occurrences of X is equal to U_1. Since there are at most k of such U_a and the sum of corresponding indices a is equal to 1, we obtain a solution to $\mathcal{I}$. $\square$

We can use similar approach to establish hardness of the VEST problem without S (or alternatively when S is the identity matrix). Recall that here the task is to obtain not necessarily a zero matrix but any matrix which contains the given vector $\mathbf{v}$ in its kernel.

Theorem 4.11. *VEST is* $\mathsf{W}[1]$*-hard for parameter k even for 2×2 integer matrices and when S is the identity matrix.*

Proof. As in the proof of Theorem 4.10, we proceed by reduction from At-Most-k-Sum with Repetitions and Target 1. Let $\mathcal{I}$ be an arbitrary instance of the problem with the set of integers A, and parameter k. We create an equivalent instance $\mathcal{I}'$ of VEST with parameter $k+1$, vector $v = (0,1)^T$ and the set of matrices $\{U_a; a \in A\} \cup \{X\}$, where U_a and X are defined same as in the proof of Theorem 4.10. We set S equal to the identity matrix.

For correctness, assume that $\mathcal{I}$ is a YES-instance and $a_1, \dots, a_\ell \in A$ are such that $\ell \leq k$ and $\sum_{i=1}^{\ell} a_i = 1$. Consider the application of following $\ell + 1$ matrices to v:

$$X \cdot \prod_{i=1}^{\ell} U_{a_i} \cdot v = XU_1 v = \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}.$$

For the other direction, assume that $\mathcal{I}'$ is a YES-instance. Let ℓ, $1 \leq \ell \leq k+1$, be the minimal integer such that $T_\ell, T_{\ell-1} \cdots T_1 v = (0,0)^T$ for some $T_1, T_2, \dots, T_\ell$ from $\{U_a; a \in A\} \cup \{X\}$. Since the determinants of all U_a are non-zero, $T_i = X$ for some $i \in \{1, 2, \dots, \ell\}$. Observe that $Xv = v$, so by minimality of ℓ we have that $T_1 \neq X$. Let i be the minimal index such that $T_i = X$, $2 \leq i \leq \ell$. Then $T_{i-1} \cdots T_1 = U_s$ for some integer s. Let us apply first i matrices to v:

$$\begin{aligned} T_i T_{i-1} \cdots T_1 v = XU_s \cdot v &= \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} 1 & s \\ 0 & 1 \end{pmatrix} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \\ &= \begin{pmatrix} 0 & 0 \\ -1 & 1 \end{pmatrix} \begin{pmatrix} s \\ 1 \end{pmatrix} = \begin{pmatrix} 0 \\ 1-s \end{pmatrix}. \end{aligned}$$

If $s \neq 1$, we get a multiple of v, which is in contradiction to minimality of ℓ. So $T_{i-1}, \dots, T_1 = U_1$, which is a product of at most k matrices of the form U_a with $a \in A$. The sum of corresponding indices a is then equal to 1, resulting in a solution to $\mathcal{I}$. □

At the end of this section, we show that VEST is equivalent to VEST without S (in other words, when $S = I_d$) and to Matrix k-Product with Repetitions resulting to Zero Matrix.

Theorem 4.12. *There is a parameterized reduction from VEST to the special case of VEST where S is the identity matrix, and the other way around.*

Proof. One direction is trivial since the case when $S = I_d$ is just a special case of VEST.

Regarding the other, let $\left(S \in \mathbb{Q}^{h \times d}, T_1, T_2, \dots, T_m \in \mathbb{Q}^{d \times d}, \mathbf{v} \in \mathbb{Q}^d, k\right)$ be an instance of VEST. First, we observe that without loss of generality we can suppose that S is a square matrix (in other words, $h = d$). Indeed, if $h < d$ then we

$$
\left(\begin{array}{c|ccc} & 0 & \dots & 0 \\ S & \vdots & \ddots & \vdots \\ & 0 & \dots & 0 \end{array}\right) \in \mathbb{Q}^{h\times h},
\left(\begin{array}{ccc|ccc} & & & 0 & \dots & 0 \\ & T_i & & \vdots & \ddots & \vdots \\ & & & 0 & \dots & 0 \\ \hline 0 & \dots & 0 & 0 & \dots & 0 \\ \vdots & \ddots & \vdots & \vdots & \ddots & \vdots \\ 0 & \dots & 0 & 0 & \dots & 0 \end{array}\right) \in \mathbb{Q}^{h\times h},
\left(\begin{array}{c} \mathbf{v} \\ \hline 0 \\ \vdots \\ 0 \end{array}\right) \in \mathbb{Q}^{h}.
$$

