Parametrized Complexity of Expansion Height

Ulrich Bauer

Department of Mathematics, Technical University of Munich (TUM),
Boltzmannstr. 3, 85748 Garching b. München, Germany
ulrich.bauer@tum.de

Abhishek Rathod

Department of Mathematics, Technical University of Munich (TUM),
Boltzmannstr. 3, 85748 Garching b. München, Germany
abhishek.rathod@tum.de

Jonathan Spreer

School of Mathematics and Statistics, The University of Sydney, NSW 2006 Australia
jonathan.spreer@sydney.edu.au

───── **Abstract** ─────

Deciding whether two simplicial complexes are homotopy equivalent is a fundamental problem in topology, which is famously undecidable. There exists a combinatorial refinement of this concept, called simple-homotopy equivalence: two simplicial complexes are of the same simple-homotopy type if they can be transformed into each other by a sequence of two basic homotopy equivalences, an elementary collapse and its inverse, an elementary expansion. In this article we consider the following related problem: given a 2-dimensional simplicial complex, is there a simple-homotopy equivalence to a 1-dimensional simplicial complex using at most p expansions? We show that the problem, which we call the *erasability expansion height*, is $\mathbf{W[P]}$-complete in the natural parameter p.

1 Introduction

Homotopy theory lies at the heart of algebraic topology. In an attempt to make the concept of homotopy equivalence more amenable to combinatorial methods, Whitehead developed what turned out to be a combinatorial refinement of the theory, called *simple-homotopy theory*. Simple-homotopy theory considers sequences of elementary homotopy equivalences defined on simplicial complexes (or, more generally, CW complexes): an *elementary collapse*, which takes a face of a complex contained only in a single proper coface and removes both faces, and its inverse operation, called an *elementary expansion*. Two simplicial complexes are then said to be of the same simple-homotopy type if one can be transformed into the other by a sequence of elementary collapses and expansions. Complexes of the same simple-homotopy type are homotopy equivalent, but the converse is not always true [18], the obstruction being an element of the *Whitehead group* of the fundamental group. However, Whitehead

proved that all homotopy-equivalent complexes with a trivial fundamental group are in fact of the same simple-homotopy type [19], and thus in this particular case the notions of simple-homotopy and homotopy coincide.

A presentation of the fundamental group can be read off from a two-dimensional complex such that the presentation is balanced and describes the trivial group if and only if this complex is contractible [12]. Since the decidability of the triviality problem for balanced presentations is open [4], the same is also true for the decidability of contractibility of 2-complexes. Hence, the decidability of the existence of a simple-homotopy equivalence from a 2-complex to a point is also open. In contrast, the problem of deciding whether a given complex has trivial fundamental group is famously undecidable already for 2-complexes through its connection to the word problem, see, for instance, [5]. It follows that sequences of elementary collapses and expansions proving simple-homotopy equivalence between a 2-complex and a point can be expected to be long, if not unbounded. Nonetheless, understanding these sequences offers a great reward: the statement that any contractible 2-complex contracts to a point using only expansions up to dimension three is equivalent to a weaker variant [13, p. 34–35] of the Andrews–Curtis conjecture [1, 20].

In this article, motivated by the aforementioned problems, we investigate the computational (parametrized) complexity of a number of variants of the problem of deciding contractibility. More precisely, we focus on the problem of deciding whether a given 2-complex admits a simple-homotopy to a 1-complex using at most p expansions, called ERASIBILITY EXPANSION HEIGHT. In addition, we consider a variant, called ORDERED ERASIBILITY EXPANSION HEIGHT, which requires that all expansions come at the very beginning of the sequence. It is worth noting that ERASIBILITY EXPANSION HEIGHT and ORDERED ERASIBILITY EXPANSION HEIGHT are equivalent for CW complexes for which one can readily swap the order of expansions and collapses [13, p. 34]. However, for simplicial complexes, the ordered and unordered expansion heights may differ.

In Section 5.1, we prove that ERASIBILITY EXPANSION HEIGHT is **W[P]**-hard, see Theorem 8. The proof uses a reduction from AXIOM SET. The same reduction also establishes **W[P]**-hardness of ORDERED ERASIBILITY EXPANSION HEIGHT. Note that a reduction from AXIOM SET is also used by the third author and others in [6] to establish **W[P]**-hardness of a parametrized version of OPTIMAL MORSE MATCHING. However, unlike in [6], the use of combinatorial and topological properties of the dunce hat is a key ingredient of the reduction used in this paper. In particular, there is only one gadget in the reduction – a subdivision of the so-called *modified dunce hat* [3], see Figure 2. In this sense the techniques used in this paper are also related to recent work by the first and second author in [3], where they show hardness of approximation for some Morse matching problems.

In Section 5.2, we show that ERASIBILITY EXPANSION HEIGHT and ORDERED ERASIBILITY EXPANSION HEIGHT are both in **W[P]**, and hence also **W[P]**-complete, see Theorems 8, 13 and 14. Both results rest on the key observation that a 2-complex is erasable if and only if greedily collapsing triangles yields a 1-dimensional complex (Proposition 11), as shown by Tancer [17, Proposition 5].

In Section 6 we show that, as a consequence of the above reduction, the problem of deciding whether a 2-complex can be shown to be simple-homotopy equivalent to a 1-complex using only 3-dimensional expansions, called ERASIBILITY 3-EXPANSION HEIGHT, is **NP**-complete, see Theorem 17.

2 Definitions and Preliminaries

2.1 Simplicial complexes

A (finite) *abstract simplicial complex* is a collection K of subsets of a finite ground set V such that if τ is an element of K, and σ is a nonempty subset of τ, then σ is an element of K. The ground set V is called the *set of vertices* of K. Since simplicial complexes are determined by their facets, we sometimes present simplicial complexes by listing their facets. A *subcomplex* of K is a subset $L \subseteq K$ which is itself a simplicial complex. Given a subset $W \subseteq V$ of the vertices of K, the *induced subcomplex on W* consists of all simplices of K that are subsets of W.

The elements of K are referred to as its *faces*. The *dimension of a face* is defined to be its cardinality minus one, and the *dimension of K* equals the largest dimension of its faces. For brevity, we sometimes refer to a d-dimensional simplicial complex as a *d-complex* and to a d-dimensional face as a *d-face*. The 0-, 1-, and 2-faces of a d-complex K are called its *vertices*, *edges*, and *triangles* respectively. Faces of K which are not properly contained in any other face are called *facets*. An $(m-1)$-face $\sigma \in K$ which is contained in exactly one m-face $\tau \in K$ is called *free*.

The *star* of a vertex v of complex K, written $\mathrm{star}_K(v)$, is the subcomplex consisting of all faces of K containing v, together with their faces. If a map $\phi : V \to W$ between the vertex sets of two simplicial complexes K and L, respectively, sends every simplex $\sigma \in K$ to a simplex $\phi(\sigma) \in L$, then the induced map $f : K \to L, \sigma \mapsto \phi(\sigma)$, is said to be *simplicial*.

