

Simplification Strategies and Extremal Examples of Simplicial Complexes

vorgelegt von
M. Sc.
Davide Lofano

von der Fakultät II - Mathematik und Naturwissenschaften
der Technischen Universität Berlin
zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften
Dr.rer.nat.
genehmigte Dissertation

Promotionsausschuss:

Vorsitzender: Prof. Dr. Wilhelm Stannat

Gutachter: Priv.-Doz. Dr. Frank H. Lutz

Gutachter: Prof. Dr. Michael Joswig

Gutachter: Professor Paweł Dłotko

Tag der wissenschaftlichen Aussprache: 7. April 2022

Berlin 2022

Acknowledgements

Moving to Berlin has not been easy. Moving to a different country, with a different language and different customs means that it takes times to get used to it and to start enjoying it. If it was not enough, after a year and a half, when I thought that I was well settled, Covid and the first of many lockdowns started. Needless to say life has changed a lot for me like for everyone else and I probably would be a very different person if the pandemic never hit. However the experience has been amazing and even with a rough start I fell in love with the city and I don't plan to move out anytime soon. For this I have to really say thanks to all my friends that have made my years here, even with all the chaos happening all around, some of the best years of my life.

I am indebted to my supervisor Frank who has always been helpful, encouraging me all the time and always ready to answer questions. It has been a pleasure to work together and hopefully we will be able to continue still in the future.

Thanks to my parents that have always been supportive and never questioned my decision to move to Berlin and pursue a Ph.D.

Thanks to Eli who has definitely made the second lockdown a very pleasant experience instead of a depressing one and who has always been at my side especially in the last stressful months.

Abstract

Since the beginning of Topology, one of the most used approaches to study a geometric object has been to triangulate it. Many invariants to distinguish between different objects have been introduced over the years, the two most important surely being homology and the fundamental group. However, the direct computation of the fundamental group is infeasible and even homology computations could become computationally very expensive for triangulations with a large number of faces without proper preprocessing. This is why methods to reduce the number of faces of a complex, without changing its homology and homotopy type, are particularly of interest. In this thesis, we will focus on these simplification strategies and on explicit extremal examples.

The first problem tackled is that of sphere recognition. It is known that 3-sphere recognition lies in NP and in co-NP, and that d -sphere recognition is undecidable for $d \geq 5$. However, the sphere recognition problem does not go away simply because it is algorithmically intractable. To the contrary, it appears naturally in the context of manifold recognition so there is a clear need to find good heuristics to process the examples. Here, we describe an heuristic procedure and its implementation in `polymake` that is able to recognize quite easily sphericity of even fairly large simplicial complexes. At the same time we show experimentally where the horizons for our heuristic lies, in particular for discrete Morse computations, which has implications for homology computations.

Discrete Morse theory generalizes the concept of collapsibility, but even for a simple object like a single simplex one could get stuck during a random collapsing process before reaching a vertex. We show that for a simplex on n vertices, $n \geq 8$, there is a collapsing sequence that gets stuck on a d -dimensional simplicial complex on n vertices, for all $d \notin \{1, n-3, n-2, n-1\}$. In contrast, for $n \leq 7$ and d arbitrary or n arbitrary and $d \in \{1, n-3, n-2, n-1\}$ it is not possible to find a collapsing sequence for the simplex on n vertices that gets stuck at dimension d . Equivalently, and in the language of high-dimensional generalizations of trees, we construct hypertrees that are anticollapsible, but not collapsible.

As for a second heuristic for space recognition, we worked on an algorithmic implementation of simple-homotopy theory, introduced by Whitehead in 1939, where

not only collapsing moves but also anticollapsing ones are allowed. This provides an alternative to discrete Morse theory for getting rid of local obstructions. We implement a specific simple-homotopy theory heuristic using the mathematical software `polymake`. This implementation has the double advantage that we remain in the realm of simplicial complexes throughout the reduction; and at the same time, theoretically, we keep the possibility to reduce any contractible complex to a single point. The heuristic algorithm can be used, in particular, to study simply-connected complexes, or, more generally, complexes whose fundamental group has no Whitehead torsion. We shall see that in several contractible examples the heuristic works very well. The heuristic is also of interest when applied to manifolds or complexes of arbitrary topology. Among the many test examples, we describe an explicit 15-vertex triangulation of the Abalone, and more generally, $(14k + 1)$ -vertex triangulations of Bing's houses with k rooms, $k \geq 3$.

One of the classes of examples on which we run the heuristic are the 3-dimensional lens spaces, which are known to have torsion in their first homology. Using this examples (minus a facet), we managed to find 2-dimensional simplicial complexes with torsion and few vertices. Studying them we constructed sequences of complexes with huge torsion on few vertices. In particular, using Hadamard matrices we were able to give a quadratic time construction of a series of 2-dimensional simplicial complexes on $5n - 1$ vertices and torsion of size $\Theta(2^{n \log(n)})$. Our explicit series of 2-dimensional simplicial complexes improves a previous construction by Speyer with torsion growth in $\Theta(2^n)$, narrowing the gap to the highest possible asymptotic torsion growth in $\Theta(2^{n^2})$ proved by Kalai via a probabilistic argument.

Contents

1	Introduction	1
2	Definitions and notations	7
2.1	Simplicial homology	8
2.2	Collapsibility	9
2.3	Simple-homotopy theory	9
2.4	Delaunay triangulation that contains a Duncce Hat	10
2.5	Discrete Morse theory	13
2.6	Monotone discrete Morse functions	17
2.7	Bistellar flips	20
3	Frontiers of sphere recognition in practice	21
3.1	Overview	21
3.2	A heuristic sphere recognition scheme	23
3.3	Experiments and runtime comparisons	31
3.3.1	Recognizing random 3-spheres with <code>polymake</code> and <code>Regina</code>	31
3.3.2	Processing a census of 4-manifolds	33
3.3.3	Higher-dimensional random spheres	34
3.3.4	A collapsible 5-manifold which is not a ball	34
3.4	Limitations	36
3.4.1	General remarks	36
3.4.2	Akbulut–Kirby spheres	37
3.4.3	Contractible but non-collapsible subcomplexes	38
3.4.4	Iterated barycentric subdivisions	42
3.4.5	Other input	42
4	The worst way to collapse a simplex	45
4.1	Overview	45
4.2	Preliminaries	47
4.2.1	Alexander duality and the top dimensions	47
4.2.2	Hypertrees	49

4.3	Constructions	51
4.3.1	Explicit constructions for $n = 8$	51
4.3.2	Induction	53
4.4	Conclusions	57
5	Random simple-homotopy theory	59
5.1	Overview	59
5.2	Pure elementary expansions	62
5.3	Implementation of random simple-homotopy	63
5.4	Bistellar flips and artefacts	65
5.4.1	Selection of expansions and simplification of complexes	68
5.5	Classical examples	69
5.5.1	The Dunce Hat	69
5.5.2	The Abalone	70
5.5.3	Bing's House with two rooms	72
5.5.4	Bing's House with k rooms	73
5.6	Experiments on various topologies and substructures	75
5.6.1	Contractible, non-collapsible complexes	75
5.6.2	Submanifolds and non-manifold substructures in manifolds	76
5.6.3	Dimensionality reduction	80
5.6.4	Akbulut–Kirby 4-spheres	81
6	Hadamard matrix torsion	83
6.1	Overview	83
6.2	Preliminaries	85
6.2.1	A first triangulation procedure	86
6.2.2	Speyer's construction	88
6.2.3	Hadamard matrices	88
6.3	An improved triangulation procedure	90
6.3.1	A modified CW disc construction	91
6.3.2	Valid sequences	92
6.3.3	Triangulations $\text{HMT}(n)$ of the Hadamard examples $H(n)$	96
7	Conclusion and outlook	101

Chapter 1

Introduction

A standard task in topology is to simplify a given presentation of a topological space. In general, this task cannot be performed algorithmically: Even the simplest homotopic property, contractibility, is undecidable.

Its combinatorial equivalent, used in computational topology, is collapsibility, first introduced by Whitehead [Whi39]. In the setting of simplicial complexes, collapsibility is a completely combinatorial notion. It is easy to check that collapsibility implies contractibility, but the converse is not true in general. The smallest such example is the Dunce Hat [Zee64], a 2-dimensional contractible complex with triangulations on eight vertices which all are not collapsible. For smaller simplicial complexes on seven or fewer vertices, a result of [BD05] shows that contractibility and collapsibility are equivalent.

By studying contractible but not collapsible complexes in Chapter 3, we were able to construct a family of examples of 2-dimensional simplicial complexes, called *saw blade complexes* that in a certain way generalize the Dunce Hat with the following properties:

Theorem A (Theorem 3.10). *The following holds for k -bladed saw blade complexes SB^k :*

- (i) *The Dunce Hat SB^1 can be triangulated with 8 vertices.*
- (ii) *SB^2 can be triangulated with 9 vertices.*
- (iii) *SB^k can be triangulated with $3k$ vertices, for $k \geq 3$.*
- (iv) *Any triangulation of a saw blade complex is contractible, but non-collapsible.*

While collapsibility and contractibility are not equivalent in general, more subtly there even exist simplicial complexes which are collapsible, but for which it is possible to choose a sequence of elementary collapses after which one gets stuck at a nontrivial complex that is not collapsible. As a first small result in Chapter 2 we show that this is possible even in the case that the starting complex is a Delaunay

triangulation, which are always collapsible [ALS19], hereby answering a question of Edelsbrunner (private communication).

Theorem B (Theorem 2.3). *There is an 8-point Delaunay triangulation in \mathbb{R}^3 that collapses to a triangulation of the Dunce Hat with eight vertices. This example is smallest possible with respect to its number of vertices.*

Moreover, it is possible to get stuck even in the case of a single simplex. It is easy to see that the simplex on n vertices, denoted Δ_{n-1} , is collapsible. Indeed, one may choose a vertex and proceed dimension-wise, from larger to smaller dimensions, collapsing at each step the free maximal faces that do not contain the chosen vertex. However, [BL13a, CEK⁺] show that there exists a sequence of elementary collapses from the 7-simplex to a triangulation of the dunce hat, so from there no further collapses are possible even though the 7-simplex is collapsible. Experimental results presented in Table 3.5 suggests that this phenomenon is quite common in higher dimensions. Trying to understand it better, in Chapter 4, we answer the following question of Lutz (private communication): Starting from the simplex Δ_{n-1} , what are the dimensions d in which a collapsing sequence may get stuck?

Theorem C (Theorem 4.1). *For $n \geq 8$ and $d \notin \{1, n-3, n-2, n-1\}$, there exists a collapsing sequence of the simplex on n vertices which gets stuck at a d -dimensional complex on n vertices. In contrast, for $n \leq 7$ and d arbitrary or n arbitrary and $d \in \{1, n-3, n-2, n-1\}$ it is not possible to find a collapsing sequence of the simplex on n vertices which gets stuck at dimension d .*

In particular, we obtain the following corollary.

Theorem D (Theorem 4.2). *For every $n \geq 8$ and $d \notin \{1, n-3, n-2, n-1\}$, there exists a contractible d -complex on n vertices with no free faces. Moreover, this is best possible; for $n \leq 7$ or $d \in \{1, n-3, n-2, n-1\}$ every contractible d -complex on n vertices has a free face.*

Having at our disposal so many examples of contractible but not collapsible complexes it is natural to ask about a combinatorial way to prove contractibility. In [Whi39] Whitehead, apart from collapsibility, also introduced a discrete version of homotopy theory, called simple-homotopy theory, where not only elementary collapses are allowed but also their inverse moves, *elementary anticollapses* or *expansions*. Simple-homotopy is strictly stronger than collapsibility; and by a famous result of Whitehead, having the simple-homotopy type of a point is equivalent to being contractible [Whi39].

In Chapter 5, we propose a space recognition heuristic based on collapses in combination with certain expansions that we call *pure elementary expansion*. Our

randomized heuristic “Random Simple-Homotopy (RSHT)” has two advantages: First, all intermediate steps are always simplicial complexes; and second, at present we do not know of a single contractible complex for which our heuristics has probability zero to succeed in recognizing contractibility.

Algorithm: Random Simple-Homotopy (Ch. 5, p. 64)

Input: simplicial complex K
Output: simplified simplicial complex
while $\dim(X) \neq 0$ **and** $i < \text{max_step}$ **do**
 while X has free faces **do**
 (C): perform a random elementary collapse
 if $\dim(X) = d \neq 0$ **and** there are induced pure d -balls on $d + 2$ vertices **then**
 (E): perform a random pure elementary $(d + 1)$ -expansion
 (CC): perform an elementary collapse deleting the newly added $(d + 1)$ -face and one of its d -faces that was already in X
 else
 (S): perform (E) + (CC) on a d -facet with $d + 1$ vertices
 $i++$
return X

Our algorithm, in low dimension, is equivalent to performing bistellar flips if the triangulation is a d -manifold.

Theorem E (Theorem 5.8). *Let X be a triangulation of a d -manifold M with $d \leq 6$. Any pure elementary $(d + 1)$ -expansion followed by collapses (as long as free faces are available) induces a bistellar flip on X .*

This implies a sort of manifold stability of our heuristic. If we run RSHT on a simplicial complex X and at some point we reach a simplicial complex X' that triangulates a d -manifold with $d \leq 6$, then from then on, whenever there are no free faces in the further run of RSHT, the respective temporary complex \tilde{X} is a d -manifold as well, and \tilde{X} is bistellarly equivalent to X' .

The algorithm can, in particular, be used to study simply-connected complexes, or, more generally, complexes whose fundamental group has no Whitehead torsion. We shall see that in several contractible examples the heuristics works very well. For example, one pure elementary 3-expansion suffices to reduce the 8-vertex triangulation of the Dunce Hat from Figure 5.3a to a single vertex, and RSHT is able to reduce the Bing’s house triangulation constructed in Chapter 5 to a point by means of five (successive) expansions. As another list of examples of contractible

but not collapsible complexes, we also constructed a generalization of Bing's house adding more rooms.

Theorem F (Theorem 5.12). *For any $k \geq 3$, Bing's House with k rooms, $BH(k)$, can be formally deformed to a point using only six pure expansions.*

We later applied the Random Simple-Homotopy algorithm to lens spaces in order to find simplicial complexes with few vertices and huge torsion. Using these experiments as a starting point, we were able in Chapter 6 to describe a triangulation procedure. Starting from any integer matrix it is possible to construct a simplicial complex with the given matrix as the boundary matrix of its first homology using a bounded number of vertices.

Theorem G (Proposition 6.8). *Given an $(m \times n)$ -matrix M and $A = (\alpha_i)$ its Smith Normal Form, there is a 2-dimensional simplicial complex K on $V(K)$ vertices with*

$$V(K) \leq 2n + m + 1 + \frac{3}{2} \sum_{i,j} |M_{ij}|.$$

Furthermore,

$$H_1(K) = \mathbb{Z}/\alpha_1\mathbb{Z} \times \cdots \times \mathbb{Z}/\alpha_r\mathbb{Z} \times \mathbb{Z}^{n-r}.$$

Focusing on Hadamard matrices as particular matrices, we were able to improve the triangulation procedure used in the proof of the previous statement.

Theorem H (Theorem 6.22). *For each $n = 2^k$, $k \geq 1$, there is a \mathbb{Q} -acyclic 2-dimensional simplicial complex $\text{HMT}(n)$ with face vector*

$$f(\text{HMT}(n)) = (5n - 1, 3n^2 + 9n - 6, 3n^2 + 4n - 4)$$

and $H_*(\text{HMT}(n)) = (\mathbb{Z}, T(\text{HMT}(n)), 0)$. The torsion in first homology is given by

$$H_1(\text{HMT}(n)) = T(\text{HMT}(n)) = (\mathbb{Z}_2)^{\binom{k}{1}} \times (\mathbb{Z}_4)^{\binom{k}{2}} \times \cdots \times (\mathbb{Z}_{2^k})^{\binom{k}{k}},$$

where $|T(\text{HMT}(n))| = n^{n/2} \in \Theta(2^{n \log n})$. Furthermore, the examples $\text{HMT}(n)$ can be constructed algorithmically in quadratic time $\Theta(n^2)$.

Our explicit series of 2-dimensional simplicial complexes with linearly many vertices with torsion growth in $\Theta(2^{n \log n})$ improves a previous quadratic time construction by Speyer [Spe] with torsion growth in $\Theta(2^n)$, narrowing the gap to the highest possible asymptotic torsion growth in $\Theta(2^{n^2})$ proved by Kalai [Kal83] via a probabilistic argument.

In 1998, Forman introduced another way to study contractibility combinatorially. His discrete Morse theory [For98, For02], a variation of the more known

smooth Morse theory, is a tool to reduce simplicial complexes using a mix of collapses and facet deletions. The advantage is that all simplicial complexes (contractible or not) can now be reduced to a vertex, possibly by using a relatively large number of facet deletions. The drawback is that even if one starts with a simplicial complex, the intermediate steps in the reduction sequence are typically non-regular CW complexes, and thus harder to handle. By only focusing on the count of facet deletions (the so-called “discrete Morse vector”) it is possible to use randomness to produce fast implementations [BL14], but at the cost of failing to recognize many contractible complexes. The algorithm is also implemented in `polymake` [GJ00] as `random_discrete_morse`. In Chapter 2, we introduce the (revised) definition of *monotone discrete Morse functions*, as asked for by Adiprasito, Benedetti and Lutz in [ABL17] as a way to describe the acyclic matchings obtained when using the heuristic.

Theorem I (Theorem 2.14). *Let X be a simplicial complex and v the lexicographically-largest among the optimal discrete Morse vectors of X . Then there exists a monotone discrete Morse function g that has v as discrete Morse vector.*

The same `random_discrete_morse` algorithm is going to be one of the main building blocks for our sphere recognition heuristic described in Chapter 3.

Algorithm: Sphere Recognition Heuristics (Ch. 3, p. 24)

Input: Hasse diagram of combinatorial d -manifold X , where $d \geq 3$

Output: Semi-Decision: Is X PL homeomorphic to \mathbb{S}^d ?

- (1) compute homology
 if *homology not spherical* **then return** NO
 for N rounds **do**
 - (2) compute random discrete Morse vector
 if *discrete Morse vector is spherical* **then return** YES
 for N' rounds **do**
 - (3) perform random bistellar flip or edge contraction
 if *boundary of simplex is reached* **then return** YES
 - (4) compute and simplify presentation of fundamental group π_1
 if *presentation is found to be trivial and $d \neq 4$* **then return** YES
 if *presentation is found to be non-trivial* **then return** NO
 return UNDECIDED
-

To tell whether a given space is homeomorphic to the sphere in a given dimension is a basic problem in computational topology. However, this problem turns out to be undecidable in dimension higher than four, lies in the complexity class

NP for 3-spheres and is still an open problem for 4-spheres. Summing up, we do not know of any efficient algorithms for d -sphere recognition in the relevant dimensions $d \geq 3$. However, the sphere recognition problem does not go away simply because it is algorithmically intractable. To the contrary it appears naturally in the context of manifold recognition so it is natural to try to find good heuristics to try and treat it. We will describe a sphere recognition heuristic and its implementation in `polymake`. The heuristic is a combination of already known methods, like discrete Morse theory and bistellar flips, combined together to create a single algorithm able to detect quite easily sphericity of even fairly large simplicial complexes.

We have run comprehensive computational experiments to show its performance and limits. In particular, we show experimentally where the horizons for our heuristic lies, especially for discrete Morse computations, which has implications to homology computations. Among the many experiments, we ran our heuristic on a census of 4-manifolds, encoded as pseudo-simplicial complexes comprising up to six maximal cells, provided by `Regina` [BBP⁺17]. Our success rate is 99.5%, with our heuristic failing on only 1,954 of these combinatorial 4-manifolds.

Theorem J (Theorem 3.8). *Among the 441,286 combinatorial types of combinatorial 4-manifolds arising from up to six maximal cells 91.5% are spheres and 8.0% are non-spheres.*

As I have tried to explain in the previous pages, the aim of this thesis is to focus on simplification strategies and on explicit extremal examples of simplicial complexes to explore the bounds and horizons of current heuristic and algorithmic approaches, as well as introducing new heuristics and implementations to expand our tool box for the computational study of topological spaces. It is now time to get started.

Chapter 2

Definitions and notations

We will explore in this chapter the concepts that are going to be central in this thesis, trying to give an adequate but succinct explanation. In particular, we start with some basic definitions from algebraic topology and we then move to simple-homotopy theory and discrete Morse theory that are going to be among our main topics of interest. Some of the definitions and explanations are taken freely from the introductory and preliminary sections of the articles [JLLT21, LN21, BLLL21, LL21] of which I was a co-author and that will be foundational for the following chapters.

The first definition that we introduce will be the building block of the entire thesis.

Definition 2.1. A *simplicial complex* X on the vertex set V is a family of subsets of V (called *simplices* or *faces* of X) such that every subset of a simplex from X is in X .

We now start with some notions from combinatorial topology, much more on these notions can be found in a standard reference such as [Hat02]. If X is a finite simplicial complex and τ is a face of X , the *dimension of* τ , denoted $\dim \tau$, is $|\tau| - 1$, and the *dimension of* X , denoted $\dim X$, is the maximum dimension of any face of X . An inclusion maximal face is called a *facet*. A complex X is said to be *pure dimensional* if all of its facets are of the same dimension. The *j -skeleton* of X , denoted $X^{(j)}$, is the subcomplex of X given by every face of X of dimension at most j ; we say that the j -skeleton of X is complete provided $X^{(j)}$ has all possible j -dimensional faces on the vertices of X .

At certain points we will also need another type of complex. A *CW complex* is constructed in an inductive way level by level. We start with a discrete set, whose points are the vertices (or 0-cells) of the complex. Inductively we add the n -dimensional cells by attaching the boundary of a new n -dimensional ball to the

$(n - 1)$ -dimensional cells already added to the complex in the previous inductive steps.

In particular, a 2-dimensional face of a CW complex can be attached to any (finite) number of 1-dimensional faces, while a 2-dimensional face of a simplicial complex is always a triangle so attached to exactly three edges.

Notation. Before going on, we introduce some notations that we use throughout the thesis:

- ▷ We use parentheses when talking about a set of vertices $V = (x_0, x_1, \dots, x_n)$.
- ▷ We use square brackets when talking about a face of a simplicial complex, e.g. $\sigma = [x_0, x_2]$. With a little abuse of notation, if $x_k \notin \sigma$, we will denote by $[x_k, \sigma]$ the face with vertex set x_k and the vertices of σ , and if τ, σ are both faces $[\tau, \sigma]$ is the face with vertex set the union of the vertices of σ and τ .
- ▷ We use curly brackets to denote a simplicial complex; we will use the same notation both to list all the faces or only the facets, which one we are using will be clear from the context; e.g. $X = \{\emptyset, [x_0], [x_1], [x_0, x_1]\}$ or $X = \{[x_0, x_1]\}$.

2.1 Simplicial homology

We are now ready to define simplicial homology, though actually we define *reduced* simplicial homology, as this makes things easier to analyze and work with in our context. We refer again, as in the previous section, to [Hat02] for a more detailed analysis.

To begin, we need a simplicial complex X and an abelian group of coefficients R ; most commonly R will be taken to be \mathbb{Z} , \mathbb{Q} , or a prime-order finite field.

For each integer i the set of i -chains, denoted C_i , is the family of formal linear combinations of faces of X of dimension i with coefficients in R (and $C_{-1} = R$). For each $i \geq 0$, one defines the i -th boundary map to be the linear map $\partial_i : C_i \rightarrow C_{i-1}$ given by sending the generator $[x_0, \dots, x_i]$ to $\sum_{j=0}^i (-1)^j [x_0, \dots, \hat{x}_j, \dots, x_i]$, where $[x_0, \dots, \hat{x}_j, \dots, x_i]$ denotes the face obtained by deleting x_j from $[x_0, \dots, x_i]$. The i -th homology group with coefficients in R , denoted $H_i(X; R)$, is defined to be the abelian group given by the quotient $\ker(\partial_i) / \text{Im}(\partial_{i+1})$ (it is routine to check that this quotient is well-defined). Without specifying the ring, throughout this thesis, the i -th homology group of X , $H_i(X)$, is assumed to be $H_i(X; \mathbb{Z})$.

A standard fact about homology is that $H_0(X) = 0$ if and only if X is path-connected. In fact, $H_0(X)$ is the free abelian group with rank equal to the number of connected components of X minus one. It is also clear from the definition that homology groups vanish in dimensions larger than the dimension of X .

A complex is said to be R -acyclic if $H_i(X; R) = 0$ for all $i \geq 0$. Regarding the cases $R = \mathbb{Q}$ and $R = \mathbb{Z}$, we have that if X is \mathbb{Z} -acyclic, then it is necessarily \mathbb{Q} -acyclic, however the converse need not hold. A \mathbb{Q} -acyclic complex will necessarily

have all homology groups finite, but they will not necessarily be trivial; this follows by the standard universal coefficient theorem for homology. For 1-dimensional complexes (i.e., for graphs), however \mathbb{Z} -acyclic and \mathbb{Q} -acyclic are equivalent and hold precisely when the graph is a tree. A complex is said to have *torsion* in homology in dimension d if its d -dimensional homology group has a finite part. The most common example is the real projective plane $\mathbb{R}P^2$ whose first homology group is \mathbb{Z}_2 , so it is said to have 2-torsion in dimension one.

2.2 Collapsibility

Another standard notion in computational topology is that of collapsibility, first introduced by Whitehead [Whi39]. In the setting of simplicial complexes, collapsibility is a completely combinatorial notion. For a simplicial complex X , a nonempty face of X is said to be *free* provided that it is a proper subset of exactly one other face in X . An *elementary collapse* of X is the process of removing a free face and the unique face properly containing it.

Generally speaking, elementary collapses preserve many topological invariants that one might want to compute for some simplicial complex. The strategy, from a computational topology perspective, therefore becomes to use elementary collapses to reduce some initial complex X to a smaller complex Y , where the relevant topological invariants may be more easily computed. The best-case scenario for such an approach is when X is *collapsible*, i.e., if there exists a sequence of elementary collapses that reduce the complex to a single vertex.

The notion of collapsibility is contrasted with the purely topological notion of contractibility. A complex is said to be *contractible* if it is homotopically equivalent to a single point. We have already said in the introduction that collapsibility implies contractibility, but the converse is not true in general. The smallest such example being the Dunce Hat [Zee64], which we will discuss in more detail below, a 2-dimensional contractible complex with triangulations on eight vertices which all are not collapsible.

2.3 Simple-homotopy theory

As a generalization of collapsibility, Whitehead introduced a discrete version of homotopy theory, called *simple-homotopy theory* [Whi39].

Recall that an *elementary collapse* is a deletion from a simplicial complex of a free face and the only facet containing it. Elementary collapses are deformation retracts and thus maintain the homotopy type; the same is true for their inverse moves, *elementary anticollapses*. Two simplicial complexes are of the same *simple-*

homotopy type if they can be transformed into one another via some sequence of collapses and anticollapses, called a *formal deformation* [Whi39].

Equivalently, two simplicial complexes are of the same simple-homotopy type if there exists a third complex that can be reduced to both the original ones via suitable sequences of elementary collapses [HAM93, p. 13]. The size of the third complex (or, using the former definition, the length of the formal deformation) cannot be bounded a priori, because the simple-homotopy type cannot be decided algorithmically. In fact, by a famous result of Whitehead, having the simple-homotopy type of a point is equivalent to being contractible [Whi39] and thus undecidable.

While collapsibility and contractibility are not equivalent in general, more subtly, there even exists simplicial complexes which are collapsible, but for which it is possible to choose a sequence of elementary collapses after which one gets stuck at a nontrivial complex that is not collapsible. Formally, we say that a collapsing sequence on a complex X gets stuck at a complex Y , if the collapsing sequence reduces X to Y and Y has no free faces. More broadly, we say that a collapsing sequence on X gets stuck at dimension d if the collapsing sequence gets stuck at some d -dimensional complex.

2.4 Delaunay triangulation that contains a Dunce Hat

We now give an example of an 8-point Delaunay triangulation that contains as a subcomplex a triangulation of the Dunce Hat and moreover collapses to it. Given a set of points P in general position in \mathbb{R}^n , a *Delaunay triangulation* is a triangulation such that no point in P is inside the circumscribed sphere of any facet of the triangulation.

Delaunay triangulations are always collapsible [ALS19]. Edelsbrunner asked, whether a Delaunay triangulation, possibly of small size, can be collapsed onto a non-collapsible complex.

The famous Dunce Hat [Zee64] is known to have triangulations with eight vertices, one of which is shown in Figure 2.1. These triangulations are contractible, but not collapsible, since there are no free edges. Having eight vertices they are smallest possible; cf. Bagchi and Datta [BD05]. Here, we show that one of the Dunce Hat triangulations with eight vertices can be found as a subcomplex of a collapsible 8-point Delaunay triangulation that collapses to it.

Benedetti and Lutz [BL13a] presented an 8-vertex triangulation of the Dunce Hat (see Figure 2.1) that can be found as a subcomplex of the 4-polytope (of abstract combinatorial type) S_{32} with eight vertices described in [GS67]. Due to

Firsching [Fir17], this polytope is known to be inscribable, meaning that it can be realized as the convex hulls of points on the unit sphere. In the table below we see a possible explicit choice of such points in \mathbb{R}^4 from the webpage [Fir].

0	-35640/102601	-17424/102601	-91476/102601	24193/102601
1	-64/105	64/105	-16/35	-23/105
2	-1690/4519	-650/4519	-1300/4519	-3931/4519
3	-378/3247	-756/3247	2058/3247	2365/3247
4	-14700/87949	-50400/87949	70560/87949	-251/87949
5	33800/76129	-27040/76129	20800/76129	-59071/76129
6	1	0	0	0
7	10164/31309	-17424/31309	-7920/31309	22597/31309

It has been known for a long time that there is a close correlation between inscribed polytopes and Delaunay triangulations. In particular, we are going to use the following theorem by Gonska.

Theorem 2.2. [Gon13, Proposition 0.3.13] *Let $d > 1$, $P \subset \mathbb{R}^d$ be a d -polytope and $V \sqcup \{N\}$ be its vertex set. Let T be the image of a vertex projection of P from N to any $(d-1)$ -dimensional subspace of \mathbb{R}^d . If P is inscribed and simplicial, then T is a Delaunay triangulation.*

Using all the above information it is now easy to prove our result:

Theorem 2.3. *There is an 8-point Delaunay triangulation in \mathbb{R}^3 that collapses to a triangulation of the Duncce Hat with eight vertices. This example is smallest possible with respect to its number of vertices.*

Proof. First of all we start with the simplicial boundary complex of the above geometric realization of S_{32} which gives a triangulation with vertices on the unit sphere S^3 in \mathbb{R}^4 . As we already know, this triangulation contains the Duncce Hat. The idea is now to add a ninth point very close to one of the vertices such that it is still possible to find an 8-point triangulation of the Duncce Hat in the newly obtained polytope as the convex hull of the nine points. If we manage to do so, then, projecting from the only vertex not contained in the Duncce Hat triangulation to the first three coordinates and using Theorem 2.2 we obtain the result.

We can add a point close to the North Pole $N = (1, 0, 0, 0)$. If we choose the point in the right direction (we have chosen the point $(0.98, 0.11, -0.11, -0.11)$), we obtain that a Duncce Hat is now contained in the subcomplex of the polytope that does not contain the North Pole. We can then use Theorem 2.2 on the set $V \sqcup \{N\}$ where V are the vertices that contains the Duncce Hat and N is the North Pole. Applying to V the stereographic projection with respect to the North Pole we obtain the following points in \mathbb{R}^3 , where the new vertex is in position (6).




Figure 2.1: Dunce Hat triangulation

0	-0.126	-0.6617	0.175
1	0.3787	-0.284	0.1361
2	-0.1047	-0.2094	-0.6331
3	0.2086	0.5677	0.6524
4	-0.491	0.6874	-0.0024
5	-0.6388	0.4914	-1.3955
6	5.5	-5.5	-5.5
7	-0.824	-0.3746	1.0687

Due to the far-away point (6), the configuration is stretched out and thus hard to zoom in and visualize. However, using the program TetGen [Si15] on these eight points we obtain their Delaunay triangulation with 17 tetrahedra in \mathbb{R}^3 :

$$\begin{aligned}
 &[0, 1, 2, 4], \quad [0, 1, 3, 4], \quad [0, 1, 3, 7], \quad [0, 1, 2, 6], \quad [0, 1, 6, 7], \\
 &[0, 2, 4, 7], \quad [0, 2, 5, 6], \quad [0, 2, 5, 7], \quad [0, 3, 4, 7], \quad [0, 5, 6, 7], \\
 &[1, 2, 3, 4], \quad [1, 2, 3, 5], \quad [1, 2, 5, 6], \quad [1, 3, 5, 6], \quad [2, 3, 4, 5], \\
 &[2, 4, 5, 7], \quad [1, 3, 6, 7].
 \end{aligned}$$

This complex contains as a subcomplex the above 8-vertices triangulation of the Dunce Hat.