Figure 4.2: A figure showing how to make all matrices square in the proof of Theorem 4.12 when $h > d$.

$$
\mathbf{v}' = \left(\begin{array}{c} \mathbf{v} \\ \hline k \\ 1 \end{array}\right), S' = \left(\begin{array}{ccc|cc} & & & 0 & 0 \\ & S & & \vdots & \vdots \\ & & & 0 & 0 \\ \hline 0 & \dots & 0 & 10 & 0 \\ 0 & \dots & 0 & 0 & 0 \end{array}\right), T_i' = \left(\begin{array}{ccc|cc} & & & 0 & 0 \\ & T_i & & \vdots & \vdots \\ & & & 0 & 0 \\ \hline 0 & \dots & 0 & 1 & -1 \\ 0 & \dots & 0 & 0 & 1 \end{array}\right).
$$

Figure 4.3: A construction forcing the matrix S' to be selected last in the proof of Theorem 4.12.

just add $d - h$ zero lines to S. If $h > d$ we add $h - d$ zero columns to S, $h - d$ zero entries to $\mathbf{v}$ and $h - d$ zero lines as well as $h - d$ zero columns to each T_i. See Figure 4.2.

Now, we add 2 dimensions: To the vector $\mathbf{v}$ we add k on the $(d+1)$-st position and 1 on the $(d+2)$-nd position. To each matrix matrix T_i we add a 2×2 submatrix which subtracts the $(d+2)$-nd component of a vector from the $(d+1)$-st. To the matrix S we add a submatrix which nullifies the $(d+2)$-nd component and multiplies the $(d+1)$-th component by 10. Let $S', T_1', T_2' \dots, T_m'$ denote the resulting $(d+2) \times (d+2)$ matrices and $\mathbf{v}'$ denote the resulting $(d+2)$-dimensional vector. See Figure 4.3. The new parameter is set to $k+1$.

If there is a solution of the original problem, that is, there are k matrices $T_{i_1}, \dots, T_{i_k}$ such that $ST_{i_k}T_{i_{k-1}} \cdots T_{i_1}\mathbf{v} = \mathbf{0}$, then $S'T_{i_k}'T_{i_{k-1}}' \cdots T_{i_1}'\mathbf{v}' = \mathbf{0}$, since 1 is k times subtracted from the $(d+1)$-st component of $\mathbf{v}'$ and the $(d+2)$-nd component is then nullified by S'.

Conversely, if there are $k+1$ matrices $Y_1, Y_2, \dots, Y_{k+1}$, where each Y_i is either S' or T_j' for some j, such that $\mathbf{r} = Y_{k+1}Y_k \cdots Y_1\mathbf{v}' = \mathbf{0}$ then Y_{k+1} must be equal to S' and the rest of the matrices are of type T_j', otherwise $\mathbf{r}_{d+1} \neq 0$ or $\mathbf{r}_{d+2} \neq 0$. Indeed, at first k matrices of type T_j' must be selected to nullify the $(d+1)$-st component: If $Y_i = S'$ for some $i \leq k$, this would increase the non-zero $(d+1)$-st component, so there would be no way to nullify it by remaining matrices $Y_{i+1}, \dots, Y_{k+1}$. At the same time, S' should be necessarily selected once to nullify the $(d+2)$-nd component, so $Y_{k+1} = S'$. Therefore, by restricting the matrices $Y_1, \dots, Y_k$ to the first d coordinates we obtain a solution to VEST with matrix S. $\square$

Theorem 4.13. *VEST and* MATRIX *k*-PRODUCT WITH REPETITIONS RESULTING TO ZERO MATRIX *are equivalent under parameterized reduction.*

Proof.

1. "Parameterized reduction from MATRIX k-PRODUCT WITH REPETITIONS RESULTING TO ZERO MATRIX to VEST"
   For each matrix $T_i \in \mathbb{Q}^{d \times d}$ we introduce a block matrix $T'_i \in \mathbb{Q}^{d^2 \times d^2}$ whose each block is T_i. We set $\mathbf{v} = (e_1, e_2, \ldots, e_d)^T \in \mathbb{Q}^{d^2}$ where each e_i is the d-dimensional unit vector with 1 on its i-th coordinate and S to the d^2-dimensional identity matrix. Therefore, $T_{i_k} T_{i_{k-1}} \cdots T_{i_1} = R$ if and only if $ST'_{i_k} \cdots T'_{i_1}\mathbf{v} = (R_{*,1}, R_{*,2}, \ldots, R_{*,d})^T$ where $R_{*,j}$ is the j-th column of the matrix R.