2.2 Simple-homotopy

We introduce the basic notions of simple-homotopy used in the present paper. The general concept of simple-homotopy can be understood independently from the notion of homotopy. In this sense this article aims to be self-contained. For further reading on homotopy theory we refer to [11].

In short, a simple-homotopy equivalence is a refinement of a homotopy equivalence. It can be described purely combinatorially with the help of the following definition.

▶ **Definition 1** (Elementary collapses and expansions)**.** *Let K_0 be a simplicial complex, and let $\tau, \sigma \in K_0$ be an m-face and an $(m-1)$-face respectively such that $\sigma \subset \tau$, and σ is free in K_0.*

We say that $K_1 = K_0 \setminus \{\tau, \sigma\}$ arises from K_0 by an elementary collapse of dimension m or elementary m-collapse, denoted by $K_0 \searrow K_1$. Its inverse, the operation $K_0 = K_1 \cup \{\tau, \sigma\}$ is called an elementary expansion of dimension m or elementary m-expansion, written $K_0 \nearrow K_1$. If the complex is implicit from context, we denote elementary collapses by $\searrow_\sigma^\tau$ and elementary expansions by $\nearrow_\sigma^\tau$. An elementary collapse or an elementary expansion is sometimes referred to as an elementary move, or simply a move.

If there exists a sequence of elementary collapses turning a complex K_0 into K_1 we write $K_0 \searrow K_1$ and say that K_0 collapses to K_1. If K_1 is one-dimensional, we say that K_0 is erasable. If K_1 is merely a point we call K_0 collapsible.

Finally, we write $K_0 \nearrow K_1$ to indicate a sequence of expansions and say that K_0 expands to K_1.

It follows that an expansion $\nearrow_\sigma^\tau$ can only be performed in a simplicial complex K if all codimension 1 faces of τ except for σ are already in K. Hence, let τ be an m-face of a simplicial complex K, and let σ be one of its $(m-1)$-faces. An *(m-dimensional) horn* $\mathbf{H}(\tau, \sigma)$ associated to the pair (τ, σ) is the simplicial complex generated by the $(m-1)$-faces of τ apart from σ.

All m-expansions and m-collapses with $m > 1$ leave the vertex set unchanged.

▶ **Definition 2** (Simple-homotopy equivalence, simple-homotopy graph)**.** *Two simplicial complexes K and L are said to be* simple homotopy equivalent *or of coinciding* simple-homotopy type*, written $K \nearrow_\searrow L$, if there exists a sequence $\mathcal{S}$ of elementary moves turning one into the other. In this case, we write $\mathcal{S} : K \nearrow_\searrow L$.*

The dimension of a simple-homotopy equivalence *is the maximum of the dimensions of K, L and of any elementary expansion or collapse in the sequence.*

The graph whose nodes are simplicial complexes, and two nodes are adjacent if their corresponding complexes are related by an elementary collapse is called simple-homotopy graph*. Naturally, its connected components are in one-to-one correspondence with simple-homotopy types.*

Two simplicial complexes of the same simple-homotopy type are homotopy equivalent, but the converse is not true, see, for instance, [18]. For simple-homotopy equivalent simplicial complexes we know the following.

▶ **Theorem 3** (Wall [17], Matveev [13, Theorem 1.3.5])**.** *Let K and L be two simplicial complexes of the same simple-homotopy type and of dimension at most $m > 2$. Then there exists a simple-homotopy equivalence of dimension at most $m + 1$, taking one to the other.*

For the case $m = 2$, Theorem 3 is still open and known as the (topological or geometric) Andrews–Curtis conjecture [2, 13, 15]. On the other hand, it is known that any contractible 2-complex is also simple-homotopy equivalent to a point [19]. Hence, any pair of contractible 2-complexes can be connected by a simple-homotopy equivalence of dimension at most four – but determining whether we can always decide if such a simple-homotopy equivalence exists is an open question [4], equivalent to the triviality problem for balanced group presentations [12].

2.3   Parametrized complexity

Parametrized complexity, as introduced by Downey and Fellows in [7], is a refinement of classical complexity theory. The theory revolves around the general idea of developing complexity bounds for instances of a problem not just based on their size, but also involving an additional *parameter*, which might be significantly smaller than the size. Specifically, we have the following definition.

▶ **Definition 4** (Parameter, parametrized problem)**.** *Let Σ be a finite alphabet.*
1. *A* parameter *of Σ^*, the set of strings over Σ, is a function $\rho : \Sigma^* \to \mathbb{N}$, attaching to every input $w \in \Sigma^*$ a natural number $\rho(w)$.*
2. *A* parametrized problem *over Σ is a pair (P, ρ) consisting of a set $P \subseteq \Sigma^*$ and a parametrization $\rho : \Sigma^* \to \mathbb{N}$.*

In this article we consider the complexity class **W[P]** for parametrized problems, following the definition by Flum and Grohe [8].

▶ **Definition 5** (Complexity Class **W[P]**)**.** *Let Σ be an alphabet and $\rho : \Sigma^* \to \mathbb{N}$ a parametrization. A nondeterministic Turing machine $\mathbb{M}$ with input alphabet Σ is called ρ-restricted if there are computable functions $f, h : \mathbb{N} \to \mathbb{N}$ and a polynomial p (with coefficients in the set of natural numbers) such that on every run with input $x \in \Sigma^*$ the machine $\mathbb{M}$ performs at most $f(k) \cdot p(|x|)$ steps, at most $h(k) \cdot \log |x|$ of them being nondeterministic, where $k := \rho(x)$. W[P] is the class of all parametrized problems (P, ρ) that can be decided by a ρ-restricted nondeterministic Turing machine.*

3 Problems

In this article we consider the following parametrized problems.

▶ **Problem 1** (ERASIBILITY EXPANSION HEIGHT)**.**

INSTANCE: A 2-dimensional simplicial complex K and a natural number p.

PARAMETER: p.

QUESTION: Is there a path in the simple-homotopy graph connecting K to a 1-complex using at most p expansions?

▶ **Problem 2** (ORDERED ERASIBILITY EXPANSION HEIGHT)**.**

INSTANCE: A 2-dimensional simplicial complex K and a natural number p.

PARAMETER: p.

QUESTION: Is there a path in the simple-homotopy graph connecting K to a 1-complex using first at most p expansions, followed by a sequence of only collapses?

In Section 5, we establish **W[P]**-completenes for ERASIBILITY EXPANSION HEIGHT and ORDERED ERASIBILITY EXPANSION HEIGHT.

The hardness proof works via a parametrized reduction using the AXIOM SET problem, which is a classical **NP**-complete problem [9, p. 263] that is well-known to be **W[P]**-complete with respect to the appropriate parameter [7, p. 473].

▶ **Problem 3** (AXIOM SET)**.**

INSTANCE: A finite set S of *sentences*, an *implication relation* R consisting of pairs (U, s) where $U \subseteq S$ and $s \in S$, and a positive integer $p \leq |S|$.

PARAMETER: p.

QUESTION: Is there a set $S_0 \subseteq S$, called an *axiom set*, with $|S_0| \leq p$ and a positive integer n such that if we recursively define

$$S_i := S_{i-1} \cup \{s \in S \mid \exists U \subseteq S_{i-1} : (U, s) \in R\}$$

for $1 \leq i \leq n$, then $S_n = S$?

