

Constructing Complicated Spheres

vorgelegt von
Master of Arts in Mathematics
Mimi Tsuruga
New York, NY, U.S.A.

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. Noemi Kurt

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

Berichter/Gutachter: Prof. John M. Sullivan

Berichter/Gutachter: Prof. Konstantin Mischaikow

Tag der wissenschaftlichen Aussprache: 1 Juli 2015

Berlin 2015

for mommy

Acknowledgements

On September 3, 2008, I arrived in Berlin utterly clueless, but excited for the adventure that lay before me. Now, nearly seven years later, I am preparing for the next adventure that awaits me in an even stranger land: California. I am not a new person now, but I am a different one. And I have many people to thank who helped me to become this changed me.

First and foremost is Frank, of course. Many people before me have thanked their thesis advisors, but I thank mine for so much more than just research advice. Frank's work—life balance; his insatiable curiosity and true love of science; his dealings with colleagues, with bureaucracy; his deep understanding of social rules and politics; and his choices to sometimes ignore them all; the time he spends doing work, the time he spends with his family; the surprisingly little time he must have not being busy. I am a better human being—and a better scientist—thanks to Frank and his guidance.

I've also had the great privilege of getting to personally know the person who guided Frank—his doctoral supervisor. It's easy to see the influence Günter has had on his students by the way Frank has been with me, and it is truly no wonder the Berlin Mathematical School is flourishing under his wing. He was my BMS mentor on paper, but also a great friend.

I also want to thank John for being a great mentor. At the times when I needed a mentor most, he helped me overcome many challenges. In a city where signs of home are few and far between, hearing him speak often brought me calm. He also gets the credit for connecting me with Frank, and for helping us get our project started (and navigating some of the bureaucracy that came with it).

But before the project started, I had a baby. And I had a problem coming back to work. It was not a problem of laziness, it was one of confidence. Carsten helped me overcome my fears of not being able to work again, having lost so much time and feeling left behind. He helped me through the toughest first weeks back from my leave; without him I may have given up.

And a great, wonderful thanks to the BMS. The Berlin Mathematical School wasn't just my source of funding. It became my new friends and family in this place I called home. Nadja and Tanja and Anja and Jürg and Marius and Domi and Chris and Shirley and Christof and Heike and Noemi and Christian and Sascha and Tim and all the rest. My fellow students—we happy few—it is for you and because of you that I found such devotion, dedication in me to give this my all. Eyal, especially, for suffering alongside me. My friends, in chronological order of my

meeting them: Natasa, Klebert, Tom, Artem, Kaie, Raman, Plamen, Josch, Matthias, Bruno, Peter, Dror, Sharad, Anna v, Margarita, Felix G, Irina, JP, Annie, Klaus, Brian, Cesar, Benjamin M, Silvia, Carla, Barbara, Emmanuel, Christoph, Faniry, Lothar, Isabella, Ella, Konstantin P, Janis, Florian, Wayne, Anna, Karim, Ahmad, Keita, Benjamin(s), Georg, Albert, Francesco, Isaac. Thank you.

Some acknowledgements for the content of this thesis: Thank you to Kontantin Mischaikow and Vidit Nanda for the first applications of our spheres. To Michael Joswig and his team of developers for all of the work related to `polymake`. A special thanks to Benjamin Lorenz for the time and care he put in to help me diagnose and repair many problems. Thanks to Javier Arsuaga and Michael Werner for inspiring the new bistellar code, which is still ongoing. Thanks to Bruno Benedetti for many fruitful conversations. And to Karim Adiprasito for many great ideas.

I also thank Google for translating my abstract to German, which was not required for the original graded submission, but apparently necessary for the library version. The library version also required that I remove many of my custom fonts, color choices, page formatting styles, etc. to conform to some standard that will likely no longer be the standard three years from now. So thank you library—and German bureaucracy—for the extra stress and headache you put me through while I prepare my family for a transatlantic move. I have surely become a more patient and more compliant person during my seven years living here; these reluctant changes to my character I hope to immediately reconvert upon my return to the US.

Lastly I want to mention some people who helped me get here. John for suggesting that I try new things. Robert for being the most understanding boss ever. George and Bettina for being amazing people I aspire to be like. All the people at Beckman and Wolfram Research who helped me work on Syzygy. Iryna and Indy and Wojtek and Jack and Howie for being my friends. Angie for being my sistah. My brother for being my brother. My husband for being my husband. My daughter for being my daughter.

Thank you ma for always supporting me. I know it's been hard for you. And there's still more ahead. I'm a better mother because of you. I'm a harder working, more ambitious woman because of you.

I'm a happier person because of you.

Berlin, August 12, 2015
Mimi Tsuruga

Preface

My daughter's default response to "what do you want to do today?" is "Lego!" We have different kinds of building blocks—wooden ones, magnetic ones, even a set of plastic tubes that you can send marbles through. But she always wants to play with those Legos. It might be because I told her that building with Legos is what I do at work.

Any parent or Lego enthusiast can tell you that Lego produces block sets that cater to two different types of customers. One customer type goes for the sets whose box cover displays a particular end product, a Millennium Falcon, perhaps; in those boxes are a step-by-step illustrated instruction manual and exactly the blocks required to build that particular version of the Millennium Falcon, plus or minus a few blocks. For those customers, finding extra blocks is like winning the lottery, while missing even one block is so frustrating he may buy a whole new box of the same set because that rare piece cannot be found elsewhere.

The other customer type does not look at the picture on the box at all, but she may pay attention to the box itself. Is the box sturdy and reusable? How many more blocks can it hold? The set comes with multiple copies of various block shapes in different colors. It is unlikely that this customer will ever count the blocks contained in this set or notice if any are missing. At some point more blocks are added to this box—like the ones left over from somebody's extra Millennium Falcon set. They will be used and reused many times over producing colorful Millennium Falcons, princess castles, or even entire zoos where each animal gets to have her own bed and kitchen. The box does contain a booklet inside, but it only makes suggestions for what one *could* build using a subset of the blocks contained in that box. Lego makes an extra effort to include a variety of suggestions that uses the blocks in different or not so obvious ways.

This thesis will show that I fall squarely into this second type of Lego consumer. The blocks we use have a particular shape and, much like the Legos, will connect to other blocks in a prescribed way. We can use as many blocks as we like, or as much as our computer's memory—our box—can hold. We include an illustrated instruction manual on how to build a particular set of colorful Millennium Falcons.

But our Millennium Falcon model is just a suggestion; we use it as an example to introduce the method we developed to build our Millennium Falcons so that anybody can build their very own Millennium Falcon, or even a whole new zoo of objects, limited only by their imagination.

Contents

Titlepage	3
1 Introduction	11
1.1 Basics	13
2 PL Sphere Recognition	17
2.1 Background	18
2.2 Tools	20
2.2.1 Preliminaries	21
2.2.2 Homology	21
2.2.3 Fundamental Group	24
2.2.4 <code>bistellar_simplification</code>	25
2.2.5 <code>Random_Discrete_Morse</code>	26
3 Akbulut–Kirby Spheres	35
3.1 Background	36
3.2 Construction	38
3.2.1 General Idea	38
3.2.2 Topological Description of Three Construction Types	40
3.2.3 Procedure for Three Construction Types	42
3.2.4 Detailed Construction	43
3.2.5 Implementation	52
3.3 Results & Experiments	58
3.3.1 Homology	59
3.3.2 <code>Random_Discrete_Morse</code>	59
3.3.3 <code>bistellar_simplification</code>	61
3.3.4 Fundamental Group	66
3.3.5 Conclusion	67
4 Mazur & Friends	69
4.1 Background	69
4.2 Construction	70
4.2.1 Mazur’s 4-manifold	70
4.2.2 A perfect non-PL 5-ball and a perfect non-PL 5-sphere	71
4.3 Results & Experiments	75
4.3.1 Mazur’s 4-manifold	75
4.3.2 A perfect non-PL 5-ball	76
4.3.3 A perfect non-PL 5-sphere	76

Chapter 1

Introduction

The term simplex comes from the same word in Latin, where it means “simple” or “plain”. We will use these simple simplices as building blocks to construct complicated spheres that turned out to be of interest:

- ▷ as hard instances for sphere recognition, Chapter 2;
- ▷ as hard instances for fundamental group computations, Chapter 3;
and
- ▷ as hard instances for homology computations [48].