Moreover, using the program DiscreteMorse [BL14], we obtain that the Delaunay triangulation actually collapses to the triangulation of the Dunce Hat in Figure 2.1, for example, through the following series of elementary collapses. Here each pair is a matching pair in the acyclic matching, so the faces of the Dunce Hat are exactly the ones of our Delaunay triangulation not appearing below.

$$\begin{aligned}
 &([3, 5, 6], [1, 3, 5, 6]), ([1, 5, 6], [1, 2, 5, 6]), ([3, 6, 7], [1, 3, 6, 7]), \\
 &([3, 4, 7], [0, 3, 4, 7]), ([2, 5, 6], [0, 2, 5, 6]), ([0, 3, 7], [0, 1, 3, 7]), \\
 &([0, 4, 7], [0, 2, 4, 7]), ([0, 2, 4], [0, 1, 2, 4]), ([0, 1, 2], [0, 1, 2, 6]),
 \end{aligned}$$

$([0, 1, 6], [0, 1, 6, 7]), ([0, 6, 7], [0, 5, 6, 7]), ([0, 5, 7], [0, 2, 5, 7]),$
 $([2, 5, 7], [2, 4, 5, 7]), ([2, 4, 5], [2, 3, 4, 5]), ([2, 3, 5], [1, 2, 3, 5]),$
 $([1, 2, 3], [1, 2, 3, 4]), ([1, 3, 4], [0, 1, 3, 4]), ([2, 3], [2, 3, 4]),$
 $([3, 6], [1, 3, 6]), ([3, 7], [1, 3, 7]).$

For the second part of the theorem, the above mentioned result of Bagchi and Datta [BD05] states that all contractible simplicial complexes with up to seven vertices are collapsible. Thus, at least 8-vertices are needed in our construction. \square

2.5 Discrete Morse theory

As another possible generalization of collapsibility, we recall here the main concepts of Forman's discrete Morse theory [For98, For02]. We follow the point of view of Chari [Cha00] and the book by Kozlov [Koz08], mainly using acyclic matchings instead of discrete Morse functions.

The *Hasse diagram* of a simplicial complex X is a directed graph G with one node per face of X and a directed arc (σ, τ) if the face σ is contained in τ and $\dim \tau = \dim \sigma + 1$.

Let P be the poset of the faces of X and let denote by E the set of edges of G . Given a subset M of E , we can orient all edges of G in the following way: an edge $(\sigma, \tau) \in E$ is oriented from σ to τ if the pair does not belong to M , in the opposite direction otherwise. Denote this oriented graph by G_M .

Definition 2.4 (Acyclic matching [Cha00]). A *matching* on G is a subset $M \subseteq E$ such that every face of X appears in at most one edge of M . A matching M is *acyclic* if the graph G_M has no directed cycle.

Given a matching M on G , an *alternating path* is a directed path in G_M such that two consecutive edges of the path do not belong both to M or both to $E \setminus M$. The faces of X that do not appear in any edge of M are called *critical* with respect to the matching M . The *discrete Morse vector* (c_0, c_1, \dots, c_d) of an acyclic matching counts the critical faces per dimension; and X is homotopy equivalent to a CW complex with c_i cells in dimension $0 \leq i \leq d$ as stated in the following theorem, which is usually considered the main theorem of discrete Morse theory.

Theorem 2.5 ([For98, Cha00]). *Let X be a simplicial complex, let P be its poset of faces, and G the Hasse diagram of P . If M is an acyclic matching on G , then X is homotopy equivalent to a CW complex X_M , called the Morse complex of M , with cells in dimension-preserving bijection with the critical cells of X . Furthermore, if the critical cells form a subcomplex X_c of X , then there exists a sequence of elementary collapses collapsing X to X_c .*




Figure 2.2: Acyclic matching in the Hasse diagram of \mathbb{RP}_6^2 obtained with `polymake` [GJ00]. The three unmatched cells are marked.

One key application of this theorem is in computational topology. Indeed it is often the case that discrete Morse theory gives us a way to find a homotopy equivalence between a given simplicial complex and a much smaller CW complex on which the homology groups are easier to compute.

Let \mathbb{F} be some field. A discrete Morse vector is \mathbb{F} -*perfect* for X if $c_i = \beta_i(X; \mathbb{F})$ for $0 \leq i \leq d$, where β_i are the Betti numbers.

Example 2.6. Figure 2.2 shows an acyclic matching for \mathbb{RP}_6^2 with three critical cells. The corresponding discrete Morse vector $(1, 1, 1)$ is \mathbb{Z}_2 -perfect. The real projective plane admits a CW-complex structure with one 0-cell, one 1-cell and one 2-cell.

Finally, we recall a standard tool to construct acyclic matchings that is going to be very useful in Chapter 4.

Theorem 2.7 (Patchwork theorem [Koz08, Theorem 11.10]). *Let P be the poset of faces of a simplicial complex X and let $\varphi: P \rightarrow Q$ be a poset map. For all $q \in Q$, assume to have an acyclic matching $M_q \subseteq E$ that involves only elements of the subposet $\varphi^{-1}(q) \subseteq P$. Then the union of these matchings is itself an acyclic matching on P .*

Another way to look at acyclic matchings is through the notion of collapsibility. As already said, an *elementary collapse* is simply the removal of two simplices σ and τ such that

- ▷ $\dim \sigma = \dim \tau - 1$,
- ▷ the only simplex containing τ is τ itself,
- ▷ the only simplices containing σ are σ and τ , in which case σ is called a *free face*.

We say that a complex is *collapsible* if, through a series of elementary collapses, it can be reduced to a single vertex. As a weaker notion, we say that a complex is *d-collapsible* provided it can be reduced to a complex of dimension less than d by a sequence of elementary collapses.

An *elementary anticollapse* [Coh12], sometimes also called *expansion*, is the dual operation, i.e., the addition of two simplices σ' and τ' such that

- ▷ $\dim \tau' = \dim \sigma' + 1$,
- ▷ σ' and τ' are not in X ,
- ▷ the only facet of τ' not contained in X is σ' .

If X is on n vertices, we say that it is *anticollapsible* if, through a series of elementary anticollapses, it can be expanded to the simplex on n vertices. As a weaker notion a complex is *d-anticollapsible* if there is a sequence of anticollapses so that the resulting complex has complete d -skeleton. It should be noted that, while it is always possible to perform an elementary anticollapse that adds a new vertex, we are prohibiting these moves while talking about anticollapsibility.

The combinatorial encoding of a set of collapses is best provided by a matching consisting of the corresponding collection of pairs of cells (σ, τ) , which it is easily seen to be an acyclic matching. In fact, a standard randomized algorithm to find an acyclic matching is the Random Discrete Morse algorithm of Benedetti and Lutz [BL14] which uses elementary collapses. This algorithm is also implemented in `polymake` [GJ00] as `random_discrete_morse`.

The idea is to perform randomly chosen elementary collapses when possible. If there are no free faces, then a top dimensional facet is chosen to be critical at random and removed from the complex. This process continues until the remaining complex is zero-dimensional. In this case, it only consists of vertices, all of which are declared critical.

The algorithm returns the number of critical cells in each dimension; i.e., the discrete Morse vector. A complex will have a discrete Morse vector $(1, 0, \dots, 0)$ with respect to some acyclic matching if and only if it is collapsible, so the `random_discrete_morse` algorithm may be used to verify collapsibility.

In at least one particular case the `random_discrete_morse` algorithm may also be used to verify that a complex is not collapsible. If a complex is d -dimensional and the `random_discrete_morse` algorithm returns a Morse vector with at least one critical cell in dimension d , then the complex is not collapsible. In fact, in such a case the complex is not even d -collapsible. This is implied by the following definition and lemma; cf. also [Tan16].

Definition 2.8. A (not necessarily pure) d -dimensional simplicial complex L is called a *core* provided that every $(d - 1)$ -dimensional face of L is contained in at least two d -dimensional faces of L .

Lemma 2.9. *Let X be a d -dimensional complex. X is d -collapsible if and only if X does not contain a d -dimensional core.*

Proof. Suppose that X is not d -collapsible, then there exists some sequence of collapses that results in a complex X' that is still d -dimensional, but has no further d -collapses. In this complex every $(d - 1)$ -dimensional face is contained in either zero d -dimensional faces, or in at least two d -dimensional faces. Thus the pure d -dimensional part of this complex, that is the subcomplex whose facets are exactly the d -dimensional faces, is a core.

Conversely, suppose that X is d -collapsible and contains a d -dimensional core Y . Then there exists a sequence of collapses of removing pairs (σ, τ) with σ of dimension $d - 1$ and τ of dimension d that reduce X to a $(d - 1)$ -dimensional complex. Let (σ, τ) be the first such pair in the sequence with $\tau \in Y$. Since Y is a subcomplex, σ also belongs to Y . At the moment we collapse (σ, τ) , σ has degree 1 in the current complex, and hence it has degree 1 in Y . But then by definition of a core, σ is contained in $\tau' \neq \tau$ so that $\tau' \in Y$. It follows, however, that τ' must have been removed from X before τ , but this contradicts the choice of τ . \square

If a complex is d -dimensional and the `random_discrete_morse` algorithm returns a discrete Morse vector with at least one critical cell in dimension d , then the complex contains a d -dimensional core. Indeed, this means that the random collapsing sequence reached a point where the complex was d -dimensional but had no free faces. By Lemma 2.9 this means the complex is not d -collapsible.

The last definition that we give here is that of a *discrete Morse function*, which is how Forman originally introduced the theory in [For98].

A map $f : X \rightarrow \mathbb{R}$ which assigns a real number to each face of X is a *discrete Morse function* if for every k -face σ of X we have

$$\begin{aligned} \#\{\tau \in X \mid f(\tau) \leq f(\sigma), \sigma \subset \tau, \dim \tau = \dim \sigma + 1\} &\leq 1 \quad \text{and} \\ \#\{\rho \in X \mid f(\rho) \geq f(\sigma), \rho \subset \sigma, \dim \rho = \dim \sigma - 1\} &\leq 1. \end{aligned} \tag{2.1}$$

Note that only one of the sets of Equation (2.1) can be non-empty; cf. [For98]. A k -face is *critical* with respect to f if both sets in Equation (2.1) are empty; the non-critical faces are *regular*, and they form an acyclic matching in the Hasse diagram of X ; cf. [Cha00]. In this sense the acyclic matchings form equivalence classes of discrete Morse functions.

2.6 Monotone discrete Morse functions

In [ABL17, Definition 3.5] *monotone discrete Morse functions* were introduced to describe functions that can be obtained using the `random_discrete_morse` algorithm. In the following, we give a slightly different definition and discuss resulting properties.

Definition 2.10. A *monotone discrete Morse function* on a simplicial complex X is a map $f : X \rightarrow \mathbb{Z}$ satisfying the following six axioms:

- (i) if $\sigma \subseteq \tau$, then $f(\sigma) \leq f(\tau)$;
- (ii) the cardinality of $f^{-1}(q)$ is at most 2 for each $q \in \mathbb{Z}$;
- (iii) if $f(\sigma) = f(\tau)$, then either $\sigma \subseteq \tau$ or $\tau \subseteq \sigma$;
- (iv*) if $\dim \sigma < \dim \tau$, then $f(\sigma) \leq f(\tau)$ or there is a face τ' with $\dim \tau' = \dim \tau$ and $f(\sigma) = f(\tau')$;
- (v) $f(X) = [0, m] \cap \mathbb{Z}$, for some $m \in \mathbb{N}$;
- (vi) for any *critical face* Δ (that is, a face such that $f(\sigma) \neq f(\Delta)$ for each face $\sigma \neq \Delta$), the complex $\{\sigma \in X \text{ s.t. } f(\sigma) \leq f(\Delta)\}$ has no free $(\dim \Delta - 1)$ -dimensional face.

Note that the set defined in condition (vi) is indeed a simplicial complex due to (i). The conditions (i)–(iii), (v), and (vi) of Definition 2.10 coincide with the respective conditions in [ABL17, Definition 3.5]. The original condition (iv) of [ABL17, Definition 3.5] is contradictory if $\sigma = \sigma' = \tau'$, but can be resolved by requiring that $\sigma \subset \tau$ instead of $\sigma \subseteq \tau$. However, in Definition 2.10 it is replaced by the slightly stronger requirement (iv*).

Definition 2.10 can further be strengthened without restricting its generality.

Definition 2.11. A *normalized monotone discrete Morse function* on a non-empty simplicial complex X is a surjective map $f : X \rightarrow \{1, \dots, m\}$ for some $m \in \mathbb{N}$ satisfying the following five axioms for $\sigma, \tau \in X$:

- (i) $f^{-1}(q) \in \{1, 2\}$ for each $q \in \{1, \dots, m\}$;
- (ii) if $\sigma \subseteq \tau$, then $f(\sigma) \leq f(\tau)$;
- (iii) if $f(\sigma) = f(\tau)$, then either $\sigma \subseteq \tau$ or $\tau \subseteq \sigma$;
- (iv) if $\dim \sigma < \dim \tau$, then $f(\sigma) \leq f(\tau)$ or there is a face τ' with $\dim \tau' = \dim \tau$ and $f(\sigma) = f(\tau')$;

- (v) for any *critical face* Δ (that is, a face such that $f(\sigma) \neq f(\Delta)$ for each face $\sigma \neq \Delta$), the complex $\{\sigma \in X \text{ s.t. } f(\sigma) \leq f(\Delta)\}$ has no free $(\dim \Delta - 1)$ -dimensional face.

The definition in particular implies that the critical faces correspond to fibers of f of cardinality one while the matching pairs correspond to the fibers of cardinality two. In the following, we will always refer to normalized monotone discrete Morse functions even if we sometimes omit the normalized for better readability.

Proposition 2.12. *For any acyclic matching in the Hasse diagram of a simplicial complex X there is a surjective function $f : X \rightarrow \{1, \dots, m\}$ that induces the same matching and in addition satisfies the first four conditions of Definition 2.11.*

Proof. Given an acyclic matching for a complex X , let us consider the oriented Hasse diagram of X , where an edge is oriented from the higher-dimensional face to the smaller one if the pair is in the acyclic matching and in the opposite direction otherwise.

The idea is to give a total order to the set $\{\text{critical faces}\} \cup \{\text{matching pairs}\}$ and then construct f by giving to the smallest element the value 1, to the next element the value 2, and so on. We will proceed by dimension d of the complex X .

First of all we consider the critical d -dimensional faces of X and we give them any total order. We then consider the Hasse diagram restricted to the d - and $(d - 1)$ -dimensional faces that we have not yet ordered. Due to the fact that the matching is acyclic it follows that there exists a $(d - 1)$ -dimensional free face that is paired in the matching with a d -dimensional face. We then impose a total order on the matching pairs following some admissible collapsing sequence, i.e., at any point a matching pair is chosen so that the corresponding $(d - 1)$ -face is free in the current complex.

The remaining not-ordered faces form a $(d - 1)$ -dimensional complex, and we repeat the procedure by first giving a total ordering to the critical $(d - 1)$ -faces and then to pairs of $(d - 2)$ - and $(d - 1)$ -faces, and so forth. It is an easy check that a function f constructed this way satisfies the first four axioms of Definition 2.11. \square

Remark 2.13. Every function that satisfies the first four conditions of Definition 2.11 is clearly a discrete Morse function. In the following we will refer to functions that satisfy these conditions as *good* discrete Morse functions. The fifth condition instead is, in general, really restrictive. See Figure 2.3 for a 6-gon with a (normalized) monotone discrete Morse function.

Adiprasito, Benedetti and Lutz asked in [ABL17, Problem 3.10] whether there is always an optimal discrete Morse function that is monotone.




Figure 2.3: A 6-gon with a (normalized) monotone discrete Morse function. The critical cells are highlighted in red.

Theorem 2.14. *Let X be a simplicial complex and v the lexicographically-largest among the optimal discrete Morse vectors of X . Then there is a monotone discrete Morse function f that has v as discrete Morse vector.*

Proof. Let $g : X \rightarrow \{1, \dots, m\}$, due to Proposition 2.12, be a good discrete Morse function on the complex X that gives rise to v as discrete Morse vector. Moreover, let T be the set of critical faces Δ such that the complex $\{\sigma \in X \text{ s.t. } g(\sigma) \leq g(\Delta)\}$ has at least one free $(\dim \Delta - 1)$ -dimensional face. If T is empty, then g satisfies also condition (v) of Definition 2.11 and it is already a monotone function. Otherwise let us choose g so that T has a minimal number of elements.

The idea is to modify g slightly without changing the Morse vector and at each step getting “closer” to a monotone function. Let Δ be the element of T with the highest value with respect to g .

We define ρ as the highest-valued (with respect to g) free $(\dim \Delta - 1)$ -dimensional face of $\{\sigma \in X \text{ s.t. } g(\sigma) \leq g(\Delta)\}$. There are two possibilities:

- ▷ ρ is a free face of Δ ;
- ▷ ρ is a free face of $\tau \neq \Delta$.

Let us first consider the second case. Given Δ , ρ , and τ as described above, we define a new function $g' : X \rightarrow \{1, \dots, m + 1\}$ in the following way:

$$\begin{aligned} g'(\sigma) &= g(\sigma) + 1 \quad \text{if } g(\sigma) > g(\Delta), \\ g'(\sigma) &= g(\Delta) + 1 \quad \text{if } \sigma = \rho \text{ or } \sigma = \tau, \\ g'(\sigma) &= g(\sigma) \quad \text{otherwise.} \end{aligned}$$

Then g' is still a discrete Morse function for X since we have chosen ρ as the highest-valued free face and g was good (we just need to check on ρ and τ). Since v is the lexicographically-largest among the optimal discrete Morse vectors of X ,

we obtain that v is also the discrete Morse vector of g' and that g and g' induce the same matching. The only other option (since v is optimal) would have been that ρ was paired before with a face of smaller dimension, but then the discrete Morse vector of g' would be lexicographically-larger than v . The first four conditions of Definition 2.11 are all clearly satisfied because they were satisfied by g so g' is still good.

Moreover the critical faces of g' that do not respect condition (v) of Definition 2.11 are a subset of that of g due to the fourth condition of 2.11 for g , but since this set was minimal they must be exactly the same.

We can now repeat the same procedure and, since the number of faces of dimension $\dim \Delta - 1$ is finite, at a certain point we will arrive at a function h and at the first case of above, i.e., ρ is a free faces of Δ .

In this case we construct the function f as equal to h everywhere except on ρ , where $f(\rho) = h(\Delta)$. Remember that ρ was a free face of Δ so the only options for ρ in the acyclic matching given by h was for ρ to be critical or to be paired below. Now instead, in the acyclic matching given by f , we have that (ρ, Δ) is a new matching pair, but this is impossible because if ρ was critical, then v was not optimal, and if ρ was paired below, then v was not the lexicographically-largest optimal discrete Morse vector. This implies that h must already be monotone. \square

2.7 Bistellar flips

To end this first chapter we put in the spotlight an additional very useful tool to reduce the number of faces of a triangulated manifold without changing its topological type.

The boundary $\partial\Delta_{d+1}$ of the $(d+1)$ -simplex is a d -dimensional simplicial complex with $d+2$ facets. A *bistellar move* is a local modification of a triangulated d -manifold X in which any subcomplex of X isomorphic to the star of a face in $\partial\Delta_{d+1}$ is replaced by its complementary facets.

To be precise, let σ be an i -face of X which is contained in exactly $d-i+1$ facets $\tau_1, \dots, \tau_{d-i+1}$ such that these facets cover exactly $d+2$ vertices. Identifying those $d+2$ vertices with the vertices of Δ_d yields $(d+2) - (d-i+1) = i+1$ complementary facets $\tau_{d-i+2}, \dots, \tau_{d+2}$ in the boundary $\partial\Delta_{d+1}$. Replacing $\tau_1, \dots, \tau_{d-i+1}$ by $\tau_{d-i+2}, \dots, \tau_{d+2}$ in X is a *candidate bistellar move of dimension $d-i$* , or a *candidate $(d-i)$ -move* for short. Let $\sigma' = \cap_{j=d-i+2}^{d+2} \tau_j$ be the complementary face to σ , where σ' is of dimension $d-i$. If σ' is not already contained in X , the move is *proper*. Applying an i -dimensional proper bistellar move reduces the f -vector of X if and only if $i > d/2$.

Two triangulated manifolds are *bistellarly equivalent* if one is obtained from the other by a finite sequence of (proper) bistellar moves.

Chapter 3

Frontiers of sphere recognition in practice

Sphere recognition is known to be undecidable in dimensions five and beyond, and no polynomial time method is known in dimensions three and four. In this chapter, we report on positive and negative computational results with the goal to explore the limits of sphere recognition from a practical point of view. An important ingredient are randomly constructed discrete Morse functions. The material of this chapter is based on a joint work with Michael Joswig, Frank H. Lutz and Mimi Tsuruga [JLLT21], which is the full version of the extended abstract [JLT14].

3.1 Overview

To tell whether a given space is homeomorphic to the sphere in a given dimension is a basic problem in computational topology. However, this is difficult in an essential way.

Theorem 3.1 (S. P. Novikov [VKF74]; cf. [CL06a]). *Given a d -dimensional finite simplicial complex X it is undecidable to check if X is homeomorphic to \mathbb{S}^d for $d \geq 5$.*

We will consider closed manifolds encoded as finite abstract simplicial complexes—but the methods and results in this chapter also hold for more general cell complexes with little modification. For a brief historical overview: d -sphere recognition is trivial in dimensions $d \leq 2$. Rubinstein [Rub95] and Thompson [Tho94] proved that 3-sphere recognition is decidable. Subsequently, Schleimer [Sch11] showed that 3-sphere recognition lies in the complexity class NP, and Lackenby [Lac21] proved that this problem also lies in co-NP; see [Lac20] for a recent survey. The complexity status of 4-sphere recognition is open. Summing up, we do not

know of any efficient algorithm for d -sphere recognition in the relevant dimensions $d \geq 3$.

Our point of departure is that the sphere recognition problem does not go away simply because it is algorithmically intractable. To the contrary it appears naturally, e.g., in the context of *manifold recognition*, which is the task of deciding whether a given simplicial complex triangulates any manifold and finding its type. In the piecewise linear (PL) category, recognizing whether a given complex triangulates some PL manifold can be reduced to *PL sphere recognition* since the links of all vertices of the given complex need to be PL spheres, sometimes also called standard PL sphere. This plays a role, e.g., for enumerating all manifolds with few vertices or facets [BK92, Bur, Bur11, CC15, Lut99, SL09]; for detecting errors in experimental topological constructions [ABL17, SK11, TL]; or for meshing [SAZ08].

In the absence of a general sphere recognition procedure the next best thing are certificates for sphericity and non-sphericity, respectively. A discrete Morse function, μ , on a finite d -dimensional abstract simplicial complex, X , may be encoded as an acyclic partial matching in the Hasse diagram of the partial ordering of the faces of X ; cf. [For98, For02] and [Cha00]. The *critical* faces are those unmatched, and (c_0, c_1, \dots, c_d) is the *discrete Morse vector* of μ , where c_k is the number of critical k -faces. We call such a discrete Morse vector *spherical* if $c_0 = c_d = 1$ and $c_k = 0$ otherwise. The relevance for our topic comes from the following key result.

Theorem 3.2 (Whitehead [Whi39]; Forman [For98, For02]). *A combinatorial d -manifold is a PL d -sphere if and only if it admits some subdivision with a spherical discrete Morse vector.*

So we propose a heuristic method for sphere recognition which navigates between Theorems 3.1 and 3.2. There are a few more obstacles though. Adiprasito and Izmistiev [AI15] showed that a sufficiently large iterated barycentric subdivision of any PL sphere is polytopal (and thus inherits a spherical discrete Morse function from linear programming). However, in view of Theorem 3.1, there cannot be any a priori bound on the number of barycentric subdivisions required to attain polytopality. Second, deciding whether a discrete Morse function with at most a fixed number k of critical cells exists is NP-hard [JP06, LLT03], intractable in the parameter k [BLPS16], and not even a polynomial approximation is available [BR19]. Finally, there are combinatorial d -spheres that do not admit any spherical discrete Morse function [BL13b, BZ11].

This chapter is the considerably expanded full version of the extended abstract [JLT14]. It is organized as follows. As our first main contribution, in Section 3.2, we present an implementation of a sphere recognition heuristic procedure in `polymake`, and demonstrate its efficiency. In the `polymake` project, Perl

and **C++** are used as programming languages; our heuristic is implemented in **C++**. It is also available through the new **Julia** interface layer `Polymake.jl` [GKT19], which supports the current `polymake` Version 4.5. Section 3.3 comprises comprehensive computational experiments which show that there are many randomly constructed, even fairly large, simplicial complexes for which deciding sphericity is surprisingly easy; this agrees with previous observations [ABL17, BL14]. Moreover, on such input our new approach proves to be superior to, e.g., the 3-sphere recognition implemented in **Regina** [BBP⁺17], which is a standard tool in computational topology. Note that **Regina**'s method is a full decision algorithm, while our heuristic may be inconclusive. However, we are not aware of any triangulation of \mathbb{S}^3 which cannot be recognized by our method. Another experiment comes from a census of 4-manifolds provided by **Regina**; here our heuristics recognizes about 99.5% of the input as spheres or non-spheres. Finally, in Section 3.4, we explore the limitations of our method. One outcome is the construction of a new family of 2-dimensional cell complexes which are contractible, but not collapsible. These *saw blade complexes* generalize the Dunce hat, and in our experiments they occur naturally as one source of difficulty for recognizing spheres. Moreover, our computer experiments show that there is a “horizon” for discrete Morse computations, along with implications to homology computations and computational topology in general.

3.2 A heuristic sphere recognition scheme

We describe our procedure for sphere recognition and its implementation in the mathematical software `polymake` [GJ00]. This is the specification:

Input: A d -dimensional (finite abstract simplicial) complex X with n vertices and m facets, where a *facet* is a face that is maximal with respect to inclusion.

Output: YES, NO, or UNDECIDED, depending whether X has been recognized as a (standard) PL d -sphere.

Our procedure features five steps, labeled (0) through (4). Discussing the trivial preprocessing Step (0) in some details allows us to introduce the basic terminology and notation. The core Steps (1), (2), (3), and (4) below together yield Algorithm A.

(0) Preprocessing

To verify whether X is a PL d -sphere, there are three elementary combinatorial checks that are useful to perform first. These checks are fast; their running time is bounded by a low-degree polynomial in the parameters d , n and m . If one of the checks fails, this will serve as the certificate that X is not a sphere.

Algorithm A: Sphere Recognition Heuristics

Input: Hasse diagram of combinatorial d -manifold X , where $d \geq 3$
Output: Semi-Decision: Is X PL homeomorphic to \mathbb{S}^d ?

```

(5) compute homology
    if homology not spherical then return NO
    for  $N$  rounds do
(6)   compute random discrete Morse vector
        if discrete Morse vector is spherical then return YES
    for  $N'$  rounds do
(7)   perform random bistellar flip or edge contraction
        if boundary of simplex is reached then return YES
(8) compute and simplify presentation of fundamental group  $\pi_1$ 
    if presentation is found to be trivial and  $d \neq 4$  then return YES
    if presentation is found to be non-trivial then return NO
    return UNDECIDED

```

More precisely, we first check if X is *pure*, i.e., each facet has exactly $d + 1$ vertices. Second, we check if each ridge is contained in exactly two facets, where a *ridge* is a face of dimension $d - 1$. Success in these two tests will assert that X is a *weak pseudo-manifold* (without boundary). Note that the 0-dimensional sphere \mathbb{S}^0 is a weak pseudo-manifold of dimension $d = 0$ with two isolated vertices.

Third, for $d \geq 1$, we check if the 1-skeleton of X is a connected graph. A connected weak pseudo-manifold X of dimension $d = 1$ is a polygon and thus triangulates \mathbb{S}^1 .

The pureness and the weak pseudo-manifold property of a simplicial complex is inherited by all face links; cf. [BD98, Rem. 8]. A (connected) weak pseudo-manifold is a *pseudo-manifold* if it is *strongly connected*, i.e., if any two of its facets can be joined by a sequence of facets for which consecutive facets share a ridge. In particular, a pseudo-manifold of dimension $d = 2$ is a triangulation of a closed surface or of a closed surface with *pinch points* (having multiple disjoint cycles as vertex links).

A d -dimensional pseudo-manifold X is a *combinatorial d -manifold* if all vertex links of X are PL homeomorphic to the boundary of the d -simplex or, equivalently, if for every proper i -face (with $0 \leq i < d$) of X its link is a PL $(d - i - 1)$ -sphere; here the case $i = d - 1$ ensures the weak pseudo-manifold property. This recursive nesting of PL spheres suggests an inductive check of the face links of X by dimension, starting with 1-dimensional links of $(d - 2)$ -faces and proceeding up until the $(d - 1)$ -dimensional links of the vertices.

A connected 2-dimensional weak pseudo-manifold X whose vertex links are

single cycles is a combinatorial 2-manifold and triangulates a closed surface. If, additionally, the Euler characteristic of X equals two, then X is \mathbb{S}^2 .

If one of the checks on the links of the overall complex X fails, then X cannot be a standard PL sphere, if one of the checks is left undecided this leaves X undecided.

After this preprocessing and an inductive check of the vertex-links we may assume that our input looks as follows:

Input (modified): Let X be a d -dimensional combinatorial manifold, for $d \geq 3$.

The subsequent four steps form the core of our sphere recognition procedure.

(1) Homology computation

Computing the simplicial homology modules of a finite simplicial complex is a standard procedure, which is implemented, e.g., in `CHomP` [M⁺a], `RedHom` [M⁺b] or `polymake` [GJ00]. The homology with field coefficients can be determined via applying Gaussian elimination to the (simplicial) boundary matrices. For finite fields of prime order or the rationals this can be achieved in polynomial time (in the size of the boundary matrices); cf. [Edm67]. Similarly, over the integers, a homology computation can be reduced to computing Smith normal forms; cf. [Mun84, Ch. 11]. Kannan and Bachem [KB79] gave the first polynomial time Smith normal form algorithm, employing modular arithmetic; see also [Ili89].

Here we employ integer coefficients throughout. A necessary condition for X to be a sphere (PL or not) is $H_d(K) \cong \mathbb{Z}$, and all other (reduced) homology groups vanish. In this case we say that X has *spherical homology*.

We recall that the *Hasse diagram* of a simplicial complex X is a directed graph with one node per face of X and a directed arc (σ, τ) if the face σ is contained in τ and $\dim \tau = \dim \sigma + 1$.

While the modular approach of [KB79] and [Ili89] is valid for matrices with arbitrary integer coefficients, simplicial boundary matrices have entries 1, -1 , and 0 only. As a consequence, in an arbitrary simplicial boundary matrix it is *always* possible to perform at least a few Gauss elimination steps. Moreover, a typical boundary matrix is sparse. If the matrix happens to stay sparse during the elimination and if, additionally, one does not run out of unit coefficients too soon (such that it is possible to continue with elimination steps), an elimination based Smith normal form algorithm can outperform the more sophisticated modular methods. This is why for computations of (simplicial) homology elimination algorithms are often preferred; cf. Dumas et al. [DHSW03] for a survey.

A (*partial*) *matching* in an arbitrary graph is a subset of the edges such that each node is covered at most once. In the Hasse diagram of X a matching corresponds to a set of non-zero coefficients in some boundary matrices. Such a

∂_2	01	02	03	04	05	12	13	14	15	23	24	25	34	35	45
012	1	-1				1									
014	1			-1				1							
023		1	-1							1					
035			1		-1									1	
045				1	-1										1
125						1			-1			1			
134							1	-1					1		
135							1		-1					1	
234										1	-1		1		
245											1	-1			1
$(\partial_1)^{\text{tr}}$	01	02	03	04	05	12	13	14	15	23	24	25	34	35	45
0	-1	-1	-1	-1	-1										
1	1					-1	-1	-1							
2		1				1				-1	-1	-1			
3			1				1			-1			-1	-1	
4				1				1			1		-1		1
5					1				1			1		1	1

Figure 3.1: Boundary matrices of \mathbb{RP}_6^2 with coefficients marked that correspond to the acyclic matching of Figure 2.2; cf. Example 3.3.

matching, μ , is called *acyclic* if reversing all arcs in μ (and keeping the arcs not in μ) still gives an acyclic graph. It is easy to see that an acyclic matching in the Hasse diagram of X yields a sequence of Gauss elimination steps that can be performed in any order without destroying the (unit) pivots required for the subsequent elimination steps.

Example 3.3. As we have already seen in Example 2.6, Figure 2.2 shows an acyclic matching, μ , in the Hasse diagram for $X = \mathbb{RP}_6^2$, which is the six-vertex triangulation of the real projective plane. Figure 3.1 shows the corresponding boundary matrices. The pivots corresponding to μ are marked. Using these pivots in an arbitrary order yields an elimination strategy for the computation of the homology modules:

$$\tilde{H}_0(\mathbb{RP}_6^2) = 0, \quad H_1(\mathbb{RP}_6^2) \cong \mathbb{Z}/2\mathbb{Z}, \quad H_2(\mathbb{RP}_6^2) = 0.$$

(2) Random discrete Morse functions

We have defined in 2.5, given any simplicial complex X , its set of *discrete Morse functions* and *discrete Morse vectors* (c_0, c_1, \dots, c_d) . We recall that X is homotopy

equivalent to a CW complex with c_i cells in dimension $0 \leq i \leq d$.