2. "Parameterized reduction from VEST to MATRIX k-PRODUCT WITH REPETITIONS RESULTING TO ZERO MATRIX"
   We first reduce VEST to the version of VEST without S as we did in the proof of Theorem 4.12. Thus, we assume that our input consists of square matrices $S', T'_1, T'_2, \ldots, T'_m$ and $\mathbf{v}$ where S' represents the original special matrix S and the parameter is $k+1$. Let us recall that S' has to be selected precisely once as the leftmost matrix otherwise the resulting vector cannot be zero by the construction from the proof of Theorem 4.12.

   Now, we make an instance of MATRIX $(k+3)$-PRODUCT WITH REPETITIONS. Let $T_{\mathbf{v}}$ be a matrix containing the vector $\mathbf{v}$ in the first column and zero otherwise. The idea is to use the matrix $T_{\mathbf{v}}$ instead of the vector $\mathbf{v}$ and force such matrix to be selected as the rightmost after S' and k matrices of type T'_i by adding some blocks. We use the construction from the proof of Theorem 4.10. Namely, we use matrices X and U_{-2} and U_{2k+1} as submatrices. By the same argument as in the proof of Theorem 4.10 the only way how to make the zero matrix by multiplying $k+3$ matrices from $\{X, U_{-2}, U_{2k+1}\}$ is to choose X twice, as the leftmost and the rightmost matrix, k-times U_{-2} and once U_{2k+1} as intermediate matrices. Therefore, we can add X to $T_{\mathbf{v}}$ and to the identity matrix as block submatrices, U_{2k+1} to S' (since S' must be selected precisely once) and U_{-2} to T'_i. It remains to force the order of $T_{\mathbf{v}}$ and the identity matrix enriched by X. For this, we add submatrices A, B such that $AB = 0$ while $BA \neq 0, AA \neq 0, BB \neq 0$. We add A to the identity matrix enriched by X, B to $T_{\mathbf{v}}$ enriched by X and identity matrices to the rest. See Figure 4.4. The following settings for A and B, respectively, work.

$$A = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 0 \end{pmatrix}, B = \begin{pmatrix} 0 & 1 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

   To sum up,

   (a) matrices of type T''_i must be chosen precisely k-times by the proof of Theorem 4.12.

   (b) The matrix S'' must be chosen precisely once and it must be on the left of the product of k matrices of type T''_i by the proof of Theorem 4.12.

$$
T_i'' = \left(\begin{array}{ccc|ccc|cc}
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\multicolumn{3}{c|}{T_i'} & \vdots & \vdots & \vdots & \vdots & \vdots \\
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & I_3 & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\
\hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{U_{-2}}
\end{array}\right),
S'' = \left(\begin{array}{ccc|ccc|cc}
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\multicolumn{3}{c|}{S'} & \vdots & \vdots & \vdots & \vdots & \vdots \\
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & I_3 & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\
\hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{U_{2k+1}}
\end{array}\right),
$$

$$
T_{\mathbf{v}}' = \left(\begin{array}{ccc|ccc|cc}
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\multicolumn{3}{c|}{T_{\mathbf{v}}} & \vdots & \vdots & \vdots & \vdots & \vdots \\
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & B & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\
\hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{X}
\end{array}\right),
H = \left(\begin{array}{ccc|ccc|cc}
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\multicolumn{3}{c|}{I_d} & \vdots & \vdots & \vdots & \vdots & \vdots \\
\multicolumn{3}{c|}{} & 0 & 0 & 0 & 0 & 0 \\
\hline
0 & \dots & 0 & & & & 0 & 0 \\
0 & \dots & 0 & & A & & 0 & 0 \\
0 & \dots & 0 & & & & 0 & 0 \\
\hline
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{} \\
0 & \dots & 0 & 0 & 0 & 0 & \multicolumn{2}{c}{X}
\end{array}\right).
$$

Figure 4.4: A figure showing what an instance of Matrix k-Product with Repetitions resulting to Zero Matrix looks like after the reduction from VEST in the proof of Theorem 4.13.

(c) The leftmost matrix has to be H by the proof of Theorem 4.10 and by the fact that $AB = 0$ while $BA \neq 0, AA \neq 0, BB \neq 0$.

(d) The rightmost matrix has to be $T_{\mathbf{v}}'$ again by the proof of Theorem 4.10 and by the fact that $AB = 0$ while $BA \neq 0, AA \neq 0, BB \neq 0$.

□

4.4 Fixed-Parameter Tractability of VEST over Finite Fields

While most of the hardness results for VEST and its variations in the previous section use constant-sized matrices, the entries of this matrices can be arbitrarily large. Here, we study the variation of the problem when all the matrices have entries from some finite field. Notice that restricting the field size by itself does not make the problem tractable: Recall the reduction from dominating set from Section 4.2 which also works over $\mathbb{Z}_2$. However, along with a bound on the matrix sizes this makes the problem tractable.