▶ Remark 6. Note that every instance of AXIOM SET can be reduced in polynomial time to an instance for which every sentence must occur in at least one implication relation: First iteratively remove all sentences from the instance which do not feature in at least one implication relation. Then, for each of them, reduce p by one (note that each of them must necessarily be an axiom). It follows that solving the reduced instance is equivalent to solving the original instance.

Similarly, note that if there exists an implication $(U, s) \in R$, $s \in U$, we can simply omit it and, if this deletes s from the instance altogether, decrease p by one.

In Section 6, we show that the following variants of the expansion height problem are **NP**-complete.

▶ **Problem 4** (ERASIBILITY 3-EXPANSION HEIGHT)**.**

INSTANCE: A finite 2-dimensional simplicial complex K and a natural number p.

QUESTION: Is there a path in the simple-homotopy graph connecting K to a 1-complex using at most p expansions, all of which are 3-expansions?

▶ **Problem 5** (ORDERED ERASIBILTY 3-EXPANSION HEIGHT)**.**

INSTANCE: A finite 2-dimensional simplicial complex K and a natural number p.

QUESTION: Is there a path in the simple-homotopy graph connecting K to a 1-complex using first at most p expansions, all of which are 3-expansions, followed by a sequence of only collapses?

4 Contractibility and collapsibility for 2-complexes

The main gadget used in the proof of our main result, Theorem 8, is based on the simplest 2-dimensional contractible complex which is not collapsible to a point – the *dunce hat*. Hence, before we describe our main gadget in detail, we start this section by briefly discussing minimal triangulations of the dunce hat, and a variant that is collapsible through a unique free edge, the *modified dunce hat*.

4.1 The dunce hat

In the category of CW complexes, the dunce hat can be obtained by identifying two boundary edges of a triangle to build a cone and then gluing the third edge along the seam of the first gluing. The resulting complex does not have a collapsible triangulation. On the other hand, the dunce hat is known to be contractible [21].

The smallest simplicial complexes realizing this construction have 8 vertices, 24 edges and 17 triangles. There are seven such minimal triangulations of the dunce hat [16]. One such triangulation, denoted by $\mathbf{D}$, is shown in Figure 1. The dunce hat $\mathbf{D}$ has two horns, namely $\mathbf{H}(\{2,7,8\},\{1,2,7,8\})$ and $\mathbf{H}(\{3,5,6\},\{1,3,5,6\})$, and hence admits two 3-expansions, namely $\nearrow_{\{2,7,8\}}^{\{1,2,7,8\}}$ and $\nearrow_{\{3,5,6\}}^{\{1,3,5,6\}}$ respectively. They are shown by the shaded areas in Figure 1.
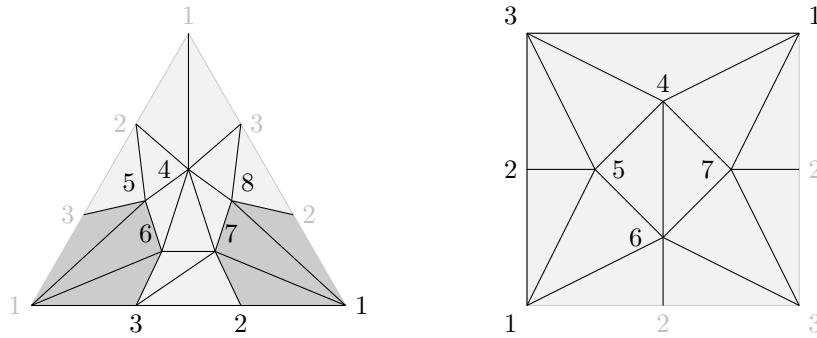

■ **Figure 1** Left: The 8-vertex triangulation $\mathbf{D}$ of the dunce hat. The two expansions turning it collapsible are highlighted. Right: The 7-vertex triangulation $\mathbf{P}$ of the modified dunce hat.

Note that after any of these two expansions we obtain a collapsible complex: After the expansion $\nearrow_{\{2,7,8\}}^{\{1,2,7,8\}}$ and the collapses $\searrow_{\{1,7,8\}}^{\{1,2,7,8\}}$, $\searrow_{\{1,7\}}^{\{1,2,7\}}$, and $\searrow_{\{1,8\}}^{\{1,2,8\}}$, the edge $\{1,2\}$ becomes free and thus $\mathbf{D}$ becomes collapsible. Similarly, starting with $\nearrow_{\{3,5,6\}}^{\{1,3,5,6\}}$, one may perform the collapse $\searrow_{\{1,5,6\}}^{\{1,3,5,6\}}$ and proceed in an analogous way. In particular, this shows that the dunce hat has the simple-homotopy type of a point, and in fact can be made collapsible by using a single expansion.

4.2 The modified dunce hat

Rather than working with the dunce hat directly, we base the construction of our gadget for the proof of Theorem 8 on the *modified dunce hat* [10]. More precisely, we "insert" a free edge into the dunce hat. For instance, Figure 1 depicts a triangulation of the modified dunce hat, which we denote by $\mathbf{P}$, with $\{1,3\}$ as the unique free edge. This particular triangulation of the modified dunce hat uses only 7 vertices, 19 edges, and 13 triangles. The modified dunce hat has previously been used as a gadget to show hardness of approximation for Morse matchings [3].

If we assume that $\mathbf{P}$ is part of a larger complex K, in which edge $\{1,3\}$ is glued to triangles not lying in $\mathbf{P}$, then $\{1,3\}$ is not free. In this case, the triangles of $\mathbf{P}$ can be collapsed away in essentially two distinct ways. Either, at some point in a simple-homotopy on K, the edge $\{1,3\}$ becomes free and thus the triangles of $\mathbf{P}$ collapse, or the triangles of $\mathbf{P}$ become collapsible by performing one of two possible 3-expansions on $\mathbf{P}$. Looking at the latter case in more detail, we have the following sequences of expansions and collapses:

$$\nearrow^{\{1,2,5,6\}}_{\{2,5,6\}}, \quad \searrow^{\{1,2,5,6\}}_{\{1,5,6\}}, \quad \searrow^{\{1,2,5\}}_{\{1,5\}}, \quad \searrow^{\{1,2,6\}}_{\{1,6\}}$$

and

$$\nearrow^{\{2,3,6,7\}}_{\{2,6,7\}}, \quad \searrow^{\{2,3,6,7\}}_{\{3,6,7\}}, \quad \searrow^{\{2,3,6\}}_{\{3,6\}}, \quad \searrow^{\{2,3,7\}}_{\{3,7\}} .$$

In the first case, the edge $\{1,2\}$ is freed, in the second case, the edge $\{2,3\}$ is freed. Both sequences can be extended to a collapsing sequence of the entire complex $\mathbf{P}$.

4.3　The main gadget

Our gadget for the proof of Theorem 8 is a subdivided version of the modified dunce hat $\mathbf{P}$ from Section 4.2. More precisely, it is determined by two positive integers m and ℓ, denoted by $\mathbf{P}_{m,\ell}$, and can be constructed from the complex $\mathbf{P}$ in essentially two steps.