Moreover, let \mathbb{F} be some field, a discrete Morse vector is \mathbb{F} -*perfect* for X if $c_i = \beta_i(K; \mathbb{F})$ for $0 \leq i \leq d$. A perfect discrete Morse vector of a sphere, which reads $(1, 0, 0, \dots, 0, 1)$, is also called *spherical*. Theorem 3.2 implies that a combinatorial d -manifold X that becomes collapsible after the removal of one facet is a PL d -sphere. In 1992, Brehm and Kühnel [BK92] used that fact to show that some 8-dimensional simplicial complex with 15 vertices is a combinatorial 8-manifold (a triangulation of the quaternionic projective plane [Gor19]).

By Theorem 3.2, the existence of an acyclic matching whose discrete Morse vector is spherical is a sufficient criterion for X to be a sphere. This gives rise to the following simple strategy: generate discrete Morse functions (or acyclic matchings) at random and check if one of them is spherical; cf. [BL14].

The `random_discrete_morse` function implemented in `polymake` has three different random strategies which we call `random-random`, `random-lex-first`, and `random-lex-last`. We will give a short outline and describe the differences among the three strategies and further differences to the original approach from [BL14].

Let X be an arbitrary d -dimensional simplicial complex, which is not necessarily a manifold. A *free face* of X is an $(i-1)$ -dimensional face that is contained in exactly one i -face, $0 < i \leq d$. To save memory, our three strategies are destructive in the sense that they keep changing the complex X . In each step we pick one of the free faces of codimension one and delete it from X together with the unique d -face containing it. This is an elementary collapse, and the two removed faces form a regular pair, which is a matching edge in the Hasse diagram. The three strategies only differ in how they pick the free face. If we run out of free faces, we pick some facet (of maximal dimension), declare it *critical* and remove it. After removing a regular pair the dimension of the resulting complex, K' , may drop to $d-1$. This process continues until K' is zero-dimensional. In this case, K' only consists of vertices, all of which are declared critical.

For the `random-random` strategy, we first find all the free faces of X and collect them in a linked list. If this list is not empty, choose a free face uniformly at random. Taking the uniform distribution means that each free face has a fair chance of being taken, but this comes at a price since the sampling itself takes time if there are many free faces to choose from. The reason is that we do not have random access to the free faces, as they are kept in a linked list. Picking a random element in a linked list takes linear time in the length of that list. If we run out of free faces, the choice of the critical d -face is again uniformly at random.

The strategy `random-random` is somehow the obvious one, but there is a much cheaper way which maintains a certain amount of randomness. Here the price is that it seems to be difficult to say something about the resulting probability distribution. The idea is to randomly relabel the vertices of X once, at the be-

ginning, and then to pick the free and critical faces in a deterministic way (which depends on the random labeling). Whenever we want to choose a free or critical face, rather than selecting one at random, we pick the first (in the case of **random-lex-first**) or the last one (in the case of **random-lex-last**) of the linked list. The **random-lex-last** strategy was called “random-revlex” in [BL14]. We changed the name here to **random-lex-last** to avoid confusion with the reverse lexicographic (term) ordering, which is different.

The cost of being fair is quite significant, with our current implementation, when dealing with large complexes. For example, running the **random-lex-first** and **random-lex-last** strategies on the fourth barycentric subdivision of $\partial\Delta^4$ took less than three minutes per run whereas the **random-random** strategy took approximately two hours per run; see Section 3.4.4. It is conceivable that there is some room for improvement here by employing a faster data structure for random sampling; we leave such an implementation for a future version of **polymake**.

Remark 3.4. In Algorithm A, the Steps (1) and (2) can also be intertwined as finding an acyclic matching results in a partial strategy for computing the homology. To this end it is most natural to process the Hasse diagram from top to bottom level by level.

(3) Random bistellar flips

If the previous tests are inconclusive, we can use a local search strategy to determine the PL type; cf. [BL00]. The boundary $\partial\Delta^{d+1}$ of the $(d+1)$ -simplex is a d -dimensional simplicial complex with $d+2$ facets. We recall that a *bistellar move* is a local modification of a combinatorial d -manifold X in which any subcomplex of X isomorphic to the star of a face in $\partial\Delta^{d+1}$ is replaced by its complementary facets.

Two simplicial complexes are *bistellarly equivalent* if one is obtained from the other by a finite sequence of (proper) bistellar moves. The following result is essential for the third step in the heuristic.

Theorem 3.5 (Pachner [Pac87]). *A d -dimensional simplicial complex is a PL d -sphere if and only if it is bistellarly equivalent to $\partial\Delta^{d+1}$.*

This is closely related to Theorem 3.2 in the following sense: Adiprasito and Izmestiev [AI15] showed that iterated barycentric subdivisions make any PL sphere polytopal; and barycentric subdivisions can be expressed as sequences of stellar subdivisions (which, by Theorem 3.5, are connected via sequences of bistellar moves). Moreover, barycentric subdivisions of polytopal spheres are polytopal, and polytopal spheres admit spherical discrete Morse vectors.

We now discuss the `polymake` implementation of the simulated annealing strategy from [BL00]. The function `bistellar_simplification` randomly applies bistellar moves to an input of type `SimplicialComplex` (required to be a combinatorial d -manifold) with the goal to lower the f -vector as much as possible. In this way the algorithm prefers moves that reduce the f -vector; this is called “cooling”. It lies in the nature of the sphere recognition problem that we may end up in a local minimum, i.e., when there are no moves to further lower the f -vector. At that point, we deliberately make moves that increase the f -vector for some number of rounds (this is called “heating”). Then we cool again, hoping that this will help jiggle us out of that local minimum. For $0 \leq i < \frac{d}{2}$ a corresponding i -move will increase the f -vector of a triangulation, while in even dimensions the f -vector is not altered by $\frac{d}{2}$ -moves. So in a heating phase we would add vertices via 0-moves, edges via 1-moves, etc., and “randomize” the triangulation by performing a (possibly large) number of $\frac{d}{2}$ -moves, before returning to cooling. As with all simulated annealing approaches, adjusting the parameters for the annealing relies on experimentation. For example, we initially may not add vertices in the heating phases via 0-moves, as this might successively increase the size of the intermediate complexes, but in case we remain in a local minimum, first adding some percentage of vertices before performing 1-moves etc. helps in some cases.

Our procedure determines all candidates for bistellar moves of X and sorts them by descending dimension. During a cooling period we first pick random d -moves, if one exists. Otherwise, we pick random candidate $(d-1)$ -moves until we find one which is proper. If this does not exist either, we continue further to dimensions $d-2$, $d-3$, etc., down to dimension $\lfloor d/2 \rfloor + 1$. Any proper move found in this way is applied immediately. Note that a bistellar move is a local operation, which is why we refrain from copying the entire complex when we apply a bistellar move. Instead, we perform the operation in place and store the reverse move in a list such that it can be undone later. Cooling continues until we get stuck with a lexicographically locally minimal f -vector. This ordering of the f -vectors is imposed indirectly by preferring higher-dimensional moves.

During a heating period, the story is slightly different. One heating strategy is to choose the dimension of the heating move at random with respect to a heuristically determined distribution. That distribution is encoded as a *heat vector* $(h_0, \dots, h_{\lfloor d/2 \rfloor})$ of integers, and we set $h := h_0 + \dots + h_{\lfloor d/2 \rfloor}$. This means, in each round of the heating period we pick the dimension k with probability h_k/h , and in that dimension we pick a random proper bistellar move. For example, the default heat vector in `polymake` for $d = 4$ is $(10, 10, 1)$. This generalizes to the default heat vector $(10, 10, \dots, 10, 1)$ in higher (even) dimensions d , while for odd d the pivot dimension k is picked uniformly at random.

Various other parameters control the precise heating behavior; and some of

them are adjusted dynamically. For instance, it is useful to heat up for more rounds if the complex is larger. Sometimes it pays off to experiment with several distributions or to use other heating schemes. E.g., as mentioned above, first add some (percentage of) or no vertices via 0-moves, then edges via 1-moves, etc.

Remark 3.6. As a speed-up for large input triangulations, we can first apply edge contractions (with admissible edges for a contraction chosen at random) as long as possible. As experienced for 3-manifold triangulations [CL06b], this eventually leads to a saturation with many edges that block further contractions. Once there is no remaining admissible edge for a contraction, we run bistellar flips to reduce the number of edges and then continue with edge contractions again. Edge contractions are useful only in an initial phase. Once a local minimum is reached for the size of the triangulation, then bistellar flips are employed to leave the local minimum.

(4) Fundamental group

A non-trivial fundamental group $\pi_1(K)$ is a certificate for not being a sphere (PL or otherwise). Conversely, there is the solution to the (PL) Poincaré Conjecture in dimensions other than four.

Theorem 3.7 (Smale [Sma61]; Perelman [Per02]). *Let X be a simply connected combinatorial d -manifold, $d \neq 4$, with spherical homology. Then X is a PL sphere.*

Freedman proved that a simply connected 4-manifold with trivial intersection form is homeomorphic to the 4-sphere [Fre82]. But his result does not say whether this also holds in the PL category. In fact, it is a major open problem whether or not “exotic” 4-spheres exist.

In [ST34, Chapter 7] Seifert and Threlfall describe how to obtain a finite presentation of $\pi_1(K)$ from any spanning tree in the 1-skeleton (with the remaining edges as generators) and all the 2-faces (as relators). However, checking if a finitely presented group is trivial is known to be undecidable [Nov55]. Discussing heuristic approaches to simplifying group presentations is beyond the scope of this chapter. In practice we rely on GAP [GAP19] which employs Tietze transformations.

Algorithm A displays our strategy in a concise form; for computational results see Sections 3.3 and 3.4 below. Notice that the ordering of the Steps (1) through (3) is arbitrary, while the “YES” answer in Step (4) is inconclusive without Step (1). Yet, there is a benefit from combining Steps (1) and (2); cf. Remark 3.4. In practice, for a complex X we suspect to be *not* a sphere, we would start with (1), while if we think that X is a PL sphere, we first try (2) as a fast routine. If we are not successful with (2), we switch to (3), which is slower but can still recognize

spheres that do not have perfect discrete Morse functions; see the discussion in Section 3.4.

If, in the case $d \neq 4$, Step (1) gives us a spherical homology vector and Step (4) a trivial presentation of the fundamental group, then the overall output is “YES”, by Theorem 3.7. In the case of spherical homology a presentation of the fundamental group with only one generator is not possible, but balanced presentations of the trivial group with two generators and two relators can already be hard to detect; see Section 3.4.2.

Clearly, when our method gives up with “UNDECIDED” this does not need to be the end of the story. For instance, in the 3-dimensional case we can feed the data into the 3-sphere recognition procedure of **Regina** [BBP⁺17]. This features a variation of the exact algorithm of Rubinstein [Rub95] and Thompson [Tho94], where, for instance, the crushing procedure (the key step of the algorithm) is dramatically simplified [Bur14]. In this way **Regina** can provide certificates for X *not* being spherical based on normal surface theory. However, we are not aware of a single triangulation of the 3-sphere for which our procedure fails.

3.3 Experiments and runtime comparisons

To find challenging input for Algorithm A is not entirely trivial. Most explicit constructions of (standard) PL spheres found in the literature are rather small and can be recognized instantaneously. All timings were taken on an AMD Phenom(tm) II X6 1090T Processor CPU (3.2 GHz, 6422 bogomips) and 8 GB RAM with openSUSE Leap 15.0 (Linux 5.1.9-5).

3.3.1 Recognizing random 3-spheres with polymake and Regina

A natural class of PL d -spheres are the boundaries of $(d + 1)$ -polytopes obtained as the convex hulls of n points chosen uniformly at random on the unit d -sphere in \mathbb{R}^{d+1} . These have been studied, e.g., in the context of the average case analysis of the simplex method of linear programming [Bor87]. Such examples can be generated with the `rand_sphere` command of **polymake**. Table 3.1 lists **polymake** and **Regina** experiments on 3-spheres with up to 100,000 vertices. For more than 15,000 vertices the convex hull computation (necessary only to construct the input) becomes a bottleneck, which is why for the larger examples (marked “*”) we used connected sums of smaller random spheres.

In all cases, the spheres were successfully recognized by each method. However, we truncated the time spent on each input to about one CPU hour, such that longer running times are omitted. The fastest method is **polymake**’s random search for a

Table 3.1: Running times (in seconds) on random polytopal 3-spheres on n vertices.

n	polymake			Regina
	Morse	bistellar	contr.+bist.	isThreeSphere
100	0.01	0.37	0.03	0.03
200	0.01	1.23	0.07	0.15
300	0.02	2.87	0.11	0.35
400	0.03	3.23	0.17	0.64
500	0.04	4.94	0.20	1.09
600	0.05	7.31	0.26	1.60
700	0.07	10.24	0.31	2.22
800	0.08	13.10	0.37	3.07
900	0.09	17.92	0.44	4.16
1000	0.10	23.03	0.49	5.23
2000	0.38	107.85	1.25	28.22
3000	0.78	281.17	2.29	74.27
4000	1.31	551.62	3.41	141.65
5000	2.26	918.09	4.82	237.42
10000	8.71	4608.71	16.48	1100.26
15000	22.11	/	39.77	2647.71
*30000	145.90	/	191.22	/
*50000	470.26	/	515.46	/
*100000	1586.41	/	2064.28	/

spherical discrete Morse function; cf. Step (2) of Algorithm A. Nearly competitive is **polymake**'s procedure of applying edge contractions, combined with random bistellar moves; cf. Step (3) of Algorithm A and Remark 3.6.

Usually, **Regina** takes 1-vertex pseudo-simplicial triangulations as input, but can also handle (abstract) simplicial complexes. In the latter case, contracting a spanning tree in the 1-skeleton yields a 1-vertex pseudo-simplicial triangulation. Conversely, the second barycentric subdivision of a pseudo-simplicial complex is a simplicial complex. In this sense these two encodings of combinatorial manifolds are similar.

Regina's recognition algorithm **isThreeSphere** runs, as a preprocessing step, the program **IntelligentSimplify** that transforms the complex into a 1-vertex triangulation and uses bistellar moves to further reduce it, similar to Step (3) of our Algorithm A. Afterwards the 3-sphere recognition procedure is employed. In this way, **Regina** is able to also find certificates for non-sphericity—which **polymake** is incapable of, beyond checking the homology. We should also point out that **IntelligentSimplify** is a heuristic designed to be an out-of-the-box first attempt to simplify a triangulation with a polynomial running time, and **Regina**'s bistellar

Table 3.2: Census of 4-manifolds with up to six pentachora.

	two pentachora		four pentachora		six pentachora	
	# sign.	percentage	# sign.	percentage	# sign.	percentage
Total:	8	100.0%	784	100.0%	440,494	100.0%
Spheres:	6	75.0%	642	81.9%	403,240	91.5%
Non-spheres:	2	25.0%	137	17.5%	35,305	8.0%
Unknown:	0	0.0%	5	0.6%	1,949	0.5%

move interface is meant to be interactive. This means that with a bit of work to build a custom-made simplification routine, the times in **Regina** could probably be improved.

The largest simplicial complex in Table 3.1, with 100,000 vertices, has 673,274 tetrahedra. The largest one successfully handled within one hour by **Regina** has 15,000 vertices and 101,088 tetrahedra.

Each row of the Tables 3.1, 3.3 and 3.4 corresponds to a single instance only. However, it is known that there is little variation of, e.g., the number of facets of the convex hull of random points on the unit sphere; cf. Reitzner [Rei05, Sec. 4]. This can also be observed experimentally; cf. [AGH⁺17, Sec. 3.5 and Fig. 6] for a closely related setup. Note that, in fixed dimension d , the expected number of facets of a random simplicial $(d + 1)$ -polytope depends linearly on the number of vertices [Bor87].

3.3.2 Processing a census of 4-manifolds

We ran our heuristics on a census of 4-manifolds provided by **Regina** [BBP⁺17]. These 4-manifolds are encoded as pseudo-simplicial complexes comprising up to six maximal cells; the 4-simplices are called *pentachora* in **Regina** for they have five facets (and five vertices). In combined form, the **tricensus** command of **Regina** generates possible facet pairings, then for each such pairing determines all possible gluing permutations, and (on the fly) reduces all gluings to isomorphism signatures that uniquely encode triangulations up to combinatorial isomorphism [Bur]. For the examples with two, four, and six pentachora there are 3, 26, and 639 facet pairings. And this yields 8, 784, and 440,494 resulting combinatorial types, respectively. Note that there are no facet pairings for an odd number of maximal cells in even dimensions. We let **Regina** expand each of these into a (proper abstract) simplicial complex via the second barycentric subdivision and pass it on to **polymake**. The simplicial complexes resulting from six pentachora have around 4,600 vertices and 86,400 facets.

Using our heuristic we found the results in Table 3.2. Each positive or negative

certificate was obtained in less than four minutes and in 90 seconds on the average. In all the cases the positive certificates arise from discrete Morse functions, while the negative certificates are provided by non-spherical homology. Taking row sums in Table 3.2 we summarize our findings as follows.

Theorem 3.8. *Among the 441,286 combinatorial types of combinatorial 4-manifolds arising from up to six pentachora 91.5% are spheres, and 8.0% are non-spheres.*

Thus, our success rate is 99.5%, with our heuristic failing on only 1,954 of these combinatorial 4-manifolds. All of these have spherical homology; to determine whether they are standard PL 4-spheres or proper combinatorial homology 4-spheres (combinatorial 4-manifolds with spherical homology, but not PL homeomorphic to the standard PL 4-sphere) is an interesting question, yet beyond the scope of this chapter. Our classification, as a list of **Regina**'s isomorphism signatures of the complexes can be found at [Lof21c].

3.3.3 Higher-dimensional random spheres

polymake can easily recognize random polytopal spheres with up to 10,000 vertices in dimension four, 1,000 vertices in dimension five, and 500 vertices in dimension six; cf. Tables 3.3 and 3.4. Again the input is constructed via uniform random sampling on the unit sphere and taking convex hulls.

Regina provides no heuristic for sphere recognition in dimension four or beyond. Yet, **Regina** can simplify a given triangulation of a 4-dimensional combinatorial manifold via contractions and bistellar moves, returning a smaller pseudo-simplicial complex. It is not immediate how to check for sphericity from that output. That implementation is deterministic; thus in each run on a fixed input it gives the same output. The running times are given in the penultimate column of Table 3.3. The last column contains the number of simplices remaining after simplification.

3.3.4 A collapsible 5-manifold which is not a ball

We consider the 5-dimensional simplicial complex C with face vector

$$f(C) = (5013, 72300, 290944, 495912, 383136, 110880)$$

constructed in [ABL17, Sec. 4]; there C is called `contractible_non_5_ball`. This is the first explicit example of a non-PL triangulation of a collapsible (and thus contractible) 5-manifold, other than the 5-ball. By construction, C is a manifold

Table 3.3: Running times (in seconds) on random polytopal 4-spheres on n vertices.

n	polymake			Regina	
	Morse	bistellar	contr.+bist.	Simplify	number of facets
100	0.04	8.22	0.46	1.52	26
200	0.13	33.50	1.28	8.44	8
300	0.30	76.62	2.63	24.97	42
400	0.54	136.85	4.77	54.51	78
500	0.82	224.67	6.17	92.25	60
600	1.21	418.50	8.06	121.24	120
700	1.64	639.45	10.95	184.11	98
800	2.28	842.94	15.38	303.16	180
900	2.88	1109.43	16.74	370.85	144
1000	3.51	1418.25	22.20	474.72	170
2000	10.86	/	40.93	2427.81	562
3000	26.40	/	219.44	/	/
5000	121.90	/	714.92	/	/
10000	594.46	/	2633.70	/	/

with boundary. To check the remaining topological properties computationally poses an interesting challenge.

First, the perfect Morse vector $(1, 0, 0, 0, 0, 0)$ for C was originally obtained in a single random discrete Morse vector search over 82 hours with a **GAP** implementation. The current implementation in **polymake** produces the same result (in most runs) in only 9 seconds with the **random-lex-first** and **random-lex-last** strategies and in about 10 minutes with the **random-random** strategy. This certifies that C is collapsible.

Second, the boundary complex ∂C with face vector

$$f(\partial C) = (5010, 65520, 212000, 252480, 100992)$$

was investigated; it was called **contractible_non_5_ball_boundary** in [ABL17]. Checking all face links for spherical discrete Morse vectors confirmed that ∂C is a combinatorial 4-manifold. For each face link a single random try sufficed. In total, the recognition of all face links took about 7.5 hours. Checking the homology reveals that ∂C is a homology 4-sphere. Finally, **GAP** identifies the fundamental group $\pi_1(\partial C)$ as the binary icosahedral group.

Table 3.4: Running times (in seconds) on random polytopal 5- and 6-spheres.

n	polymake, $d = 5$		polymake, $d = 6$	
	Morse	contr.+bist.	Morse	contr.+bist.
100	0.33	10.29	10.01	301.62
200	1.86	40.52	86.78	2634.24
300	5.78	102.32	387.90	/
400	11.49	169.86	967.55	/
500	21.95	340.08	1788.44	/
600	35.31	515.86	/	/
700	55.55	820.34	/	/
800	78.48	1120.07	/	/
900	104.28	1441.08	/	/
1000	133.34	2016.53	/	/

3.4 Limitations

In the previous section we saw that many, even fairly large, simplicial spheres can be recognized easily, despite Theorem 3.1. Here we explore the limitations of our heuristic. The combination of our (positive and negative) experiments may serve as a description of a “horizon” within which we can hope for effective recognition results.

3.4.1 General remarks

We refrain from a detailed comparison of simplicial homology computations. However, standard implementations, as CHomP [M⁺a], RedHom [M⁺b], Perseus [MN13], and polymake [GJ00], employ elimination schemes for computing the integer homology, which are equivalent to finding discrete Morse functions with few critical cells. In this sense, the horizon within which we can compute the simplicial homology is essentially the same as the horizon for the discrete Morse Step (2). There are more software systems to compute simplicial homology, but many, including, e.g., Dionysus [Mor] and PHAT [BKR], are restricted to \mathbb{Z}_2 -coefficients. Finding an optimal discrete Morse function is NP-hard; cf. [JP06, LLT03]. Recently Bauer and Rathod established that we may not even hope for polynomial approximability [BR19].

In the subsequent we will exhibit several scenarios in which finding a spherical discrete Morse function for a simplicial sphere may fail in practice. An obvious impediment is the lack of any spherical discrete Morse function. The smallest known example is an 18-vertex triangulation of \mathbb{S}^3 , constructed from a triple trefoil knot supported on three edges [BL13b].

Table 3.5: Collapsing the d -simplex.

d	Rounds	Non-perfect	Percentage
7	10^{10}	0	0.0%
8	10^9	12	0.0000012%
9	10^8	2	0.000002%
10	10^7	3	0.00003%
11	10^7	12	0.00012%
12	10^6	4	0.0004%
13	10^6	6	0.0006%
14	10^5	4	0.004%
15	10^5	8	0.008%
16	10^4	4	0.04%
17	10^4	10	0.10%
18	10^3	2	0.2%
19	10^3	6	0.6%
20	10^3	13	1.3%
21	10^3	62	6.2%
22	10^3	153	15.3%
23	10^2	35	35%
24	10^2	67	67%
25	$5 \cdot 10^1$	46	92%

In dimension three, the known recognition *algorithms* for the 3-sphere make use of normal surface theory. As a consequence of Theorem 3.7, a trivial fundamental group suffices to show that a 3-manifold is, in fact, the 3-sphere.

3.4.2 Akbulut–Kirby spheres

A family of standard PL \mathbb{S}^4 -triangulations, $\text{AK}(r)$, for $r \geq 3$, of the *Akbulut–Kirby spheres* [AK85] has been constructed in [TL]. In fact, Akbulut and Kirby [AK85] gave handlebody decompositions of a family of 4-manifolds, in the hope of obtaining exotic 4-spheres. Yet, later Gompf [Gom91] and Akbulut [Akb10] showed that these manifolds are PL homeomorphic to the standard 4-sphere \mathbb{S}^4 . We have

$$f(\text{AK}(r)) = (176 + 64r, 2390 + 1120r, 7820 + 3840r, 9340 + 4640r, 3736 + 1856r).$$

The triangulated Akbulut–Kirby spheres $\text{AK}(r)$ so far constitute the single explicit family of simplicial spheres that we could not recognize easily by our heuristic. More precisely, Step (2) failed on all complexes $\text{AK}(r)$ for all $r \geq 3$. Step (3) worked for $r = 3$, but failed for $r \geq 4$. Steps (2) and (3) are particularly relevant in dimension four—in all other dimensions, as a consequence of Theorem 3.7, they

can conceptually be replaced by the combination of Steps (1) and (4). We do not know if the spheres $\text{AK}(r)$ admit perfect discrete Morse vectors. E.g., the smallest one we found for $r = 5$ is $(1, 2, 4, 2, 1)$, possibly reflecting that we used two winded up 1-handles in the construction.

What makes the PL 4-spheres $\text{AK}(r)$ difficult to recognize is that the original handlebody decomposition of Akbulut and Kirby [AK85] is based on the non-trivial presentation

$$G(r) = \langle x, y \mid xyx = yxy; x^r = y^{r-1} \rangle$$

of the trivial group built in as the fundamental group, i.e., $\pi_1(\mathbb{S}^4)$, of $\text{AK}(r)$. In 100 out of 450 runs we found **GAP** to be able to identify $\pi_1(\text{AK}(4))$ as the trivial group. However, in dimension four (and knowing that the homology is spherical) this only shows that the input is a topological 4-sphere, without yielding any information on the PL type. In one run with $r = 5$ we obtained the initial presentations $G(r)$ as the *output* of **GAP**'s simplification procedure. For $r \geq 4$, none of the Steps (1)–(4) was conclusive to determine that $\text{AK}(r)$ is the standard PL 4-sphere \mathbb{S}^4 .

3.4.3 Contractible but non-collapsible subcomplexes


A simplicial sphere X admits a perfect discrete Morse function (respectively vector) if and only if there is a facet σ of X such that $X - \sigma$ is a collapsible ball. In this way a key difficulty in finding a perfect discrete Morse vector for $X - \sigma$ stems from subcomplexes that are contractible, but not collapsible. The most prominent such example is the 2-dimensional *dunce hat* which can be obtained from a single triangle by identifying its three boundary edges in a non-coherent way [Zee64].

Crowley et al. [CEK⁺] showed that the 7-simplex with 8 vertices contains in its 2-skeleton an 8-vertex triangulation of the dunce hat onto which it collapses. (This result can easily be verified by running the random discrete Morse Step (2) on the 7-simplex, but not allowing the triangles of an 8-vertex triangulation of the dunce hat be used as free faces.) While the dunce hat has triangulations with 8 vertices [BL13a], every contractible complex with fewer vertices is collapsible [BD05].

This leads us to our next experiment, where we compute random discrete Morse vectors for d -simplices, $7 \leq d \leq 25$; cf. Table 3.5. For instance, in dimension seven every one of the 10^{10} runs that we tried gave a perfect discrete Morse function, i.e., a collapsing sequence. With increasing dimension that success rate drops slowly until $d = 20$, where we get stuck with a discrete Morse vector which is not perfect in 1.3% of all tries. Going to even higher dimensions shows a rapid decline of the probability to find a perfect discrete Morse vector. From this we conclude that dimension 25 marks a “horizon” for Step (2), even for a single simplex. Note also that the implementation of the algorithm requires to store the entire Hasse diagram, which is memory expensive; e.g., the Hasse diagram

Table 3.6: Discrete Morse vectors for 10^9 runs on the 8-simplex.

Discrete Morse vector	Count
(1 0 0 0 0 0 0 0)	999999988
(1 1 1 0 0 0 0 0)	4
(1 0 1 1 0 0 0 0)	7
(1 0 0 1 1 0 0 0)	1

**Figure 3.2:** The simplicial complexes D , SB_a^2 , SB_b^3 and SQ (clockwise from top left)

of the 25-simplex needs around 200 GB of RAM. The data structure underlying the Hasse diagram is optimized for speed with respect to enumerating all faces of a simplicial complex (and the covering relations of the inclusion poset) from the facets. That algorithm, which is linear in the size of the output, builds on a very general method of Ganter [Gan87] for closure systems, and the implementation details in `polymake` are discussed in [HJS19].

It is instructive to look at the subcomplexes which arise from non-perfect discrete Morse vectors in our experiments for a d -simplex. For $d = 8$, we found four examples of 2-dimensional contractible and non-collapsible complexes on nine vertices which we call D , SB_a^2 , SB_b^3 , and SQ ; cf. Figure 3.2 and (the second line of) Table 3.6. The first one, D , is a triangulation of the dunce hat. The following concept is derived from scrutinizing the complexes in Figure 3.2.

Definition 3.9. The k -bladed saw blade complex SB^k is the 2-dimensional CW complex obtained from a polygonal disk with $3k$ edges

$$a_1, a_1^{-1}, a_1, a_2, a_2^{-1}, a_2, \dots, a_k, a_k^{-1}, a_k$$




Figure 3.3: The sawblade complex SB^3

by identifying a_1, a_1^{-1}, a_1 as well a_2, a_2^{-1}, a_2 and so on until a_k, a_k^{-1}, a_k , for $k \geq 1$.

In particular, for $k = 1$ we obtain a triangle whose three edges are identified in the order a_1, a_1^{-1}, a_1 , i.e., SB^1 is the dunce hat. More generally, SB^k consists of k vertices, k edges, and a single disk; so the Euler characteristic equals one; cf. Figure 3.3, which explains the name. We use notation like SB_x^k and SB_y^k to denote specific triangulations.

Theorem 3.10. *The following holds for k -bladed saw blade complexes SB^k :*

- (i) *The dunce hat SB^1 can be triangulated with 8 vertices.*
- (ii) *SB^2 can be triangulated with 9 vertices.*
- (iii) *SB^k can be triangulated with $3k$ vertices, for $k \geq 3$.*
- (iv) *Any triangulation of a saw blade complex is contractible, but non-collapsible.*

Proof: We first prove (iii): For $k \geq 3$, assume that the identified boundary of the polygonal disk reads $1-2-1-2-3-2-3-4-3-4-\dots-k-1-k-1$; see Figure 3.3 in the case $k = 3$. In the interior of the disk we place a cycle with $2k$ vertices and connect the cycle vertices with the boundary cycle vertices. More precisely, we connect every other cycle vertex to the beginning vertex of a blade and its two neighbors, and we connect the remaining cycle vertices to the two middle vertices of a blade. Finally, the interior $2k$ -gon can be triangulated arbitrarily without additional vertices.

(ii)+(i): In the case of two blades, we start with $1-2-1-2-3-1-3-2-3-1$ as the identified boundary cycle. The extra vertex, say 3, is needed to avoid unwanted additional identifications. In the interior we then place a hexagon and connect

Table 3.7: Discrete Morse vectors for 10^3 runs on the 20-simplex.

Discrete Morse vector	Count
(1, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	987
(1, 0, 0, 0, 6, 26, 59, 87, 61, 13, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 3, 30, 111, 158, 132, 82, 24, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 8, 34, 80, 126, 155, 126, 61, 27, 10, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 14, 27, 24, 13, 3, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 30, 117, 278, 409, 393, 213, 39, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 2, 25, 110, 236, 305, 175, 19, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 3, 5, 9, 34, 85, 134, 109, 33, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 19, 82, 150, 161, 90, 15, 0, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 3, 18, 51, 118, 196, 264, 207, 57, 0, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 11, 107, 243, 366, 463, 450, 261, 54, 0, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 1, 5, 30, 95, 160, 163, 124, 72, 27, 7, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 6, 48, 182, 377, 657, 876, 801, 493, 170, 22, 0, 0, 0, 0, 0, 0, 0)	1
(1, 0, 0, 0, 0, 0, 0, 8, 14, 13, 14, 7, 0, 0, 0, 0, 0, 0, 0)	1

its vertices similar to before. For an 8-vertex triangulation of the dunce hat SB^1 see [BL13a].

(iv): The dunce hat SB^1 is contractible, and none of its triangulations is collapsible [Zee64]. For $k \geq 2$, the saw blade complex SB^k has k vertices that (by labeling appropriately) appear in order $1-2-1-2-3-2-3-4-3-4-\dots-k-1-k-1$ on the boundary of the original $3k$ -gon. We cut the $3k$ -gon along an interior arc into two polygonal disks with identifications on the boundary, $1-2-1-2$ -(interior arc) and the remainder $2-3-2-3-4-3-4-\dots-k-1-k-1$ -(interior arc). Both parts are contractible CW complexes (that retract to the paths $1-2$ and $2-3-4-\dots-k-1$) glued along a contractible subcomplex (the interior arc). We conclude that their union SB^k is contractible; cf. Hachimori [Hac08] for a similar decomposition of a contractible 2-complex.

Now we consider any triangulation X of the saw blade complex SB^k , for $k \geq 1$. There is no free 2-face as all edges in such a triangulation either have degree two or three. It follows that X is non-collapsible. \square

Saw blade triangulations with different numbers of blades are combinatorially non-isomorphic. These simplicial complexes and their quotients provide 2-dimensional contractible, but non-collapsible simplicial complexes on which we can get stuck when trying to randomly collapse a simplex. For instance, the simplicial complex SQ from Figure 3.2 is obtained as a quotient from identifying two vertices in some triangulation of SB^3 . Similar examples and higher-dimensional analogs exist in abundance.

Tables 3.6 and 3.7 give the actual discrete Morse vectors found for the 8-simplex and the 20-simplex, respectively. We observe that we can get stuck (i.e., run out of free faces) in different dimensions; see [LN21] for an analysis of this phenomenon.