Lemma 4.14. *Computing M_k for a given instance of VEST over finite field $\mathbb{F}$ is* FPT *when parameterized by the size of $\mathbb{F}$ and the size of matrices.*

Proof. Let $\mathcal{M}_{\mathbb{F}}^d$ be the set of all $d \times d$ matrices with entries from $\mathbb{F}$, then $|\mathcal{M}_{\mathbb{F}}^d| = |\mathbb{F}|^{d^2}$. For every $X \in \mathcal{M}_{\mathbb{F}}^d$ and every integer $i \in \{1, \ldots, k\}$ we will compute a value $a_X^i \in \mathbb{N}_0$ equal to the number of sequences of i matrices from the input such that their product is equal to X. In particular, this allows to obtain

$$M_k = \sum_{X \in \mathcal{M}_{\mathbb{F}}^d; SX\mathbf{v}=\mathbf{0}} a_X^k.$$

For $i = 1$, the computation can be done simply by traversing the input matrices. Assume that a_X^i have been computed for all the matrices X and all $i \in \{1, \ldots, j\}$. We initiate by setting $a_X^{j+1} = 0$ for every $X \in \mathcal{M}_{\mathbb{F}}^d$. Then, for every pair (X, q), where $X \in \mathcal{M}_{\mathbb{F}}^d$ and $q \in \{1, \ldots, m\}$, we increment $a_{XT_q}^{j+1}$ by a_X^j. In the end, we will then have a correctly computed value

$$a_Y^{j+1} = \sum_{q=1}^{m} \sum_{X; XT_q=Y} a_X^j.$$

□

Our next step is to consider the matrices of unbounded size, but with at most p first rows containing non-zero entries. In particular, if $\mathbb{F} = \mathbb{Z}_2$, we can associate to every such matrix T a graph with the vertex set $\{1, \ldots, d\}$ such that there exists an edge between the vertices i and j, $i \leq j$, if and only if $T^{i,j} = 1$. Conversely, a graph with the vertex set $\{1, ldots, d\}$ can be represented by such a matrix if and only if the vertices in $\{1, \ldots, p\}$ form it's vertex cover.

Observe that every matrix T with at most p first non-zero rows has the following form:

$$T = \left(\begin{array}{c|ccc} A & & B & \\ \hline 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{array}\right),$$

where A is $p \times p$ matrix and B is $p \times (d-p)$ matrix. Further, we will denote matrices of this form by $A|B$. Consider the product of two such matrices $T_1 = A_1|B_1$ and $T_2 = A_2|B_2$:

$$\left(\begin{array}{c|ccc} A_1 & & B_1 & \\ \hline 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{array}\right) \left(\begin{array}{c|ccc} A_2 & & B_2 & \\ \hline 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{array}\right) = \left(\begin{array}{c|ccc} A_1A_2 & & A_1B_2 & \\ \hline 0 & 0 & \ldots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \ldots & 0 \end{array}\right) = (A_1A_2)|(A_1B_2).$$

Corollary 4.15. *$\prod_{i=1}^{k}(A_i|B_i) = (\prod_{i=1}^{k} A_i)|(\prod_{i=1}^{k-1} A_i \cdot B_k) = (XA_k)|(XB_k)$, where $X = \prod_{i=1}^{k-1} A_i$. In particular, the product does not depend on B_i for $i < k$.*

Theorem 4.5. *Given an instance of VEST and $k \in \mathbb{N}$, computing M_k is* FPT *when parameterized by $|\mathbb{F}|$ and p, if all but the first p rows of the input matrices are zeros.*

Proof. We slightly modify the definition of $a_X^i \in \mathbb{N}_0$ from the proof of Theorem 4.14. Now, for every $i \in [k]$ and every matrix $X \in \mathcal{M}_{\mathbb{F}}^p$, let a_X^i be the number of sequences of i matrices $T_j = A_j|B_j$ from the input such that corresponding product of A_j is equal to X.