A handlebody decomposition (detailed in Chapter 3) is yet another way to build manifolds using building blocks—in this case, topological balls—which connect to each other in a prescribed way. An “ i -handle” is a d -ball which “attaches” to a d -manifold. They attach to each other along each other’s boundary—but only a certain portion of the boundary. The number i indicates what portion of their boundaries will be attached. For example, if $i = 0$, then none of the i -handle’s boundary will be attached, it is just a new ball that does not attach to any part of the existing d -manifold. If $i = d$, then the entire boundary of the i -handle will attach to some appropriate portion of the d -manifold’s boundary. There is also an “attaching map” which gives directions to where an i -handle will attach to a d -manifold. One way to define an attaching map is using finite group presentations (see Section 2.1 and Section 2.2.3).

The Akbulut–Kirby spheres are a family of 4-manifolds that have a handlebody decomposition described by a certain family of groups that are nontrivial presentations of the trivial group. They each have one 0-handle, two 1-handles, two 2-handles, and one 4-handle. The attaching maps for the two 2-handles is described by:

$$G_{AK}(r) = \langle x, y \mid xyx = yxy; x^r = y^{r-1} \rangle \text{ for } r \in \mathbb{Z}.$$

The (smooth) 4-manifolds having these handlebody decompositions are known to be diffeomorphic to the standard 4-sphere [5]. The objects we construct, called simplicial complexes, are explicit triangulations of the Akbulut–Kirby spheres.

We built triangulations of the Akbulut–Kirby spheres for the following reasons:

1. In dimension 4, the categories $\text{PL}=\text{DIFF}$ coincide. In particular, it follows that a PL triangulation of the Akbulut–Kirby spheres is PL homeomorphic to the boundary of a 5-simplex. A finite PL triangulation can be read into a computer as a combinatorial object. Using topological software, we can then compute topological properties of such simplicial complexes, including (heuristically) whether or not a given complex is PL homeomorphic to a standard sphere.
2. In dimension 4, the sphere recognition problem is open, that is, the existence of an algorithm that can decide whether a given triangulated 4-manifold is a PL 4-sphere is unknown.
3. In dimension 4, the PL (and, therefore, also the smooth) Poincaré conjecture is still open. Computing the homology and fundamental group alone (see also reason 4) is not sufficient to say that a PL 4-manifold is a PL sphere.
4. As heuristics to recognize spheres, we use bistellar flips and random searches for discrete Morse functions. These triangulated Akbulut–Kirby spheres may be helpful test examples for improving the implementation of these heuristic algorithms.
5. The Akbulut–Kirby spheres were once exotic candidates for the smooth Poincaré conjecture in dimension 4. That is, since the smooth Poincaré conjecture is still open in dimension 4, it is not known whether there are 4-manifolds that are homeomorphic but not diffeomorphic to the standard 4-sphere. If the (combination of) heuristic algorithms we use can successfully recognize the triangulated Akbulut–Kirby spheres to be PL 4-spheres, we could potentially use the same method to triangulate current exotic candidates and again try the heuristic algorithms to recognize them to be PL 4-spheres.

The remainder of this chapter will go over much of the nomenclature and basic notions that is used throughout this thesis. Chapter 2 discusses the current state of the sphere recognition problem and, in particular, the heuristic algorithms being used in the topological software `polymake` and their limitations. Chapter 3 is dedicated to the Akbulut–Kirby spheres: their history, how we construct their triangulations, and the results of experiments run on these triangulations. Chapter 4 demonstrates the versatility of the construction method we describe in Chapter 3 by building Mazur’s 4-manifold (whose boundary is a homology sphere that is not the standard 3-sphere) and a non-PL 5-ball and non-PL 5-sphere that have perfect discrete Morse vectors.

1.1 Basics

This section provides the basic terms and concepts used throughout this thesis. We take heavily from many of the classic texts from piecewise linear manifold theory including Hudson [36], Kirby [41], Munkres [56], Stillwell [69], Whitehead [77], Zeeman [79], Zomorodian [81] and papers by Benedetti [9], Björner–Lutz [14], Lickorish [46]. Keep in mind that all of the combinatorial notions described here can be implemented as computer programs; they, in fact, are implemented in distributed software packages such as `polymake` [30] and Regina [16], but we wrote many custom functions on our own using Mathematica [78], GAP [29], and `polymake` scripts.

We start with the simplex. A(n abstract) **simplex** is a finite non-empty set. A simplex with one element is often called a vertex; a simplex with two elements, an edge; with three elements, a triangle; and with four elements, a tetrahedron. More generally, a **d -simplex** is a finite non-empty set of exactly $d + 1$ distinct elements; the order of the elements is arbitrary. A subset τ of a simplex σ is its **face** and we write $\tau \prec \sigma$. An i -face τ of a simplex σ is an i -simplex with $\tau \prec \sigma$. Our convention for notation may refer directly to the simplex, like σ , or write out the vertices of the simplex $[v_0 \ v_1 \ \cdots \ v_d]$. We often abuse notation and refer to a vertex $[v_0]$ by its vertex label v_0 , that is, we may refer to the element of a 0-face of a simplex as its vertex.

The **boundary** $\partial\sigma$ of a d -simplex σ is the collection of all i -faces of σ with $i < d$. Two simplices may share faces; then we say that they **intersect** and refer to the highest dimensional face they share as their intersection. If two simplices $\sigma = [v_0 \ \cdots \ v_\nu]$ and $\tau = [w_0 \ \cdots \ w_\mu]$ do not intersect, we can define their **join** as the $(\nu + \mu + 1)$ -simplex $[v_0 \ \cdots \ v_\nu \ w_0 \ \cdots \ w_\mu]$ and write $\sigma * \tau$.

A(n abstract) **simplicial complex**, is a finite non-empty collection of simplices and all their faces. A simplicial d -complex is a simplicial complex whose simplices have a maximum dimension of d . A simplex of a d -complex K who is not a face of any other simplex in K is a **facet** of K . If every facet of a d -complex K is of dimension d , then K is said to be **pure**. From this point forth we will only work with pure complexes.

A **subcomplex** J of a d -complex K is a subcollection of K that is itself a simplicial complex, denoted $J \subset K$. The boundary ∂K of K , for example, is a subcomplex of K . For $k \leq d$, the collection of all i -simplices of a d -complex K with $0 \leq i \leq k$ is a subcomplex called the **k -skeleton** of K and write $K^{(k)}$. The **star** of a simplex τ in a d -complex K is $\text{star}(\tau, K) = \{\beta \prec \sigma \mid \tau \prec \sigma, \sigma \in K\}$, a d -dimensional subcomplex whose facets are the d -simplices σ of K such that τ is a face of σ . The **link** of an i -simplex τ in a d -complex K is $\text{link}(\tau, K) = \{\sigma \in K \mid \tau * \sigma \in K\}$, a $(d - (i + 1))$ -dimensional subcomplex of simplices of K whose join with τ is a simplex of K . Notice that $\text{link}(\tau, K) \subset \text{star}(\tau, K)$.