While in the case of the 8-simplex at most two extra critical cells are picked up (see Table 3.6), the discrete Morse vector

$$(1, 0, 6, 48, 182, 377, 657, 876, 801, 493, 170, 22, 0, 0, 0, 0, 0, 0, 0, 0)$$

for the 20-simplex in Table 3.7 contains 3,632 extra critical cells. Thus, in higher dimensions, not only do we get stuck with non-collapsible, contractible subcomplexes more often, but when we get stuck, the resulting discrete Morse vectors will also be larger. This may be seen as empirical evidence for the non-approximability of perfect Morse function; cf. [BR19].

3.4.4 Iterated barycentric subdivisions

As already seen in Section 3.3, it is sometimes rather easy to find perfect discrete Morse vectors, even in fairly large simplicial complexes, provided that the complexes are nicely structured; cf. also [BL14]. Adiprasito and Izmestiev [AI15] showed that sufficiently large iterated barycentric subdivisions of any PL sphere admit spherical discrete Morse functions. Yet, the average number of critical cells for random discrete Morse vectors grows exponentially with the number of barycentric subdivisions [ABL17]. Here we try our sphere recognition heuristic on higher barycentric subdivisions of boundaries of simplices.

For the third barycentric subdivision $\text{sd}^3 \partial\Delta_4$ of the boundary of the 4-simplex with face vector (12600, 81720, 138240, 69120) the perfect discrete Morse vector (1, 0, 0, 1) was found in 994 out of 1000 runs of the **random-lex-last** version (cf. [ABL17] and Section (8)) of the random discrete Morse search; see Table 3.8. For $\text{sd}^4 \partial\Delta_4$ with face vector (301680, 1960560, 3317760, 1658880) the (same) perfect discrete Morse vector was found in only 844 out of 1000 runs. This suggests that the 4th barycentric subdivision is still within the “horizon” for computations with the version **random-lex-last**, while the 5th barycentric subdivision was too large to fit into the main memory of the machine we used for the experiments. The **random-lex-first** strategy behaved slightly better than **random-lex-last**; the strategy **random-random** was always successful.

3.4.5 Other input

Except for the Akbulut–Kirby spheres all the examples studied so far arise from easy to understand procedures. Searching for entirely different triangulations

Table 3.8: Discrete Morse vectors for iterated barycentric subdivisions of the 3-sphere $\partial\Delta_4$.

random-random		random-lex-first		random-lex-last	
sd ³ ∂Δ ₄ with f = (12600, 81720, 138240, 69120)					
(1,0,0,1):	1000	(1,0,0,1):	999	(1,0,0,1):	994
		(1, 1, 1, 1):	1	(1, 1, 1, 1):	6
sd ⁴ ∂Δ ₄ with f = (301680, 1960560, 3317760, 1658880)					
(1,0,0,1):	100	(1,0,0,1):	829	(1,0,0,1):	844
		(1, 1, 1, 1):	143	(1, 1, 1, 1):	107
		(1, 2, 2, 1):	19	(1, 2, 2, 1):	30
		(2, 3, 2, 1):	3	(1, 3, 3, 1):	9
		(2, 5, 4, 1):	2	(1, 4, 4, 1):	4
		(1, 3, 3, 1):	2	(2, 5, 4, 1):	2
		(1, 4, 4, 1):	1	(1, 5, 5, 1):	2
		(1, 5, 5, 1):	1	(2, 3, 2, 1):	1
				(2, 7, 6, 1):	1

of \mathbb{S}^3 , we started out with $\partial\Delta_4$ with five vertices. Then we added 525 vertices via random 0-moves, followed by 50,000 random 1-moves, followed by another 10^6 rounds of random bistellar moves where we allowed both 1- and 2-moves. This resulted in a “random” triangulation of the 3-sphere with face vector $f = (530, 50474, 99888, 49944)$, which we fed into our heuristic. The smallest discrete Morse vector found was $(1, 2192, 2192, 1)$ —far away from the perfect vector $(1, 0, 0, 1)$. This means that Step (2) fails on such input. Yet, applying bistellar moves again quickly gives back the initial $\partial\Delta_4$. We also used **GAP** to actually find a trivial presentation for the fundamental group of the example, which took 16 hours for the simplification. It could be interesting to further investigate this or similar classes of examples.

Acknowledgments

We are indebted to Witte for his original **polymake** implementation of the bistellar simplification algorithm. Further, we are grateful to Benedetti, Kerber, Mischaikow, Nanda, and five anonymous reviewers of the journal and conference version of this chapter for their valuable comments. In particular, we would like to thank one of the reviewers, who suggested to us to process the 4-manifold census of **Regina**. This lead to the results of Section 3.3.2.

Chapter 4

The worst way to collapse a simplex

In general, a contractible simplicial complex need not be collapsible. Moreover, there exist complexes which are collapsible but even so admit a collapsing sequence where one “gets stuck”, that is one can choose the collapses in such a way that one arrives at a nontrivial complex which admits no collapsing moves as we have already seen in Table 3.5. Here we examine this phenomenon in the case of a simplex. In particular we characterize all values of n and d so that the n -simplex may collapse to a d -complex from which no further collapses are possible. Equivalently, and in the language of high-dimensional generalizations of trees, we construct hypertrees that are anticollapsible, but not collapsible. The material of this chapter is based on a joint work with Andrew Newman [LN21].

4.1 Overview

Collapsibility and contractibility are not equivalent in general, more subtly there even exists simplicial complexes which are collapsible, but for which it is possible to choose a sequence of elementary collapses after which one gets stuck at a nontrivial complex that is not collapsible, as we have already seen in Example 2.4. Formally, we say that a collapsing sequence on a complex X gets stuck at a complex Y , if the collapsing sequence reduces X to Y and Y has no free faces. More broadly, we say that a collapsing sequence on X gets stuck at dimension d if the collapsing sequence gets stuck at some d -dimensional complex.

There are examples of this even in the case of a simplex. Recall that the simplex on n vertices, denoted Δ_{n-1} is the simplicial complex on ground set $\{1, \dots, n\}$ where every subset of $\{1, \dots, n\}$ is a face. It is easy to see that Δ_{n-1} is collapsible. Indeed one may choose a vertex and proceed dimension-wise, collapsing at each

step the free maximal faces that do not contain the chosen vertex. However, [BL13a] shows that there exists a sequence of elementary collapses from the 7-simplex to a triangulation of the Dunce Hat, so from there no further collapses are possible even though the 7-simplex is collapsible. This notion is central in the implicit question in the title of this chapter, what is the worst way to collapse a simplex? That is starting from the simplex Δ_{n-1} what is the maximal dimension d in which a collapsing sequence may get stuck? We answer this question with the following theorem:

Theorem 4.1. *For $n \geq 8$ and $d \notin \{1, n-3, n-2, n-1\}$, there exists a collapsing sequence of the simplex on n vertices which gets stuck at a d -dimensional complex on n vertices. In contrast, for $n \leq 7$ and d arbitrary or n arbitrary and $d \in \{1, n-3, n-2, n-1\}$ it is not possible to find a collapsing sequence of the simplex on n vertices which gets stuck at dimension d .*

Theorem 4.1 implies the following corollary.

Corollary 4.2. *For every $n \geq 8$ and $d \notin \{1, n-3, n-2, n-1\}$, there exists a contractible d -complex on n vertices with no free faces. Moreover this is best possible, for $n \leq 7$ or $d \in \{1, n-3, n-2, n-1\}$ every contractible d -complex on n vertices has a free face.*

Strictly speaking only the first part of Corollary 4.2 is implied by Theorem 4.1 and not the second part. However our proof of the second part of Theorem 4.1 proceeds by showing that for n arbitrary and $d \in \{1, n-3, n-2, n-1\}$ every contractible complex on n vertices of dimension d has a free face and for $n \leq 7$ we cite the previously mentioned result of [BD05].

Another way to state Theorem 4.1 is in terms of *anticollapsibility*. An elementary anticollapse is the reverse of an elementary collapse, we defined this formally in Section 2.5, and we say that a complex X on n vertices is anticollapsible provided that there exists a sequence of anticollapsing moves from X to Δ_{n-1} . Equivalently, there is a standard notion of Alexander duality for simplicial complexes, which we also discuss below, and X is anticollapsible if and only if its Alexander dual is collapsible. In terms of anticollapsibility, Theorem 4.1 tells us that for $n \geq 8$ and $d \notin \{1, n-3, n-2, n-1\}$, there exists a d -dimensional simplicial complex on n vertices which is anticollapsible, but has no free faces.

Within this framework of constructing complexes which satisfy some specified nonempty subset of the conditions of contractibility, collapsibility, and anticollapsibility, we are really considering properties of higher-dimensional generalizations of trees. For 1-dimensional complexes contractibility, collapsibility, and anticollapsibility are all equivalent and hold exactly for trees. Each of these properties may be generalized to higher dimensions to describe high-dimensional

trees, but the properties are no longer equivalent. In general either collapsibility or anticollapsibility will imply contractibility, but no other implications hold. High-dimensional trees have been studied from various viewpoints by, for example [LP19, Kal83, DKM09, KSS84, ABL17, BL00, KLP18]. In Section 4.2.2 we overview some of the literature on this topic, but for now we point out that our results in particular characterize values of n and d so that there are d -dimensional trees on n vertices which are anticollapsible but not collapsible.

4.2 Preliminaries

4.2.1 Alexander duality and the top dimensions

Given a simplicial complex X , there is a natural way to define an Alexander dual X^* . Here we give the definition and main theorem for this duality as described in [BT09].

Let X be a simplicial complex having vertex set V . Given a subset $\sigma \subseteq V$ let $\sigma^c = V \setminus \sigma$ denote the complementary vertex set.

Definition 4.3. The Alexander dual of X on V is the simplicial complex defined by

$$X^* := \{\sigma \subseteq V \mid \sigma^c \notin X\}.$$

It is easy to see that $X^{**} = X$. Furthermore, for simplicial complexes we have the following notion of combinatorial Alexander duality similar to classic Alexander duality for more general topological spaces.

Theorem 4.4 (Combinatorial Alexander duality [Kal83]). *Let X be a simplicial complex on n vertices. Then*

$$H_i(X) \cong H^{n-i-3}(X^*).$$

where $H^{n-i-3}(X^*)$ is the $(n - i - 3)$ rd cohomology group of X .

We don't formally define (reduced) cohomology here, but it suffices for our purposes to just mention that if X is acyclic then all its cohomology groups vanish and that $H^0(X) = 0$ if and only if X is connected.

The Alexander dual behaves exceptionally well with respect to collapsibility, indeed the dual of an elementary collapse in X is an elementary anticollapse in X^* . This is standard to check, but we prove it in Proposition 4.5.

Proposition 4.5. *If X collapses to Y then X^* anticollapses to Y^* . In particular if X is collapsible then X^* is anticollapsible.*

Proof. Let X be a simplicial complex on the ground set V with $|V| = n$. We show that an elementary collapse on X corresponds to an elementary anticollapse on X^* . Suppose that $\tau = [x_0, \dots, x_k]$ is free in X with unique coface $\tau' = [x_0, \dots, x_k, x_{k+1}]$ and we perform the elementary collapse removing τ and τ' to arrive at X' . We show that $(X')^*$ is obtained from X^* by an elementary anticollapse. The claim will then follow by induction.

Since $X' = X \setminus \{\tau, \tau'\}$, we have that $(X')^* = \{\sigma \subseteq V \mid \sigma^c \notin X\} \cup \{\tau^c, (\tau')^c\} = X^* \cup \{\tau^c, (\tau')^c\}$. Thus $(X')^*$ is obtained from X^* by adding $\tau^c = V \setminus [x_0, \dots, x_k]$ and $(\tau')^c = V \setminus [x_0, \dots, x_{k+1}]$. Therefore $(\tau')^c$ is an $(n - k - 3)$ -simplex and τ^c is an $(n - k - 2)$ -simplex with $(\tau')^c \subseteq \tau^c$. Moreover since $\tau, \tau' \in X$, we have that $(\tau')^c, \tau^c$ do not belong to X^* . Thus we only have to check that all of the facets of τ^c different from $(\tau')^c$ are contained in X^* . Let $\sigma \subseteq \tau^c$ and suppose that $\sigma \notin X^*$, then $\sigma^c \in X$ and $\tau \subseteq \sigma^c$, but since τ is free we have that $\sigma^c = \tau'$. Thus $(X')^*$ is obtained from X^* by an elementary anticollapse. \square

Remark 4.6. To be completely precise when discussing duality and collapsibility we have to allow for the trivial collapse of the empty set as a free face of a simplicial complex with only one vertex. Typically, this case is not considered when discussing collapsible complexes, but observe that the dual of the simplex on n vertices is the empty simplicial complex on ground set $[n] = \{1, \dots, n\}$. Nonetheless, even if we do not allow the trivial collapse the second part of Proposition 4.5 remains true as the dual of a complex on the ground set $[n]$ with only one vertex is the boundary of the simplex on n vertices with a single $(n - 2)$ -dimensional face removed, and this anticollapses in one step to the simplex, with this step dual to the trivial collapse.

As an application of Alexander duality we see the following proof which is a simple generalization of what was done in [BD05] in the case of 7 vertices.

Proposition 4.7. *Any contractible simplicial complex on n vertices of dimension larger or equal to $n - 3$ must have at least one free face.*

Proof. The only simplicial complex of dimension $n - 1$ on n vertices is the $(n - 1)$ -simplex, which clearly has every one of its $(n - 2)$ -dimensional faces free. And of course there are no complexes of dimension bigger than $n - 1$ on n vertices.

Let us assume that X is a contractible simplicial complex on n vertices and of dimension $n - 2$. Then any $(n - 3)$ -dimensional face has $n - 2$ vertices, so can be contained only in 0, 1 or 2 $(n - 2)$ -dimensional faces of X . If all the $(n - 3)$ -dimensional faces are contained in 0 or 2 $(n - 2)$ -dimensional faces then the union of all the $(n - 2)$ -dimensional faces yields a cycle in the degree $(n - 2)$ homology group with $\mathbb{Z}/2\mathbb{Z}$ -coefficients which is impossible since the complex is contractible. Then we have at least one free $(n - 3)$ -dimensional face.

The remaining case is when the complex X is $(n - 3)$ -dimensional. In this case we can look at the Alexander dual X^* of X . By Combinatorial Alexander duality 4.4, X^* will be a connected complex on n vertices since X is contractible. Moreover, since X is $(n - 3)$ -dimensional, some edges of X^* is missing. Thus there exists vertices x , y , and z in X^* so that $[x, z], [y, z] \in X^*$, but $[x, y] \notin X^*$. But this implies that the simplex $[x, y, z]^c$ is a free $(n - 4)$ -dimensional face of X . \square

Corollary 4.8. *Given a simplicial complex X on n vertices there does not exist a collapsing sequence which gets stuck at dimension at least $n - 3$.*

4.2.2 Hypertrees

In the 1-dimensional case, a tree is characterized as a graph which is connected and has no cycles. Thus in the language of homology, a graph G is a tree if and only if $H_1(G) = H_0(G) = 0$. This definition was extended by Kalai in [Kal83] to the higher-dimensional notion of a \mathbb{Q} -acyclic complex; a \mathbb{Q} -acyclic complex X , also called a hypertree, is a simplicial complex so that $H_i(X; \mathbb{Q}) = 0$ for all $i \geq 0$.

We point out that in Kalai's original formulation a d -dimensional \mathbb{Q} -acyclic complex was defined to have complete $(d - 1)$ -skeleton. This will not be a requirement for our complexes, though in many of our constructions it will hold and gives an easy way to check that homology vanishes in degrees below $d - 1$.

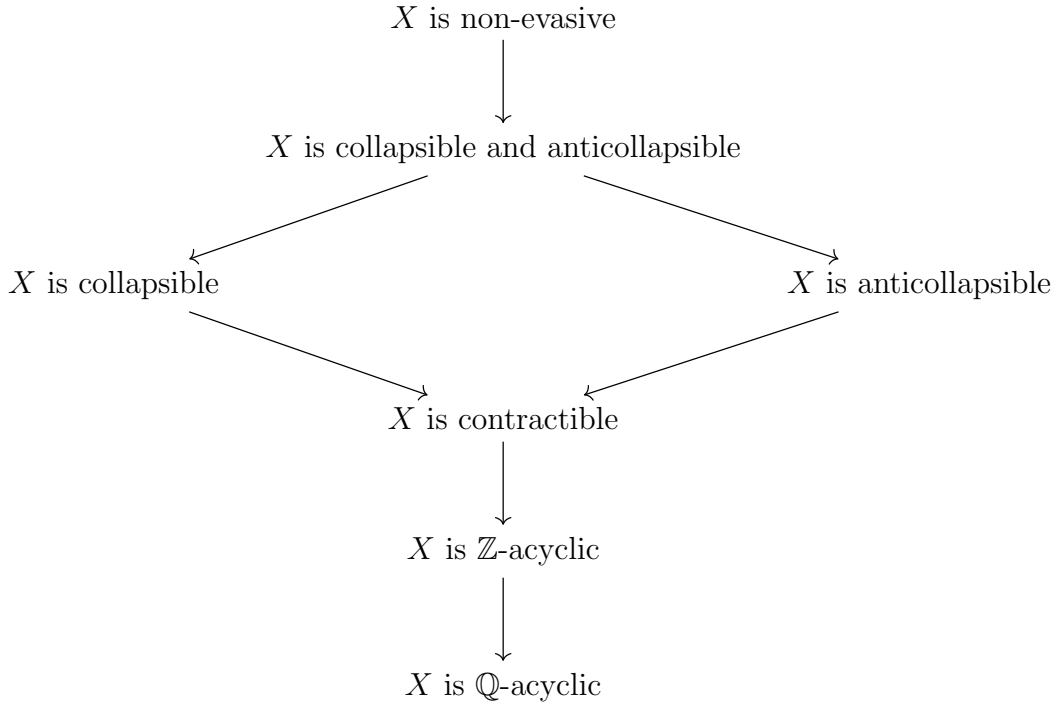
One way in which d -dimensional hypertrees are more interesting than 1-dimensional trees is in the ways that certain equivalent properties for trees generalize to nonequivalent properties for hypertrees.

One such property, non-evasiveness, which we have not yet defined, was first described in [KSS84], and, among other definitions, has the following nice inductive one:

- ▷ A single vertex is non-evasive.
- ▷ X is non-evasive if and only if there exists a vertex v of X so that both $\text{link}(v, X)$ and $\text{del}(v, X)$ are non-evasive, where

$$\begin{aligned}\text{link}(v, X) &= \{\sigma \in X \mid v \notin \sigma, [v, \sigma] \in X\}, \\ \text{del}(v, X) &= \{\sigma \in X \mid v \notin \sigma\}.\end{aligned}$$

In particular there is the following chain of implications for properties of a hypertree X .



The inductive definition of non-evasiveness may be used to show the first implication and the others are obvious. Moreover, any tree is non-evasive as any one of its leaves has single vertex link and in the $d = 1$ case a \mathbb{Q} -acyclic complex is a tree. So it is clear that all the above properties are equivalent for $d = 1$ and hold exactly for trees.

On the other hand, for $d \geq 2$, none of the implications are reversible in general. For the bottom implication in the chain, recall by the universal coefficient theorem that a complex is \mathbb{Z} -acyclic if and only if it is \mathbb{Q} -acyclic and $\mathbb{Z}/q\mathbb{Z}$ acyclic simultaneously for every prime q . Thus, for example, any triangulation of the projective plane is \mathbb{Q} -acyclic but not $\mathbb{Z}/2\mathbb{Z}$ -acyclic, so hence not \mathbb{Z} -acyclic. The standard such triangulation is given by identifying antipodal faces of the icosahedron to produce a triangulation of the projective plane with 6 vertices, 15 edges, and 10 triangles.

Continuing from bottom to top, [BL00] gives an example of a \mathbb{Z} -acyclic complex which is not contractible. Any triangulation of the Dunce Hat gives an example of a contractible, but not collapsible 2-complex. Such triangulations are given by [BL13a, Zee64]. The dual of a triangulated Dunce Hat gives an example of a contractible, but not anticollapsible complex. The example of [BL13a], that one can get stuck in dimension 2 when collapsing the 7-simplex, gives an example of an anticollapsible complex which is not collapsible. The dual of such a complex shows also that there are collapsible complexes that are not anticollapsible. Finally, [ABL17] gives an example of a complex which is evasive, but is nonetheless anticollapsible and collapsible.

4.3 Constructions

4.3.1 Explicit constructions for $n = 8$

Probably the best known example of a contractible but non-collapsible complex is the Dunce Hat [Zee64], which is known to have triangulations with 8 vertices. Benedetti and Lutz [BL13a] presented an 8-vertex triangulation of the Dunce Hat, see Figure 2.1, that can be found as a subcomplex of, and anticollapses to, a non-evasive ball with 8 vertices. This in particular implies that this triangulation is anticollapsible.

By Proposition 4.7 we know that any contractible simplicial complex on 8 vertices in dimension bigger than four has at least one free face. For $d = 3$ and $d = 4$ we considered the dual problem. We looked for 3-dimensional hypertrees and 2-dimensional hypertrees which are collapsible but have no anticollapsing moves possible and found the following examples given as a list of facets.

$$Y_8^2 := \{ [1, 2, 3], [1, 3, 4], [1, 4, 5], [1, 5, 6], [1, 3, 8], [1, 6, 8], \\ [1, 7, 8], [2, 3, 7], [3, 4, 6], [2, 4, 6], [2, 5, 8], [2, 6, 7], [2, 7, 8], \\ [3, 4, 7], [3, 5, 7], [3, 5, 8], [4, 5, 8], [4, 6, 8], [4, 7, 8], [5, 6, 7], \\ [1, 2, 6] \};$$

$$Y_8^3 := \{ [4, 6, 7, 8], [2, 5, 7, 8], [1, 5, 7, 8], [3, 4, 7, 8], [2, 4, 7, 8], \\ [2, 3, 7, 8], [1, 3, 7, 8], [2, 5, 6, 8], [3, 4, 6, 8], [1, 4, 6, 8], [2, 3, 6, 8], \\ [1, 3, 6, 8], [3, 4, 5, 8], [2, 4, 5, 8], [1, 3, 5, 8], [1, 2, 5, 8], [2, 3, 4, 8], \\ [1, 2, 4, 8], [4, 5, 6, 7], [3, 5, 6, 7], [2, 5, 6, 7], [1, 4, 6, 7], [1, 3, 6, 7], \\ [1, 2, 6, 7], [1, 4, 5, 7], [1, 3, 4, 7], [1, 2, 4, 7], [3, 4, 5, 6], [1, 4, 5, 6], \\ [2, 3, 5, 6], [2, 3, 4, 6], [1, 2, 4, 6], [1, 3, 4, 5], [1, 2, 3, 5], [1, 2, 3, 4] \};$$




Figure 4.1: An intermediate step in the construction of Y_2^8 .

The example Y_8^2 was constructed by hand. For Y_8^3 we implemented the following generalization of Kruskal's algorithm to generate d -dimensional hypertrees on n vertices and checked collapsibility and d -anticollapsibility:

1. Begin with the complete $(d - 1)$ -dimensional complex on n vertices.
2. While there are fewer than $\binom{n-1}{d}$ faces do:
 - (a) Pick σ uniformly at random from among all d -dimensional faces in Δ_{n-1} not yet considered. If σ does not complete a cycle in the top homology group with \mathbb{Q} -coefficients of the complex so far, add it to the complex.
 - (b) Otherwise, do not add σ to the complex.
3. Return the complex.

This algorithm is the higher-dimensional analogue of Kruskal's algorithm for finding a minimal-weight spanning tree in an Erdős–Rényi random graph process with weights indexing the random order in which the edges are added. Even though Kruskal's algorithm classically refers to an algorithm for finding a spanning tree, here it makes sense to consider it as an algorithm for generating a random tree. It is important to note that this algorithm in general will, even in the 1-dimensional case, *not* return a uniform spanning hypertree.

The complex Y_8^3 was found by running 10,000 trials of Kruskal's algorithm. It was one of two examples generated which was collapsible but not anticollapsible.

A later attempt using Kruskal's algorithm with 100,000 runs with $n = 8$ and $d = 2$ also yielded an example of a collapsible but not anticollapsible hypertree which could be used for the base case in place of Y_8^2 .

The collapsibility of Y_8^2 can be proved by hand. The simplicial complex in Figure 4.1 is a subcomplex of Y_8^2 and is clearly collapsible since the edge $[3, 6]$ is free and after we enter the square we can easily collapse away all the triangles. The only thing left to check is that Y_8^2 collapse to this complex but this can also easily be done in only six collapsing steps. We leave the details to the reader.

Alternatively, the reader may use, for example, the Random Discrete Morse algorithm implementation in `polymake` [GJ00] to verify that Y_8^2 and Y_8^3 are collapsible, but are dual to non-collapsible complexes, in particular not a single anti-collapsing move is possible.

The duals of these two examples and the triangulation of the Dunce Hat in Figure 2.1 gives us a proof of the following:

Proposition 4.9. *There exist simplicial complexes with 8 vertices in dimension 2, 3 and 4 that are anticollapsible, and so in particular are contractible, but with no free faces.*




Figure 4.2: Construction of Definition 4.10 applied to the Duncce Hat triangulation of Figure 2.1 with 0 as the special vertex.

4.3.2 Induction

We now want to prove the inductive step. That is, given a d -dimensional on n vertices simplicial complex X , which is anticollapsible and non-collapsible, we want to construct X' which is $(d + 1)$ -dimensional on $(n + 1)$ vertices while still being anticollapsible and non-collapsible. To do so we need the following construction.

Definition 4.10. Let X be a simplicial complex of dimension d on n vertices (x_1, \dots, x_n) and let $x := x_i$ be one of them. Given a label a we will denote by $X_{x,a}$ the simplicial complex X where the vertex x is labeled by a . We then define:

$$X_x^1 = \{[a]\} * X_{x,b} \cup \{[b]\} * X_{x,a} \quad (4.1)$$

Where $*$ is the join of two complexes. The faces of the join are the union of a face of the first complex and a face of the second one.

X_x^1 is a simplicial complex on $n + 1$ vertices $(a, b, x_1, \dots, x_{i-1}, x_{i+1}, \dots, x_n)$ of dimension $d + 1$.

This construction has also a nice presentation in the Alexander dual, in particular there is a bijection between the facets of X^* and the facets of X_x^{1*} , where a cell σ is sent to itself if it does not contain x , otherwise if it is of the form $\sigma = [x, \sigma']$ it is sent to $[a, b, \sigma']$.

We are now going to show that many interesting properties are preserved while going from X to X_x^1 , especially those we are interested in: contractibility, non-collapsibility, and anticollapsibility.

Lemma 4.11. *Let X be a d -dimensional simplicial complex with no free faces, then for any vertex x of X , we have that X_x^1 has no free faces.*

Proof. Let τ be a d -dimensional face of X_x^1 . There are then three possible cases:

- ▷ $a \in \tau$, then $\tau = [a, \tau']$ and τ' is a $(d - 1)$ -dimensional face of $X_{x,b}$, in particular it is contained in at least two facets σ and σ' , which implies that τ is contained in $[a, \sigma]$ and $[a, \sigma']$.
- ▷ $b \in \tau$, which is exactly the same as above.
- ▷ $a, b \notin \tau$, but this clearly implies that τ is contained in $[a, \tau]$ and $[b, \tau]$, so is not a free face.

□

We turn now to discrete Morse theory and, given an acyclic matching on a simplicial complex X , we would like to lift it to X_x^1 . We will do this in two steps. First, recall that, by definition, we have that $\text{link}(a, X_x^1) = X_{x,b}$. Then, since $X_{x,b}$ is combinatorially isomorphic to X , we start by lifting the entire matching to the cells that contain a ; i.e. given a matching pair (σ, τ) in $X_{x,b}$ we add the pair $([a, \sigma], [a, \tau])$ to our newly defined matching in X_x^1 . We could now be tempted to do the same with respect to b , but it can be easily seen that in this way we will obtain something not well defined. Instead what we do is to look at the restriction of the initial matching to $\text{del}(x, X)$ and lift it to the cells that do not contain a . We describe this construction formally below.

Construction of a matching on X_x^1 . Given an acyclic matching M on X and a vertex x , we will call by M_b the same matching on $X_{x,b}$. We then construct a matching M_x^1 on X_x^1 in the following way.

Let $(\sigma, \tau) \in M_b$ be a matching pair with $\sigma \subset \tau$, then:

- ▷ $([a, \sigma], [a, \tau]) \in M_x^1$,
- ▷ If $b \notin \sigma$ then $(\sigma, \tau) \in M_x^1$,
- ▷ If $b \notin \tau$ then $([b, \sigma], [b, \tau]) \in M_x^1$.

Lemma 4.12. *The matching defined above is acyclic and, if the critical cells of the matching on X forms a subcomplex Y then, the critical cells of the lifted matching on X_x^1 are exactly the cells of Y_x^1 .*

Proof. First of all, by construction, we immediately obtain that the collection of edges defined above is a matching.

The fact that it is acyclic follows from the Patchwork Theorem 2.7 where $Q = \{0, 1\}$ and the poset map is the map that sends a cell to 1 if it contains a and to 0 otherwise. This is clearly a well-defined poset map and the matching can be restricted to the fibers, so proving that our matching is acyclic is equivalent to proving that the matching restricted to each fiber is acyclic. The matching on the fiber of 1 is clearly acyclic because it is equivalent to the starting matching on

X . We need now to check that the matching on the fiber of 0, i.e. the matching restricted to the cells that do not contain a , is acyclic. We are going to prove this by contradiction.

Let $\tau_0 \searrow \sigma_0 \nearrow \tau_1 \searrow \sigma_1 \cdots \nearrow \tau_k = \tau_0$ be a cycle in the directed Hasse diagram of X_x^1 , i.e. for each i , (σ_i, τ_i) is a pair in the matching while (σ_i, τ_{i+1}) is not a pair in the matching, but σ_i is a face of τ_{i+1} .

Let τ'_i and σ'_i be the restrictions of these cells to the vertices different from b . By construction we obtain that for each i the pair (σ'_i, τ'_i) is a matched pair in X or $\sigma'_i = \tau'_i$ and equivalently σ'_i is a face of τ'_{i+1} and is not paired with it or the two cells are equal.

Then the restriction is still a cycle in X . But since the matching on X is acyclic we must have that all the restrictions are equal to τ'_0 which is impossible.

Let us now suppose that the critical cells of the matching on X forms a sub-complex Y .

Let σ be a cell of X_x^1 , we will show that σ is critical if and only if it belongs to Y_x^1 . To do so we need to analyze various cases separately.

▷ $a \in \sigma$. Let us then write $\sigma = [a, \sigma']$. The following chain of implications is true:

$$\sigma \text{ is critical in } M_x^1 \Leftrightarrow \sigma' \text{ is critical in } M_b \Leftrightarrow \sigma' \in Y_{x,b} \Leftrightarrow \sigma \in Y_x^1.$$

▷ $a \notin \sigma, b \in \sigma$. As before let us write $\sigma = [b, \sigma']$ and we obtain the exact same chain:

$$\sigma \text{ is critical in } M_x^1 \Leftrightarrow \sigma' \text{ is critical in } M_b \Leftrightarrow \sigma' \in Y_{x,b} \Leftrightarrow \sigma \in Y_x^1.$$

▷ $a, b \notin \sigma$. This last case again follow from a simple chain of implications:

$$\sigma \text{ is critical in } M_x^1 \Leftrightarrow \sigma \text{ is critical in } M_b \Leftrightarrow \sigma \in Y_{x,b} \Leftrightarrow \sigma \in Y_x^1.$$

□

We should notice that, while lifting the matching, we do not really need for x to be a vertex of X or Y . In the case $x \notin Y$ by Y_x^1 we mean, with a slight abuse of notation, the double cone over Y on the new vertices a and b , and the same if $x \notin Y$. This can be observed if we recall that Y_x^1 is a union of $\{a\} * Y_{x,b}$ and $\{b\} * Y_{x,a}$ where $Y_{x,a}$ is the labeled complex resulting from relabeling vertex x as a , and likewise for $Y_{x,b}$. The previous lemma is still true in these special cases.

Using this newly constructed matching and simple homotopy theory we are now able to show that our construction preserves contractibility.

Corollary 4.13. *Given a contractible simplicial complex X and any $x \in X$ we have that X_x^1 is contractible.*

Proof. By [Thm. 21 [Whi39]] any contractible simplicial complex can be reduced to a point by a sequence of elementary collapses and anticollapses. Let $X = X_0 \downarrow Y_1 \uparrow X_1 \downarrow \dots \downarrow Y_k = \{v\}$ be one of such sequences where by $X_{i-1} \downarrow Y_i$ we mean that X_{i-1} collapses to Y_i , while by $Y_i \uparrow X_i$ that Y_i anticollapses to X_i . Each step of these sequences is in particular an acyclic matching on a X_i . We can then use the lifting of the matching defined above and obtain that for any $x \in X$, X_x^1 is homotopy equivalent to $(Y_k)_x^1$. If $v = x$ then $(Y_k)_x^1$ is a segment; otherwise, with the same abuse of notation already discussed, $(Y_k)_x^1$ is the union of two segments attached at one vertex. In both cases $(Y_k)_x^1$ is contractible, which means that X_x^1 is contractible. \square

Remark 4.14. It can be shown that for any nonempty complex X , the complex X_x^1 satisfies for all $i \geq 0$, $H_{i+1}(X_x^1) = H_i(X)$ and $H_0(X_x^1) = 0$. Recall that $X_x^1 = \{[a]\} * X_{x,b} \cup \{[b]\} * X_{x,a}$. Therefore X_x^1 is a union of two contractible complexes whose intersection is $\text{del}(x, X) \cup (\{a, b\} * \text{link}(x, X))$. From this a routine Mayer–Vietoris argument may be applied to show the shift in homology from X to X_x^1 . We omit the proof as it isn't necessary to the discussion of contractible complexes.