The values of a_X^i for every $i \in \{1, \ldots, k-1\}$ can be computed same as in the proof of Theorem 4.14. Given this information, we can count the sequences of length k that nullify $\mathbf{v}$. Indeed, by Corollary 4.15, the number of such sequences with the last matrix $T_j = A_j|B_j$ is precisely

$$b_j = \sum_{X \in \mathcal{M}_{\mathbb{F}}^p; S\cdot(XA_j)|(XB_j)\cdot\mathbf{v}=\mathbf{0}} a_X^{k-1},$$

and M_k is then equal to $\sum_{j=1}^m b_j$. $\square$

We remark that the algorithm for computing a_X^i from the last proof can be exploited to determine minimal k such that $M_k > 0$, or to report that there is no such k. For this, let us run the algorithm with $k = 1$, then with $k = 2$ and so on. If after some iteration $k = j + 1$ we obtain that $M_i = 0$ for all $i \in \{1, \ldots, j\}$ and there is no $X \in \mathcal{M}_{\mathbb{F}}^p$ such that $a_X^1 = \cdots = a_X^j = 0$ and $a_X^{j+1} \neq 0$, we may conclude that $M_k = 0$ for all $k \in N$, since every product of length more than j can be obtained as a product of length at most j, and none of the latter nulify $\mathbf{v}$. Otherwise, there exists at least one $X \in \mathcal{M}_{\mathbb{F}}^p$ such that $a_X^1 = \cdots = a_X^j = 0$ and $a_X^{j+1} \neq 0$. Note that every $X \in \mathcal{M}_{\mathbb{F}}^p$ can play this role only for one value of k. Therefore, it always suffices to make $|\mathcal{M}_{\mathbb{F}}^p|$ iterations of the algorithm.

4.5 Undecidability of VEST with no parameter

In this section, we show that determining whether there exists $k \in \mathbb{N}$ such that $M_k > 0$ for an instance of VEST is an undecidable problem. The reduction is from Post's Correspondence Problem which is known to be undecidable. See [Pos47].

(Binary) Post's Correspondence Problem	
Input:	m pairs $(v_1, w_2), (v_2, w_2), \ldots, (v_m, w_m)$ of words over alphabet $\{0, 1\}$.
Question:	Is possible to choose k pairs $(v_{i_1}, w_{i_k}), (v_{i_2}, w_{i_2}) \ldots, (v_{i_k}, w_{i_k})$, for some $k \in \mathbb{N}$, such that $v_{i_1}v_{i_2}\cdots v_{i_k} = w_{i_1}w_{i_2}\cdots w_{i_k}$? (Where $v_{i_1}v_{i_2}\cdots v_{i_k}$ and $w_{i_1}w_{i_2}\cdots w_{i_k}$ denote concatenation of the words $v_{i_1}, v_{i_2}, \cdots, v_{i_k}$ and $w_{i_1}, w_{i_2}, \cdots, w_{i_k}$, respectively.)

For a word $v \in \{0, 1\}^*$ let $|v|$ be its length and let $(v)_2$ be the integer value of v interpreting it as a binary number. Let us define the following matrix for a binary word v.

$$T_v = \begin{pmatrix} 2^{|v|} - (v)_2 & (v)_2 \\ 2^{|v|} - (v)_2 - 1 & (v)_2 + 1 \end{pmatrix}$$

The following holds.

Lemma 4.16. *Let v, w be binary words. Then, $T_vT_w = T_{wv}$ where wv is the concatenation of w and v.*

Note that the construction of T_v is a based on [Cla71][Satz 28, p. 157] which we are aware of thanks to Günter Rote.

Proof of Lemma 4.16. First of all, note that $2^{|v|}(w)_2 + (v)_2 = (wv)_2$. Using this observation, we compute all the entries of the matrix

$$T_v T_w = \begin{pmatrix} 2^{|v|} - (v)_2 & (v)_2 \\ 2^{|v|} - (v)_2 - 1 & (v)_2 + 1 \end{pmatrix} \begin{pmatrix} 2^{|w|} - (w)_2 & (w)_2 \\ 2^{|w|} - (w)_2 - 1 & (w)_2 + 1 \end{pmatrix}.$$