The **join** $K_1 * K_2$ of a d_1 -complex K_1 and a d_2 -complex K_2 , where K_1 and K_2 are disjoint, is the collection of the join of each of the simplices of K_1 with each of the simplices of K_2 ; it is a $(d_1 + d_2 + 1)$ -complex. In particular, $\text{star}(\tau, K) = \tau * \text{link}(\tau, K)$.

The join is very different from the union. The union $K_1 + K_2$ of two d -complexes K_1, K_2 is just the usual set union, the collection of both of their simplices; it is a d -complex. [It is not recommended to take the union of complexes of different dimension as there is a chance you end up with a complex that is not pure.] Notice that in the case of the join of two complexes, the simplices of K_1 and K_2 may not have vertices in common to be able to take the join of them, whereas for a union it is possible (and likely) that the complexes have simplices that share vertices. For set difference, we write $K_1 - K_2$.

The **cone** of (or over) a d -complex K is $\text{cone}(K)$, a $(d + 1)$ -complex obtained by taking the join $a * K$ for a new vertex a ; this a is called the **apex** of the cone. A **stellar subdivision** of a d -complex K on a simplex $\sigma \in K$ is $\Psi_\sigma(K) = K - \text{star}(\sigma, K) + \text{cone}(\partial\sigma * \text{link}(\sigma, K))$. A **combinatorial d -ball** is a pure simplicial d -complex that has a common refinement by stellar subdivisions as a d -simplex. A **combinatorial d -sphere** is a pure simplicial d -complex that has a common refinement by stellar subdivisions as the boundary of a $(d + 1)$ -simplex.

A d -complex K is a **combinatorial manifold** (with or without boundary) if the link of every vertex of K is a combinatorial $(d - 1)$ -sphere or $(d - 1)$ -ball. A $(d - 1)$ -simplex in a combinatorial d -manifold is called a **ridge**. A ridge τ of a d -complex K can be a face of at most 2 facets of K ; if τ is a face of exactly 2 facets of K , it is said to be an **interior face** of K ; if τ is in exactly 1 facet of K , it is a **boundary face** of K . The **boundary** ∂K of a combinatorial d -manifold K is a subcomplex of K whose facets are the boundary faces of K . The vertices whose links are combinatorial balls are the vertices of the boundary ∂K of K .

In fact, any combinatorial d -manifold can be equipped with a compatible PL manifold structure; and conversely, every PL d -manifold can be triangulated as a combinatorial d -manifold [36, 79]. However, if $d \geq 5$, not all triangulations of a topological d -manifold need to be combinatorial triangulations. According to Kirby and Siebenmann [42], for $d \geq 5$, there are topological manifolds that don't admit combinatorial triangulations at all, that is, all their triangulations are non-PL triangulations. By the double suspension theorem of Edwards [25] and Cannon [17], it is also possible to have non-PL triangulations of all PL d -manifolds for $d \geq 5$, see for example the discussion in [14] and references therein.

A **bistellar i -move** on a (closed) combinatorial d -manifold K of an i -simplex $\sigma \in K$ is $\Phi_\sigma(K) = K - \text{star}(\sigma, K) + \tau * \partial\sigma$, where $\tau \notin K$ is a $(d - i)$ -simplex and $\partial\tau = \text{link}(\sigma, K)$. Note that a bistellar i -move on σ cannot be made unless there is an appropriate τ satisfying these strict restrictions. A stellar subdivision, on the other hand, is always possible. A bistellar d -move on a d -complex K is a stellar subdivision of a facet of K . Two combinatorial d -manifolds K, L are **bistellarly equivalent** if there exists a sequence of bistellar moves from one to the other. Two PL manifolds are bistellarly equivalent if and only if they are PL homeomorphic [58]. For example, a combinatorial d -sphere is bistellarly equivalent to the boundary of a $(d + 1)$ -simplex.

A **discrete vector field** \mathcal{V} on a simplicial complex K is a collection of pairs $\{\tau \prec \sigma\}$ of simplices of K such that τ is an $(i-1)$ -face of an i -simplex σ and each simplex is in at most one pair of \mathcal{V} . A **\mathcal{V} -path** is a sequence $\tau_0 \prec \sigma_0 \succ \tau_1 \prec \sigma_1 \succ \cdots \prec \sigma_{t-1} \succ \tau_t$ where each $\{\tau_n \prec \sigma_n\}$ is a pair in \mathcal{V} and for every n , $\tau_n \prec \sigma_{n-1} \neq \sigma_n$. If $\tau_t = \tau_0$ then the \mathcal{V} -path is said to be **closed**. A **gradient vector field** is a discrete vector field \mathcal{V} that has no nontrivial closed \mathcal{V} -paths. An i -simplex not contained in a gradient vector field \mathcal{V} is called a **critical cell** of dimension i . The **discrete Morse vector** of a gradient vector field \mathcal{V} is the vector (c_0, c_1, \dots, c_d) that counts the number of critical cells of each dimension. See [27] for more details.

A combinatorial manifold that has a discrete Morse vector having exactly one critical cell is a combinatorial d -ball [27, 77]. A discrete Morse vector is **spherical** if it reads $(1, 0, \dots, 0, 1)$. A combinatorial manifold (without boundary) that has a discrete Morse vector having exactly two critical cells is a combinatorial d -sphere. If K is a combinatorial ball or sphere, then we simply say their respective discrete Morse vectors are **perfect**.

Notation

d -simplex	$\sigma = [v_0 \ v_1 \ \cdots \ v_d]$
τ is a face of σ	$\tau \prec \sigma$
boundary of a simplex σ	$\partial\sigma$
join of simplices	$\sigma * \tau$
(abstract) simplicial complex	$K = \{\alpha, \beta, \dots, \sigma, \tau\}$
subcomplex J of K	$J \subset K$
star of a simplex τ in K	$\text{star}(\tau, K) = \{\beta \prec \sigma \mid \tau \prec \sigma, \sigma \in K\}$
link of a simplex τ in K	$\text{link}(\tau, K) = \{\sigma \in K \mid \tau * \sigma \in K\}$
join of complexes	$J * K = \{\sigma * \tau \mid \sigma \in J, \tau \in K\}$
cone over complex K	$\text{cone}(K) = a * K, \ a \notin K$
stellar subdivision	$\Psi_\sigma(K) = K - \text{star}(\sigma, K)$ $+ \text{cone}(\partial\sigma * \text{link}(\sigma, K))$

Chapter 2

PL Sphere Recognition

One fundamental question in topology is the homeomorphism problem. It asks whether a space X_1 and another space X_2 can be continuously transformed from one to the other, and back again. We can answer this question either by finding an explicit homeomorphism between the two spaces, or by showing that there exists a topological invariant that they do not share. Homology and the fundamental group are examples of topological invariants. Poincaré famously used these two invariants and conjectured whether they alone are sufficient to determine whether a topological manifold is homeomorphic to a sphere—and indeed they are.

A related problem is to ask whether one can algorithmically recognize certain spaces. To do so would require to first input the spaces into a computer. One option (if possible) is to discretize the topological space by triangulating it; that way we can work combinatorially on the abstract simplicial complexes that induced the triangulation. To check whether two triangulations are then (PL) homeomorphic to each other, we must first verify that they are both (PL) manifolds.