1. Subdivide the edge $\{1,3\}$ of $\mathbf{P}$ $(m-1)$ times, thereby introducing vertices $x_1, \ldots x_{m-1}$. Relabel $1 \to x_0$ and $3 \to x_m$ to obtain m free edges $f_i = \{x_{i-1}, x_i\}$, $1 \leq i \leq m$.
2. Remove the edge $\{4,6\}$ and place ℓ vertex-disjoint copies of the disk

$$\{\{c_j, a_j, y_j\}, \{c_j, y_j, z_j\}, \{c_j, z_j, b_j\}, \{d_j, a_j, y_j\}, \{d_j, y_j, z_j\}, \{d_j, z_j, b_j\}\},$$

   $1 \leq j \leq \ell$, inside the 4-gon in the center of $\mathbf{P}$ bounded by 4, 5, 6, and 7. Triangulate the remaining space in the interior of the 4-gon. This creates edges $e_j = \{y_j, z_j\}$, $1 \leq j \leq \ell$, with pairwise vertex disjoint stars disjoint to 4, 5, 6, and 7 (now $\{4,6\}$ reappears as a path from 4 to 6, and thus $\mathbf{P}_{m,\ell}$ is in fact a proper subdivision of $\mathbf{P}$). See Figure 2 for an illustration.
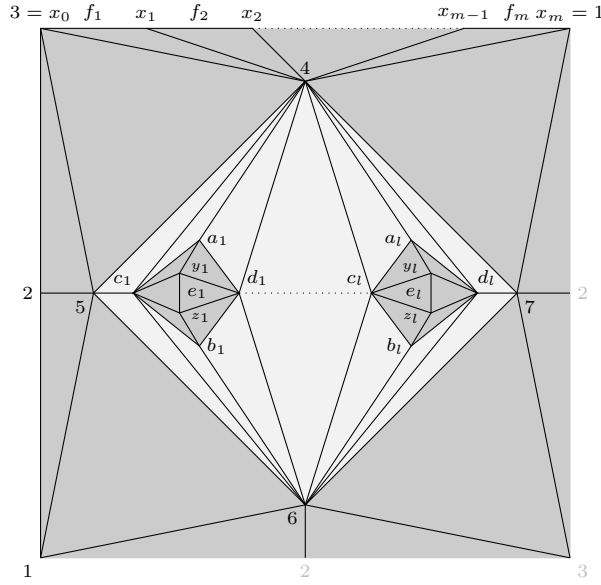
One key property of $\mathbf{P}_{m,\ell}$ is that we do not subdivide any faces of $\mathbf{P}$ near to the two available 3-expansions. As a result, again, assuming that $\mathbf{P}_{m,\ell}$ is part of a larger complex K where all free edges of $\mathbf{P}_{m,\ell}$ are glued to other triangles of K outside of $\mathbf{P}_{m,\ell}$ and thus are not free, the triangles of $\mathbf{P}_{m,\ell}$ can be collapsed according to the following observation:

▶ **Remark 7.** Let K be a two-dimensional simplicial complex such that $\mathbf{P}_{m,\ell}$ is a subcomplex whose vertices do not span any other faces of K (i.e., $\mathbf{P}_{m,\ell}$ is an induced subcomplex of K), and $K \sear013 L$ where L is a 1-complex. Then, at least one of the following three statements holds true at some point in $K \searrow L$, enabling us to eventually collapse away all the triangles of $\mathbf{P}_{m,\ell}$.

1. one of the edges $f_i \in \mathbf{P}_{m,\ell}$ becomes free;
2. one of two 3-expansions on $\mathbf{P}_{m,\ell}$: $\nearrow^{\{1,2,5,6\}}_{\{2,5,6\}}$ or $\nearrow^{\{2,3,6,7\}}_{\{2,6,7\}}$ is performed;
3. multiple expansions result in a complex in which all the triangles of $\mathbf{P}_{m,\ell}$ can be collapsed.

In other words, if one of the edges $f_i \in \mathbf{P}_{m,\ell}$ does not become free at some point in $K \searrow L$, then one is forced to use 3-expansions (either directly on $\mathbf{P}_{m,\ell}$, or after performing additional expansions) to collapse away the triangles of $\mathbf{P}_{m,\ell}$.

In Section 5.1 we use this gadget to reduce an instance $A = (S, R, p)$ of Axiom Set to Erasibility Expansion Height: Every sentence $s \in S$ is associated with one copy of $\mathbf{P}_{m,\ell}$, the edges f_i correspond to implications $(U, s) \in R$, and the edges e_j correspond to whenever $s \in U$ for some implication $(U, u) \in R$.

Figure 2 The main gadget of the construction $\mathbf{P}_{m,\ell}$.

5   Parametrized complexity of Erasibility Expansion Height

In this section, we first prove that ERASIBILITY EXPANSION HEIGHT and ORDERED ERASIBILITY EXPANSION HEIGHT are **W[P]**-hard by a reduction from AXIOM SET, a problem known to be **W[P]**-complete. We then show that the two problems are also contained in **W[P]**.

5.1   W[P]-hardness of expansion height problems

▶ **Theorem 8.** ERASIBILITY EXPANSION HEIGHT *and* ORDERED ERASIBILITY EXPANSION HEIGHT *are* **W[P]**-*hard problems.*

The following lemma is used to assemble the gadgets in our reduction into a simplicial complex K.

▶ **Lemma 9** (Munkres, [14, Lemma 3.2]). *Let C be a finite set, let K be a simplicial complex with set of vertices V, and let $f : V \to C$ be a surjective map associating to each vertex of K a color from C. The coloring f extends to a simplicial map $g : K \to K_f$ where K_f has vertex set C and is obtained from K by identifying vertices with equal color.*

If for all pairs $v, w \in V$, $f(v) = f(w)$ implies that their stars $\mathrm{star}_K(v)$ and $\mathrm{star}_K(w)$ are vertex disjoint, then, for all faces $\tau, \sigma \in K$ we have that
- *τ and $g(\tau)$ have the same dimension, and*
- *$g(\tau) = g(\sigma)$ implies that either $\tau = \sigma$ or τ and σ are vertex disjoint in K.*

Lemma 9 provides a way of gluing faces of a simplicial complex by a simplicial quotient map obtained from vertex identifications, and tells us when this gluing does not create unwanted identifications.

Proof of Theorem 8. We want to reduce AXIOM SET to ERASIBILITY EXPANSION HEIGHT.

Fix an instance $A = (S, R, p)$ of AXIOM SET such that every sentence $s \in S$ is subject to at least one implication $(U, s) \in R$ and such that $(U, s) \in R$ implies $s \notin U$. By Remark 6, this is not a restriction, since every instance of AXIOM SET can be reduced to such an

instance in polynomial time. For every sentence $s \in S$, take a copy $\mathbf{P}_s$ of the gadget $\mathbf{P}_{m,\ell}$ to model s, where $\ell \geq 0$ is the number of implications $(U, u) \in R$ with $s \in U$ and $m \geq 1$ is the number of implications $(U, s) \in R$. Thus, for all values $\ell \geq 0$ and $m \geq 1$ the gadget $\mathbf{P}_{m,\ell}$ is a simplicial complex without any unintended identifications. Denote the free edges of $\mathbf{P}_s$ by $f_i^s = \{x_{i-1}^s, x_i^s\}$, $1 \leq i \leq m$, and its edges of type e_j by $e_j^s = \{y_j^s, z_j^s\}$, $1 \leq j \leq \ell$.