Many other properties of X are also preserved by X_x^1 , for example non-evasiveness. We do not prove it here, but it is easy to check.

Lemma 4.15. *Let X be a simplicial complex on n vertices of dimension d without free faces. If X anticollapses to Δ_{n-1} , then for any vertex $x \in X$, X_x^1 is a simplicial complex on $n + 1$ vertices of dimension $d + 1$ without free faces that anticollapses to the simplex Δ_n .*

Proof. The statement follows immediately from the previous lemmas. Notice that, for any $n \in \mathbb{N}$, and any vertex $x \in \Delta_{n-1}$, $(\Delta_{n-1})_x^1$ is combinatorially isomorphic to Δ_n . \square

Lemma 4.15 is the main inductive tool to prove Theorem 4.1. Namely, Lemma 4.15 tells us that if Theorem 4.1 holds for (n, d) then it holds for $(n + 1, d + 1)$.

We almost have the full proof of Theorem 4.1. To finish the proof we will show that if the theorem holds for $(n, 2)$ then it holds also for $(n + 1, 2)$. Fortunately this can be easily accomplished by the following proposition.

Proposition 4.16. *If X is a simplicial complex on n vertices of dimension d without free faces that anticollapses to the simplex Δ_{n-1} , then Y obtained from X by deleting a facet and adding the cone over its boundary is anticollapsible to the simplex Δ_n and has no free faces.*

Proof. It is obvious that the complex Y still has no free faces.

We now check anticollapsibility. Let v be the new vertex of Y and σ the facet of X that we have deleted. By construction we can perform the elementary anticollapse $([\sigma], [v, \sigma])$ and call Y' the new complex obtained. We now have that $\text{del}(v, Y') = X$, and since X anticollapses to Δ_{n-1} we can perform the same anticollapsing moves to Y' obtaining a new complex Y'' . Now $\text{del}(v, Y'') = \Delta_{n-1}$ and $\text{link}(v, Y'') = \sigma$ which are both non-evasive. In particular Y'' is non-evasive and therefore anticollapsible. \square

From this we immediately obtain Theorem 4.1 as Proposition 4.16 implies that if Theorem 4.1 holds for $(n, 2)$ then it holds $(n + 1, 2)$, and the Duncie Hat on 8 vertices gives the base case to this induction, and the other cases have already been proved.

Remark 4.17. The change to a complex described in the proof of Proposition 4.16 is called a bistellar-0 flip or a stacking move. Moves of this type are described in [Pac87, Pac91].

4.4 Conclusions

A motivation to consider the topic discussed here comes from computational experiments of the previous chapter. Where, examining a *random* collapse procedure on simplices of increasing dimension and quite surprisingly it seems, at least empirically, that the probability to get stuck increases exponentially as $n \rightarrow \infty$. For example the experiments in table 3.5 showed that in only 12 out of one billion attempts did a random collapsing sequence fail to reach a single vertex on the 8-simplex, but on the 16-simplex 4 trials out of 10,000 failed to reach a single vertex. In the largest example, 46 out of 50 collapsing sequences on the 25-simplex failed. Here we characterize where a collapsing sequence can get stuck, but the expected behavior of a random collapsing procedure remains unclear.

Questions about random collapses of the simplex are also related to the questions about hypertree enumeration. Rather than analyzing random collapses, one could take a uniform distribution over all d -complexes that the n -simplex can collapse to and ask how many are collapsible. This is a special case of the problem of enumerating different types of hypertrees. Questions of this type appear to be quite difficult. For $d = 1$, there is the well known enumeration commonly called Cayley's formula that show that there are n^{n-2} labeled trees on n vertices. However, for larger values of d it is not even known how many d -dimensional hypertrees on n vertices there are. The closest we have to an enumeration formula is the following classic result of Kalai giving a weighted enumeration formula.

Theorem 4.18 (Kalai [Kal83]). *For $n, d \geq 1$ let $\mathcal{T}_{n,d}$ denote the collection of d -dimensional \mathbb{Q} -acyclic complexes on n vertices with complete $(d - 1)$ -skeleton,*

then

$$\sum_{X \in \mathcal{T}_{n,d}} |H_{d-1}(X)|^2 = n^{\binom{n-2}{d}}.$$

This weighted enumeration formula was later extended by Duval, Klivans, and Martin [DKM09] to the case where the $(d-1)$ -skeleton need not be complete in analogy to how Kirchhoff's Matrix Tree theorem generalizes Cayley's formula. Unweighted enumeration remains an open problem, with the best known bounds given by Linial and Peled [LP19]. In [LP19] the authors conjecture that almost all d -dimensional hypertrees, with complete $(d-1)$ -skeleton, are not d -collapsible for $d \geq 2$. Similar conjectures could be made for other properties of hypertrees. For instance is it true that almost all collapsible d -dimensional hypertrees fail to be anticollapsible for $d \geq 2$?

Our constructions give examples of complexes that are anticollapsible but not collapsible. Naturally, one could ask for contractible complexes that are *neither* anticollapsible nor collapsible. The example C_3^8 below is such a complex. This example was found by using Kruskal's algorithm to generate 500,000 examples of 3-dimensional hypertrees on 8 vertices. We point out that for $d \geq 3$ any d -dimensional \mathbb{Z} -acyclic hypertree with complete 2-skeleton is necessarily contractible by the Whitehead theorem. In our 500,000 runs C_3^8 was one of three examples that was neither 3-collapsible nor 3-anticollapsible.

$$\begin{aligned} C_3^8 := \{ & [1, 5, 7, 8], [3, 4, 5, 8], [1, 2, 6, 7], [1, 2, 3, 5], [1, 3, 4, 6], \\ & [2, 4, 7, 8], [4, 5, 6, 7], [2, 3, 7, 8], [1, 3, 5, 6], [2, 4, 5, 8], \\ & [1, 3, 4, 8], [2, 3, 4, 5], [1, 2, 4, 6], [2, 4, 6, 7], [2, 4, 5, 7], \\ & [1, 3, 5, 7], [1, 3, 4, 5], [2, 3, 6, 7], [3, 5, 7, 8], [3, 4, 5, 7], \\ & [1, 3, 4, 7], [2, 3, 6, 8], [2, 3, 4, 6], [1, 3, 7, 8], [1, 5, 6, 7], \\ & [2, 5, 6, 8], [4, 6, 7, 8], [1, 5, 6, 8], [2, 3, 5, 6], [1, 2, 3, 8], \\ & [3, 4, 6, 8], [1, 2, 5, 7], [1, 2, 4, 8], [5, 6, 7, 8], [3, 4, 6, 7] \} \end{aligned}$$

Non-collapsibility and non-anticollapsibility for this complex may be verified using the implementation of the Random Discrete Morse algorithm in **polymake**. For both C_3^8 and $(C_3^8)^*$ the Random Discrete Morse algorithm returns a Morse vector with a critical cell in the top dimension.

Even without this example, one can establish that contractible complexes which are neither collapsible nor anticollapsible must exist. Indeed all possible collapsible sequences from a complex on n vertices are finite, but deciding if a complex is contractible is undecidable as it requires deciding if the complex has trivial fundamental group [Tan16].

While C_3^8 was one of only three examples out of 500,000 randomly generated 3-dimensional trees to be contractible but neither anticollapsible nor collapsible, we suspect that for large n almost all contractible complexes are neither collapsible nor anticollapsible.

Chapter 5

Random simple-homotopy theory

After having spent some time in the previous chapter on theoretical facts about collapsibility, we now go back to computational aspects. In particular, we implement an algorithm, *RSHT* (*Random Simple-Homotopy*) to study the simple-homotopy types of simplicial complexes, with a special focus on contractible spaces and on finding substructures in higher-dimensional complexes. The algorithm combines elementary simplicial collapses with *pure elementary expansions*. For triangulated d -manifolds with $d \leq 6$, we show that RSHT reduces to (random) bistellar flips.

Among the many examples on which we test RSHT, we describe an explicit 15-vertex triangulation of the Abalone, and more generally, $(14k + 1)$ -vertex triangulations of Bing's houses with k rooms, $k \geq 3$, which all can be deformed to a point using only six pure elementary expansions. The material of this chapter is based on a joint work with Bruno Benedetti, Crystal Lai and Frank H. Lutz [BLLL21]

5.1 Overview

We propose a simple randomized algorithm to modify triangulations while keeping the simple-homotopy type of a space. The algorithm can be used as a heuristic to study simply-connected complexes, or, more generally, complexes whose fundamental group has no Whitehead torsion. We shall see that in several contractible examples the heuristics works very well. The algorithm is also of interest when applied to manifolds or complexes of arbitrary topology, as we discuss below.

Our work builds on that of Whitehead, who in 1939 introduced a discrete version of homotopy theory, called *simple-homotopy theory* [Whi39]. We recall that an *elementary collapse* is a deletion from a simplicial complex of a free face, along with the face it is contained in. Elementary collapses are deformation retracts, and thus maintain the homotopy type; the same is true for their inverse moves, *ele-*

mentary anticollapses. Two simplicial complexes are of the same *simple-homotopy type* if they can be transformed into one another via some sequence of collapses and anticollapses, called a *formal deformation* [Whi39]. By a famous result of Whitehead, having the simple-homotopy type of a point is equivalent to being contractible [Whi39] and thus undecidable.

In contrast, it is possible to decide algorithmically whether a given complex is *collapsible*, even if this decision problem is NP-complete [Tan16]. The advantage of the collapsibility notion is that all intermediate steps in the reduction are simplicial complexes of smaller and smaller size, hence very easy to encode and work with. The drawback is that collapsibility is strictly stronger than contractibility: Many “elementary” contractible complexes, like the Dunce Hat [Zee64] or Bing’s House with two rooms [Bin64], are not collapsible.

In 1998, Forman introduced a second way to study contractibility combinatorially. His Discrete Morse theory [For98, For02] is a tool to reduce simplicial complexes using a mix of collapses and facet deletions. The advantage is that all simplicial complexes (contractible or not) can now be reduced to a vertex, possibly by using a relatively large number of facet deletions. The drawback is that even if one starts with a simplicial complex, the intermediate steps in the reduction sequence are typically non-regular CW complexes, and thus harder to handle. By only focusing on the count of facet deletions (the so-called “discrete Morse vector”) it is possible to use randomness to produce fast implementations [BL14], but at the cost of failing to recognize many contractible complexes.

In this chapter, we go back to Whitehead’s original idea, and propose a third simplification method based on collapses in combination with certain expansions. Our randomized heuristic *Random Simple-Homotopy (RSHT)*; see Section 5.3) has two advantages: First, all intermediate steps are indeed simplicial complexes; and second, at the moment we do not know of a single contractible complex for which our heuristics has probability zero to succeed in recognizing contractibility.

Here is the idea. We perform elementary collapses until we get stuck. Then we select a top dimensional face ϱ uniformly at random, and for all d -faces ϱ' adjacent to ϱ via a $(d - 1)$ -dimensional ridge, we check if the subcomplex induced on the $d + 2$ vertices of $\varrho \cup \varrho'$ is a *pure d -dimensional ball*. This test is very fast. If for some ϱ' the answer is positive, we glue onto our complex the full $(d + 1)$ -simplex σ on the vertices of $\varrho \cup \varrho'$. If for several ϱ' s the answer is positive, we simply choose one uniformly at random.

This glueing step is called a *pure elementary $(d + 1)$ -expansion*, and it is also classical from the topological perspective, compare [HAM93, Chapter I]. After this step, we may collapse away the newly introduced $(d + 1)$ -simplex σ together with any d -face τ of it. To avoid undoing the pure elementary expansion, we must select a τ that was already present in the complex we got stuck at before the pure

elementary expansion. This first elementary collapse after the pure elementary expansion is called “(CC) step” below (see Section 5.3). The combination “pure elementary expansion + (CC) step”, known in the topological literature as “transient move” [HAM93], maintains both the dimension and the simple-homotopy type: In fact, any pure elementary expansion can be viewed as a composition of back-to-back elementary anticollapses.

Whitehead proved that for every contractible complex there is a formal deformation that reduces it to a single point [Whi39]. It is not known if there is also a formal deformation to a point in which one performs anticollapses or expansions “only when stuck”, i.e., only to intermediate complexes without free faces. If this is true, then indeed any contractible complex would have a positive probability to be recognized by our heuristics. Of course, we cannot in any case expect any universal upper bound on the number of elementary anticollapses needed, or else we would have found an algorithm that recognizes contractibility.

However, we shall see in Sections 5.5 and 5.6 that in many key examples the number of pure elementary expansions needed is relatively low. As a benchmark series, we build *Bing’s house with k rooms*, a one-parameter generalization of the aforementioned Bing’s house with two rooms. For all $k \geq 3$, we prove that Bing’s house with k rooms can be collapsed by adding only six further tetrahedra, cleverly chosen (**Theorem 5.12**). Of course, since our algorithm is randomized, there is no guarantee that precisely those tetrahedra will be selected as expansions. But even with a quick attempt consisting of 10^4 runs, our algorithm is able to reduce Bing’s house with seven rooms (which is a 2-complex on 99 vertices) to a point by adding only 41 tetrahedra; see Table 5.1.

Random Simple-Homotopy (RSHT) works with simplicial complexes of arbitrary dimension, but it is of particular interest when applied to triangulations of low-dimensional manifolds. When $d \leq 6$, we show (in **Theorem 5.8**) that on any (closed) d -manifold RSHT has basically the same effect of performing *bistellar flips*, also known as *Pachner moves*, which are the standard ergodic moves that allow to transform into one another any two PL homeomorphic triangulations of the same manifold [Pac87].

In Section 5.6, we discuss how RSHT can be used to reach interesting small (or even vertex-minimal) triangulations and subcomplexes “hidden” inside triangulated manifolds. For the sake of applications, one should declare right away that RSHT is designed to focus on the (simple-)homotopy, and not the homeomorphism type. So in case we start with a collection of points in 10-space, say, which all lie “approximately” on a Möbius strip, the effect of performing RSHT on the Čech complex of the point set would be to detect an S^1 , and not a Möbius strip. Yet, RSHT is capable of detecting, for example, closed surfaces or higher-dimensional closed manifolds in data, beyond just determining their homologies.

It takes considerable effort to build examples of contractible complexes for

which RSHT does *not* practically succeed in revealing contractibility, if interrupted after a million steps, say. Respective examples, showcased in the last Section 5.6.4 of this chapter, are based on the Akbulut–Kirby 4-spheres [AK85] and triangulations thereof [TL] already used in previous chapters. The homeomorphism type of these “tangled” triangulations of S^4 is notoriously difficult to recognize.

5.2 Pure elementary expansions

Any two simple-homotopy equivalent complexes are homotopy equivalent. The converse is true for complexes whose fundamental group has trivial *Whitehead group* (see [Coh12] or [Mne14] for the definition), but false in general: Counterexamples in dimension two can be obtained by triangulating the cell complexes in [Lus91], while counterexamples in dimension three or higher had been known to exist long before [Mil66].

It is an easy consequence of the theory of Gaussian elimination for integer matrices that the Whitehead group of the trivial group is trivial. Therefore, any two homotopy-equivalent simply-connected complexes are also simple-homotopy equivalent. More generally, it is known that the Whitehead group of a group G is trivial if G is

- ▷ \mathbb{Z} [Hig40], $\mathbb{Z} \oplus \mathbb{Z}$ [BHS64], and more generally, any free Abelian group [BHS64],
- ▷ any of the cyclic groups $\mathbb{Z}_2, \mathbb{Z}_3, \mathbb{Z}_4, \mathbb{Z}_6$ [Coh12],
- ▷ any subgroup of the braid group B_n [FR00], or of any Artin group of type $A_n, D_n, F_4, G_2, I_2(p), \tilde{A}_n, \tilde{B}_n, \tilde{C}_n$, or $G(de, e, r)$ ($d, r \geq 2$) [Rou20];
- ▷ any free product of groups listed above, so in particular $\mathbb{Z} * \mathbb{Z}$ or any free group [Sta65];
- ▷ and further cases [GB07]; in fact, the Farrell–Jones conjecture implies that any torsion-free group should appear in the present list [LRRV17].

Any two homotopy-equivalent complexes whose fundamental group appears in the list above are of the same simple-homotopy type.

Whitehead’s work allows us to be more specific on the dimension (although not on the number) of the intermediate complexes involved in the definition of simple-homotopy equivalence, as follows. An elementary collapse is called an *i-collapse* if the dimension of the two faces removed are $i - 1$ and i . Similarly, an *i-anticollapse* is one that adds a pair of faces of dimension $i - 1$ and i . The *order* of a formal deformation will be the maximum i for which *i*-collapses or *i*-anticollapses are involved in the sequence.

Theorem 5.1 (Whitehead [Whi39, Theorems 20 & 21]). *Let X and Y be two homotopy-equivalent simplicial complexes. If the Whitehead group of their fundamental group is trivial, then there is a formal deformation from X to Y whose*

order does not exceed $\max\{\dim X, \dim Y\} + 2$. If, in addition, Y is a deformation retract of X , and $\dim X > 2$, then there is a formal deformation from X to Y whose order does not exceed $\dim X + 1$.

The conjecture that the last statement of the previous theorem might hold also for the case $\dim X = 2$ goes under the name of “Generalized Andrews–Curtis conjecture”, and represents a major open problem in combinatorial topology. It is, however, generally believed to be false [HAM93].

Based on Whitehead’s work, we would now like to perform “random anticollapses”. Yet, if we wish to add a $(d + 1)$ -dimensional face σ to X in an elementary anticollapse, then all d -faces of σ need to be present in X already, except for a single d -face τ . However, it is not difficult to construct contractible d -complexes X that do not allow any $(d + 1)$ -anticollapses; cf. [LN21]. In these cases, lower-dimensional faces need to be added first. Computationally, this brings an extra difficulty to the introduction of a random model. To bypass this difficulty, we adopt a different set of moves.

Definition 5.2. Let X be a d -dimensional complex. A *pure elementary $(d + 1)$ -expansion* is the gluing of a $(d + 1)$ -dimensional simplex σ to X in case σ intersects X in a d -ball.

A pure elementary $(d + 1)$ -expansion combines together in a single move one $(d + 1)$ -anticollapse plus all the lower-dimensional anticollapses that have to be performed first. Hence a sequence of pure elementary expansions and elementary collapses can be rewritten as a formal deformation. Whenever we run out of collapsing steps, we perform exactly one pure elementary $(d + 1)$ -expansion, and then switch back to elementary collapses. When the complex is reduced to a point, we stop.

5.3 Implementation of random simple-homotopy

Algorithm RSHT provides a description of the Random Simple-Homotopy procedure in pseudocode. The actual implementation can be found on GitHub at [Lof21b] as a `polymake` [GJ00] extension. It is based on the two different types of basic operations: random *collapses* (C) and random *pure elementary expansions* (E) plus collapsing steps (CC) that ensure that a pure elementary expansion is not undone immediately by the next regular collapsing step (C). The step (S) allows facet *subdivisions* in case no other pure elementary expansions are available.

Random collapses (C) are discussed as part of Random Discrete Morse Theory in [BL14]. A fast implementation of random collapses in `polymake` is described in [JLLT21]. Hence, it remains to implement random pure elementary expansions (E).

Algorithm RSHT A: Random Simple-Homotopy

Input: simplicial complex X **Output:** simplified simplicial complex**while** $\dim(X) \neq 0$ **and** $i < \text{max_step}$ **do** **while** X has free faces **do**

| (C): perform a random elementary collapse

if $\dim(X) = d \neq 0$ **and** there are induced pure d -balls on $d + 2$ vertices **then** | (E): perform a random pure elementary $(d + 1)$ -expansion | (CC): perform an elementary collapse deleting the newly added $(d + 1)$ -face and one of its d -faces that was already in X **else** | (S): perform (E) + (CC) on a d -facet with $d + 1$ vertices | $i++$ **return** X

While collapses in `polymake` can be carried out fast in the Hasse diagram of X , there is no explicit implementation in `polymake` to expand the Hasse diagram of X to include the faces of a $(d + 1)$ -simplex σ that is added in a pure elementary expansion. Thus, for every pure expansion we recompute the Hasse diagram for $X + \sigma$ and then proceed with random collapses in the new Hasse diagram of $X + \sigma$. For various input examples of non-collapsible, contractible complexes, relatively few pure expansions are needed (see Sections 5.5 and 5.6); thus the extra cost for recomputing Hasse diagrams stays low.

Remark 5.3. Pure elementary expansions are not the only operations to modify a given complex X by expanding it. Another more general possibility would be to glue additional $(d + 1)$ -simplices to X along induced contractible subcomplexes (of mixed dimension). This provides more options to modify X , but at the price of having to check subcomplexes for contractibility. As we experienced after running various experiments, this seems expensive without any advantage. We therefore decided to stick to pure elementary expansions. In fact, checking whether an induced subcomplex on $d + 2$ vertices is a pure d -ball is very fast: It can be achieved by a lookup in the Hasse diagram.

Remark 5.4. By Whitehead's Theorem 5.1, we might be forced to first go up by two dimensions (and not just by one as we do in Algorithm RSHT) to find a formal deformation from a complex X to some homotopy equivalent complex L . This could be incorporated in the algorithm by performing not only single pure elementary $(d + 1)$ -expansions followed immediately by collapses, but by allow-




Figure 5.1: Expansions as bistellar flips.

ing sequences of pure elementary $(d+1)$ -expansions followed by pure elementary $(d+2)$ -expansions before switching back to collapses. In principle, this generalized procedure could be set up in a simulated annealing fashion, in a completely analogous way to what we do here; but for the examples we study in the subsequent Sections 5.5 and 5.6, we shall restrict ourselves to the basic algorithm RSHT, as this already works well.

5.4 Bistellar flips and artefacts

Pure elementary $(d+1)$ -expansions have (at least for d -manifolds in low dimensions $d \leq 6$) a clear interpretation in terms of bistellar flips. In fact, let X be a d -complex. In a pure elementary $(d+1)$ -expansion, some $(d+1)$ -simplex σ is glued to X along a d -ball consisting of $1 \leq k \leq d+1$ of the d -faces of σ ; let r be the intersection of these k faces. If r is contained in no further d -face of X , then after adding σ , collapsing it away with one of the k d -faces, and collapsing further lower-dimensional faces, we are left with a complex X' that is obtained from X via a bistellar move; cf. [BL00]. If instead r is contained in more than k d -faces of X , then in passing from X to X' the facet degree of r is decreased by one.

Example 5.5. If we glue a tetrahedron σ to a 2-complex X along a 2-disk in $\partial\sigma$, the disk can either consist of 1, 2, or 3 triangles. In the first case, the complex X' resulting after the collapses is a subdivision of X . (The triangle τ of X is subdivided using the unique vertex of σ not in X ; see Figure 5.1, left.) In the second case, if r is the edge common to the two triangles of $\partial\sigma$ in which σ intersects X and r is contained in exactly these two triangles of X , then r is flipped to yield X' ; see Figure 5.1. In the third case, the transition from X to X' “undoes” a subdivision.

Example 5.6. Let X be 2-dimensional and let σ be a tetrahedron glued to X along two triangles whose intersection is r , and suppose that this r is contained in exactly three triangles of X . Then after the addition of σ and its removal, r will be contained in two of the triangles of X' ; see Figure 5.2.

If, in a pure elementary 3-expansion, some tetrahedron is glued on top of two adjacent triangles ϱ_1, ϱ_2 of a triangulated 2-manifold, then, after collapsing away




Figure 5.2: Reduction of the face degree.

the tetrahedron together with ϱ_1 , the resulting triangulation will still contain ϱ_2 and (as a free face) the edge $e = \varrho_1 \cap \varrho_2$. This edge e is thus the only free (1-)face; hence, it will be selected in the incoming (C) step of RSHT. As a result, the combination (E) + (CC) + (C) is a proper bistellar flip—and the diagonal of the two initial triangles gets flipped. In the case of a subdivision, the combination (E) + (CC) is a proper bistellar flip as well. Thus, it remains to inspect the case when a subdivision is undone. After the addition of a tetrahedron (E) and the deletion of one of the initial three triangles along with the tetrahedron in the (CC) step, the other two initial triangles remain, and we have (two) free edges for two further (C) steps. In contrast to the previous cases, after the two (C) steps, the resulting triangulation is not a surface yet—as we still have the intersection vertex of the three initial triangles as a free vertex that is connected to the modified triangulated surface by an edge. That is, the result of (E) + (CC) + (C) + (C) is a triangulated surface with an additional edge sticking out. This edge is then collapsed away in another (C) step.

This situation generalizes as follows:

Lemma 5.7. *Let X be a triangulation (without free faces) of a d -manifold M and suppose that the $(d+1)$ -simplex $\sigma = [0, 1, \dots, d+1]$ intersects X in a pure d -ball B with $1 \leq k \leq d+1$ d -facets on the $d+2$ vertices $0, 1, \dots, d+1$ of σ so that (w.l.o.g.) $B = [0, 1, \dots, d-k+1] * \partial[d-k+2, d-k+1, \dots, d+1]$. We add σ (and its faces) to X and, by step (CC) of RSHT, ban those facets of σ as free faces that do not contain $[0, 1, \dots, d-k+1]$.*

- ▷ *If $k \leq 7$, then running RSHT on $X \cup \sigma$ until no further free faces are available yields a triangulation $X' = X - B + B'$ of M with*

$$B' = \partial[0, 1, \dots, d-k+1] * [d-k+2, d-k+1, \dots, d+1],$$

i.e., X' is obtained from X by a bistellar flip.

- ▷ *If $k > 7$ (which can occur for $d > 6$ only), then running RSHT on $X \cup \sigma$ until no further free faces are available might terminate in a non-pure simplicial complex X'' that is the union of a triangulation of M with a contractible, but non-collapsible lower-dimensional complex on the vertices $d-k+2, d-k+1, \dots, d+1$.*

Proof. Step (CC) of RSHT implies that our first collapsing move will remove a facet of B along with the added $(d+1)$ -simplex σ . At any consecutive collapsing step (C), the faces involved in the collapses will be of the form $[0, 1, \dots, d-k+1] * \tau$, where $\tau \in \partial[d-k+2, d-k+1, \dots, d+1]$ (because our starting complex X had no free faces). The restriction of the collapsing sequence to $\partial[d-k+2, d-k+1, \dots, d+1]$ gives us a valid collapsing sequence of the simplex $[d-k+2, d-k+1, \dots, d+1]$, where the first collapsing move is induced by the initial step (CC). Now:

- ▷ If $k \leq 7$, the simplex $[d-k+2, d-k+1, \dots, d+1]$ has at most seven vertices; and by [BD05] every contractible simplicial complex with $k \leq 7$ vertices is collapsible, i.e., the collapsing sequence induced by RSHT on $[d-k+2, d-k+1, \dots, d+1]$ will terminate at a single point. It follows that $X' = X - B + B'$.
- ▷ If $k > 7$, then the collapsing sequence on $[d-k+2, d-k+1, \dots, d+1]$ might get stuck on a contractible, but non-collapsible subcomplex of dimension at least two [CEK⁺, LN21], and thus the resulting complex X'' need not be pure. \square

Note that in the special case when $d = 7$ and $k = 8$ we might get stuck on a subcomplex $[0] * D \subseteq [0] * \partial[1, \dots, 8]$, with D an 8-vertex triangulation of the Dunce Hat; cf. [BL13a]. The resulting complex $X' = X - B + B' + [0] * D$ then deviates from the modification via a bistellar flip, $X - B + B'$, by the additional cone $[0] * D$ with apex $[0]$ over the (contractible) Dunce Hat D in the 2-skeleton of $\partial[1, \dots, 8]$. The complex $X' = X - B + B' + [0] * D$ deformation retracts to $X - B + B'$, but has no free faces.

Theorem 5.8 (Reduction of pure elementary $(d+1)$ -expansions to bistellar flips). *Let X be a triangulation of a d -manifold M with $d \leq 6$. Any pure elementary $(d+1)$ -expansion followed by collapses (as long as free faces are available) induces a bistellar flip on X .*

Proof. The statement follows from Lemma 5.7 and the fact that the maximum number of facets of a pure d -ball on $d+2$ vertices is $d+1$. \square

Corollary 5.9 (Manifold stability). *Let X be a (not necessarily pure) simplicial complex. If we run RSHT on X and at some point reach a simplicial complex X' that triangulates a d -manifold with $d \leq 6$, then from then on, whenever there are no free faces in the further run of RSHT, the respective temporary complex \tilde{X} is a d -manifold as well, and \tilde{X} is bistellarly equivalent to X' .*

To avoid lower-dimensional artefacts $[0, 1, \dots, d-k+1] * N \subseteq [0, 1, \dots, d-k+1] * \partial[d-k+2, d-k+1, \dots, d+1]$ in the modification $X' = X - B + B' + [0, 1, \dots, d-k+1] * N$ of a triangulated manifold X , involving a contractible, non-collapsible complex N for $d \geq 7$ and $k \geq 8$, we should switch to bistellar flips

$X' = X - B + B'$ once we know that X is a manifold. Quite often, this is not clear a priori—in fact, testing whether X is a manifold is an undecidable problem for $d \geq 6$; cf. [JLLT21].

In practice [JLLT21], on a 7-simplex it is nearly impossible to get stuck with random collapses. On the 8-simplex, only about 0.0000012% of the runs of random collapses get stuck. But in higher dimensions, the situation changes dramatically: For example, for the 25-simplex, contractible but non-collapsible substructures are encountered in 92% of the runs.

Another option to deal with the artefacts would be to run RSHT on lower-dimensional parts to “melt away” the artefacts. However, in our experiments in Sections 5.5 and 5.6 we only focus on top-dimensional pure elementary expansions, since the terminal triangulations of the examples we consider are all of dimension $d \leq 6$.

In case a general complex X has no free faces and is not a manifold, then a sequence (E) + (CC) + (C) + ... + (C) until no further collapses are possible might reduce X in dimension or can reduce (or increase) the degree of a face in X , as we have seen in Example 5.6 and Figure 5.2. In the latter case, we can regard the sequence as a *generalized bistellar flip*. These generalized operations give flexibility in the modification of a given complex X .

5.4.1 Selection of expansions and simplification of complexes

We next discuss in more detail how the pure elementary expansions are selected and why Algorithm RSHT has a tendency to simplify simplicial complexes to yield small or even vertex-minimal triangulations. First, we note that RSHT, apart from temporarily adding $(d + 1)$ -faces in the pure elementary expansion steps (E), never increases the dimension of the complex.

As outlined in the introduction, for any d -facet of a d -dimensional complex X , chosen uniformly at random, we can check for all neighboring d -facets whether the induced subcomplexes on the combined $d + 2$ vertices are pure d -dimensional. From the collection of all available such pure induced d -balls on $d + 2$ vertices, we pick one uniformly at random for a pure elementary d -expansion step (E). However, in general, such pure induced d -balls on $d + 2$ vertices need not exist. For example, in the case of neighborly triangulations of surfaces, the induced subcomplexes on the four vertices of two adjacent triangles are the two triangles *plus* the opposite diagonal edge; such subcomplexes are not contractible. In such a case, the only possible pure elementary expansion is by picking a facet (uniformly at random) as a pure d -ball and initiating a subdivision (S). An example of a triangulated 3-sphere on 16 vertices that allows no bistellar flips (apart from subdivisions of tetrahedra) is given in [DFM04].

Lemma 5.10. *Let X be a triangulated circle S^1 with $n > 3$ vertices. Then X is reduced by Algorithm RSHT to the boundary of a triangle in $n - 3$ pure elementary expansion steps (E), each followed by two collapsing steps (CC) + (C).*

In the case of triangulations of S^2 with $n > 4$ vertices, there always are admissible edge flips, and thus Algorithm RSHT never adds a vertex in a subdivision step (S). A vertex can get removed in the reversal of a subdivision once the current triangulation has a vertex of degree 3. However, the boundary of the octahedron has all of its vertices of degree 4; in fact, there are infinitely many triangulations of S^2 with all vertex degrees at least four. In any such example, the removal of a vertex is not immediately possible. But after a suitably long sequence of random edge flips, eventually vertices of degree 3 show up, and the three incident triangles to such a vertex have the chance to get chosen for an induced pure 2-ball to remove the vertex of degree 3.

Similarly, general complexes X are simplified and reduced in size by collapsing away collapsible parts and by reversing subdivisions to reduce the number of vertices—but without a universal guarantee for success (as contractibility is undecidable).