For a simplicial complex to be a PL manifold, we need to verify that all its vertex links are PL spheres. Here we can invoke the PL Poincaré conjecture, which is true in all dimensions except dimension 4. For dimensions $d \neq 4$, given a PL d -manifold K , if the homology groups $H_*(K) \cong (\mathbb{Z}, 0, \dots, 0, \mathbb{Z})$ and $\pi_1(K) \cong 1$, then K must be PL homeomorphic to the boundary of a $d + 1$ -simplex. There are (heuristic) algorithms for computing the homology and fundamental group for the vertex links. But to use the PL Poincaré conjecture, we need (to not be in dimension 4 and) to verify that those vertex links are also PL manifolds, which requires us to check if *their* vertex links are PL spheres. In this way, we recursively check for PL spheres—taking extra care in dimension 4.

In this chapter, we will discuss our strategy for recognizing PL spheres. We begin with some background in Section 2.1. We will then discuss the tools we use and some of their limitations in Section 2.2. The material presented here is a portion of joint work with Michael Joswig and Frank Lutz; more details can be found in our preprint [38].

2.1 Background

Before we dive into the history, we have to distinguish the different problems at hand. There is the Poincaré conjecture, of course. There is also the word and isomorphism problem, the manifold recognition problem, and the PL sphere recognition problem. For a thorough survey of the results and history of many of these related problems, we refer the reader to [53] and [69]. Here we only mention some of the highlights.

TOP and PL and DIFF

The topological category, or TOP, is a category whose objects are topological spaces with some manifold structure and morphisms are continuous maps. The piecewise linear category (PL) and smooth category (DIFF) place different conditions on the structure and maps. In dimensions 2 and 3, the categories TOP=DIFF=PL coincide [55, 63]. In dimensions 4–6, PL=DIFF coincide, but TOP does not [64]. In dimensions 7 and above, TOP≠PL≠DIFF.

The Poincaré Conjecture

The original statement of this famous conjecture is that a compact 3-dimensional topological manifold that is simply connected (i.e., has trivial fundamental group) must be homeomorphic to the standard 3-sphere S^3 . The Poincaré conjecture in higher dimensions, where there are different categories to consider, is referred to as the generalized Poincaré conjecture and states that a simply connected d -manifold with the homology of the d -sphere is the d -sphere. The topological Poincaré conjecture is true in all dimensions, $d = 3$: [59], $d = 4$: [28], $d \geq 5$: [68]. The PL Poincaré conjecture is true in all $d \neq 4$, $d = 3$ [59], $d \geq 5$: [68]. The smooth Poincaré conjecture is true in $d = 2, 3, 5, 6$ and open in $d = 4$. In dimension $d \geq 7$ exotic spheres, which are homeomorphic but not diffeomorphic to a standard d -sphere, are known to exist [54].

Poincaré introduced the fundamental group of a manifold in [61] connecting the geometric study of spaces with groups. His conjecture, then, motivated the next set of problems in group theory.

The Word, Conjugacy, and Isomorphism Problems

All three problems were first formally posed by Dehn [24]. The three problems are about finite group presentations. A finitely presented group

$$G = \langle x_1, \dots, x_n \mid \varphi_1, \dots, \varphi_m \rangle$$

has finitely many n generators x_i and finitely many m relators φ_i . The relators are written as words formed by the x_i 's. A finite presentation is said to be balanced if $n = m$.

Given a finitely presented group G , the word problem is a decision problem which asks whether an arbitrary word is trivial in G . The conjugacy problem, similarly, asks whether two words define conjugate

elements of G . The isomorphism problem asks, when given two arbitrary finite presentations, whether the two groups they present are isomorphic. In geometric terms, if we think of the group as a presentation of the fundamental group of a manifold, the word problem corresponds to determining whether loops are contractible; the conjugacy problem determines whether two loops are homotopic; and the isomorphism problem determines whether two manifolds have the same fundamental group.

The word and (therefore also) conjugacy problems are unsolvable in general [57]. However, in a probabilistic sense, almost all finitely presented groups have solvable word and conjugacy problems [32]. The isomorphism problem is unsolvable in general [3, 62].

The Manifold Recognition Problem

The manifold recognition problem asks whether there exists an algorithm that can decide whether one manifold is homeomorphic to another. The manifold recognition problem is often referred to as the homeomorphism problem in the literature. We will make a distinction here: the homeomorphism problem is the topological problem, the manifold recognition problem is the corresponding algorithmic decision problem (assuming the manifolds are presented in some finite encoding).

A.A. Markov [49] showed that for all $d \geq 4$, the PL d -manifold recognition problem is unsolvable, see also [21]. [15] extended it to the smooth category. In Markov's proof, he builds a manifold M_0 which cannot be recognized; [74] later showed that M_0 is the connected sum of several copies of $S^2 \times S^{d-3}$.

The Sphere Recognition Problem

The sphere recognition problem asks whether there exists an algorithm that can decide whether a given PL manifold is PL homeomorphic to the standard PL sphere. While sphere recognition is trivial in dimensions $d \leq 2$, Sergey P. Novikov showed that the problem is undecidable when $d \geq 5$; his proof can be found in the last chapter of [75], see also [21].

Since then most of the research focused on $d = 3$. [75] gave and tested¹ an algorithm for recognizing 3-spheres which was proven [35] to only be valid for input manifolds with Heegaard splittings of genus 2. In 1992, Rubinstein presented the existence of an algorithm for recognizing 3-spheres using normal surface theory during a workshop in Haifa and later published the results [65]. Thompson [72] simplified Rubinstein's proof. The proof was further generalized from simplicial complexes to handlebodies by [51] in 1995.

Once 3-sphere recognition was known to be possible, the question was then to determine its complexity. The first result in this direction was in 2001 when [37] showed that at least in a special class which he called Q-triangulations, the problem is in NP (and irreducibly co-NP). In 2011

¹They used a Soviet computer BESM 6. Only 355 of these machines were built. It ran at about 10 MHz and had 192kB of memory.

Schleimer [66] proved that generally 3-sphere recognition lies in NP. Hass and Kuperberg are now preparing a proof that 3-sphere recognition is in co-NP assuming the Generalized Riemann hypothesis (announced in [33]).

2.2 Tools

The task of recognizing higher dimensional spheres (or manifolds) seems to be doomed given S.P. Novikov’s non-recognizability result. Still, in many situations sphere recognition can be solved (easily) even for huge instances with few exceptions. In this section, we take a look at a few heuristic tools that have been developed and the success we observed in practice using these tools for recognizing explicit examples. We also discuss some of their limitations. Figures 2.1 and 2.2 describe our strategy for sphere recognition.

Key ideas for PL sphere recognition

Given a simplicial d -complex K , determine whether K is a PL d -sphere. See Figures 2.1 and 2.2 for a flow chart.

- In any dimension d : If K is pure and its ridges are contained in exactly 2 facets, then K is a pseudo-manifold.
- In dimension $d = 1$: If K is a pseudo-manifold and connected, then K is a polygon (which is a PL 1-sphere).
- In dimension $d = 2$: If K is a pseudo-manifold, has Euler characteristic $\chi(K) = 2$, and vertex links are single circles, then K is a PL 2-sphere.
- In any dimension d : If K is a pseudo-manifold and all its vertex links are PL $(d - 1)$ -spheres, then K is a combinatorial d -manifold.
- In dimension $d \neq 4$: If K is a combinatorial d -manifold, homology $H_*(K, \mathbb{Z}) = (\mathbb{Z}, 0, \dots, 0, \mathbb{Z})$, and the fundamental group $\pi_1(K) = 1$, then K is a PL d -sphere.
- In any dimension d : If K is a combinatorial d -manifold and admits a spherical discrete Morse vector, then K is a PL d -sphere.
- In any dimension d : If K is bistellarly equivalent to the boundary of a $(d + 1)$ -simplex, then K is a PL d -sphere.