For a fixed $s \in S$, endow the set of implications $(U, s) \in R$ of s with an arbitrary order $(U_1, s), \ldots, (U_m, s)$. Similarly, for a fixed $u \in S$, order the set of implications in R containing u arbitrarily as $(U^1, s^1), \ldots, (U^\ell, s^\ell)$. Now for every $(U_i, s) \in R$ and every $u \in U_i = U^j$ (i.e., $s^j = s$), glue the edge $f_i^s = \{x_{i-1}^s, x_i^s\}$ of the gadget $\mathbf{P}_s$ to the edge $e_j^u = \{y_j^u, z_j^u\}$ of $\mathbf{P}_u$ by identifying x_{i-1}^s with y_j^u and x_i^s with z_j^u.

Performing these identifications for all implications in R yields a complex, which we denote by K. Note that, fixing $s \in S$, and $0 \leq i \leq m$, the only vertices to which x_i^s can possibly be identified to in K are y_j^u, z_j^u $(1 \leq j \leq \ell)$.

More precisely, using the orderings from above for vertex x_i^s, let (U_i, s) and (U_{i+1}, s) be the i-th and $(i+1)$-st implication of s in R (if $i \in \{0, m\}$ there is only one implication to consider) and denote their sentences by $u_1, \ldots, u_r$ and $u^1, \ldots, u^t$, where $r = |U_i|$ and $t = |U_{i+1}|$. Moreover, let U_i (resp. U_{i+1}) be the j_f-th (resp. j^g-th) implication where the sentence u_f (resp. u^g) occurs, for $1 \leq f \leq r$ (resp. $1 \leq g \leq t$). Then x_i^s is identified with $z_{j_f}^{u_f}$ $(1 \leq f \leq r)$ and $y_{j^g}^{u^g}$ $(1 \leq g \leq t)$.

Now since every fixed edge of type e_j^u is only identified with one edge of type f_i^s, those vertices are not identified with any other vertices of K. Since, by construction, the set of the vertex stars of $z_{j_f}^{u_f}$ $(1 \leq f \leq r)$, $y_{j^g}^{u^g}$ $(1 \leq g \leq t)$, and x_i^s are pairwise vertex disjoint, we can apply Lemma 9 to ensure that no unwanted identifications occur in building up K. In particular, every gadget $\mathbf{P}_s$ is a subcomplex of K via the canonical isomorphism given by the gluing map.

We now show that the following three statements are equivalent for our complex K:

(a) there exists a simple-homotopy equivalence turning K into a 1-dimensional complex using first at most p expansions, followed by a sequence of only collapses,

(b) there exists a simple-homotopy equivalence $K \nearrow\!\!\searrow L$ turning K into a 1-dimensional complex L using at most p expansions, and

(c) there exists an axiom set $S_0 \subset S$ for $A = (S, R, p)$ using at most p elements.

We trivially have that (a) $\Longrightarrow$ (b).

In order to show that (c) $\Longrightarrow$ (a), assume that there exists an axiom set $S_0 \subset S$ of size p, and perform one 3-expansion on each gadget $\mathbf{P}_u$ with $u \in S_0$. As described in Section 4.3, these expansions admit all triangles of these gadgets to collapse. This, in turn, frees all edges f_i^s where $s \in S$ has an implication $(U, s) \in R$ with $U \subset S_0$. Consequently, all triangles of such gadgets $\mathbf{P}_s$ can be collapsed. Since S_0 is an axiom set, repeating this process eventually collapses away all tetrahedra and triangles, leaving a 1-complex.

In order to show that (b) $\Longrightarrow$ (c), we start with a few definitions. For a complex K' with $K \nearrow\!\!\searrow K'$, we say that a gadget $\mathbf{P}_s \subseteq K$ is *touched* with respect to a simple-homotopy sequence $\mathcal{S} : K \nearrow\!\!\searrow K'$ if at some point in the simple-homotopy sequence one of the triangles of $\mathbf{P}_s$ is removed. Otherwise $\mathbf{P}_s$ is said to be *untouched*. Note that even if all triangles of $\mathbf{P}_s$ are present in K', $\mathbf{P}_s$ might still be touched. Being touched or untouched is a property of the sequence $\mathcal{S} : K \nearrow\!\!\searrow K'$, not of the complex K'.

We build the axiom set $S_0 \subset S$ for $A = (S, R, p)$ in the following way: A sentence $s \in S$ is in S_0 if and only if a triangle in $\mathbf{P}_s$ is removed by a 3-collapse of the given simple-homotopy sequence $\mathcal{S} : K \nearrow\!\!\searrow L$. We first inductively prove a claim about

$$S^k = \{s \in S \mid \mathbf{P}_s \text{ is touched by a 3-collapse in the first } k \text{ moves of } K \nearrow\!\!\searrow L\}.$$

We first inductively prove a claim about

$$S^k = \{s \in S_0 \mid \mathbf{P}_s \text{ is by a 3-collapse in the first } k \text{ moves of } K \searrow L\}.$$

▷ **Claim 10.** For $s \in S$, if the gadget $\mathbf{P}_s$ is touched by the first k elementary moves in $\mathcal{S} : K \searrow L$, then s is implied by sentences in S^k.

Proof. First note that all gadgets are untouched in K and $S^0 = \emptyset$.

By induction hypothesis, if a gadget $\mathbf{P}_s$ is touched by one of the first $k - 1$ moves in $\mathcal{S} : K \searrow L$, then s is implied by sentences in S^{k-1}.

The induction claim is trivially true if $\mathbf{P}_s$ is touched in the first $k - 1$ moves, or if $\mathbf{P}_s$ is touched by a 3-collapse in the k-th move ($s \in S^k \setminus S^{k-1}$), causing the sentence s to be included in S_0.