Thus,

$$\begin{aligned}
(T_v T_w)^{1,1} =& \left(2^{|v|} - (v)_2\right)\left(2^{|w|} - (w)_2\right) + (v)_2\left(2^{|w|} - (v)_2 - 1\right) \\
=& 2^{|wv|} - 2^{|v|}(w)_2 - 2^{|w|}(v)_2 + (v)_2(w)_2 + 2^{|w|}(v)_2 - (v)_2(w)_2 - (v)_2 \\
=& 2^{|wv|} - 2^{|v|}(w)_2 - (v)_2 \\
=& 2^{|wv|} - (wv)_2, \\
(T_v T_w)^{1,2} =& \left(2^{|v|} - (v)_2\right)(w)_2 + (v)_2\left((w)_2 + 1\right) \\
=& 2^{|v|}(w)_2 - (v)_2(w)_2 + (v)_2(w)_2 + (v)_2 \\
=& 2^{|v|}(w)_2 + (v)_2 \\
=& (wv)_2, \\
(T_v T_w)^{2,1} =& \left(2^{|v|} - (v)_2 - 1\right)\left(2^{|w|} - (w)_2\right) + ((v)_2 + 1)\left(2^{|w|} - (w)_2 - 1\right) \\
=& 2^{|wv|} - 2^{|v|}(w)_2 - 2^{|w|}(v)_2 + (v)_2(w)_2 - 2^{|w|} + (w)_2 \\
& + 2^{|w|}(v)_2 - (v)_2(w)_2 - (v)_2 + 2^{|w|} - (w)_2 - 1 \\
=& 2^{|wv|} - 2^{|v|}(w)_2 - (v)_2 - 1 \\
=& 2^{|wv|} - (wv)_2 - 1, \\
(T_v T_w)^{2,2} =& \left(2^{|v|} - (v)_2 - 1\right)(w)_2 + ((v)_2 + 1)((w)_2 + 1) \\
=& 2^{|v|}(w)_2 - (v)_2(w)_2 - (w)_2 + (v)_2(w)_2 + (v)_2 + (w)_2 + 1 \\
=& 2^{|v|}(w)_2 + (v)_2 + 1 \\
=& (wv)_2 + 1.
\end{aligned}$$

□

Reduction Given an instance of Post's Correspondence Problem we describe what an instance of VEST may look like. For each pair (v, w) we define

$$T_{(v,w)} = \left(\begin{array}{c|c} T_v & \begin{matrix} 0 & 0 \\ 0 & 0 \end{matrix} \\ \hline \begin{matrix} 0 & 0 \\ 0 & 0 \end{matrix} & T_w \end{array}\right),$$

we set the initial vector $\mathbf{v} := (0, 1, 0, 1)^T$ and $S := (1, 0, -1, 0)$. The undecidability of VEST then follows from the following lemma.

Lemma 4.17. *Let* $(v_{i_1}, w_{i_1}), (v_{i_2}, w_{i_2}) \ldots, (v_{i_k}, w_{i_k})$ *be* k *pairs of binary words. Then*

$$ST_{(v_{i_k},w_{i_k})}T_{(v_{i_{k-1}},w_{i_{k-1}})} \cdots T_{(v_{i_1},w_{i_1})}\mathbf{v} = \mathbf{0}$$

if and only if

$$v_{i_1}v_{i_2} \cdots v_{i_k} = w_{i_1}w_{i_2} \cdots w_{i_k}.$$

Proof. By Lemma 4.16, $T_{(v_{i_k},w_{i_k})}T_{(v_{i_{k-1}},w_{i_{k-1}})} \cdots T_{(v_{i_1},w_{i_1})} = T_{(v_{i_1}v_{i_2}\cdots v_{i_k},w_{i_1}w_{i_2}\cdots w_{i_k})}$. The vector $\mathbf{v}$ selects the second column of the submatrix $T_{v_{i_1}v_{i_2}\cdots v_{i_k}}$ and the second column of the submatrix $T_{w_{i_1}w_{i_2}\cdots w_{i_k}}$. In other words, the result is equal to

$$\Big((v_{i_1}v_{i_2} \cdots v_{i_k})_2\,, (v_{i_1}v_{i_2} \cdots v_{i_k})_2 + 1, (w_{i_1}w_{i_2} \cdots w_{i_k})_2\,, (w_{i_1}w_{i_2} \cdots w_{i_k})_2 + 1\Big)^T.$$

The final result after multiplying S with the vector above is the following 1-dimensional vector $(v_{i_1}v_{i_2} \cdots v_{i_k})_2 - (w_{i_1}w_{i_2} \cdots w_{i_k})_2$. $\square$

Bibliography

[AB09] S. Arora and B. Barak. *Computational complexity: A modern approach*. Cambridge University Press, Cambridge, 2009.

[AB17] K. A. Adiprasito and B. Benedetti. Subdivisions, shellability, and collapsibility of products. *Combinatorica*, 37(1):1–30, 2017.

[ALW14] Amir Abboud, Kevin Lewi, and Ryan Williams. Losing weight by gaining edges. In Andreas S. Schulz and Dorothea Wagner, editors, *Algorithms - ESA 2014 - 22th Annual European Symposium, Wroclaw, Poland, September 8-10, 2014. Proceedings*, volume 8737 of *Lecture Notes in Computer Science*, pages 1–12. Springer, 2014.

[Ani85] D. J. Anick. Diophantine equations, Hilbert series, and undecidable spaces. *Annals of Mathematics*, 122:87–112, 1985.

[Ani87] D. J. Anick. Generic algebras and CW complexes. *Algebraic topology and algebraic K-theory*, pages 247–321, 1987.