Checking whether K is a combinatorial d -manifold, as mentioned earlier, means we have to recursively check for PL spheres. This is equivalent to checking for all face links to be PL spheres for faces of K of every dimension. Despite the fact that exact methods for $d = 3$ exist [65, 72], employing a heuristic approach even in the 3-dimensional case often turns out to be effective.

2.2.1 Preliminaries

Before running any other tests, we will first run some elementary combinatorial checks:

- (1) Is K pure?
- (2) Are all ridges of K contained in exactly 2 facets?
- (3) Is (the 1-skeleton of) K connected?

These checks are fast; their running time is bounded by a low-degree polynomial in the number of dimensions, number of facets and number of vertices. If one of these preliminary tests fail, this will serve as the certificate that K is not a sphere. Tests (1)–(3) always work and run in polynomial time; they are used to discard simplicial complexes that for obvious reasons are not manifolds.

One important remark for these first steps is that it requires us to build the Hasse diagram. The Hasse diagram will be used in all of the other tests. In fact, some programs were written in the `polymake` system specifically to take advantage of `polymake`’s fast Hasse diagram computation. Details on how `polymake` computes the Hasse diagram of a simplicial complex can be found in [38].

2.2.2 Homology

A necessary condition for a combinatorial d -manifold K 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 K has *spherical homology*. Computing the simplicial homology groups is a fairly standard procedure. For field coefficients this reduces to Gauss elimination being applied to the (simplicial) boundary matrices. Over the integers the computation of the Smith normal forms of the boundary matrices serves as a replacement; see [56, § 11]. Usually, we compute with integer coefficients. Current algorithms for computing the Smith normal form, though polynomial [40], are too slow for explicit computations on large examples. However, computation time can be dramatically improved by employing a topological preprocessor, e.g., CHomP [22], RedHom [18], Perseus [60]. For example, Perseus reduces the input complex by what they call iterated discrete Morse reductions. The idea here, very simply stated, is to find the “best” discrete Morse function by computing all discrete Morse functions (in the combinatorial sense as described in Section 1.1) on the complex and keeping track of the faces that appear relatively frequently. This discrete Morse function is an instruction for how one can reduce a complex to a complex that has fewer faces. One can then iterate this process until we reach a “smallest possible” reduced complex, which we can then compute the homology of directly.

We will not discuss homology here. For most of our tests, we used the homology client in `polymake` to compute the homology directly since the complexes we worked on were relatively small. All of our test examples had spherical homology, as expected.

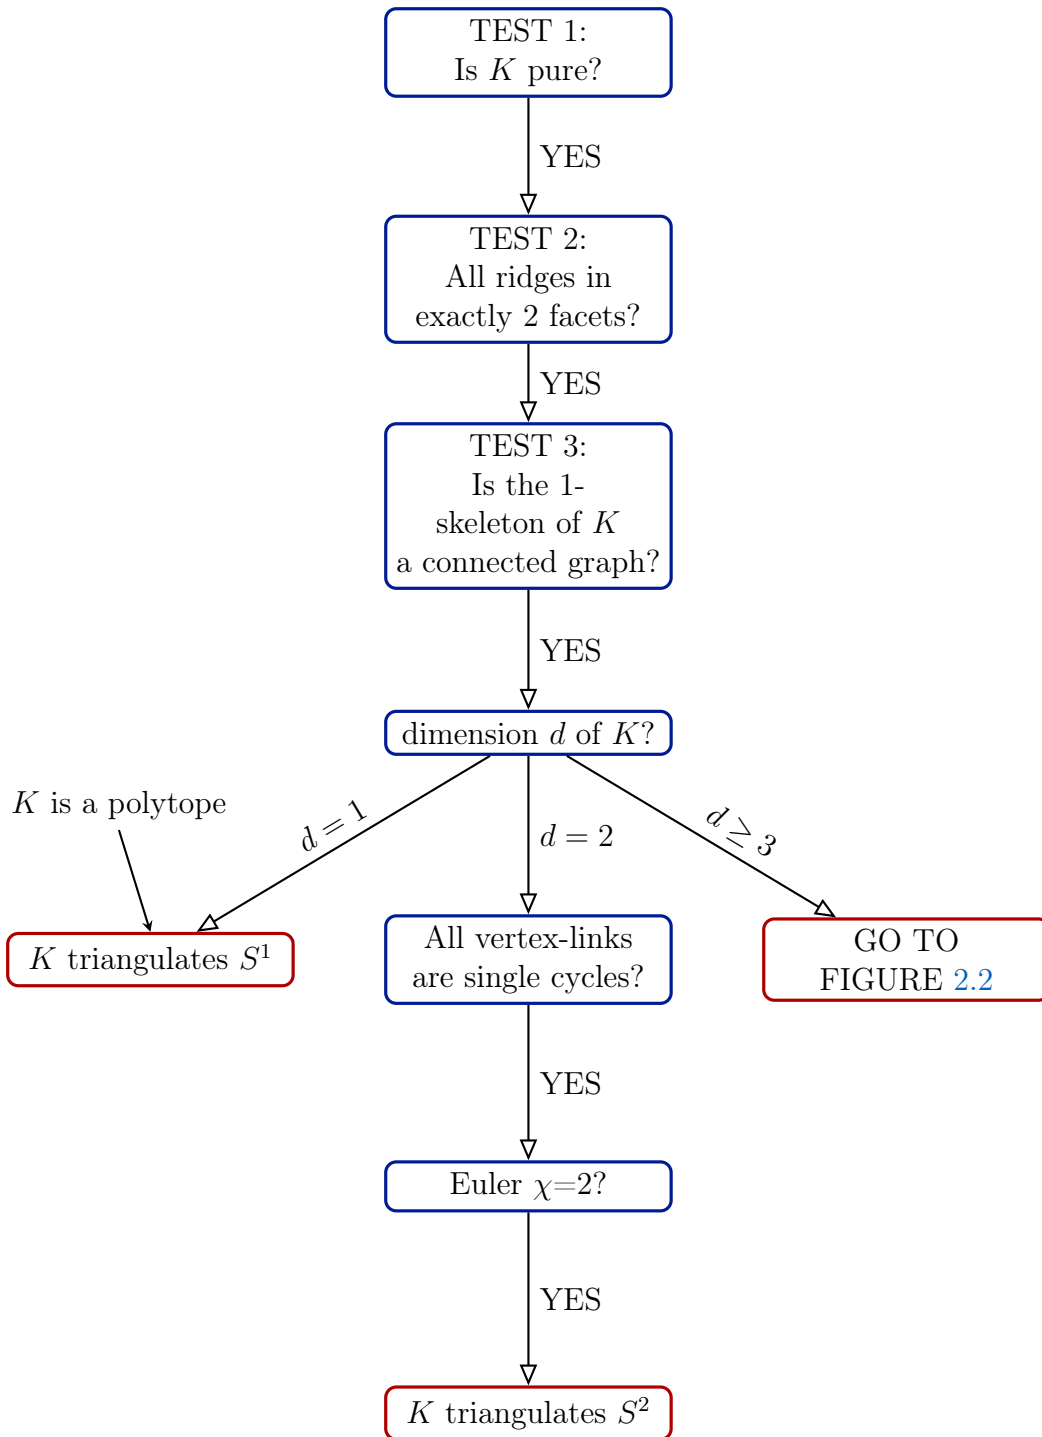


Figure 2.1: For a simplicial complex K .