So, suppose that this is not the case. That is, suppose that $\mathbf{P}_s$ is untouched in the length $k-1$ prefix $\mathcal{S}' : K \searrow K'$ of $\mathcal{S} : K \searrow L$ and touched by a 2-collapse in the k-th move. This implies that $S^k = S^{k-1}$ and that $\mathbf{P}_s$ is a subcomplex of K' and one of the edges f_i^s must be free in K'. Now let $\mathbf{P}_{u_1}, \mathbf{P}_{u_2}, \ldots, \mathbf{P}_{u_q}$ be the set of other gadgets containing triangles glued to f_i^s in the original complex K (that is, there is an implication $(\{u_1, u_2, \ldots, u_q\}, s) \in R$). Since none of these triangles are present in K', all gadgets $\mathbf{P}_{u_1}, \mathbf{P}_{u_2}, \ldots, \mathbf{P}_{u_q}$ must be touched in $\mathcal{S}' : K \searrow K'$. Thus, either they were touched by a 3-collapse and their corresponding sentences are part of S^{k-1}, or they were touched by a 2-collapse and, by the induction hypothesis, their corresponding sentences are implied by sentences in S^k. It follows that s is implied by sentences in $S^k = S^{k-1}$, proving the claim. ◁

By assumption, K is simple homotopy equivalent to a 1-complex L. That is, $\mathcal{S}' : K \searrow L$ eventually removes all triangles from K. Hence, every sentence $s \in S$ is touched as a result of a 2-collapse or a 3-collapse. Let m be the number of elementary moves needed to reach L starting from K. Then, by Claim 10, $S^m = S_0$ is the desired axiom set. Also, since a sentence s is included in S_0 only if a triangle belonging to gadget $\mathbf{P}_s$ is removed as part of a 3-collapse, and since a triangle belonging to gadget $\mathbf{P}_s$ does not belong to any other gadget $\mathbf{P}_u$ for $u \neq s$, S_0 cannot contain more elements than the number of 3-collapses (and hence 3-expansions).

Finally, we infer **W[P]**-hardness of ERASIBILITY EXPANSION HEIGHT and ORDERED ERASIBILITY EXPANSION HEIGHT from the above equivalence and the **W[P]**-hardness of AXIOM SET [7, p. 473]. ◀

5.2   W[P]-membership of Erasibility Expansion Height

We now show that ORDERED ERASIBILITY EXPANSION HEIGHT and ERASIBILITY EXPANSION HEIGHT are in **W[P]** by describing suitable nondeterministic algorithms for deciding both problems. We begin with a well-known fact about checking collapsibility of 2-complexes.

▶ **Proposition 11** (Tancer [17], Proposition 5). *Let K be a 2-complex that collapses to a 1-complex L and to another 2-complex M. Then M also collapses to a 1-complex.*

▶ Remark 12. The proposition above implies that we can collapse an input 2-complex K greedily until no more 2-collapses are possible, and if K collapses to a 1-complex L, the algorithm is guaranteed to terminate with a 1-complex as well.

▶ **Theorem 13.** ORDERED ERASIBILITY EXPANSION HEIGHT *is in W[P].*

Proof. Let K be a simplicial complex with n simplices. First, note that if there exists a simple homotopy sequence $\mathcal{S}$ taking K to a 1-complex with p expansions that all come at the beginning of the sequence, then there also exists a simple homotopy sequence $\mathcal{S}_{\mathbb{M}}$ taking K to a 1-complex where p expansions are followed by collapses such that, for each d, all collapses of dimension $d+1$ are executed before collapses of dimension d. This follows from observing that, for any two d-collapses $\searrow_\sigma^\tau$ and $\searrow_\alpha^\beta$, if the d-collapse $\searrow_\sigma^\tau$ is executed before the d-collapse $\searrow_\alpha^\beta$ in $\mathcal{S}$, then the same can be carried out in $\mathcal{S}_{\mathbb{M}}$. Also, in any simple homotopy sequence $\mathcal{S}$ that takes K to a 1-complex, for every $d > 2$, the number of d-expansions equals the number of d-collapses in $\mathcal{S}$. This follows from a simple inductive argument starting with highest dimensional moves.

Denoting the total number of d-collapses, $d > 2$, in $\mathcal{S}$ by $q \leq p$, it follows that, if there exists a simple homotopy sequence $\mathcal{S}$ with p expansions that come at the beginning, then there exists a simple homotopy sequence $\mathcal{S}_{\mathbb{M}}$ with p expansions in the beginning followed by q collapses that gives rise to a 2-complex K' with $O(n^3)$ faces. The faces can be as many as $O(n^3)$ since $\mathbb{M}$ does not guess any 2-collapses. Furthermore, if $\mathcal{K}$ is erasable through the simple homotopy sequence $\mathcal{S}$, then K' is also erasable, once again, because the 2-collapses of $\mathcal{S}$ can be carried out in the same order in $\mathcal{S}_{\mathbb{M}}$. Hence, the non-deterministic Turing machine $\mathbb{M}$ can now be described as follows:

1. Guess p expansions and q collapses non-deterministically to obtain a complex K'.
2. Deterministically check if K' is erasable.

By Remark 12, erasability of K' can be deterministically checked in time polynomial in n.

Since any simplex in the desired simple homotopy sequence has at most $n + p$ vertices, the number of bits required to encode a single vertex is $O(\log(n + p))$. Also, because the dimension of the faces involved in expansions and collapses is certainly in $O(p)$, and since an expansion or a collapse can be fully described by a pair of simplices, the number of bits required to encode an expansion or a collapse is $O(p \log(n + p))$. Hence, in order to guess $p + q$ moves, it suffices for $\mathbb{M}$ to guess $O(p^2 \cdot \log(n + p))$ bits in total since $q \leq p$. Now, assuming $n, p \geq 2$, we have

$$p^2 \log(n + p) \leq p^2 \log(np) = p^2 \log(n) + p^2 \log(p) \leq p^2 (1 + \log(p)) \log(n).$$

Hence, for sufficiently large n and p, the number of bits guessed by $\mathbb{M}$ is bounded by a function of the form $f(p) \log n$. Thus, $\mathbb{M}$ is a p-restricted Turing machine, and ORDERED ERASIBILITY EXPANSION HEIGHT is in **W[P]**. ◀

▶ **Theorem 14.** ERASIBILITY EXPANSION HEIGHT *is in* **W[P]**.

Proof. Assume that there exists a simple homotopy that takes K to a 1-complex using no more than p expansions. The Turing machine $\mathbb{M}$ needs to generate one such sequence. Below, we show that, in order to achieve this, $\mathbb{M}$ does not have to guess an entire simple homotopy sequence, but only a subsequence, and the remaining part of the sequence can be found deterministically by $\mathbb{M}$.

Given a 2-dimensional complex K with n faces, $\mathbb{M}$ first nondeterministically guesses p expansions and p collapses, and the order in which they are to be executed. These moves are referred to in the following as *prescribed moves*. While these moves are meant to appear in a specified order, they need not appear consecutively. The moves that are not prescribed are computed deterministically by $\mathbb{M}$. A simple homotopy sequence of K, that takes K to a 1-complex, in which all the prescribed moves occur as a subsequence, is called a sequence *compatible* with the prescribed moves. By assumption, there exists a set of prescribed moves for which a compatible sequence exists.

In order to give a description of $\mathbb{M}$, we introduce some additional terminology. Let $\mathcal{S}_q^j$ be an intermediate simple homotopy sequence computed by $\mathbb{M}$, such that the first j prescribed moves guessed by $\mathbb{M}$ form a subsequence of $\mathcal{S}_q^j$, q is the total number of moves in $\mathcal{S}_q^j$, and $\mathcal{S}_q^j$ is a prefix of a set of sequences $\mathcal{S}$ compatible with the prescribed set of moves. Let $K \diagdown K_q^j$ be the complex obtained by executing the moves in $\mathcal{S}_q^j$. Then, a collapse $\diagdown_\sigma^\tau$ in K_q^j is *valid* for this prefix if appending the collapse still leaves a compatible prefix. That is, $\mathcal{S}_q^j$ appended with the collapse $\diagdown_\sigma^\tau$ (giving $\mathcal{S}_{q+1}^j$) continues to be a prefix of at least one compatible sequence $\mathcal{S} \in \mathcal{S}$. A collapse that is not valid is said to be *forbidden*.