5.5 Classical examples

In this section, we test how the Algorithm RSHT performs on the Duncce Hat, on Bing’s House with two rooms, and on similar, “classical” examples of contractible complexes. It turns out that the number of pure elementary expansions needed to reduce these complexes to a single vertex is conveniently low: one pure elementary expansion suffices for an 8-vertex triangulation of the Duncce Hat; five pure elementary expansions suffice for a simplicial version of Bing’s house with two rooms; and in general, six tetrahedra are sufficient to collapse Bing’s house with k rooms (**Theorem 5.12**). Triangulations of these examples can be found online at the “Library of Triangulations” [BL21].

5.5.1 The Duncce Hat

The Duncce Hat [Zee64] is the most famous example of a contractible, but non-collapsible complex; compare [BL13a]. It is obtained by glueing together the three edges of a single triangle in a non-coherent way. The Duncce Hat can be triangulated as a simplicial complex with eight vertices (see Figure 5.3a); and eight vertices is fewest possible, as every contractible simplicial complex on seven vertices is collapsible [BD05]. No triangulation of the Duncce Hat is collapsible, since there are no free edges to start with.




Figure 5.3: A formal deformation of the Dunce Hat.

The Dunce Hat of Figure 5.3a admits two anticollapsing moves, the addition of the tetrahedron 1245 or alternatively the addition of the tetrahedron 1367. In Figure 5.3b we added 1367. All of the triangles in 1367 are free, since this is now the only tetrahedron present. If we collapse away the triangle 367, we recover the initial complex of Figure 5.3a. If instead we choose to delete the free triangle 136, we obtain the triangulation displayed in Figure 5.3c. This triangulation has a free edge, 16, that allows us to get rid of the triangle 167. After this elementary collapse, the edge 17 becomes free, allowing us to remove the triangle 137. But now the edge 13 is free, and it can easily be seen that the deletion of the triangle 138 paves the way to a full collapse down to a single vertex.

Lemma 5.11. *One pure elementary 3-expansion suffices to reduce to a vertex the 8-vertex triangulation of the Dunce Hat from Figure 5.3a.*

In 10^4 runs, RSHT used on average 2.4145 pure elementary 3-expansions to reduce the 8-vertex Dunce Hat to a point; see Section 5.6.1 and Table 5.1.

5.5.2 The Abalone

The Abalone [HAM93], sometimes called *Bing's House with one room*, is another example of a contractible but non-collapsible complex. We are not aware of any triangulation of this space in the literature, so we present one, **Abalone**, with 15 vertices:

1 2 7	1 2 9	1 3 8	1 3 9	1 4 7	1 4 8	1 4 9	2 3 7
2 3 15	2 9 15	3 7 8	3 9 14	3 14 15	4 5 7	4 5 8	4 6 7
4 6 9	5 6 9	5 6 10	5 7 10	5 8 9	6 7 11	6 10 11	7 8 10
7 8 11	8 9 12	8 9 13	8 10 12	8 11 13	8 12 13	9 12 14	9 13 15
10 11 12	11 12 13	12 13 14	13 14 15.				

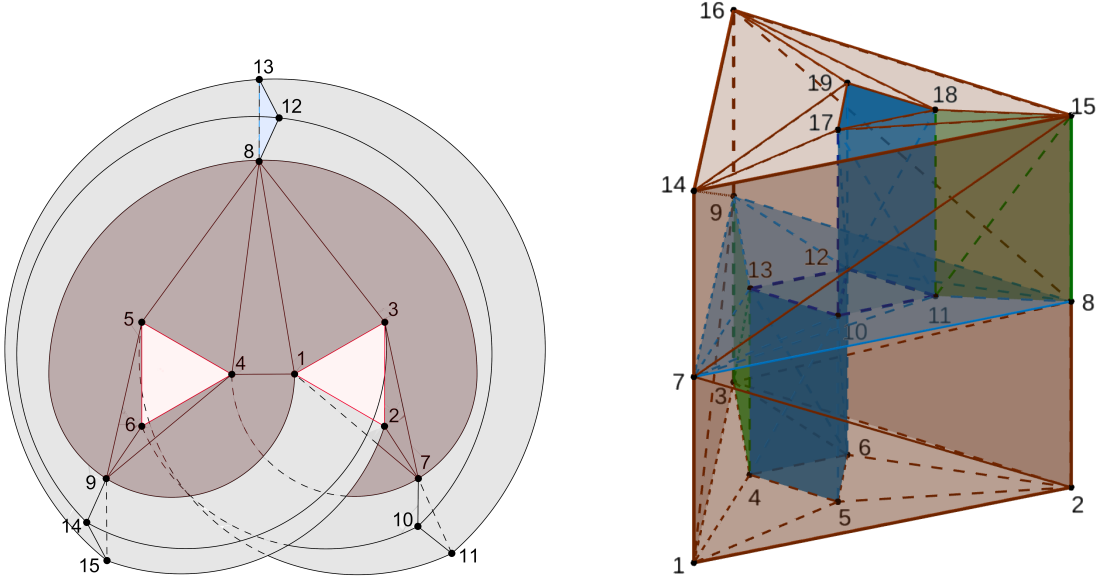


Figure 5.4: Triangulations **Abalone** of the Abalone (left) and **BH** of Bing's house with two rooms (right).

Figure 5.4 displays this triangulation, although some diagonals have been omitted for reasons of pictorial clarity. Essentially, the triangulation consists of a membrane (in dark) from which two prismatic tunnels (in light) originate at the two empty triangles 123 and 456; and the tunnels are separated by the highlighted triangle 81213. The Abalone is contractible as can be seen by filling in the two tunnels.

RSHT can reduce the Abalone to a point using only three expansions. One way to do so is to free the edge 89 of Figure 5.4 by first adding the three tetrahedra 891213, 9121314, and 9131415, in this order, as anticollapsing moves. The resulting complex is then collapsible. This can either be verified by hand, or via the `random_discrete_morse` algorithm implemented in `polymake` [BL14]: The three tetrahedra fill in the prism between the triangle 81213 and the (formerly empty) triangle 91415. By collapsing away this prism, the edge 89 becomes free so that the (dark) membrane around the empty triangle 456 can be collapsed away, which frees the tunnel originating at this empty triangle. Its removal then allows to collapse the remaining disk.

We can interpret the anticollapsing moves followed by collapses as operations that move the walls of the tunnel so that eventually the obstruction to collapsibility vanishes.

5.5.3 Bing's House with two rooms

Bing's House with two rooms [Bin64] is an early example of a contractible space no triangulation of which is collapsible. For our purposes, we triangulate Bing's House as a triangular prism with two floors, two tunnels to reach the floors, and all rectangular walls subdivided into two triangles each. Figure 5.4 displays the following (small) triangulation BH with $f = (19, 65, 47)$ (with the list of facets also available online as example BH at [BL21]):

1 2 5	1 2 7	1 3 4	1 3 9	1 4 5	1 7 9	2 3 6
2 3 8	2 5 6	2 7 8	3 4 6	3 4 13	3 8 9	3 9 13
4 5 10	4 6 13	4 10 13	4 12 13	5 6 10	6 10 12	6 12 13
7 8 11	7 8 15	7 9 13	7 9 14	7 10 11	7 10 13	7 14 15
8 9 12	8 9 16	8 11 12	8 11 15	8 15 16	9 12 13	9 14 16
10 11 17	10 12 17	11 12 18	11 15 18	11 17 18	12 17 19	12 18 19
14 15 17	14 16 19	14 17 19	15 16 18	15 17 18	16 18 19.	

RSHT is able to reduce Bing's house to a point by means of five (successive) expansions (in the upper room, each followed by collapses so that the outer walls of Bing's house are moved towards the upper tunnel). Here is a possible strategy. By successively adding five tetrahedra in the upper room of our Bing's House triangulation, we fill in a cubical prism between the horizontal square 7–8–11–10 of the medium floor and the square 14–15–18–17 of the ceiling. The first two tetrahedra 7 8 11 15 and 11 15 17 18 can be added independently, and their addition are proper anticollapsing steps. The third tetrahedron 7 11 15 17 is a pure expansion, and the addition of the two final tetrahedra 7 10 11 17 and 7 14 15 17 are again anticollapsing steps. The newly introduced cubical prism connects the outer vertical square 7–8–15–14 with the vertical square 10–11–18–17 of the upper tunnel. The resulting complex is collapsible; an explicit collapsing sequence proving this claim is detailed below.

We start from the outside, by perforating the back square 7–8–15–14. Then we entirely remove the interior of the cubical prism along with the two triangles 7 8 15 and 7 14 15 of the back square and the two triangles 14 15 17 and 15 17 18 of the top square. The result is an indented Bing's House triangulation with two new side triangles 7 10 17 and 7 14 17. But now the edge 1 7 18 has been freed, and we can use it to collapse away the subdivided squares of the triangulation one by one. First the square 10–11–18–17 is collapsed away, which frees the edge 10 11. This edge in turn can be used to remove the horizontal square 7–8–11–10, thus freeing the edge 7 8. Next, we remove the squares 1–2–8–7, 2–3–9–8, 1–3–9–7, the vertical wall 3–4–13–9, then all triangles of the lower floor, then the lower tunnel, to end up with the indented upper room with empty triangle 10 12 13. This remaining complex is a triangulated disc and thus collapsible.




Figure 5.5: Ground floor (left) and room R_2 (right) of Bing's House $BH(3)$ with three rooms.

5.5.4 Bing's House with k rooms

A recent example of a non-collapsible, contractible complex is *Bing's House with three rooms (and thin walls)* by Tancer [Tan16]. He introduced the example as a gadget to prove that the problem of recognizing collapsible complexes is NP-complete. The basic layout of the example can be found in [Tan16]. Here, we give an explicit triangulation $BH(3)$; and extend this construction to k rooms, $BH(k)$, $k \geq 3$.

The starting point for the construction of $BH(3)$ is to have a ground floor with three triangular holes as depicted in Figure 5.5. The floor has the following triangles:

1 2 5	1 2 15	1 4 6	1 4 10	1 5 7	1 6 7	1 9 11	1 9 14
1 10 12	1 11 12	1 14 16	1 15 16	2 3 5	2 13 15	3 4 6	3 5 6
4 8 10	8 9 11	8 10 11	9 13 14	13 14 15.			

Onto the ground floor, we glue three rooms in a coherent way. Room R_1 is glued onto the two regions A and B and uses nine additional vertices from 17 to 25. Room R_2 , depicted in Figure 5.5, is glued onto the regions B and C and uses the nine vertices from 26 to 34. Finally, room R_3 is glued onto the regions C and A with further nine vertices ranging from 35 to 43. The rooms R_2 and R_3 are cyclic copies of the room R_1 , where 9 and 18 are added to the vertex-labels 17 to 25 of room R_1 , respectively. Concretely, the triangles of room R_1 are

1 2 17 1 9 17 2 3 18 2 5 18 2 17 18 3 4 19 3 18 19
4 8 20 4 19 20 5 6 21 5 7 21 5 18 21 6 7 22 6 21 22
7 21 23 7 22 23 8 9 24 8 20 24 9 17 25 9 24 25 17 18 21
17 20 22 17 20 24 17 21 23 17 21 23 17 24 25 18 19 21 19 20 22
19 21 22.

Those of room R_2 are

1 2 26 1 4 26 2 13 33 2 26 34 2 33 34 4 8 27 4 10 27
4 26 27 8 9 28 8 27 28 9 13 29 9 28 29 10 11 30 10 12 30
10 27 30 11 12 31 11 30 31 12 30 32 12 31 32 13 29 33 26 27 30
26 29 31 26 29 33 26 30 32 26 31 32 26 33 34 27 28 30 28 29 31
28 30 31,

and those of room R_3 are

1 4 35 1 9 35 2 3 38 2 13 37 2 37 38 3 4 42 3 38 42
4 35 43 4 42 43 9 13 36 9 14 36 9 35 36 13 36 37 14 15 39
14 16 39 14 36 39 15 16 40 15 39 40 16 39 41 16 40 41 35 36 39
35 38 40 35 38 42 35 39 41 35 40 41 35 42 43 36 37 39 37 38 40
37 39 40.

The three rooms R_1 , R_2 , and R_3 are then all glued to the upper side of the ground floor. Since the vertices of the upper layer of a room are distinct from the vertices of the upper layers of the other two rooms, there is no conflict for the chosen gluing to the same side. To enter the interior of a room, one has to first pass through the tunnel from above of the room to the left, before the room itself can be entered from below through the lower left empty triangle.

The previous construction can be generalized to create Bing's Houses $BH(k)$ with k rooms, $k \geq 3$. Instead of just three regions, start with k regions that have a triangular hole each, cyclically arranged around a central vertex 1 on the ground floor, and attach to it k rooms, R_1, \dots, R_k , in a coherent way, as before. The resulting triangulation has face vector

$$f = (14k + 1, 50k, 36k).$$

A C++-implementation `BH_k.cc` to generate the examples $BH(k)$ along with explicit triangulations `BH_3`, `BH_4`, and `BH_5` can be found online at [\[BL21\]](#).

Our next result highlights that in terms of simple-homotopy theory, $BH(k)$ is easy to understand.

Theorem 5.12. *For any $k \geq 3$, Bing's House with k rooms, $BH(k)$, can be formally deformed to a point using only six pure expansions.*

Proof. Since the rooms R_1, \dots, R_k are all identical, we extend to $BH(k)$ the labelling scheme that we used for the ground floor and the rooms of $BH(3)$. First

we do all the expansions in room R_1 . By adding the following six tetrahedra

$$2\,3\,5\,18, \, 3\,5\,18\,19, \, 5\,18\,19\,21, \, 3\,5\,6\,19, \, 5\,6\,19\,21, \, 6\,19\,21\,22$$

we fill in the cubical prism between the horizontal square on the vertices 2–3–6–5 of the main floor and the horizontal square on the vertices 18–19–22–21 of room R_1 's ceiling. We may now start the collapsing sequence from the outside. We perforate the back square 2–3–19–18 and then remove the whole interior of the prism, along with the back square 2–3–19–18 and the horizontal square 18–19–22–21 of the ceiling. Now the edge 21 22 is free. Thus, we can proceed exactly as for Bing's House with two rooms: We collapse away the squares 5–6–22–21 and 2–3–5–6, in this order. But now the edge 23 is free; so we can use it to collapse away room R_k . By induction, we can thus collapse all the rooms one by one. \square

How does this compare with the experimental results? In 10^4 runs, RSHT was always able to reduce Bing's house with three rooms, $BH(3)$, to a point, using on average about 148 additional tetrahedra. In the “best run”, only 12 additional tetrahedra were used. For Bing's house with k rooms, $BH(k)$, $4 \leq k \leq 7$, in 10^4 runs, even in the best case, RSHT tends to perform a growing number of expansions; see Table 5.1. This growing number of used tetrahedra is not surprising, due to the probabilistic model that we used: When selecting from more rooms, the number of options for possible expansions gets larger. So if we keep the number of rounds fixed, the chances to pick the cleverest sequence of pure expansions will get thinner.

5.6 Experiments on various topologies and substructures

In this section, we explore how our algorithm RSHT performs for further interesting simplicial complexes, whether contractible or not. All timings were taken on an Intel(R) Core(TM) i7-4720HQ CPU with 2.60 GHz and 16 GB RAM.

5.6.1 Contractible, non-collapsible complexes

Table 5.1 lists the number of expansions used for the Dunce Hat and Bing's Houses described in the previous section, as well as for the contractible complex `two_optima` of [ABL17] and for some knotted balls [Lut04, BL13b]. Furch's knotted 3-ball is the only example in this set for which the runtime is not negligible. In fact, due to the large number of expansions required, it took an average of 85 seconds to complete one round of the algorithm for this 3-ball.

Table 5.1: RSHT run for a selection of contractible, non-collapsible complexes from [BL21].

complex	f -vector	rounds	# expansions (minimum)	# expansions (mean)
Dunce hat	(8, 24, 17)	10^4	1	2.4145
Abalone	(15, 50, 36)	10^4	3	32.4156
Bing's House	(19, 65, 47)	10^4	5	58.0964
Bing's House 3 rooms, $BH(3)$	(43, 150, 108)	10^4	12	147.9727
$BH(4)$	(57, 200, 144)	10^4	15	167.7727
$BH(5)$	(71, 250, 180)	10^4	27	195.8890
$BH(6)$	(85, 300, 216)	10^4	34	221.2596
$BH(7)$	(99, 350, 252)	10^4	41	244.5763
two_optima	(106, 596, 1064, 573)	10^3	1	7.050
Knotted balls				
Furch's knotted ball	(380, 1929, 2722, 1172)	10^3	1459	1949.950
double_trefoil_ball	(15, 93, 145, 66)	10^3	1	29.600
triple_trefoil_arc	(17, 127, 208, 97)	10^3	6	94.678

The explanation of Table 5.2 is as follows. If one starts with a single d -simplex, with $8 \leq d \leq 15$, and one tries to collapse it down to a point, sometimes one gets stuck in contractible, non-collapsible complexes of intermediate dimension [LN21]. For each initial d -simplex we recorded 10 such examples, and on each one of these 10 examples we let RSHT run for 10^3 rounds. In each of the rounds, RSHT was able to reduce the respective examples to a point: In columns three and four of Table 5.2, we recorded the minimal and average numbers of expansions used. With the increase of the dimension, the runtime started to become an issue. For the largest examples, with $d = 15$, it took on average around 25 seconds to complete one round.

5.6.2 Submanifolds and non-manifold substructures in manifolds

If we remove a facet from a triangulation of the d -dimensional sphere S^d , the resulting simplicial complex is a triangulated d -ball, and thus has the simple-homotopy type of a point by Whitehead's Theorem 5.1. In case the initial d -manifold M^d is not a sphere, the removal of a simplex from a triangulation yields a simplicial complex that, depending on M^d , may deform to a submanifold or to a non-manifold substructure in M^d . Table 5.3 provides results for some classical

Table 5.2: RSHT run for contractible, non-collapsible complexes obtained when trying to collapse a d -simplex.

d	# examples \times # rounds	# expansions (minimum)	# expansions (mean)
8	10×10^3	1.0	1.9310
9	10×10^3	1.0	3.6845
10	10×10^3	1.0	8.4502
11	10×10^3	1.1	7.6552
12	10×10^3	2.5	29.4564
13	10×10^3	1.2	38.3988
14	10×10^3	7.9	174.7835
15	10×10^3	36.3	205.1362


Table 5.3: RSHT run for manifold triangulations minus a facet.

initial complex	initial f -vector	resulting complex	resulting f -vector
$\mathbb{R}P^3$ [Wal70]	(11, 51, 80, 40)	$\mathbb{R}P^2$	(6, 15, 10)
$\mathbb{R}P^4$ [Lut99, Ch. 3]	(16, 120, 330, 375, 150)	$\mathbb{R}P^3$	(11, 51, 80, 40)
$\mathbb{C}P^2$ [KB83]	(9, 36, 84, 90, 36)	S^2	(4, 6, 4)
$\mathbb{H}P^2$ [BK92]	(15, 105, 455, 1365, 3003, 4515, 4230, 2205, 490)	S^4	(6, 15, 20, 15, 6)
Poincaré 3-sphere [BL00]	(16, 106, 180, 90)	\mathbb{Z} -acyclic 2-complex	(10, 40, 31)


examples: Starting with the vertex-minimal triangulation of $\mathbb{R}P^3$ with 11 vertices, and removing a facet, in 10^4 runs of RSHT it took on average 25.2510 expansions to reach the 6-vertex triangulation of $\mathbb{R}P^2$. From $\mathbb{R}P^4$ to $\mathbb{R}P^3$ it took 885.5957 expansions. From $\mathbb{C}P^2$ to S^2 no expansions were used around half of the times; the average number of expansions needed was 2.3543. Finally, it took 30.0784 expansions to reach S^4 from $\mathbb{H}P^2$. For the Poincaré homology 3-sphere [BL00], the RSHT algorithm found a 2-dimensional \mathbb{Z} -acyclic 2-complex on 10 vertices (the boundary of the identified dodecahedron) using 2031.732 expansions in less than two minutes per run.

The 3-dimensional lens spaces $L(p, q)$, introduced by Tietze [Tie08], are well-known topological spaces with torsion in homology. Starting from triangulations of the 3-manifolds $L(p, 1)$ [BS93, Lut03a] for $p \geq 3$, we aimed for small triangulations of 2-dimensional simplicial complexes that still have p -torsion. (The case $p = 2$ has been already considered, since $L(2, 1) = \mathbb{R}P^3$.) The table in the top left of Figure 5.6 gives the f -vectors of these smaller complexes; Figure 5.6 (a)–


torsion	f -vector
\mathbb{Z}_3	(8, 24, 17)
\mathbb{Z}_4	(8, 26, 19)
\mathbb{Z}_5	(9, 32, 24)
\mathbb{Z}_6	(9, 33, 25)
\mathbb{Z}_7	(9, 34, 26)
\mathbb{Z}_8	(9, 35, 27)
\mathbb{Z}_9	(9, 36, 28)
\mathbb{Z}_{10}	(9, 36, 28)
\mathbb{Z}_{11}	(10, 42, 33)
\mathbb{Z}_{12}	(10, 42, 33)
\mathbb{Z}_{13}	(10, 43, 34)
\mathbb{Z}_{14}	(11, 50, 40)
\mathbb{Z}_{15}	(11, 50, 40)



(a) Complex d2_n8_3torsion with 3-torsion.



(b) Complex d2_n8_4torsion with 4-torsion.



(c) Complex d2_n8_5torsion with 5-torsion.

Figure 5.6: Small substructures with p -torsion of the lens spaces $L(p, 1)$.

(c) shows resulting small triangulations d2_n8_3torsion, d2_n8_4torsion, and d2_n8_5torsion (with facets lists available at [BL21]) with torsion \mathbb{Z}_3 , \mathbb{Z}_4 , and \mathbb{Z}_5 , respectively. The example d2_n8_3torsion has the combinatorial symmetry $(2, 3)(4, 8)(6, 7)$; the example d2_n8_4torsion has symmetry $(1, 2)(4, 6)(7, 8)$. In (b), the obtained complex is the union of an 8-vertex triangulation of the projective plane and a Möbius band. The complex d2_n8_5torsion originates from a triangulated disk by identifications highlighted in blue and red.

The following natural problem is open for $p \geq 3$:

Question 5.13. What is the minimal number of vertices $n_{\min}(p)$ for a simplicial 2-complex with p -torsion?

An earlier construction of a 2-dimensional simplicial complex with 3-torsion as

Table 5.4: RSHT run for triangulations of sphere products minus a facet.

initial complex	initial f -vector	resulting complex	size of the resulting complex
$S^2 \times S^1$	(12, 48, 72, 36)	$S^2 \vee S^1$	$\partial\Delta_3 \cup X^1(3.5382)$
$S^3 \times S^1$	(15, 75, 150, 150, 60)	$S^3 \vee S^1$	$\partial\Delta_4 \cup X^1(7.7617)$
$S^2 \times S^2$	(16, 84, 216, 240, 96)	$S^2 \vee S^2$	$\partial\Delta_3 \cup \partial\Delta_3$
$S^3 \times S^2$	(20, 130, 420, 710, 600, 200)	$S^3 \vee S^2$	$\partial\Delta_4 \cup X^2(11.5460)$

a *sum complex* on eight vertices is by Linial, Meshulam and Rosenthal [LMR10]. Their example is based on the following collection of subsets of \mathbb{Z}_8 :

$$X_{\{0,1,3\}} = \{ \sigma \subset \mathbb{Z}_8 : |\sigma| = 3, \sum_{x \in \sigma} x \equiv 0, 1 \text{ or } 3 \pmod{8} \}.$$

This complex has complete 1-skeleton and face vector $f = (8, 28, 21)$. Three edges of the complex are free, and after collapsing the respective triangles we reach a 2-complex with $f = (8, 25, 18)$, which still has one triangle and one edge more than the example `d2_n8_3torsion`. By running RSHT on the triangulation with 18 triangles repeatedly, we again reach `d2_n8_3torsion`—or a second non-isomorphic triangulation with the same f -vector that is obtained from `d2_n8_3torsion` by flipping the edge 1–5.

Conjecture 5.14. *The examples `d2_n8_3torsion` and `d2_n8_4torsion` have componentwise minimal f -vectors for complexes with 3- and 4-torsion, respectively.*

In the description of the torus $S^1 \times S^1$ as a square with opposite edges identified, the removal of the interior of the identified square yields the wedge product $S^1 \vee S^1$ of two circles S^1 that are glued together at a point. In general, if we remove a facet from a triangulation of a sphere product, the resulting complex is simple-homotopy equivalent to the wedge product of the constituting spheres. In the case of $S^2 \times S^1$, the wedge product $S^2 \vee S^1$ is of mixed dimension. Since in the implementation of RSHT our focus is on the top-dimensional faces, RSHT is not further touching lower-dimensional parts once these are reached via collapses. Thus, the resulting triangulations of $S^2 \vee S^1$ are of the form $\partial\Delta_3 \cup X^1$, consisting of the vertex-minimal triangulation of S^2 as the boundary complex $\partial\Delta_3$ of a 3-simplex Δ_3 union a 1-dimensional complex X^1 .

Depending on the intersection of X^1 with $\partial\Delta_3$, X^1 either is a path (a 1-dimensional ball) or a loop (a 1-sphere S^1). For a unified description in Table 5.4, we write $X^1(4.5382)$ to point out that X^1 has (in 10^4 runs of RSHT) on average 4.5382 edges. Table 5.4 gives results for further sphere products, where for the lower-dimensional parts the average number of facets are listed. The initial triangulations

Table 5.5: RSHT run for triangulations of products of surfaces.

initial complex	initial f -vector	resulting complex	resulting smallest f -vector
$T \times I$	(77, 511, 854, 420)	T	(7, 21, 14)
$g_2 \times I$	(121, 929, 1586, 780)	g_2	(9, 32, 24)
$g_5 \times I$	(253, 2183, 3782, 1860)	g_5	(12, 60, 40)
$g_6 \times I$	(297, 2601, 4514, 2220)	g_6	(13, 69, 46)
$g_{10} \times I$	(473, 4273, 7442, 3660)	g_{10}	(18, 108, 72)
$g_{50} \times I$	(2233, 20993, 36722, 18060)	g_{50}	(51, 683, 534)

of the sphere products in Table 5.4 are produced via product triangulations of boundaries of simplices [Lut03b].

In a separate experiment, we started with a triangulation of S^1 with 10 vertices and with a triangulation of S^2 with 100 vertices as the boundary complex of a random simplicial 3-polytope, for which 100 points on the round 2-dimensional sphere were chosen randomly via the `rand_sphere` client of the software system `polymake` [GJ00]). The initial triangulation of $S^2 \times S^1$ has face-vector $f = (1000, 6880, 11760, 5880)$. It took RSHT an average of 1108.23 expansions, in 10^2 runs, to reduce the triangulation (minus a facet) to a triangulation $\partial\Delta_3 \cup X^1(21.76)$ of the wedge product $S^2 \vee S^1$. We repeated the same experiment, but this time applying 200,000 preliminary random bistellar edge flips to the 100-vertex triangulation of S^2 , before taking the sphere product. The results of this experiment are similar to the one before (though with a slightly higher average number of expansions). This suggests that RSHT may be reliable even for larger complexes.

5.6.3 Dimensionality reduction

“Finding meaningful low-dimensional structures hidden in their high-dimensional observations” [TDSL00] is a major theme in analyzing higher-dimensional data of various origins. Usually, the data is given as a finite set of points in some Euclidean or metric space and is then often transformed to (higher-dimensional) simplicial complexes via taking Čech complexes or Vietoris–Rips complexes. Here, we did not start with explicit data sets, but instead “hid” a (closed) surface in a higher-dimensional product as another model to test RSHT on.

Starting with the standard 7-vertex triangulation T of the torus, we first took connected sums of T to create surfaces of higher genus g_k , $k \geq 2$. Then we took the cross product $g_k \times I$ of g_k with an interval (subdivided into 10 edges on 11 vertices), and reduced the resulting triangulation of the cross product with RSHT.

In every single one out of 10^2 runs, the product $g_k \times I$ gets reduced back to a small or even vertex-minimal triangulation of the original surface of genus g_k , as displayed in Table 5.5. In a second experiment, we performed 200,000 random edge flips to “randomize” the surfaces g_k ; then, we took cross products with the 10-edge interval I . Again, in 10^2 runs of RSHT, we always achieved the respective f -vectors of Table 5.5.

In a final experiment, we started with the triangulation of the surface g_{50} from before, but this time we added 100 vertices in subdivision steps before performing the 200,000 random edge flips. We then took again the cross product with the interval I to get a randomized triangulation of $g_{50} \times I$ with $f = (2728, 24278, 42212, 20760)$. We then took another cross product of this 3-manifold with boundary with the 4-simplex Δ_4 . The resulting complex is 7-dimensional with around 34 million faces and face vector

$$f = (13420, 386630, 2446620, 6910210, 10432052, 8786210, 3909060, 718200).$$

In less than an hour and by using a few thousand expansions, in each out of 10^2 runs of RSHT, we were able to reduce this complex back to a triangulation of the 2-dimensional orientable surface of genus 50 with fewer than 60 vertices. In some cases we were even able to reach the same f -vector with 51 vertices as in Table 5.6c. Due to memory constraints that come from the computation of the Hasse diagram of the starting complex (requiring around 10 GB of RAM for this example), this was the largest complex that we were able to study.

5.6.4 Akbulut–Kirby 4-spheres

As stated early on, contractibility is, in general, undecidable. However, it takes considerable effort to pose challenges to RSHT. A notoriously hard series of complexes is given by the triangulations of the Akbulut–Kirby 4-dimensional spheres [TL]. These PL-triangulated standard 4-dimensional spheres are built in an intricate way via non-trivial presentations of the trivial group as their fundamental group [AK85]. By Pachner’s theorem, these examples are bistellarly equivalent to the boundary of the 5-simplex, and by Whitehead’s theorem, the examples minus a facet are simple-homotopy equivalent to a single vertex. However, establishing connecting sequences of bistellar flips failed in [TL], beyond the first easy examples of the series. Indeed, here RSHT made no progress either, even when we set $\text{max_step} = 1,000,000$ and waited for a total runtime of 60 hours.

Chapter 6

Hadamard matrix torsion

In this chapter, we construct in quadratic time $\Theta(n^2)$ a series $\text{HMT}(n)$ of 2-dimensional simplicial complexes with torsion $H_1(\text{HMT}(n)) = (\mathbb{Z}_2)^{\binom{k}{1}} \times (\mathbb{Z}_4)^{\binom{k}{2}} \times \cdots \times (\mathbb{Z}_{2^k})^{\binom{k}{k}}$, $|H_1(\text{HMT}(n))| = |\det(H(n))| = n^{n/2} \in \Theta(2^{n \log n})$, where the construction is based on the Hadamard matrices $H(n)$ for $n \geq 2$ a power of 2, i.e., $n = 2^k$, $k \geq 1$. The examples have linearly many vertices, their face vector is $f(\text{HMT}(n)) = (5n - 1, 3n^2 + 9n - 6, 3n^2 + 4n - 4)$.

Our explicit series with torsion growth in $\Theta(2^{n \log n})$ improves a previous construction by Speyer [Spe] with torsion growth in $\Theta(2^n)$, narrowing the gap to the highest possible asymptotic torsion growth in $\Theta(2^{n^2})$ proved by Kalai [Kal83] via a probabilistic argument. The material of this chapter is based on a joint work with Frank H. Lutz [LL21].

6.1 Overview

The most elementary way to build a two-dimensional CW complex with torsion \mathbb{Z}_r in the first integer homology starts out with a single polygonal disc with r edges, $r \geq 2$, that are all oriented in the same direction and jointly identified. The resulting CW complex has one vertex, one edge, and one two-dimensional face—it has homology $H_* = (\mathbb{Z}, \mathbb{Z}_r, 0)$.

Definition 6.1 (Matrix Disc Complexes). Let $M = (M_{ij})$ be an $(m \times n)$ -matrix with integer entries. Let the set of two-dimensional **matrix disc complexes** $DC(M)$ associated with M comprise the CW complexes constructed level-wise in the following way:

- ▷ Every complex in $DC(M)$ has a single 0-cell (with label 0 in the following).
- ▷ The 1-skeleton of a complex in $DC(M)$ has an edge cycle a_j for every column index $j \in \{1, \dots, n\}$ of the matrix M .




Figure 6.1: Klein bottle on the left and a pinched $\mathbb{R}P^2$ on the right, both obtained from the matrix $M = \begin{pmatrix} 2 & 2 \end{pmatrix}$.

- ▷ Every row i of M with row sum $s_i = |M_{i1}| + \dots + |M_{in}|$, $i \in \{1, \dots, m\}$, contributes a polygonal disc with s_i edges. For every positive entry M_{ij} , M_{ij} edges of the disc are oriented coherently and are assigned with the label a_j . In the case of a negative entry, the direction of the corresponding edges is reversed; in the case of a zero-entry, the respective edge does not occur.

Example 6.2. The (1×1) -matrix $M = (r)$ yields a single disk with the r identified edges all oriented in the same direction—the elementary construction from above.

Example 6.3. Let us now consider $M = \begin{pmatrix} 2 & 2 \end{pmatrix}$ with one row. The examples in $DC(M)$ have two cycles, a_1 and a_2 , and a single disc with edges $2a_1 + 2a_2$. There are two choices for the edge sequences of the disc, either $a_1a_1a_2a_2$ or $a_1a_2a_1a_2$. In the first case, the resulting CW complex is the Klein bottle (Figure 6.1, left), in the second case, we obtain a pinched real projective plane $\mathbb{R}P^2$ (Figure 6.1, right). While the Klein bottle is a manifold, the latter example is not—thus the two examples in $DC(M)$ are not homeomorphic to each other.