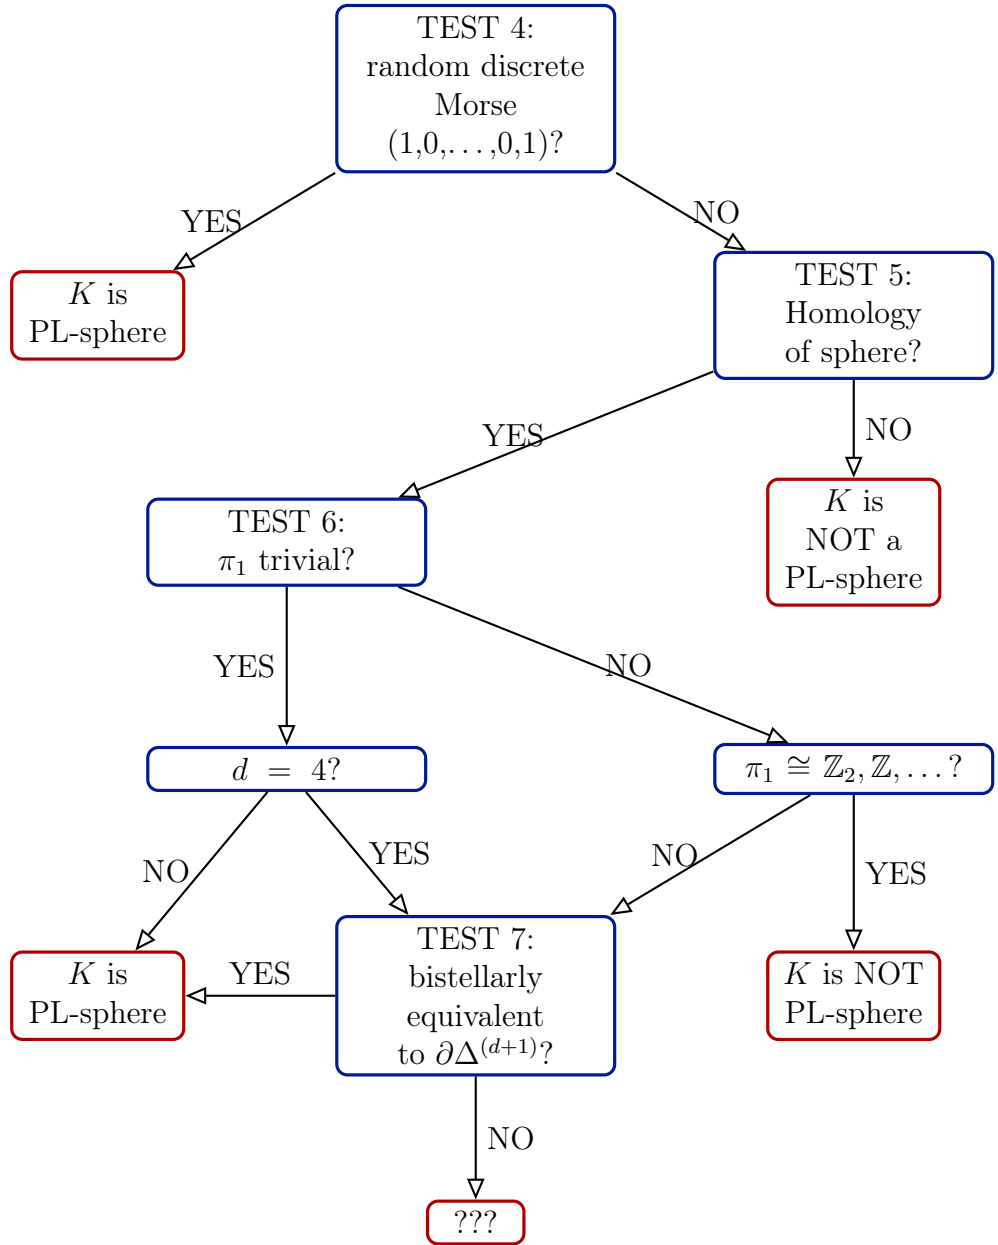


Figure 2.2: For a combinatorial d -manifold K of dimension $d \geq 3$.

2.2.3 Fundamental Group

A few years before Dehn formally posed the three (word, conjugacy, and isomorphism) problems, Tietze [73], in his habilitation, introduced four operations on group presentations which can be used to check whether the presentations are isomorphic, which we now know as Tietze moves. Given a finite group presentation $\langle x_1, \dots, x_n \mid \varphi_1, \dots, \varphi_m \rangle$:

- Add a relation φ_{m+1} which is a consequence of $\varphi_1, \dots, \varphi_m$.
- Its inverse: Remove a redundant relation, if one exists.
- Add a generator x_{m+1} along with a relation $\varphi_{m+1} = \omega(x_1, \dots, x_m)x_{m+1}^{-1}$ where $\omega(x_1, \dots, x_m)$ is a word composed of the alphabet $\{x_1, \dots, x_m\}$.
- Its inverse: Find a generator as a word in other generators and replace it in all relators.

Tietze showed that any two finite presentations of a group are convertible into each other by a finite sequence of Tietze moves. He even stated ² that checking for isomorphisms of group presentations is unsolvable decades before the concept of unsolvability was formalized. Indeed the isomorphism problem was later shown to be unsolvable [3, 62].

One example of a non-trivial presentation of the trivial group with only two generators and two relators that are not obviously trivial (using Andrews–Curtis moves; see Chapter 3) is

$$G_{AK}(r) = \langle x, y \mid xyx = yxy; x^r = y^{r-1} \rangle \text{ for } r \in \mathbb{Z}.$$

To see that $G_{AK}(r)$ is trivial, rewrite the first relator as

$$y = x^{-1}y^{-1}xyx.$$

Then take it to the r -th power and

$$\begin{aligned} y^r &= (x^{-1}y^{-1}xyx)(x^{-1}y^{-1}xyx) \cdots (x^{-1}y^{-1}xyx) \\ &= x^{-1}y^{-1}x^r yx \\ &= x^{-1}y^{r-1}x \\ &= x^r \end{aligned}$$

by substituting the second relator twice. This yields

$$y^r = y^{r-1}$$

or $y = e$ and therefore $x = e$. Thus $G_{AK}(r)$ is indeed isomorphic to the trivial group. This group is associated with the Akbulut–Kirby spheres.

In our tests, we derive a presentation of the fundamental group as an edge-path group of the complex [67] using `polymake` and then use `GAP` [29] to

²“Während sich nämlich die Gleichheit von zwei Zahlenreihen stets feststellen läßt, ist die Frage, ob zwei Gruppen isomorph seien, nicht allgemein lösbar.”

simplify the resulting presentation. GAP uses known techniques to simplify group presentations using Tietze transformations [34]. One choice made in their simplification strategy is that it prefers moves which shorten relations. One can see easily that the proof shown above of the triviality of the group G cannot be used for large r (since y^r produces a long word). However, it seems to be particularly difficult to construct a triangulation of a simply connected manifold for which GAP has been unable to recognize the presentation of its fundamental group to be trivial. The first such examples (known to us) are our triangulations of the Akbulut–Kirby spheres, see Chapter 3.