Note that labelling vertices of a complex C by natural numbers determines a lexicographic order on the simplices of C. The lexicographic order $<_C$ on simplices of C can be extended to a lexicographic order $\prec$ on collapses as follows: If $(\diagdown_\sigma^\tau), (\diagdown_\alpha^\beta)$ are two collapses in C, then $(\diagdown_\sigma^\tau) \prec (\diagdown_\alpha^\beta)$ if $\sigma <_C \alpha$.

The Turing machine $\mathbb{M}$ for deciding ERASIBILITY EXPANSION HEIGHT can be described as follows:

1. Guess $2p$ prescribed moves non-deterministically.
2. Execute 2-collapses in lexicographic order until no more 2-collapses are valid.
3. Repeat until all prescribed moves have been executed:
   a. Execute the next prescribed move.
   b. Execute 2-collapses in lexicographic order until no more collapses are valid.

Let $\mathcal{S}$ be a sequence compatible with an ordered set of prescribed moves X (of cardinality $2p$), and let $\mathcal{S}_\mathbb{M}$ be a simple homotopy sequence computed by $\mathbb{M}$ as above such that X is a subsequence of $\mathcal{S}_\mathbb{M}$. Now, let σ be a free edge associated with a 2-collapse $\diagdown_\sigma^\tau$ in $\mathcal{S}_q^j$ for some j and q, where $\mathcal{S}_q^j$ is a subsequence of $\mathcal{S}_\mathbb{M}$. If there exist future prescribed moves including cofaces of σ, then the next prescribed move including cofaces of σ that is not an expansion involving σ is denoted by m_1. Similarly, if there exist future prescribed moves including cofaces of τ, then the next prescribed move including cofaces of τ that is not an expansion involving τ is denoted by m_2. Note that m_2 cannot come before m_1 but we may have $m_1 = m_2$. Then, the 2-collapse $\diagdown_\sigma^\tau$ is forbidden if and only if m_1 exists and is not preceded by a future prescribed expansion involving σ or m_2 exists, and is not preceded by a future prescribed expansion involving τ. It follows that, for each free edge, checking if a collapse is forbidden (or valid) can be done deterministically in time polynomial in p. To see this note that the most expensive atomic operation is to check if a simplex (of dimension 1 or 2) is a face of a simplex that is at most p dimensional, and the number of prescribed moves is at most $2p$. Altogether, the set of valid collapses can be computed in time polynomial in n and p, which can also be lexicographically ordered in polynomial time.

Finally, let K' denote the complex obtained from K by the sequence $\mathcal{S}_\mathbb{M}$. Then, the following claim establishes the effectiveness of the greedy strategy employed by $\mathbb{M}$.

▷ **Claim 15.** If there exists a simple homotopy sequence with at most p expansions that takes K to a 1-complex, then there exists an execution branch of the Turing machine that terminates successfully, i.e., the complex K' obtained by $\mathbb{M}$ is a 1-complex.

Proof. Let $\mathcal{S}$ be a simple homotopy sequence with p expansions that takes K to a 1-complex. Let X_e be the ordered set of expansions in $\mathcal{S}$. Thus, $|X_e| = p$. Moreover, let X_e^+ (X_c^+) denote the d-expansions (d-collapses) in $\mathcal{S}$ with $d > 2$. As in Theorem 13, by a simple inductive argument starting from the highest dimension it can be shown that $|X_c^+| = |X_e^+|$. To the p expansions of $\mathcal{S}$, we associate a set X_c of collapses of $\mathcal{S}$ as follows: If $|X_c^+| < p$, then let X_c^- be an arbitrary set of d-collapses in $\mathcal{S}$ with $d \le 2$, and $|X_c^-| = p - |X_c^+|$. Now, let $X_c = X_c^+ \cup X_c^-$, so that $|X_c| = p$. Finally, let the ordered set X of prescribed moves be the set containing all elements of $X_c \cup X_e$ seen as a subsequence of $\mathcal{S}$.

We assume that the non-deterministic Turing machine $\mathbb{M}$ correctly guesses the specified sequence of prescribed moves X. It now suffices to show the following claim about the sequence $\mathcal{S}_{\mathbb{M}}$.   ◁

▷ **Claim 16.** $\mathcal{S}_{\mathbb{M}}$ is compatible with the prescribed moves X, and $\mathcal{S}_{\mathbb{M}}$ takes K to a 1-complex if $\mathcal{S}$ takes K to a 1-complex.

Proof. Let $K \searrow K^j$ denote the complex obtained from $\mathcal{S}$ after executing the j-th prescribed move in $\mathcal{S}$. We show that there exists a complex $K_{\mathbb{M}}^j$ obtained from $\mathcal{S}_{\mathbb{M}}$ after executing the j-th prescribed move in $\mathcal{S}_{\mathbb{M}}$. Also, let T^j ($T_{\mathbb{M}}^j$) denote the set of 2-simplices of K^j ($K_{\mathbb{M}}^j$).

First observe that, $K_{\mathbb{M}}^0 = K$ exists and that $T_{\mathbb{M}}^0 = T^0$. We now show that the following claim is inductively true: $T_{\mathbb{M}}^j \subset T^j$ for all $j \in [1, 2p]$. Suppose we make the induction hypothesis that $T_{\mathbb{M}}^{j-1} \subset T^{j-1}$ for some $j \in [1, 2p]$. Then, the set of forbidden collapses for $\mathcal{S}$ and $\mathcal{S}_{\mathbb{M}}$ are the same until the j-th move in X can be reached. Let τ_1 be the first 2-face of K^{j-1} that is removed as part of a 2-collapse after $j-1$ prescribed moves have been executed in $\mathcal{S}$. Without loss of generality, assume that the 2-collapse that removes τ_1 is non-prescribed. Then, there exists an edge $\sigma_1 \subset \tau_1$ such that τ_1 is the unique coface of σ_1 in K^{j-1}. By induction hypothesis, since $T_{\mathbb{M}}^{j-1} \subset T^{j-1}$ the same is also true for $K_{\mathbb{M}}^{j-1}$. Since $\mathbb{M}$ greedily removes every valid collapse it can (in lexicographic order), at some appropriate lexicographic index, τ_1 is also removed from $K_{\mathbb{M}}^{j-1}$ (possibly along with σ_1). Now, let $\tau_1, \tau_2, \ldots, \tau_{q-1}$ be the first $q - 1$ 2-faces removed from K^{j-1} (as part of non-prescribed collapses). Assume that $\tau_1, \tau_2, \ldots, \tau_{q-1}$ have also been removed from $K_{\mathbb{M}}^{j-1}$. By the same reasoning as before, if τ_q is the q-th face to be removed from $K_{\mathbb{M}}^{j-1}$ (as part of non-prescribed collapses), then τ_q may also be removed from $\mathcal{S}_{\mathbb{M}}$ as part of a valid collapse. Hence, by induction, $T_{\mathbb{M}}^j \subset T^j$ for all $j \in [1, 2p]$.