[Ani89] D. J. Anick. The computation of rational homotopy groups is $\#\wp$-hard. Computers in geometry and topology, Proc. Conf., Chicago/Ill. 1986, Lect. Notes Pure Appl. Math. 114. pages 1–56, 1989.

[Bar41] R. Barkley. Explicit bounds for some functions of prime numbers. *American Journal of Mathematics*, 63(1):211–232, 1941.

[BKS23] C. Brand, V. Korchemna, and M. Skotnica. Deterministic Constrained Multilinear Detection. In J. Leroux, S. Lombardy, and D. Peleg, editors, *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*, volume 272 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:14, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.

[BKSS23] C. Brand, V. Korchemna, K. Simonov, and M. Skotnica. Counting vanishing matrix-vector products, 2023. arXiv: 2309.13698.

[BL13] B. Benedetti and F. H. Lutz. Knots in collapsible and non-collapsible balls. *Electron. J. Combin.*, 20(3):Paper 31, 29, 2013.

[BMS02] G. Baumslag, A. G. Myasnikov, and V. Shpilrain. Open problems in combinatorial group theory. Second edition. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 1–38. Amer. Math. Soc., Providence, RI, 2002.

[Boo59] W. W. Boone. The word problem. *Annals of mathematics*, 70:207–265, 1959.

[Bor20] E. Borghini. Some remarks on PL collapsible covers of 2-dimensional polyhedra. *Osaka Journal of Mathematics*, 57(1):39–49, 2020.

[Bro57] E. H. Brown. Finite computability of Postnikov complexes. *Annals of Mathematics*, 65:1, 1957.

[BW83] A. Björner and M. Wachs. On lexicographically shellable posets. *Trans. Amer. Math. Soc.*, 277(1):323–341, 1983.

[ČKM+14a] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Extendability of continuous maps is undecidable. *Discrete & Computational Geometry*, 51(1):24–66, 2014.

[ČKM+14b] M. Čadek, M. Krčál, J. Matoušek, L. Vokřínek, and U. Wagner. Polynomial-time computation of homotopy groups and Postnikov systems in fixed dimension. *SIAM Journal on Computing*, 43(5):1728–1780, 2014.

[Cla71] V. Claus. *Stochastische Automaten.* Vieweg+Teubner Verlag, 1971. (in German).

[CLOT03] O. Cornea, G. Lupton, J. Oprea, and D. Tanré. *Lusternik-Schnirelmann category*, volume 103 of *Mathematical Surveys and Monographs.* American Mathematical Society, Providence, RI, 2003.

[DF95] R. G. Downey and M. R. Fellows. Fixed-parameter tractability and completeness II: On completeness for W[1]. *Theoretical Computer Science*, 141(1):109–131, 1995.

[Far03] M. Farber. Topological complexity of motion planning. *Discrete Comput. Geom.*, 29(2):211–221, 2003.

[Far04] M. Farber. Instabilities of robot motion. *Topology Appl.*, 140(2-3):245–266, 2004.

[FG04a] J. Flum and M. Groge. *Parameterized Complexity Theory.* Springer, 2004.

[FG04b] J. Flum and M. Grohe. The parameterized complexity of counting problems. *SIAM Journal on Computing*, 33(4):892–922, 2004.

[FM20] M. Farber and S. Mescher. On the topological complexity of aspherical spaces. *J. Topol. Anal.*, 12(2):293–319, 2020.

[GPP+19] X. Goaoc, P. Paták, Z. Patáková, M. Tancer, and U. Wagner. Shellability is NP-complete. *J. ACM*, 66(3):21:1–21:18, 2019.

[Hac08] M. Hachimori. Decompositions of two-dimensional simplicial complexes. *Discrete Math.*, 308(11):2307–2312, 2008.

[HAMS93] C. Hog-Angeloni, W. Metzler, and A. J. Sieradski. *Two-Dimensional Homotopy and Combinatorial Group Theory.* London Mathematical Society Lecture Note Series. Cambridge University Press, 1993.

[Hat02] A. Hatcher. *Algebraic topology*. Cambridge Univ. Press, Cambridge, 2002.

[HT02] J. Herzog and Y. Takayama. Resolutions by mapping cones. *Homology Homotopy Appl.*, 4(2, part 2):277–294, 2002. The Roos Festschrift volume, 2.

[Hud69] J. F. P. Hudson. *Piecewise linear topology*. University of Chicago Lecture Notes prepared with the assistance of J. L. Shaneson and J. Lees. W. A. Benjamin, Inc., New York-Amsterdam, 1969.

[Koz97] D. N. Kozlov. General lexicographic shellability and orbit arrangements. *Ann. Comb.*, 1(1):67–90, 1997.