2.2.4 bistellar_simplification

The `bistellar_simplification` client implemented in `polymake` by Nikolaus Witte follows the simulated annealing strategy in [14]. It uses a local search strategy to determine the PL type of a simplicial complex. See Section 1.1 for the definition of bistellar moves [58].

The goal of the algorithm is to make these local changes, the bistellar moves, to the input simplicial complex to lower its f -vector (lexicographically) as much as possible. Naturally, the algorithm prefers moves that lower the f -vector (“cooling”). Unfortunately, we may fall into a local minimum, i.e., when there are no moves to further improve the f -vector. At that point, we deliberately make moves that increase the f -vector for some number of rounds (“heating”) then cool again, hoping that this will help jiggle us out of that local minimum. The threshold for the number of rounds we tolerate before heating and the amount we heat is changed dynamically throughout the computation (greater when the complex is large and smaller when the complex is small).

Since each move only makes local changes, the algorithm only updates the list of possible moves instead of manipulating and carrying the entire Hasse diagram, which would be computationally very costly. More specifically, there is a subset `raw_options` of all the i -dimensional faces of the complex that are contained in exactly $d - i + 1$ facets, $0 \leq i \leq d$. The `raw_options` are updated after each move. But only some of the `raw_options` are ‘proper’ options, i.e., potential moves, which must satisfy the additional condition that performing the move does not introduce a face that is already present in the complex. (The reason we do not update the list of proper options is because updating the raw options then checking that list is faster in the long run.)

The list of `raw_options` is grouped by dimension. The options that introduce new faces of lower dimension are heating moves and moves that introduce faces of higher dimension are cooling moves. During a cooling period we begin by checking for 0-moves, which will remove a vertex (0-face) from the complex (and replace its star with a new d -face, where d is the dimension of the complex). To select the 0-move at random, we start with the list of `raw_options` of dimension 0. Create a random permutation of that list. Then we check each option in that permuted order until we find

one that is proper. If none were proper, we move on to look for 1-moves, and continue to dimension $\lfloor \frac{d}{2} \rfloor$ until we come across a proper option.

During a heating period, the story is slightly different. All raw options of dimension d (of a d -complex) are proper, i.e., all facets can be stellarly subdivided. For the heating strategy, the dimension of the heating move is chosen at random respecting a chosen distribution. For example, say we input a 4-dimensional complex. The default heat distribution for dimension 4 is $[10, 10, 1]$; the distribution vector is of length $\lfloor \frac{d}{2} \rfloor + 1$; the first entry refers to $\lfloor \frac{d}{2} \rfloor$ -moves and the last entry refers to d -moves. So if the amount we heat `heating=210`, then the algorithm will heat for 210 rounds of which 10 (on average) are 4-moves (adds a vertex), 100 (on average) are 3-moves (adds an edge) and 100 (on average) are 2-moves (no change to f -vector), but the order is chosen uniformly at random. In higher dimensions ($d > 3$) choosing appropriate parameters for options, like the heating distribution, is non-trivial.

The triangulations of the Akbulut–Kirby 4-spheres $AK(r)$ in Chapter 3 are the only explicit examples of combinatorial spheres we know of for which `bistellar_simplification` did not find the boundary of a simplex.

2.2.5 Random_Discrete_Morse

A randomized search for small discrete Morse vectors was introduced in [10]. This approach traverses the Hasse diagram level-wise from top to bottom. The free faces for elementary collapses are chosen at random; if there are no free faces, a face of the current maximal dimension is chosen at random, marked critical, and removed.

In testing for whether a simplicial complex is a PL sphere, we use the following result.

THEOREM 2.1 (Whitehead [77], Forman [26, 27])

A combinatorial d -manifold (without boundary) is PL-homeomorphic to the standard PL d -sphere S^d if and only if it has some subdivision which admits a *spherical* discrete Morse vector, i.e., a discrete Morse vector with exactly one critical 0-cell and exactly one critical d -cell.

So this means that a given combinatorial manifold K is a PL sphere if `Random_Discrete_Morse` is able to find the discrete Morse vector $\langle 1, 0, \dots, 0, 1 \rangle$. However, this algorithm must overcome a number of major difficulties.

- Deciding whether a discrete Morse function with at most a fixed number of critical cells exists is NP-hard [39, 45].
- There are combinatorial d -spheres that do not admit any spherical discrete Morse function [12, 13] (however, there must be some subdivisions of them which do).

- Adiprasito and Izmestiev recently showed that a sufficiently large iterated barycentric subdivision of any PL sphere is polytopal [2] (and thus admits a spherical discrete Morse function). However, as pointed out above, for $d \geq 5$, recognizing the PL d -sphere is undecidable. This implies that a priori there is no bound on the number of barycentric subdivisions required to admit a spherical discrete Morse function.
- In iterated barycentric subdivisions, finding a spherical discrete Morse function quickly becomes cumbersome [1].

Below we will demonstrate that despite these drawbacks, finding optimal discrete Morse functions is often surprisingly easy, even for large input; see also [1, 10]. Not surprising, however, is that there are also other kinds of input for which our methods fail. A thorough analysis of the reasons for failure reveals several interesting families of simplicial complexes.

The `Random_Discrete_Morse` client implemented in `polymake` has three random strategies which we call *random-random*, *random-lex-first*, and *random-lex-last*. Here we outline our implementation and describe the differences between the three strategies.

Let K be a d -dimensional simplicial complex, which is not necessarily a manifold. A *free face* of K is a $(d-1)$ -dimensional face that is contained in exactly one d -face. In each step we try to pick one of the free faces and delete both it and the unique facet containing it from K . 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 differ in how they pick the free face. If we run out of free faces we pick some facet, declare it critical and remove it. In both cases, after removing a regular pair or after removing a critical face 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 K and collect them in a set or array data type. 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; see Knuth [44, §3.4.2]. If we run out of free faces the choice of the critical d -face is again uniformly at random.

Random-random is somehow the obvious strategy but there is a much cheaper way which maintains a certain amount of randomness. The idea is to randomly relabel the vertices of K once, at the beginning, and then to pick the free and critical faces in a deterministic way (which depends on the resulting labeling). Whenever a free or critical face is chosen, rather than selecting one at random, we pick the first (in the case of random-lex-first) or the last (in the case of random-lex-last) one.

The cost of being fair is quite significant when dealing with large complexes; for example, running the random-lex-first and random-lex-last strategies on the 4th barycentric subdivision of the boundary of a 4-simplex took less than 3 minutes per round whereas the random-random strategy took approximately 2 hours per round.

d	rounds	non-perfect	percentage
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	0.1%
18	10^3	2	0.2%
19	10^3	6	0.6%
20	10^3	13	1.3%

Table 2.1: Collapsing the d -simplex.

The random-lex-last strategy was called “random-revlex” in [10]. We changed this here to random-lex-last to avoid confusion with (the term) reverse lexicographic ordering, which is different.

Limitations for Random_Discrete_Morse

Dimension of the Input

The main difficulty we face is that once we start collapsing a triangulated sphere (after the removal of an initial critical facet), we might encounter subcomplexes that are contractible, but non-collapsible. The most prominent example of a non-collapsible, contractible complex is the 2-dimensional *dunce hat* [80] which can be obtained from a single triangle by identifying, in a non-coherent way, its three boundary edges. The dunce hat can be triangulated as a simplicial complex with 8 vertices (see [11]), while every contractible complex with fewer vertices is collapsible [8].