Finally, since by assumption, K^{2p} collapses to a 1-complex, by applying arguments analogous to the induction above, the same is true for $K_{\mathbb{M}}^{2p}$ since $T_{\mathbb{M}}^{2p} \subset T^{2p}$.   ◁

Since given a 2-complex with n faces, $\mathbb{M}$ non-deterministically guesses $2p$ moves, as in Theorem 13, the number of bits guessed by $\mathbb{M}$ is bounded by $f(p) \log(n)$, where $f(p) = O(p^2(1 + \log(p)))$. Hence, $\mathbb{M}$ is a p-restricted Turing machine, and Erasibility Expansion Height is in **W[P]**.   ◀

6   NP-completeness of Erasibility 3-Expansion Height

Note that the parametrized reduction from Axiom Set to Erasibility Expansion Height (and Ordered Erasibility Expansion Height) is also a polynomial-time reduction (or Karp reduction) from Axiom Set to Erasibility 3-Expansion Height (and Ordered Erasibilty 3-Expansion Height), since the complexity of reduction is independent of the parameter p and depends only on the size of the input complex. This observation leads us to the following result.

▶ **Theorem 17.** *The decision problems* Erasibility 3-Expansion Height *and* Ordered Erasibilty 3-Expansion Height *are* **NP**-*hard.*

Proof. Since the Axiom Set problem is known to be **NP**-hard [9], it follows that Erasibility 3-Expansion Height and Ordered Erasibilty 3-Expansion Height are also **NP**-hard.   ◀

For the rest of the section, we assume that K is a 2-complex K with n faces and m vertices. The total number of simplices that one can encounter in any simple homotopy

sequence of K using only 3-expansions is at most $M = O(m^4)$. (Note that the ground set of K is fixed since we do not allow 1-expansions). Hence, the total number of *elementary moves* that may be available at any given point in the sequence is bounded by $O(M)$. That is, p itself is bounded by $O(M)$.

▶ **Theorem 18.** ERASIBILITY 3-EXPANSION HEIGHT *is in* ***NP***.

Proof. The non-deterministic algorithm $\mathbb{M}$ for deciding ERASIBILITY 3-EXPANSION HEIGHT first guesses at each point in the simple homotopy sequence starting with K, one elementary move (out of at most $O(M)$ available moves), and constructs a new complex from the move. The total number of moves made by $\mathbb{M}$ is bounded by $(\frac{n+2p-1}{2})$. Finally, $\mathbb{M}$ checks if the final complex is a 1-complex. ◀

▶ **Theorem 19.** ORDERED ERASIBILTY 3-EXPANSION HEIGHT *is in* ***NP***.

Proof. The non-deterministic algorithm $\mathbb{M}$ for deciding ORDERED ERASIBILTY 3-EXPANSION HEIGHT first guesses at most p 3-expansions followed by an equal number of 3-collapses, resulting in a 2-complex K' with n faces. From Remark 12, the erasability of K' can be deterministically checked in time polynomial in n, proving the claim. ◀

───── **References** ─────

1 J. J. Andrews and M. L. Curtis. Free groups and handlebodies. *Proc. Amer. Math. Soc.*, 16:192–195, 1965. `doi:10.2307/2033843`.

2 Jonathan A. Barmak. *Algebraic topology of finite topological spaces and applications*, volume 2032 of *Lecture Notes in Mathematics*. Springer, Heidelberg, 2011. `doi:10.1007/978-3-642-22003-6`.

3 Ulrich Bauer and Abhishek Rathod. Hardness of approximation for Morse matching. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2663–2674. SIAM, Philadelphia, PA, 2019. `doi:10.1137/1.9781611975482.165`.

4 Gilbert Baumslag, Alexei G. Myasnikov, and Vladimir 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. `doi:10.1090/conm/296/05067`.

5 V. V. Borisov. Simple examples of groups with unsolvable word problem. *Mat. Zametki*, 6:521–532, 1969.

6 Benjamin A. Burton, Thomas Lewiner, João Paixão, and Jonathan Spreer. Parameterized complexity of discrete Morse theory. *ACM Trans. Math. Software*, 42(1):Art. 6, 24, 2016. `doi:10.1145/2738034`.

7 R. G. Downey and M. R. Fellows. *Parameterized complexity*. Monographs in Computer Science. Springer-Verlag, New York, 1999. `doi:10.1007/978-1-4612-0515-9`.

8 J. Flum and M. Grohe. *Parameterized complexity theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer-Verlag, Berlin, 2006.

9 Michael R. Garey and David S. Johnson. *Computers and intractability*. W. H. Freeman and Co., San Francisco, Calif., 1979. A guide to the theory of NP-completeness, A Series of Books in the Mathematical Sciences.

10 Masahiro Hachimori. *Combinatorics of constructible complexes*. PhD thesis, Tokyo University, 2000. URL: `https://pdfs.semanticscholar.org/81ee/09ed08ca6c7487e8e18639a7a05c110781c3.pdf`.

11 Allen Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.

12 Cynthia Hog-Angeloni and Wolfgang Metzler, editors. *Two-dimensional homotopy and combinatorial group theory*, volume 197 of *London Mathematical Society Lecture Note Series*. Cambridge University Press, Cambridge, 1993. `doi:10.1017/CBO9780511629358`.

13 Sergei Matveev. *Algorithmic topology and classification of 3-manifolds*, volume 9 of *Algorithms and Computation in Mathematics.* Springer, Berlin, second edition, 2007.

14 James R. Munkres. *Elements of algebraic topology.* Addison-Wesley Publishing Company, Menlo Park, CA, 1984.

15 Alexei D. Myasnikov, Alexei G. Myasnikov, and Vladimir Shpilrain. On the Andrews-Curtis equivalence. In *Combinatorial and geometric group theory (New York, 2000/Hoboken, NJ, 2001)*, volume 296 of *Contemp. Math.*, pages 183–198. Amer. Math. Soc., Providence, RI, 2002. `doi:10.1090/conm/296/05074`.

16 J. Paixão and J. Spreer. Random collapsibility and 3-sphere recognition, 2015. Preprint, 18 pages, 6 figures. `arXiv:1509.07607`.

17 Martin Tancer. Recognition of collapsible complexes is NP-complete. *Discrete Comput. Geom.*, 55(1):21–38, 2016. `doi:10.1007/s00454-015-9747-1`.

18 J. H. C. Whitehead. On incidence matrices, nuclei and homotopy types. *Ann. of Math. (2)*, 42:1197–1239, 1941. `doi:10.2307/1970465`.

19 J. H. C. Whitehead. Simple homotopy types. *Amer. J. Math.*, 72:1–57, 1950. `doi:10.2307/2372133`.

20 Perrin Wright. Group presentations and formal deformations. *Trans. Amer. Math. Soc.*, 208:161–169, 1975. `doi:10.2307/1997282`.

21 E. C. Zeeman. On the dunce hat. *Topology*, 2:341–358, 1964. `doi:10.1016/0040-9383(63)90014-4`.