Lemma 6.4. *The examples in $DC(M)$ all have the same integer homology H_* .*

Proof. The examples $DC(M)$ all have a single vertex only and therefore are connected, respectively. For every representative $C \in DC(M)$ it thus follows that $H_0(C) = \mathbb{Z}$. Each edge of C is a cycle, i.e., the first homology group $H_1(C)$ of C is determined by the n rows of M as relations, and therefore H_1 coincides for all the examples in $DC(M)$. Further, the second homology H_2 of any representative C is simply the kernel of the matrix M . \square

Example 6.5 (Projective plane). For the (2×2) -matrix $H(2) = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ we have two discs with edges a_1a_2 and $a_1a_2^{-1}$, as in Figure 6.2 on the left. If we glue together the two discs along the common edge a_2 , as in Figure 6.2 on the right, we obtain a single disc with two (identified) edges a_1a_1 , the standard scheme for the real projective plane $\mathbb{R}P^2$ with homology $H_* = (\mathbb{Z}, \mathbb{Z}_2, 0)$.




Figure 6.2: \mathbb{RP}^2 as a CW complex with two disks.

A simple consequence of the proof of Lemma 6.4 is that if M is a square matrix with $\det(M) \neq 0$, then $|H_1(C)| = |\det(M)|$ for any example $C \in DC(M)$. Our goal in the following is to construct triangulations of CW-complexes $C \in DC(M)$ as abstract simplicial complexes with few vertices, in particular, for square matrices M with large determinant, yielding simplicial complexes with huge torsion.

Kalai proved in [Kal83] that, asymptotically, there are \mathbb{Q} -acyclic simplicial complexes on n vertices with torsion growth in $\Theta(2^{n^2})$, and that this is the maximal possible growth. Recently, Newman [New19] showed, via a randomized construction, that any abelian group G can be obtained as torsion of a simplicial complex with $\Theta(\log(|G|)^{\frac{1}{2}})$ vertices. Explicit classes of \mathbb{Q} -acyclic simplicial complexes were provided by Linial, Meshulam and Rosenthal [LMR10], called sum complexes, however, without control on the torsion growth.

The overall approach in this chapter is inspired by a construction of Speyer [Spe] of 2-dimensional \mathbb{Q} -acyclic simplicial complexes on $\Theta(n)$ vertices with exponential torsion growth $\Theta(2^n)$, corresponding to particular square matrices M of size $\Theta(n) \times \Theta(n)$; see Section 6.2.2. For general matrices M , a first (elementary) triangulation approach for associated complexes $C \in DC(M)$ is given in Section 6.2.1.

Starting with Hadamard matrices $H(n)$, we give an explicit construction in quadratic time $\Theta(n^2)$ that achieves torsion growth $\Theta(2^{n \log n})$; see Section 6.3.

6.2 Preliminaries

In this section, we first give a (straight-forward) procedure to triangulate a general complex $C \in DC(M)$ for some given integer $(m \times n)$ -matrix $M = (M_{ij})$. Afterwards, we discuss basic facts about Hadamard matrices $H(n)$ that will be used in the main Section 6.3, where we introduce an improved triangulation scheme for the complexes corresponding to these square matrices to obtain our torsion growth bound $\Theta(2^{n \log n})$.

6.2.1 A first triangulation procedure

Given some $(m \times n)$ -matrix $M = (M_{ij})$, $i \in \{1, \dots, m\}$, $j \in \{1, \dots, n\}$, the complexes $C \in DC(M)$ are 2-dimensional CW complexes with a single vertex (with label 0), m cycles a_j , $j \in \{1, \dots, n\}$, corresponding to the column indices, and a disc for every non-zero row of M . In the following, we assume that M has no zero columns and rows. (In case rows appear multiple times, we saw by Example 6.3 that different choices for the edge sequences can yield different gluings of the discs.)

When we subdivide a CW complex C to obtain a triangulation of it as an abstract simplicial complex K , we have to ensure that K has no loops and no parallel edges:

- ▷ The only vertex 0 of C is kept as vertex 0 in K .
- ▷ Each cycle a_j , $j \in \{1, \dots, n\}$, is triangulated by using two extra vertices, v_j^1 and v_j^2 .
- ▷ We next triangulate the m discs corresponding to the m rows of M . For each row $i \in \{1, \dots, m\}$, the respective disc is an s_i -gon for $s_i = |M_{i1}| + \dots + |M_{in}|$. Into each s_i -gon we place a $\lceil \frac{3}{2}s_i \rceil$ -gon using $\lceil \frac{3}{2}s_i \rceil$ new vertices c_i^k , $k \in \{1, \dots, \lceil \frac{3}{2}s_i \rceil\}$. The inside of the $\lceil \frac{3}{2}s_i \rceil$ -gon can be triangulated arbitrarily by consecutively adding diagonals. The annulus between the inner $\lceil \frac{3}{2}s_i \rceil$ -gon and the outer s_i -gon is triangulated by connecting any inner vertex with three consecutive vertices of the s_i -gon, creating a cone, and then filling the remaining gaps with triangles. In particular, we choose a starting vertex on the s_i -gon and a direction, and we connect c_i^1 with the starting vertex and the two consecutive ones. Next, we connect c_i^2 with the last vertex to which we connected c_i^1 and then to the consecutive two vertices. We continue like this till we reach the last vertex of the $\lceil \frac{3}{2}s_i \rceil$ -gon which will be connected with the starting vertex again (be careful, this last vertex could be connected with only two vertices of the outer s_i -gon instead of three). See Figure 6.3 for an example on how to fill the inside.

Remark 6.6. In many examples, we could actually triangulate the polygonal discs with fewer vertices than according to the above procedure. However, improvements will depend on the concrete entries of the given matrix.

Since all the polygonal discs are triangulated in a way such that no two discs share an interior edge we have that the homology of the constructed simplicial complex K is the same as that of the original CW complex C .

In the following proposition on triangulations of matrix disc complexes we use the Smith Normal Form of a matrix M , so before the proposition, let us remember the definition.

Definition 6.7. (Smith [Smi61]) Let M be a nonzero $(m \times n)$ -matrix over a principal ideal domain. There exist invertible $(m \times m)$ - and $(n \times n)$ -matrices S and T , respectively, such that $M = SAT$ and

$$A = \begin{pmatrix} \alpha_1 & 0 & \dots & & 0 \\ 0 & \alpha_2 & 0 & \dots & 0 \\ \vdots & 0 & \ddots & 0 & \dots & 0 & \vdots \\ & \vdots & 0 & \alpha_r & & & \\ & & & & 0 & & \\ & & & & & \ddots & \\ 0 & 0 & \dots & & & & 0 \end{pmatrix},$$

where $\alpha_i \mid \alpha_{i+1}$ for all $1 \leq i < r$. A is called the *Smith Normal Form* of M .

Proposition 6.8. Given an $(m \times n)$ -matrix M and $A = (\alpha_i)$ its Smith Normal Form, there is a 2-dimensional simplicial complex K on $V(K)$ vertices with

$$V(K) \leq 2n + m + 1 + \frac{3}{2} \sum_{i,j} |M_{ij}|.$$

Furthermore,

$$H_1(K) = \mathbb{Z}/\alpha_1\mathbb{Z} \times \dots \times \mathbb{Z}/\alpha_r\mathbb{Z} \times \mathbb{Z}^{n-r}.$$

Proof. Let M be an $(m \times n)$ -matrix, $C \in \text{DC}(M)$, K a triangulation of C as described above, and $s_i = |M_{i1}| + \dots + |M_{in}|$. We easily see that

$$\sum_{i=1}^m \left\lceil \frac{3}{2} s_i \right\rceil \leq m + \frac{3}{2} \sum_{i,j} |M_{ij}|.$$

For the number of vertices $V(K)$ of K we obtain the bound

$$V(K) = 1 + 2n + \sum_{i=1}^m \left\lceil \frac{3}{2} s_i \right\rceil \leq 2n + m + 1 + \frac{3}{2} \sum_{i,j} |M_{ij}|,$$

with 1 vertex for the original vertex of C , $2n$ vertices to triangulate the n cycles, and $\sum_{i=1}^m \left\lceil \frac{3}{2} s_i \right\rceil$ interior vertices in the m discs. Since M is the boundary matrix for the homology of K , the homology of K is represented by the Smith Normal Form A of M . \square

Remark 6.9. Given a square matrix M with $\det(M) \neq 0$, it follows by Proposition 6.8 that for the simplicial complex K associated with M , $|H_1(K)| = |\det(M)|$.

6.2.2 Speyer's construction

Speyer [Spe] provided a construction to produce 2-dimensional \mathbb{Q} -acyclic simplicial complexes for which the size of the torsion grows exponentially in the number of vertices.

Let $k \geq 2$ be an integer and $k = \gamma_m 2^m + \gamma_{m-1} 2^{m-1} + \dots + \gamma_0 2^0$ be its binary expansion, with leading coefficient $\gamma_m = 1$ and otherwise $\gamma_i \in \{0, 1\}$ for all $0 \leq i \leq m-1$. An $((m+1) \times (m+1))$ -matrix $M(k)$ is constructed in the following way:

- ▷ The first row contains the entries $(-1)^i \gamma_{m-i}$ for $i \in \{0, \dots, m\}$.
- ▷ The lower part of the matrix $M(k)$ is an $(m \times (m+1))$ -matrix with 1's on the first diagonal followed by 2's on the diagonal to the right, and all other entries equal to zero. It is then easy to see that $\det(M(k)) = k$.

Using Proposition 6.8, there is a 2-dimensional simplicial complex K on $V(K) \leq 3(m+1) + 1 + \frac{3}{2}(3m + (m+1)) \leq 9m + 6$ vertices corresponding to $M(k)$ that has torsion of size k . Since the number of vertices is linear in m , the torsion k grows exponentially in the size $n = m+1$ of the matrix, i.e., $k \in \Theta(2^n)$.

Example 6.10. For an explicit example, we consider $k = 11 = 8 + 2 + 1$ with

$$M(11) = \begin{pmatrix} \text{[redacted]} \\ 1 & 2 & 0 & 0 \\ 0 & 1 & 2 & 0 \\ 0 & 0 & 1 & 2 \end{pmatrix}.$$

Figure 6.3 displays a triangulation $K(11)$ of a representative $C \in \text{DC}(M(11))$ with 29 vertices that has torsion \mathbb{Z}_{11} .

Remark 6.11. In case k is prime, the torsion of complexes corresponding to the matrix $M(k)$ is cyclic, but in general it is a product of the factors in the Smith Normal Form of $M(k)$.

As already pointed out in Remark 6.6, also in this particular example $K(11)$ we could save on the number of vertices necessary for the triangulation of the interiors of the polygonal discs. However, it is open what the minimum number of vertices is for a re-triangulation of $K(11)$.

6.2.3 Hadamard matrices

In our quest to construct 2-dimensional simplicial complexes with few vertices but high torsion via matrices, Hadamard matrices are of particular interest. In fact, Hadamard matrices were used earlier in geometry for extremal constructions; see for example [HKL96] and [Zie00].




Figure 6.3: The four subdivided triangles of the complex $K(11)$ of Example 6.10. In the upper left part of $K(11)$ we highlighted the five cones of inner vertices with respect to the subsequences of three boundary vertices.

Definition 6.12. A Hadamard matrix H is a square $(n \times n)$ -matrix whose entries are either $+1$ or -1 and whose rows are mutually orthogonal.

Hadamard matrices H have in common that their determinants attain the Hadamard bound, $|\det(H)| = n^{\frac{n}{2}}$. In general, $|\det(M)| \leq n^{\frac{n}{2}}$ for any integer $(n \times n)$ -matrix $M = (M_{ij})$ with $|M_{ij}| \leq 1$, and a matrix M attains the bound if and only if it is a Hadamard matrix [Had93].

It is known that the order of a Hadamard matrix must be 1, 2, or a multiple of 4, but it is open whether Hadamard matrices exist for all multiples of 4. However, a very nice construction of Sylvester [Syl67] tells us that if H is a Hadamard matrix of order n , then the matrix $\begin{pmatrix} H & H \\ H & -H \end{pmatrix}$ is a Hadamard matrix of order $2n$.

The following sequence of matrices $H(n)$, $n = 2^k$, $k \geq 1$, also called Walsh matrices [Wal23], are Hadamard matrices:

$$H(1) = (1),$$

$$H(2^k) = \begin{pmatrix} H(2^{k-1}) & H(2^{k-1}) \\ H(2^{k-1}) & -H(2^{k-1}) \end{pmatrix}, \text{ for } k \geq 1.$$

In the following, we will always refer to this sequence when we talk about Hadamard matrices. The matrix $H(2)$ was discussed before in Example 6.5 above.

Lemma 6.13. *Let $n = 2^k$, $k \geq 0$, and let $A = (\alpha_i)$ be the Smith Normal Form of $H(n)$, with $H(n) = SAT$ for integral invertible matrices S and T .*

Then $\text{diag}(A) = (\alpha_i) = (2^0, 2^1, \dots, 2^1, \dots, 2^j, \dots, 2^j, \dots, 2^{k-1}, \dots, 2^{k-1}, 2^k)$, where each 2^j appears $\binom{k}{j}$ times.

Proof. We prove the statement by induction, where the base case for $n = 1$ is clear, since $H(1) = (1)$.

Let $A = S^{-1}H(n)T^{-1}$ be the Smith Normal Form of $H(n)$. For the following two invertible matrices $\tilde{S} = \begin{pmatrix} S^{-1} & 0 \\ S^{-1} & -S^{-1} \end{pmatrix}$ and $\tilde{T} = \begin{pmatrix} T^{-1} & -T^{-1} \\ 0 & T^{-1} \end{pmatrix}$ we have:

$$\begin{aligned} \tilde{S}H(2n)\tilde{T} &= \begin{pmatrix} S^{-1} & 0 \\ S^{-1} & -S^{-1} \end{pmatrix} \begin{pmatrix} H(n) & H(n) \\ H(n) & -H(n) \end{pmatrix} \begin{pmatrix} T^{-1} & -T^{-1} \\ 0 & T^{-1} \end{pmatrix} \\ &= \begin{pmatrix} S^{-1}H(n) & S^{-1}H(n) \\ 0 & 2S^{-1}H(n) \end{pmatrix} \begin{pmatrix} T^{-1} & -T^{-1} \\ 0 & T^{-1} \end{pmatrix} \\ &= \begin{pmatrix} S^{-1}H(n)T^{-1} & 0 \\ 0 & 2S^{-1}H(n)T^{-1} \end{pmatrix} = \begin{pmatrix} A & 0 \\ 0 & 2A \end{pmatrix}. \end{aligned}$$

The resulting matrix is the Smith Normal Form of $H(2n)$ —after a suitable reordering of the diagonal elements to ensure that $\alpha_i | \alpha_{i+1}$ for all $1 \leq i \leq 2n - 1$. The lemma then follows from the known relation $\binom{k+1}{j} = \binom{k}{j-1} + \binom{k}{j}$ of Pascal's triangle. \square

6.3 An improved triangulation procedure

Using Proposition 6.8, we can triangulate a complex $C \in \text{DC}(H(n))$ with $\Theta(n^2)$ vertices to produce torsion of size $n^{\frac{n}{2}} \in \Theta(2^{n \log n})$. Asymptotically, this is a worse bound than what can be achieved by considering the Speyer matrices $M(k)$ of Section 6.2.2 that yield triangulations with $\Theta(n)$ vertices and torsion of size $\Theta(2^n)$. The inferior torsion growth is due to using linearly many additional vertices inside each of the n discs for the examples $C \in \text{DC}(H(n))$, compared to the constant number of vertices needed for the discs associated with the rows of the lower part

of the Speyer matrices $M(k)$. Our aim in this section is to provide a modified construction for complexes associated with the Hadamard matrices $H(n)$ that requires in total only linearly many vertices, $5n - 1$, but still yields torsion of size $n^{\frac{n}{2}}$.

6.3.1 A modified CW disc construction

Instead of triangulating complexes $C \in \text{DC}(H(n))$ in the way outlined in Section 6.2.1, in the following we consider CW complexes \tilde{C} that are derived from the matrices $H(n)$ in a modified way, yet keeping the homotopy type, i.e. $\tilde{C} \simeq C$. Our goal is to use exactly one interior vertex for each of the polygonal discs (see the square discs in the upper part of Figure 6.5).

The vertex 0 appears n times along the boundary of each of the polygonal discs and therefore has to be shielded off when we triangulate the interiors of the discs—to avoid unwanted identifications of interior edges of the discs.

The easiest way to shield off the special vertex 0 is by connecting its two neighbors, at every occurrence along the boundary of the discs, by a diagonal. However, we have to ensure that each such diagonal is used only once for all of the discs. To meet this requirement, we modify our CW disc construction:

- ▷ A complex associated with $H(n)$ has a single 0-cell only.
- ▷ To each column j of $H(n)$ we associate two different edge cycles a_j^+ and a_j^- .
- ▷ To each row i of $H(n)$ we associate a polygonal disc, as before in Definition 6.1. In addition we consider n discs that each connect the two cycles a_j^+ and a_j^- .

For every complex $C \in \text{DC}(H(n))$, we then define an *augmented* complex \tilde{C} , so that in each disc of C we revert negatively oriented edges by using the connecting digon pieces; see Figure 6.4.

The boundaries of the polygonal discs still represent the defining rows of the matrix $H(n)$, and since the digon connectors are homotopy equivalent to single cycles, it follows that $\tilde{C} \simeq C$ for the modified complexes.

Towards the interior of the discs, we have positively oriented edges only—which will allow (see Section 6.3.2) to choose an ordering of the cycles on the boundary of each polygonal disc so that each diagonal that shields the vertex 0 occurs only once in \tilde{C} .

Remark 6.14. An augmented complex \tilde{C} has a straight-forward representation by an augmented $(2n \times 2n)$ -matrix $\tilde{H}(n)$ so that $\tilde{C} \in \text{DC}(\tilde{H}(n))$. Here, we obtain $\tilde{H}(n)$ from $H(n)$ by splitting every column into two columns, where the first copy contains the positive entries of the original column and the second copy contains




Figure 6.4: A polygonal disc in the construction of \tilde{C} .

the absolute values of the negative entries of the original column. We further add n new rows that represent the inserted digons.

E.g., the matrix $H(2) = \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ is augmented to

$$\tilde{H}(2) = \begin{pmatrix} \text{blue } 1 & \text{green } 1 & \text{blue } 0 & \text{green } 0 \\ \text{red } 1 & \text{red } 0 & 0 & 0 \\ 0 & 0 & \text{red } 1 & \text{red } 0 \end{pmatrix},$$

where the positive entries of $H(2)$ are listed in the blue columns, whereas the green columns represent the negative entries of $H(2)$. The connecting digons are highlighted in red. By subtracting the second column from the first one and the

fourth column from the third, we obtain the matrix $\begin{pmatrix} 1 & 0 & 1 & 0 \\ 1 & 0 & -1 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}$. In particular,

it follows that $|\det(\tilde{H}(2))| = |\det(H(2))|$.

6.3.2 Valid sequences

To avoid identified diagonals, we next choose suitable orderings of the edge boundaries of the discs to select representatives $\tilde{C}(n) \in DC(\tilde{H}(n))$ that can be triangu-

lated with linearly many vertices. In particular, we ensure that two consecutive cycles occur exactly once along the boundaries of the discs.

Definition 6.15. Let $M = (M_{ij})$ be an $(n \times n)$ -matrix with ± 1 -entries. A *valid sequence* (τ_i^j) is a sequence of orderings of the positively oriented n boundary edges a_j^\pm , $j \in \{1, \dots, n\}$, of the first n discs, $i \in \{1, \dots, n\}$, of an augmented complex $\tilde{C} \in DC(\tilde{M})$, associated with M that satisfy:

1. For each disc $i \in \{1, \dots, n\}$, (τ_i^*) is a permutation of the numbers $\{1, \dots, n\}$, always starting with 1,
2. For all distinct disc indices i_1, i_2 and all edge indices j_1, j_2 such that if two consecutive edge labels coincide, $\tau_{i_1}^{j_1} = \tau_{i_2}^{j_2}$ and $\tau_{i_1}^{j_1+1} = \tau_{i_2}^{j_2+1}$, then at least one pair of corresponding matrix entries differs in sign, $M_{i_1, \tau_{i_1}^{j_1}} \neq M_{i_2, \tau_{i_2}^{j_2}}$ or $M_{i_1, \tau_{i_1}^{j_1+1}} \neq M_{i_2, \tau_{i_2}^{j_2+1}}$. It is assumed that if j_1 is equal to n , then $j_1 + 1$ is equal to 1, and the same for j_2 .

Example 6.16. The sequences $((1))$ and $((1, 2), (1, 2))$ are the unique valid sequences for $H(1)$ and $H(2)$, respectively.

Example 6.17. The sequence

$$(\tau_i^j) = ((1, 3, 2, 4), (1, 2, 4, 3), (1, 3, 2, 4), (1, 4, 3, 2))$$

is a valid sequence for the Hadamard matrix

$$H(4) = \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -1 & 1 & -1 \\ 1 & 1 & -1 & -1 \\ 1 & -1 & -1 & 1 \end{pmatrix}.$$

The first permutation $(1, 3, 2, 4)$ gives the ordering of the edges of the first disc that is associated to the first row of the matrix $H(4)$. As all entries of this first row are positive, we have all corresponding edges of the first disc with forward orientation.

In this example, the first and the third permutation are identical. In particular, the first parts $(1, 3)$ of the two permutations coincide, but $H(4)_{13} \neq H(4)_{33}$. Another consecutive pair that appears in the second and the fourth permutation is $(4, 3)$, but $H(4)_{24} \neq H(4)_{44}$ and $H(4)_{23} \neq H(4)_{43}$. If we compare all further consecutive edge pairs, we can easily check that the given sequence of permutations is a valid sequence.

Using the inductive definition

$$H(2n) = \begin{pmatrix} H(n) & H(n) \\ H(n) & -H(n) \end{pmatrix}$$

for the Hadamard matrices $H(n)$, with the above sequence for $H(1)$ as the base for the induction, we next provide a procedure to obtain a valid sequence $(\tilde{\tau}_i^j)$ for $H(2n)$ from any valid sequence (τ_i^j) for $H(n)$.

For any fixed $i \in \{1, \dots, n\}$, we create two permutations of the numbers $\{1, \dots, 2n\}$ for $H(2n)$, starting from the permutation τ_i^* of $H(n)$:

$$\begin{aligned}\beta_i^* &= (\tau_i^1, \dots, \tau_i^n; \tau_i^1 + n, \dots, \tau_i^n + n), \\ \gamma_i^* &= (\tau_i^1, \tau_i^2 + n, \tau_i^3, \tau_i^4 + n, \dots, \tau_i^n + n; \tau_i^1 + n, \tau_i^2, \tau_i^3 + n, \tau_i^4, \dots, \tau_i^n),\end{aligned}$$

i.e., β_i^* is obtained from τ_i^* by first taking a copy of τ_i^* followed by another copy of τ_i^* , where n is added to each entry of the second copy; and γ_i^* is obtained by adding n to all even positions of a copy of τ_i^* followed by doing so for all odd positions of a second copy of τ_i^* .

Proposition 6.18. *Given a valid sequence (τ_i^j) for $H(n)$, the sequence $(\tilde{\tau}_i^j)$, inductively defined by*

$$\tilde{\tau}_i^j = \begin{cases} \beta_i^j, & \text{if } i \leq n, \\ \gamma_{i-n}^j, & \text{if } i > n, \end{cases} \quad \text{for } i, j \in \{1, \dots, 2n\},$$

is a valid sequence for $H(2n)$.

In the definition of $\tilde{\tau}_i^j$ in Proposition 6.18, we first take all n permutations β_i^j of length $2n$ and then (after a shift of the row index by n) all n permutations γ_i^j of length $2n$ (in the given order, respectively).

Example 6.19. Starting with the valid sequence for $H(2)$ from Example 6.16, by the construction of Proposition 6.18 we obtain the following valid sequence for $H(4)$:

$$(\tilde{\tau}_i^j) = ((1, 2, 3, 4), (1, 2, 3, 4), (1, 4, 3, 2), (1, 4, 3, 2)),$$

which is different from the sequence of Example 6.17, showing that valid sequences are not necessarily unique.

Example 6.20. Starting with the valid sequence for $H(4)$ from Example 6.19, we obtain the following valid sequence $(\tilde{\tau}_i^j)$ for $H(8)$:

$$\begin{aligned} & ((1, 2, 3, 4, 5, 6, 7, 8), (1, 2, 3, 4, 5, 6, 7, 8), (1, 4, 3, 2, 5, 8, 7, 6), (1, 4, 3, 2, 5, 8, 7, 6), \\ & (1, 6, 3, 8, 5, 2, 7, 4), (1, 6, 3, 8, 5, 2, 7, 4), (1, 8, 3, 6, 5, 4, 7, 2), (1, 8, 3, 6, 5, 4, 7, 2)). \end{aligned}$$

Proof of Proposition 6.18. It is clear from the construction that for every i , $\tilde{\tau}_i^*$ is a permutation of the numbers $\{1, \dots, 2n\}$ starting with 1, so we only need to check the second condition to prove that $(\tilde{\tau}_i^j)$ is a valid sequence for $H(2n)$.

The sequence $(\tilde{\tau}_i^j)$ for $H(2n)$ is obtained from the valid sequence (τ_i^j) for $H(n)$ in an inductive way. Let $i_1, i_2, j_1, j_2 \in \{1, \dots, 2n\}$, $i_1 \neq i_2$, be such that $\tilde{\tau}_{i_1}^{j_1} = \tilde{\tau}_{i_2}^{j_2}$ and $\tilde{\tau}_{i_1}^{j_1+1} = \tilde{\tau}_{i_2}^{j_2+1}$, i.e., in the permutations corresponding to the two different discs i_1, i_2 two consecutive edge labels are the same. Let $\hat{i}_1 = i_1$ for $1 \leq i_1 \leq n$ and $\hat{i}_1 = i_1 - n$ for $n+1 \leq i_1 \leq 2n$ be the original disc indices in the initial valid sequence (τ_i^j) . Same for \hat{i}_2 and analogously for \hat{j}_1, \hat{j}_2 . It then follows by construction that $\tau_{\hat{i}_1}^{\hat{j}_1} = \tau_{\hat{i}_2}^{\hat{j}_2}$ and $\tau_{\hat{i}_1}^{\hat{j}_1+1} = \tau_{\hat{i}_2}^{\hat{j}_2+1}$, as in the transition from (τ_i^j) to $(\tilde{\tau}_i^j)$ either n was added to the entries or not. There are then four different cases to consider, where we compare pairs of β - and γ -permutations:

- ▷ $[\beta-\beta]$ $i_1, i_2 \leq n$: In this case $H(2n)_{i_1, \tilde{\tau}_{i_1}^{j_1}} = H(n)_{\hat{i}_1, \tau_{\hat{i}_1}^{\hat{j}_1}}$, and the same correlation holds for the other three matrix entries that appear in the definition of a valid sequence. By assumption, (τ_i^j) is a valid sequence for $H(n)$, which implies $H(n)_{\hat{i}_1, \tau_{\hat{i}_1}^{\hat{j}_1}} \neq H(n)_{\hat{i}_2, \tau_{\hat{i}_2}^{\hat{j}_2}}$ or $H(n)_{\hat{i}_1, \tau_{\hat{i}_1}^{\hat{j}_1+1}} \neq H(n)_{\hat{i}_2, \tau_{\hat{i}_2}^{\hat{j}_2+1}}$. By the correlation between the entries of $H(n)$ and $H(2n)$ as written before, it follows that $H(2n)_{i_1, \tau_{i_1}^{j_1}} \neq H(2n)_{i_2, \tau_{i_2}^{j_2}}$ or $H(2n)_{i_1, \tau_{i_1}^{j_1+1}} \neq H(2n)_{i_2, \tau_{i_2}^{j_2+1}}$.
- ▷ $[\gamma-\gamma]$ $i_1, i_2 > n$: This time $H(2n)_{i_1, \tilde{\tau}_{i_1}^{j_1}} = \pm H(n)_{\hat{i}_1, \tau_{\hat{i}_1}^{\hat{j}_1}}$ with “+” for column indices $1 \leq \tilde{\tau}_{i_1}^{j_1} \leq n$ and “-” otherwise. But this does not change anything, since $\tilde{\tau}_{i_1}^{j_1} = \tilde{\tau}_{i_2}^{j_2}$. Thus both $H(2n)_{i_1, \tilde{\tau}_{i_1}^{j_1}}$ and $H(2n)_{i_2, \tilde{\tau}_{i_2}^{j_2}}$ will keep the sign or change it, so we can use the same argument as before saying that since (τ_i^j) was a valid sequence for $H(n)$, we have a pair of matrix entries that are not equal.
- ▷ $[\beta-\gamma]$ $i_1 \leq n$ and $i_2 > n$: As in the definition of β_i^* and γ_i^* for some permutation τ_i^* , the number n was added to all entries of the second half to obtain β_i^* , but added in an alternating fashion to the entries for obtaining γ_i^* , only the following boundary cases are possible for which consecutive edge labels $\tilde{\tau}_{i_1}^{j_1} = \tilde{\tau}_{i_2}^{j_2}$ and $\tilde{\tau}_{i_1}^{j_1+1} = \tilde{\tau}_{i_2}^{j_2+1}$ can agree: positions $j_1 = n$ or $j_1 = 2n$ in β_i^* or $j_2 = n$ or $j_2 = 2n$ in γ_i^* . If $j_1 = n$, then $\tilde{\tau}_{i_1}^{j_1+1} = n + \tau_{i_1}^1 = n + 1 = n + \tau_{i_2}^1 = \tilde{\tau}_{i_2}^{n+1}$. In this case, the only possibility is that $j_2 = n$, but by construction, we have that $\tilde{\tau}_{i_1}^{j_1} \neq \tilde{\tau}_{i_2}^{j_2}$. The same argument can be used in the three other cases.
- ▷ $[\gamma-\beta]$ $i_1 > n$ and $i_2 \leq n$: This case is analogous to the previous one.

In all four cases we showed that the second condition for having a valid sequence is satisfied, so we proved that $(\tilde{\tau}_i^j)$ is a valid sequence for $H(2n)$. \square

As a consequence of Proposition 6.18 and the Examples 6.16 and 6.17 we have that for every $n = 2^k$, $k \geq 2$, there is a (not necessarily unique) valid sequence for the Hadamard matrix $H(n)$.

Via the recursive procedure of Proposition 6.18, it takes quadratic time to obtain a valid sequence for $H(n)$ from a valid sequence for $H(\frac{n}{2})$, as the $(\frac{n}{2})^2$ numbers in the valid sequence for $H(\frac{n}{2})$ have to be read and copied in modified form four times. This gives us a recursive formula, $t(n) = 4n^2 + t(\frac{n}{2})$, for the total number of steps $t(n)$ to obtain a valid sequence for $H(n)$ from the unique valid sequence for $H(1)$. It follows that for $n = 2^k$, $t(n) \in O(\sum_{i=0}^k (2^i)^2)$, where $\sum_{i=0}^k (2^i)^2 = 2^{2k} \sum_{i=0}^k (2^{-i})^2$ gives us that $t(n) \in O(n^2)$.

6.3.3 Triangulations $HMT(n)$ of the Hadamard examples $H(n)$

Let (τ_i^j) be the valid sequence for the $(n \times n)$ -Hadamard matrix $H(n) = (H_{ij})$ that is derived from the unique valid sequence of $H(1)$ by the above procedure.