[Koz08] D. Kozlov. *Combinatorial algebraic topology*, volume 21 of *Algorithms and Computation in Mathematics*. Springer, Berlin, 2008.

[Lic91] W. B. R. Lickorish. Unshellable triangulations of spheres. *European J. Combin.*, 12(6):527–530, 1991.

[Mat03] J. Matoušek. *Using the Borsuk-Ulam theorem*. Universitext. Springer-Verlag, Berlin, 2003.

[Mat13] J. Matoušek. Computing higher homotopy groups is W[1]-hard. *arXiv preprint arXiv:1304.7705*, 2013.

[MF08] R. Malgouyres and A. R. Francés. Determining whether a simplicial 3-complex collapses to a 1-complex is NP-complete. In *Discrete geometry for computer imagery*, volume 4992 of *Lecture Notes in Comput. Sci.*, pages 177–188. Springer, Berlin, 2008.

[MMS22] J. Minařík, S. Moran, and M. Skotnica. How expressive are friendly school partitions?, 2022. arXiv: 2203.10772.

[MST21] T. Magnard, M. Skotnica, and M. Tancer. Shellings and sheddings induced by collapses. *SIAM Journal on Discrete Mathematics*, 35(3):1978–2002, 2021.

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

[Nov55] P. S. Novikov. On the algorithmic unsolvability of the word problem in group theory. *Trudy Mat. Inst. Steklov*, 44:1–143, 1955. (in Russian).

[PB80] J. S. Provan and L. J. Billera. Decompositions of simplicial complexes related to diameters of convex polyhedra. *Math. Oper. Res.*, 5(4):576–594, 1980.

[Pos47] E. L. Post. A variant of a recursively unsolvable problem. *Journal of Symbolic Logic*, 12(2):55–56, 1947.

[Roo79] J.-E. Roos. Relations between the poincaré-betti series of loop spaces and of local rings. In *Séminaire d'Algèbre Paul Dubreil*, pages 285–322. Springer, 1979.

[RS82] C. P. Rourke and B. J. Sanderson. *Introduction to piecewise-linear topology*. Springer Study Edition. Springer-Verlag, Berlin-New York, 1982. Reprint.

[SGW21] A. Santamaría-Galvis and R. Woodroofe. Shellings from relative shellings, with an application to NP-completeness. *Discrete Comput. Geom.*, 66(2):792–807, 2021.

[Sko19] M. Skotnica. No-three-in-line problem on a torus: Periodicity. *Discrete Mathematics*, 342(12):111611, 2019.

[Sko22] Michael Skotnica. VEST is W[2]-hard, 2022. arXiv: 2209.09788.

[ST23] M Skotnica and M Tancer. NP-hardness of computing PL geometric category in dimension 2. *SIAM Journal on Discrete Mathematics*, 37(3):2016–2029, 2023.

[Tan16] M. Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016.

[VKF74] I. A. Volodin, V. E. Kuznecov, and A. T. Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Uspehi Mat. Nauk*, 29(5(179)):71–168, 1974. Appendix by S. P. Novikov.

[Whi39] J. H. C. Whitehead. Simplicial spaces, nuclei and m-groups. *Proc. London Math. Soc. (2)*, 45(4):243–327, 1939.

List of publications

The papers related to this thesis are marked by $\star$.

Published papers

1. M. Skotnica. No-three-in-line problem on a torus: Periodicity. *Discrete Mathematics*, 342(12):111611, 2019

2. $\star$ T. Magnard, M. Skotnica, and M. Tancer. Shellings and sheddings induced by collapses. *SIAM Journal on Discrete Mathematics*, 35(3):1978–2002, 2021

3. C. Brand, V. Korchemna, and M. Skotnica. Deterministic Constrained Multilinear Detection. In J. Leroux, S. Lombardy, and D. Peleg, editors, *48th International Symposium on Mathematical Foundations of Computer Science (MFCS 2023)*, volume 272 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:14, Dagstuhl, Germany, 2023. Schloss Dagstuhl – Leibniz-Zentrum für Informatik

4. $\star$ M Skotnica and M Tancer. NP-hardness of computing PL geometric category in dimension 2. *SIAM Journal on Discrete Mathematics*, 37(3):2016–2029, 2023

Preprints

5. J. Minařík, S. Moran, and M. Skotnica. How expressive are friendly school partitions?, 2022. arXiv: 2203.10772

6. $\star$ Michael Skotnica. VEST is W[2]-hard, 2022. arXiv: 2209.09788

7. $\star$ C. Brand, V. Korchemna, K. Simonov, and M. Skotnica. Counting vanishing matrix-vector products, 2023. arXiv: 2309.13698