From now on, we assume $n \geq 2$, as for $n = 1$ there is only one disc, which can easily be triangulated as a single triangle (hereby subdividing the boundary loop). Further, we will work with the augmented CW complex $\tilde{C}(n) \in DC(\tilde{H}(n))$ as defined earlier, i.e., in the polygonal disc corresponding to the row i the order of the edges is given by τ_i^* . In the following, we will construct a triangulation $HMT(n)$ of the complex $\tilde{C}(n)$:

- ▷ For each column j of $H(n)$ we subdivide the two corresponding cycles a_j^+ and a_j^- by using two additional vertices each, in particular, the subdivision of a_j^+ will be of the form $0-v_j^1-v_j^2-0$ and the subdivision of a_j^- will be of the form $0-w_j^1-w_j^2-0$, using a total of $4n + 1$ vertices for the $2n$ cycles and the initial vertex 0 .
- ▷ For each row i , we triangulate the interior of the corresponding disc (that has only positively oriented edges). Each copy of the special vertex 0 is shielded off by a triangle that contains the respective copy and its two adjacent vertices, which together uses n triangles. Then a single additional vertex c_i is placed at the center of the disc and connected to the $2n$ subdivision vertices of the loops using $2n$ triangles. In total, this gives $3n$ triangles and a single additional vertex to triangulate one of the discs.
- ▷ We triangulate each of the n digons between the loops a_j^+ and a_j^- using the following four triangles:

$$\{\{0, v_j^1, w_j^2\}, \{v_j^1, v_j^2, w_j^2\}, \{v_j^2, w_j^1, w_j^2\}, \{0, v_j^2, w_j^1\}\}, \quad \text{for } j \in \{1, \dots, n\}.$$

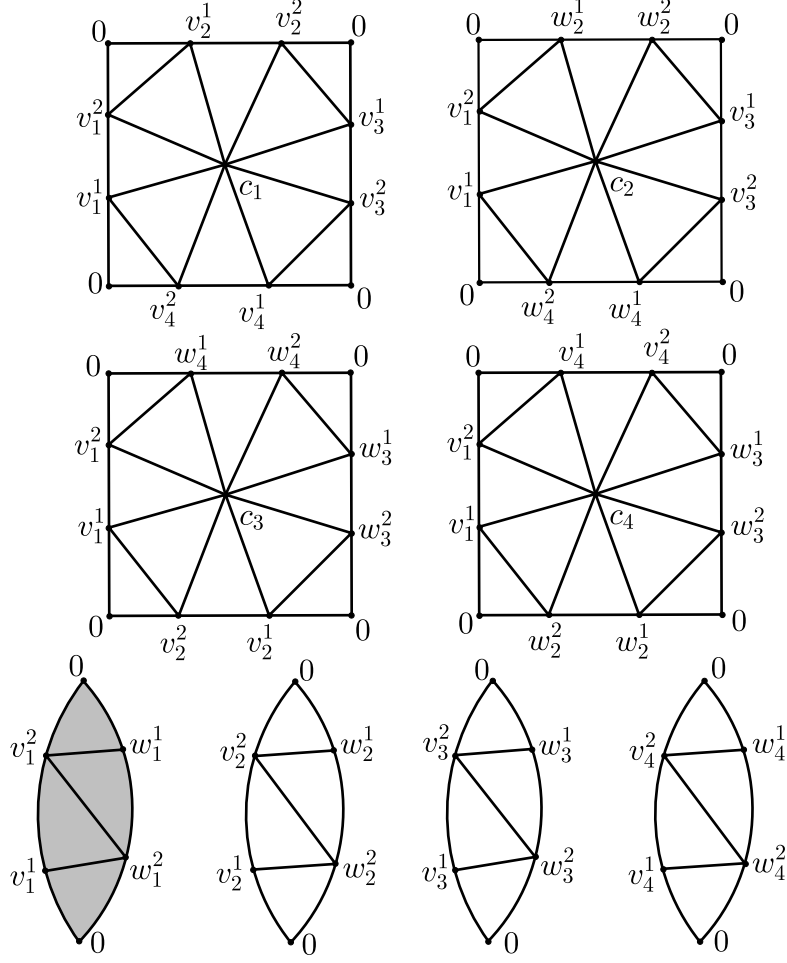


Figure 6.5: The simplicial complex $\text{HMT}(4)$ associated with the Hadamard matrix $H(4)$.

The triangulation corresponding to the valid sequence described in Example 6.19 for $H(4)$ is drawn in Figure 6.5.

By the definition of a valid sequence, we easily see that the interiors of the n -gons do not share a single edge. This implies that $\tilde{H}(n)$ is the boundary matrix of the first homology of our newly constructed simplicial complex.

Remark 6.21. We actually never use the (negative) cycle a_1^- in the construction of the n -gons, since for every n , the first column of $H(n)$ has only $+1$'s. We therefore delete from our construction the digon $a_1^+ a_1^-$ (marked in grey in Figure 6.5) and save two vertices, w_1^1 and w_1^2 , without modifying the homotopy type of our complex. In total, we then need $5n - 1$ vertices in our construction.

Since in Section 6.3.2 we showed the existence of valid sequences and given that we know the determinant and the Smith Normal Form of the Hadamard matrices $H(n)$, thanks to Lemma 6.13, we obtain the following:

Theorem 6.22. *For each $n = 2^k$, $k \geq 1$, there is a \mathbb{Q} -acyclic 2-dimensional simplicial complex $\text{HMT}(n)$ with face vector*

$$f(\text{HMT}(n)) = (5n - 1, 3n^2 + 9n - 6, 3n^2 + 4n - 4)$$

and $H_*(\text{HMT}(n)) = (\mathbb{Z}, T(\text{HMT}(n)), 0)$. The torsion in first homology is given by

$$H_1(\text{HMT}(n)) = T(\text{HMT}(n)) = (\mathbb{Z}_2)^{\binom{k}{1}} \times (\mathbb{Z}_4)^{\binom{k}{2}} \times \cdots \times (\mathbb{Z}_{2^k})^{\binom{k}{k}},$$

where $|T(\text{HMT}(n))| = n^{n/2} \in \Theta(2^{n \log n})$. Furthermore, the examples $\text{HMT}(n)$ can be constructed algorithmically in quadratic time $\Theta(n^2)$.

Proof. For any $k \geq 1$, due to Proposition 6.18, we have a valid sequence for $H(n)$, where $n = 2^k$. By Remark 6.21, $\text{HMT}(n)$ has $5n - 1$ vertices. For the edges, we have that each of the $2n - 1$ cycles (not $2n$, because one is not used according to Remark 6.21) has three distinct edges. By construction we are adding exactly $3n$ distinct edges in the interior of each of the n polygonal discs and three additional edges for each of the $n - 1$ digons, yielding a total of $3n^2 + 9n - 6$ edges. Again by construction, we are triangulating each of the n polygonal discs with $3n$ triangles and each of the $n - 1$ digons with four additional triangles, which gives a total of $3n^2 + 4n - 4$ triangles.

The simplicial complex $\text{HMT}(n)$ is clearly connected, so $H_0(\text{HMT}(n)) = \mathbb{Z}$, and the statement about $H_1(\text{HMT}(n))$ follows by Lemma 6.13 on the Smith Normal Form of the Hadamard matrices $H(n)$. The Euler characteristic of the simplicial complex $\text{HMT}(n)$ is 1 and we do not have a free part in the first homology, so we immediately obtain that $H_2(\text{HMT}(n)) = 0$.

Our construction of a valid sequence for $H(n)$ takes quadratic time in n . Given the valid sequence, building the triangulation $\text{HMT}(n)$ takes linear time (in the size of the valid sequence), i.e., quadratic time in n to output the triangulation with $3n^2 + 4n - 4$ triangles. Thus, in total, the examples $\text{HMT}(n)$ are constructed in quadratic time $\Theta(n^2)$. \square

An implementation `HMT.py` of our construction in `python` is available on GitHub at [Lof21a]. Triangulations of the examples $\text{HMT}(2^k)$, $2 \leq k \leq 5$, can be found online at the “Library of Triangulations” [BL21].

Remark 6.23. The construction above is not necessarily producing vertex-minimal triangulations of the disc complexes associated with the Hadamard matrices. For example, with the Random Simple Homotopy heuristic RSHT from [BLLL21] we were able to reduce the examples $\text{HMT}(2)$, $\text{HMT}(4)$, and $\text{HMT}(8)$ to smaller triangulations with 6, 11, and 34 vertices, respectively.

A general strategy for a possible asymptotic reduction of the number of vertices needed is by Newman [New19] via the *pattern complex* for a given complex. In particular, Newman used a revisited and randomized Speyer’s construction that yields linearly many edges in the number of vertices for complexes associated with a given abelian group G as their torsion group to obtain pattern complexes with $\Theta(\log(|G|)^{\frac{1}{2}})$ vertices—thus achieving Kalai’s asymptotic growth $\Theta(2^{n^2})$ via a randomized construction. In our construction, however, the number of edges is quadratic in the number of vertices, and therefore the effect of taking the pattern complex gives at most a linear improvement, which is not changing the asymptotic growth $\Theta(2^{n \log n})$ of the torsion size for our series of triangulations.

Remark 6.24. We have been made aware (by an anonymous reviewer of an earlier conference submission of this chapter as a self-standing article) of another folklore, but unpublished approach to achieve torsion that is exponential in n^2 . The idea of that construction is to use Steiner systems $S(2, 6, n)$. Such a system consists of a family of 6-element subsets (called blocks) of an n -element set, with the property that every pair of vertices is contained in exactly one block. On every block, one can construct a 6-vertex triangulation of $\mathbb{R}P^2$. This gives a complex with $\Theta(n^2)$ edge-disjoint $\mathbb{R}P^2$ ’s whose torsion is therefore $\mathbb{Z}_2^{\Theta(n^2)}$. The existence of such systems has recently been proved by Keevash in [Kee18], though in a non-constructive way. A deterministic construction that asymptotically yields high torsion can still be obtained, using the fact that it is enough to build on a partial system containing $\Theta(n^2)$ 6-element blocks where each pair is contained in *at most* one block. These can be generated via a polynomial-time derandomization of Rodl’s Nibble [Gra96], though the procedure seems to be impractical and we are not aware of an implementation.

Chapter 7

Conclusion and outlook

To wrap things up, we have seen in the previous four chapters different approaches to the problem of topological spaces' simplifications and type recognition. However, the problem in its generality is far from being solved and there are still many open questions that need to be addressed. I will give here a very brief overview of some of the open problems that have arisen in these years and that, unfortunately, we have not been able to solve yet.

In Chapter 3, we applied random collapses to a single simplex; cf. Table 3.5. The results strongly suggest that a simplex is nearly never randomly collapsible while going up in dimension. A theoretical proof of this however, is still missing—it would be a very nice addition and would greatly contribute to our understanding of the horizon for collapsibility. I must admit that I have spent quite a considerable amount of time in trying to find one but some small detail was always lacking. Hopefully someone will come soon and show me a nice and clean solution so that I can finally sleep better.

Another interesting challenge arising in the same chapter is how to implement computational topology tools to show that the Abkolut–Kirby spheres are indeed standard PL spheres. As we already said, this has been proved in the PL setting but till now all our heuristics and, as far as we know, every sphere recognition heuristic fails on these examples.

Finally in Section 3.3.2 we analysed the census of 4-dimensional complexes with up to six pentachora. In doing so we divided the census in three parts, spheres, non-spheres, and unknown examples. We believe that both the lists of non-spheres and of unknown examples could be a very useful library of interesting to recognise 4-dimensional manifolds; we encourage everyone to try their hands on them. The lists can be found on GitHub at [Lof21c].

In the following Chapter 4, we looked at contractible complexes that are not collapsible. We focused on finding explicit examples but of course another very interesting question would be to ask what is the probability of a contractible

complex to be collapsible for a large number of vertices n . We suspect that for large n almost all contractible complexes are neither collapsible nor anticollapsible but we don't have any proof of it.

The Random-Simple-Homotopy heuristic of Chapter 5, while performing incredibly well, is also still open for improvement and is far from being perfect especially on the technical side. In the future, we hope it would actually become a useful tool as a preprocessing step in homology computations.

Finally, as already said in Chapter 6, our explicit series of examples with high torsion in homology is still far from the upper bound by Kalai. Thus, would be very interesting to see concrete series that reach the bound.

We feel that the field of combinatorial and computational topology is in demand of more interesting examples and explicit constructions. I have tried in this thesis to provide some, but of course more are very welcome.

The above are only some of the uncountable open problems that are still present in the field and that hopefully will be solved in the upcoming years.

Bibliography

- [ABL17] Karim A. Adiprasito, Bruno Benedetti, and Frank H. Lutz, *Extremal examples of collapsible complexes and random discrete Morse theory*, Discrete Comput. Geom. **57** (2017), no. 4, 824–853.
- [AGH⁺17] Benjamin Assarf, Evgenij Gawrilow, Katrin Herr, Michael Joswig, Benjamin Lorenz, Andreas Paffenholz, and Thomas Rehn, *Computing convex hulls and counting integer points with `polymake`*, Math. Program. Comput. **9** (2017), no. 1, 1–38.
- [AI15] Karim A. Adiprasito and Ivan Izmistiev, *Derived subdivisions make every PL sphere polytopal*, Israel J. Math. **208** (2015), no. 1, 443–450.
- [AK85] Selman Akbulut and Robion Kirby, *A potential smooth counterexample in dimension 4 to the Poincaré conjecture, the Schoenflies conjecture, and the Andrews–Curtis conjecture*, Topology **24** (1985), no. 4, 375–390.
- [Akb10] Selman Akbulut, *Cappell–Shaneson homotopy spheres are standard*, Ann. of Math. (2) **171** (2010), no. 3, 2171–2175.
- [ALS19] Dominique Attali, André Lieutier, and David Salinas, *When convexity helps collapsing complexes*, 35th International Symposium on Computational Geometry (SoCG 2019), 2019, 15 pages.
- [BBP⁺17] Benjamin A. Burton, Ryan Budney, William Pettersson, et al., *Regina: Software for low-dimensional topology*, <http://regina-normal.github.io/>, 1999–2017.
- [BD98] Bhaskar Bagchi and Basudeb Datta, *A structure theorem for pseudomanifolds*, Discrete Math. **188** (1998), 41–60.
- [BD05] ———, *Combinatorial triangulations of homology spheres*, Discrete Math. **305** (2005), no. 1-3, 1–17.
- [BHS64] Hyman Bass, Alex Heller, and Richard G. Swan, *The Whitehead group of a polynomial extension*, Publications Mathématiques de l’Institut des Hautes Études Scientifiques **22** (1964), no. 1, 61–79.

- [Bin64] Rudolf H. Bing, *Some aspects of the topology of 3-manifolds related to the Poincaré conjecture*, Lectures on Modern Mathematics, Vol. II, Wiley, New York, 1964, pp. 93–128.
- [BK92] Ulrich Brehm and Wolfgang Kühnel, *15-vertex triangulations of an 8-manifold*, Math. Ann. **294** (1992), no. 1, 167–193.
- [BKR] Ulrich Bauer, Michael Kerber, and Jan Reininghaus, *PHAT*, <https://bitbucket.org/phant-code/phant/src>.
- [BL00] Anders Björner and Frank H. Lutz, *Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the Poincaré homology 3-sphere*, Experiment. Math. **9** (2000), no. 2, 275–289.
- [BL13a] Bruno Benedetti and Frank H. Lutz, *The dunce hat in a minimal non-extendably collapsible 3-ball*, Electronic Geometry Models No. 2013.10.001 (2013), <http://www.eg-models.de/2013.10.001>.
- [BL13b] ———, *Knots in collapsible and non-collapsible balls*, Electron. J. Combin. **20** (2013), no. 3, Paper 31, 29 pages.
- [BL14] ———, *Random discrete Morse theory and a new library of triangulations*, Exp. Math. **23** (2014), no. 1, 66–94.
- [BL21] ———, *Library of triangulations*, http://page.math.tu-berlin.de/~lutz/stellar/library_of_triangulations, 2013–2021.
- [BLLL21] Bruno Benedetti, Crystal Lai, Davide Lofano, and Frank H. Lutz, *Random simple-homotopy theory*, arXiv:2107.09862 (2021).
- [BLPS16] Benjamin A. Burton, Thomas Lewiner, Joao Paixao, and Jonathan Spreer, *Parameterized complexity of discrete Morse theory*, ACM Transactions on Mathematical Software (TOMS) **42** (2016), no. 1, 1–24.
- [Bor87] Karl-Heinz Borgwardt, *The Simplex Method, a probabilistic analysis*, Algorithms and Combinatorics: Study and Research Texts, vol. 1, Springer-Verlag, Berlin, 1987.
- [BR19] Ulrich Bauer and Abhishek Rathod, *Hardness of approximation for Morse matching*, Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SIAM, Philadelphia, PA, 2019, pp. 2663–2674.
- [BS93] Ulrich Brehm and Jacek Świątkowski, *Triangulations of lens spaces with few simplices*, SFB 288 Preprint NO. 59, TU Berlin, 26 pages, 1993.
- [BT09] Anders Björner and Martin Tancer, *Combinatorial Alexander duality—a short and elementary proof*, Discrete Comput. Geom. **42** (2009), no. 4, 586–593.

-
- [Bur] Benjamin A. Burton, *tricensus* — *Form a census of triangulations*, <https://regina-normal.github.io/docs/man-tricensus.html>.
 - [Bur11] ———, *Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations*, ISSAC 2011—Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation, ACM, New York, 2011, pp. 59–66.
 - [Bur14] ———, *A new approach to crushing 3-manifold triangulations*, Discrete Comput. Geom. **52** (2014), no. 1, 116–139.
 - [BZ11] Bruno Benedetti and Günter M. Ziegler, *On locally constructible spheres and balls*, Acta Math. **206** (2011), no. 2, 205–243.
 - [CC15] Maria Rita Casali and Paola Cristofori, *Cataloguing PL 4-manifolds by gem-complexity*, Electron. J. Combin. **22** (2015), no. 4, Paper 4.25, 25.
 - [CEK⁺] K. Crowley, A. Ebin, H. Kahn, P. Reyfman, J. White, and M. Xue, *Collapsing a simplex to a noncollapsible simplicial complex*, Preprint, 2003, 7 pages.
 - [Cha00] Manoj K. Chari, *On discrete Morse functions and combinatorial decompositions*, Discrete Math. **217** (2000), no. 1-3, 101–113, Formal power series and algebraic combinatorics (Vienna, 1997).
 - [CL06a] A. V. Chernavsky and V. P. Leksine, *Unrecognizability of manifolds*, Ann. Pure Appl. Logic **141** (2006), no. 3, 325–335.
 - [CL06b] Péter Csorba and Frank H. Lutz, *Graph coloring manifolds*, Algebraic and geometric combinatorics, Contemp. Math., vol. 423, Amer. Math. Soc., Providence, RI, 2006, pp. 51–69.
 - [Coh12] Marshall M. Cohen, *A Course in Simple-homotopy Theory*, vol. 10, Springer Science & Business Media, 2012.
 - [DFM04] Randall Dougherty, Vance Faber, and Michael Murphy, *Unflippable tetrahedral complexes*, Discrete Comput. Geom. **32** (2004), no. 3, 309–315.
 - [DHSW03] Jean-Guillaume Dumas, Frank Heckenbach, David Saunders, and Volkmar Welker, *Computing simplicial homology based on efficient Smith normal form algorithms*, Algebra, geometry, and software systems, Springer, Berlin, 2003, pp. 177–206.
 - [DKM09] Art Duval, Caroline Klivans, and Jeremy Martin, *Simplicial matrix-tree theorems*, Trans. Amer. Math. Soc. **361** (2009), no. 11, 6073–6114.
 - [Edm67] Jack Edmonds, *Systems of distinct representatives and linear algebra*, J. Res. Nat. Bur. Standards Sect. B **71B** (1967), 241–245.

- [Fir] Moritz Firsching, *Classifications and realizations: combinatorial spheres, polytopes and oriented matroids*, <http://page.mi.fu-berlin.de/moritz/polytopes>.
- [Fir17] ———, *Realizability and inscribability for simplicial polytopes via nonlinear optimization*, *Mathematical Programming* **166** (2017), no. 1-2, 273–295.
- [For98] Robin Forman, *Morse theory for cell complexes*, *Adv. Math.* **134** (1998), no. 1, 90–145.
- [For02] ———, *A user’s guide to discrete Morse theory*, *Sém. Lothar. Combin.* **48** (2002), Art. B48c, 35.
- [FR00] F. Thomas Farrell and Sayed K. Roushon, *The Whitehead groups of braid groups vanish*, *International Mathematics Research Notices* **2000** (2000), no. 10, 515–526.
- [Fre82] Michael H. Freedman, *The topology of four-dimensional manifolds*, *J. Differ. Geom.* **17** (1982), no. 3, 357–453.
- [Gan87] Bernhard Ganter, *Algorithmen zur formalen Begriffsanalyse*, *Beiträge zur Begriffsanalyse* (Bernhard Ganter, Rudolf Wille, and Karl Erich Wolff, eds.), Bibliographisches Inst., Mannheim, 1987, pp. 241–254.
- [GAP19] The GAP Group, *GAP – Groups, Algorithms, and Programming, Version 4.10.1*, 2019.
- [GB07] Nicolas Grenier-Boley, *On the triviality of certain Whitehead groups*, *Mathematical Proceedings of the Royal Irish Academy*, JSTOR, 2007, pp. 183–193.
- [GJ00] Ewgenij Gawrilow and Michael Joswig, *polymake: a framework for analyzing convex polytopes*, *Polytopes—combinatorics and computation* (Oberwolfach, 1997), DMV Sem., vol. 29, Birkhäuser, Basel, 2000, pp. 43–73.
- [GKT19] Sebastian Gutsche, Marek Kaluba, and Sascha Timme, *Polymake.jl, version 1.0*, <https://github.com/oscar-system/Polymake.jl/>, 2019.
- [Gom91] Robert E. Gompf, *Killing the Akbulut-Kirby 4-sphere, with relevance to the Andrews-Curtis and Schoenflies problems*, *Topology* **30** (1991), no. 1, 97–115.
- [Gon13] Bernd Gonska, *Inscribable polytopes via delaunay triangulations*, Ph.D. thesis, FU Berlin, 2013.
- [Gor19] Denis Gorodkov, *A 15-vertex triangulation of the quaternionic projective plane*, *Discrete Comput. Geom.* **62** (2019), no. 2, 348–373.
- [Gra96] David A. Grable, *Nearly-perfect hypergraph packing is in NC*, *Information Processing Letters* **60** (1996), no. 6, 295–299.

-
- [GS67] Branko Grünbaum and Vadakekkara P. Sreedharan, *An enumeration of simplicial 4-polytopes with 8 vertices*, Journal of Combinatorial Theory **2** (1967), no. 4, 437–465.
 - [Hac08] Masahiro Hachimori, *Decompositions of two-dimensional simplicial complexes*, Discrete Math. **308** (2008), no. 11, 2307–2312.
 - [Had93] Jacques Hadamard, *Resolution d’une question relative aux determinants*, Bull. des sciences math. **2** (1893), 240–246.
 - [HAM93] Cynthia Hog-Angeloni and Wolfgang Metzler (eds.), *Two-Dimensional Homotopy and Combinatorial Group Theory*, London Mathematical Society Lecture Note Series, vol. 197, Cambridge University Press, 1993.
 - [Hat02] Allen Hatcher, *Algebraic Topology*, Cambridge University Press, 2002.
 - [Hig40] Graham Higman, *The units of group-rings*, Proceedings of the London Mathematical Society **2** (1940), no. 1, 231–248.
 - [HJS19] Simon Hampe, Michael Joswig, and Benjamin Schröter, *Algorithms for tight spans and tropical linear spaces*, J. Symbolic Comput. **91** (2019), 116–128.
 - [HKL96] Matthew Hudelson, Victor Klee, and David Larman, *Largest j -simplices in d -cubes: Some relatives of the Hadamard maximum determinant problem*, Linear Algebra and its Applications **241-243** (1996), 519–598, Proceedings of the Fourth Conference of the International Linear Algebra Society.
 - [Ili89] Costas S. Iliopoulos, *Worst-case complexity bounds on algorithms for computing the canonical structure of finite abelian groups and the Hermite and Smith normal forms of an integer matrix*, SIAM J. Comput. **18** (1989), no. 4, 658–669.
 - [JLLT21] Michael Joswig, Davide Lofano, Frank H. Lutz, and Mimi Tsuruga, *Frontiers of sphere recognition in practice*, arXiv:1405.3848 (2021).
 - [JLT14] Michael Joswig, Frank H. Lutz, and Mimi Tsuruga, *Heuristic for sphere recognition*, Mathematical software—ICMS 2014, Lecture Notes in Comput. Sci., vol. 8592, Springer, Heidelberg, 2014, pp. 152–159.
 - [JP06] Michael Joswig and Marc E. Pfetsch, *Computing optimal Morse matchings*, SIAM J. Discrete Math. **20** (2006), no. 1, 11–25.
 - [Kal83] Gil Kalai, *Enumeration of \mathbf{Q} -acyclic simplicial complexes*, Israel J. Math. **45** (1983), no. 4, 337–351.
 - [KB79] Ravindran Kannan and Achim Bachem, *Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix*, SIAM J. Comput. **8** (1979), no. 4, 499–507.

- [KB83] Wolfgang Kühnel and Thomas F. Banchoff, *The 9-vertex complex projective plane*, The Mathematical Intelligencer **5** (1983), no. 3, 11–22.
- [Kee18] Peter Keevash, *Counting Steiner triple systems*, European Congress of Mathematics, 2018, pp. 459–481.
- [KLNP18] Matthew Kahle, Frank H. Lutz, Andrew Newman, and Kyle Parsons, *Cohen–Lenstra heuristics for torsion in homology of random complexes*, Experimental Mathematics (2018).
- [Koz08] Dmitry Kozlov, *Combinatorial Algebraic Topology*, Algorithms and Computation in Mathematics, vol. 21, Springer, Berlin, 2008.
- [KSS84] Jeff Kahn, Michael Saks, and Dean Sturtevant, *A topological approach to evasiveness*, Combinatorica **4** (1984), no. 4, 297–306.
- [Lac20] Marc Lackenby, *Algorithms in 3-manifold theory*, arXiv:2002.0217 (2020).
- [Lac21] ———, *The efficient certification of knottedness and Thurston norm*, Adv. Math. **387** (2021), Paper No. 107796, 142.
- [LL21] Davide Lofano and Frank H. Lutz, *Hadamard matrix torsion*, arXiv:2109.13052 (2021).
- [LLT03] Thomas Lewiner, Hélio Lopes, and Geovan Tavares, *Optimal discrete Morse functions for 2-manifolds*, Comput. Geom. **26** (2003), no. 3, 221–233.
- [LMR10] Nathan Linial, Roy Meshulam, and Mishael Rosenthal, *Sum complexes—a new family of hypertrees*, Discrete Comput. Geom. **44** (2010), no. 3, 622–636.
- [LN21] Davide Lofano and Andrew Newman, *The worst way to collapse a simplex*, Israel J. Math. **244** (2021), 625–647.
- [Lof21a] Davide Lofano, *Hadamard torsion program*, <https://github.com/davelofa/HadamardTorsion>, 2021.
- [Lof21b] ———, *Random homotopy extension*, <https://github.com/davelofa/RandomHomotopyExt>, 2021.
- [Lof21c] ———, *Topological types of the 4-manifolds with up to 6 pentachora*, <https://github.com/davelofa/Census6Pentachora>, 2021.
- [LP19] Nati Linial and Yuval Peled, *Enumeration and randomized constructions of hypertrees*, Random Structures & Algorithms **55** (2019), no. 3, 677–695.
- [LRRV17] Wolfgang Lück, Holger Reich, John Rognes, and Marco Varisco, *Algebraic K-theory of group rings and the cyclotomic trace map*, Advances in Mathematics **304** (2017), 930–1020.

-
- [Lus91] Martin Lustig, *Nielsen equivalence and simple-homotopy type*, Proceedings of the London Mathematical Society **3** (1991), no. 3, 537–562.
 - [Lut99] Frank H. Lutz, *Triangulated Manifolds with Few Vertices and Vertex-Transitive Group Actions*, Ph.D. thesis, TU Berlin, 1999.
 - [Lut03a] ———, *Seifert*, <http://page.math.tu-berlin.de/~lutz/stellar/SEIFERT>, 2003.
 - [Lut03b] ———, *Triangulated manifolds with few vertices: Geometric 3-manifolds*, arXiv:math/0311116 (2003).
 - [Lut04] ———, *Small examples of nonconstructible simplicial balls and spheres*, SIAM J. Discrete Math. **18** (2004), no. 1, 103–109.
 - [M⁺a] Konstantin Mischaikow et al., *CHomP: Computational homology project*, <http://chomp.rutgers.edu>.
 - [M⁺b] Marian Mrozek et al., *CAPD:RedHom — simplicial and cubical homology*, <http://redhom.ii.uj.edu.pl>.
 - [Mil66] John Milnor, *Whitehead torsion*, Bull. Amer. Math. Soc. **72** (1966), 358–426.
 - [MN13] Konstantin Mischaikow and Vidit Nanda, *Morse theory for filtrations and efficient computation of persistent homology*, Discrete Comput. Geom. **50** (2013), no. 2, 330–353.
 - [Mne14] Pavel Mnev, *Lecture notes on torsions*, arXiv:1406.3705 (2014).
 - [Mor] Dmitriy Morozov, *Dionysus*, <https://mrzv.org/software/dionysus2/>.
 - [Mun84] James R. Munkres, *Elements of Algebraic Topology*, Addison-Wesley Publishing Company, Menlo Park, CA, 1984.
 - [New19] Andrew Newman, *Small simplicial complexes with prescribed torsion in homology*, Discrete Comput. Geom. **62** (2019), no. 2, 433–460.
 - [Nov55] Pyotr S. Novikov, *On the algorithmic unsolvability of the word problem in group theory*, Izdat. Akad. Nauk SSSR, Moscow, 1955, Trudy Mat. Inst. Steklov. no. 44.
 - [Pac87] Udo Pachner, *Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulationen kompakter semilinearer Mannigfaltigkeiten*, Abh. Math. Sem. Univ. Hamburg **57** (1987), 69–86.
 - [Pac91] ———, *P.L. homeomorphic manifolds are equivalent by elementary shellings*, European J. Combin. **12** (1991), no. 2, 129–145.

- [Per02] Grisha Perelman, *The entropy formula for the Ricci flow and its geometric applications*, arXiv:math/0211159 (2002).
- [Rei05] Matthias Reitzner, *The combinatorial structure of random polytopes*, Adv. Math. **191** (2005), no. 1, 178–208.
- [Rou20] Sayed K. Roushon, *A certain structure of Artin groups and the isomorphism conjecture*, Canadian Journal of Mathematics (2020), 1–18.
- [Rub95] Joachim H. Rubinstein, *An algorithm to recognize the 3-sphere*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 601–611.
- [SAZ08] Emil Saucan, Eli Appleboim, and Yehoshua Y. Zeevi, *Sampling and reconstruction of surfaces and higher dimensional manifolds*, J. Math. Imaging Vision **30** (2008), no. 1, 105–123.
- [Sch11] Saul Schleimer, *Sphere recognition lies in NP*, Low-dimensional and symplectic topology, Proc. Sympos. Pure Math., vol. 82, Amer. Math. Soc., Providence, RI, 2011, pp. 183–213.
- [Si15] Hang Si, *Tetgen, a delaunay-based quality tetrahedral mesh generator*, ACM Transactions on Mathematical Software (TOMS) **41** (2015), no. 2, 11.
- [SK11] Jonathan Spreer and Wolfgang Kühnel, *Combinatorial properties of the K3 surface: simplicial blowups and slicings*, Exp. Math. **20** (2011), no. 2, 201–216.
- [SL09] Thom Sulanke and Frank H. Lutz, *Isomorphism-free lexicographic enumeration of triangulated surfaces and 3-manifolds*, European J. Combin. **30** (2009), no. 8, 1965–1979.
- [Sma61] Stephen Smale, *Generalized Poincaré’s conjecture in dimensions greater than four*, Ann. of Math. (2) **74** (1961), 391–406.
- [Smi61] Henry J. S. Smith, *Xv. on systems of linear indeterminate equations and congruences*, Philosophical transactions of the royal society of london (1861), no. 151, 293–326.
- [Spe] David E. Speyer, *Small simplicial complexes with torsion in their homology?*, MathOverflow, <https://mathoverflow.net/q/17852> (version: 2010-04-20).
- [ST34] Herbert Seifert and William Threlfall, *Lehrbuch der Topologie*, B. G. Teubner, Leipzig, 1934.
- [Sta65] John Stallings, *Whitehead torsion of free products*, Ann. of Math. (1965), 354–363.

-
- [Syl67] James J. Sylvester, *Thoughts on inverse orthogonal matrices, simultaneous sign successions, and tessellated pavements in two or more colours, with applications to Newton's rule, ornamental tile-work, and the theory of numbers*, The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science **34** (1867), no. 232, 461–475.
- [Tan16] Martin Tancer, *Recognition of collapsible complexes is NP-complete*, Discrete Comput. Geom. **55** (2016), no. 1, 21–38.
- [TDSL00] Joshua B. Tenenbaum, Vin De Silva, and John C. Langford, *A global geometric framework for nonlinear dimensionality reduction*, Science **290** (2000), no. 5500, 2319–2323.
- [Tho94] Abigail Thompson, *Thin position and the recognition problem for S^3* , Math. Res. Lett. **1** (1994), no. 5, 613–630.
- [Tie08] Heinrich Tietze, *Über die topologischen Invarianten mehrdimensionaler Mannigfaltigkeiten*, Monatshefte für Mathematik und Physik **19** (1908), no. 1, 1–118.
- [TL] Mimi Tsuruga and Frank H. Lutz, *Constructing complicated spheres*, EuroCG 2013, pp. 29–32.
- [VKF74] I. A. Volodin, V. E. Kuznetsov, and A. T. Fomenko, *The problem of discriminating algorithmically the standard three-dimensional sphere*, Russ. Math. Surv. **29** (1974), no. 5, 71–172.
- [Wal23] Joseph L. Walsh, *A closed set of normal orthogonal functions*, American Journal of Mathematics **45** (1923), no. 1, 5–24.
- [Wal70] David W. Walkup, *The lower bound conjecture for 3- and 4-manifolds*, Acta Mathematica **125** (1970), 75–107.
- [Whi39] John H. C. Whitehead, *Simplicial spaces, nuclei and m -groups*, Proc. London Math. Soc. (2) **45** (1939), no. 4, 243–327.
- [Zee64] E. Christopher Zeeman, *On the dunce hat*, Topology **2** (1964), 341–358.
- [Zie00] Günter M. Ziegler, *Lectures on 0/1-polytopes*, Polytopes—Combinatorics and Computation, Springer, 2000, pp. 1–41.