Once an initial critical facet is removed from a manifold candidate, our goal is to show that the remaining simplicial complex collapses to a vertex and thus is a PL ball (see Theorem 2.1). The most basic example of a d -dimensional PL ball is a single d -dimensional simplex. [23] gave an explicit collapsing sequence from the 7-simplex with 8 vertices to an 8-vertex triangulation of the dunce hat. Our examples below are similar in spirit.

discrete Morse vectors	count
(1 0 0 0 0 0 0 0 0)	999999988
(1 1 1 0 0 0 0 0 0)	4
(1 0 1 1 0 0 0 0 0)	7
(1 0 0 1 1 0 0 0 0)	1

Table 2.2: Spectrum for 10^9 runs on the 8-simplex.

discrete Morse vectors	count
(1, 0, 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, 0)	1
(1, 0, 3, 30, 111, 158, 132, 82, 24, 0, 0, 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, 0)	1
(1, 0, 1, 14, 27, 24, 13, 3, 0, 0, 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, 0)	1
(1, 0, 2, 25, 110, 236, 305, 175, 19, 0, 0, 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, 0, 0)	1
(1, 0, 1, 19, 82, 150, 161, 90, 15, 0, 0, 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, 0, 0)	1
(1, 0, 1, 11, 107, 243, 366, 463, 450, 261, 54, 0, 0, 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, 0, 0)	1
(1, 0, 6, 48, 182, 377, 657, 876, 801, 493, 170, 22, 0, 0, 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, 0, 0)	1

Table 2.3: Spectrum for 10^3 runs on the 20-simplex.




Figure 2.3: The four 2-complexes that were left over from attempted collapses of the 8-simplex.

When we collapse the faces of a d -simplex using strategies random-lex-first or random-lex-last, we will always reach a single vertex (since a d -simplex is a cone and the respective collapses are towards an apex). However, if free faces are chosen randomly for a d -simplex, $d \geq 7$, we might run into the dunce hat, like in the Crowley example, or other contractible, but non-collapsible subcomplexes.

Table 2.1 displays the results of some of our random experiments. In dimension 7, all 10^9 rounds of sequences of random collapses were perfect. From dimension 8 on, we see a clear increase in the number of non-perfect discrete Morse vectors encountered, which means that randomly finding perfect discrete Morse vectors for the d -simplex beyond dimension 20 becomes increasingly more difficult.

Although we have deterministic strategies to determine optimal discrete Morse vectors for various complexes, such as shellable complexes (which include single simplices), we cannot expect that running random collapses on a general ‘random’ input can perform better than on a single simplex.

Tables 2.2 and 2.3 give the actual discrete Morse vectors we found for the 8-simplex and the 20-simplex, respectively. We observe that we can get stuck (i.e., run out of free faces at a dimension $d > 0$) in different dimensions, as we already see for the 8-simplex in Table 2.2. While in the case of the 8-simplex we at most picked up two extra critical cells, the discrete Morse vector

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

for the 20-simplex in Table 2.3 contains 3632 extra critical cells. Thus, in higher dimensions (experimentally), not only do we get stuck with non-collapsible, contractible subcomplexes more often, but when we do get stuck, the resulting discrete Morse vectors typically are huge. The first time we get stuck, the remaining complex is still homotopy equivalent to the d -simplex we started with. In particular, this reduced complex is *contractible*,




Figure 2.4: A saw blade labeling. See Definition 2.2.

but non-collapsible. This means our heuristic procedure can also be used to search for such complexes.

In high dimensions, randomly finding spherical discrete Morse vectors for even the boundary d -sphere of a $(d + 1)$ -simplex is difficult, so testing combinatorial d -spheres that have more vertices and more facets will be even harder. Thus, a random search for spherical discrete Morse vectors will fail with high probability for *all* high-dimensional simplicial complexes, which renders `Random_Discrete_Morse` useless in high dimensions.

Let us return to the four examples of 2-dimensional contractible complexes we got stuck with when collapsing the 8-simplex; we call these four complexes D , S_a , S_b , and SQ . The first of these complexes D is a dunce hat as displayed on the top left in Figure 2.3. The other three complexes S_a , S_b , and SQ are not dunce hats, and they have inspired us to define a nice family of contractible, but non-collapsible 2-dimensional complexes that generalize the dunce hat.

DEFINITION 2.2

A *saw blade complex*, see Figure 2.4, is a 2-dimensional simplicial complex obtained from a triangulated disk by identifying vertices on the boundary of the disk such that

- the identification of the boundary vertices may not induce identifications of interior edges,
- the identification of the boundary vertices is 3-to-1, and
- the identification of the boundary edges is 3-to-1.

And after identification

- every boundary edge appears exactly twice with the same orientation and once with opposite orientation, and
- all identified boundary edges form a cycle.

The complexes D, S_a, S_b in Figure 2.3 are saw blade complexes; also every triangulation of the dunce hat is a saw blade complex. Saw blade

complexes with a different number of blades are combinatorially non-homeomorphic complexes. The complex SQ is obtained as a quotient of the sawblade complex S_b . Notice that the edge $[0\ 4]$ occurs four times on the boundary of SQ (the four colored edges in Figure 2.3). If we glue the two blue edges together and the two pink edges together, then the boundary of the new 2-disc we obtain is exactly the boundary of S_b . We can see there are an abundance of 2-dimensional contractible, non-collapsible simplicial complexes on which we can get stuck when randomly searching for simplicial collapses — and there will be similar constructions in higher dimensions. See [38] for a more detailed discussion.

Iterated barycentric subdivisions

As remarked earlier, it is often rather easy to find optimal discrete Morse vectors even for huge (nicely structured) complexes, see [10]. However, for higher-dimensional simplices or for higher barycentric subdivisions in fixed dimension $d \geq 4$ of boundaries of $(d + 1)$ -simplices [1] we might get stuck in substructures like the dunce hat. The probability to indeed encounter a dunce hat or a similar contractible, but non-collapsible subcomplex is extremely small in low dimension, as we have seen in our experiments above.

Adiprasito and Izmistiev [2] showed that a sufficiently large iterated barycentric subdivision of any PL sphere admits a spherical discrete Morse function. Yet, the average number of critical cells for random discrete Morse vectors grows exponentially with the number of barycentric subdivisions [1]. We ran our implementation on higher barycentric subdivisions of boundaries of simplices. For the 3rd barycentric subdivision of the boundary of the 4-simplex with $f = (12600, 81720, 138240, 69120)$ the optimal discrete Morse vector $(1, 0, 0, 1)$ was found in 994 out of 1000 runs of the random-lex-last version [1] of the random discrete Morse search; see Table 2.4. For the 4-th barycentric subdivision of the boundary of the 4-simplex with face vector $f = (301680, 1960560, 3317760, 1658880)$ the optimal discrete Morse vector $(1, 0, 0, 1)$ was found in only 844 out of 1000 runs, which may indicate that the horizon for computations lies near the 5-th barycentric subdivision.

random-random	random-lex-first	random-lex-last
sd_3_bd_delta_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_bd_delta_4 with $f = (301680, 1960560, 3317760, 1658880)$		
(1, 0, 0, 1): 20	(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

Table 2.4: Distribution of discrete Morse vectors for barycentric subdivisions.

Chapter 3

Akbulut–Kirby Spheres

The Akbulut–Kirby spheres are one of the first examples that appear in Kirby’s lecture notes “The Topology of 4-manifolds” [41]. They are given as explicit (and interesting) examples of a handlebody. A **handlebody** is a (compact PL¹) d -manifold \mathcal{M} with a handlebody decomposition. A **handlebody decomposition** of a compact PL manifold \mathcal{M} is a sequence

$$\mathcal{M}_0 \subset \mathcal{M}_1 \subset \cdots \subset \mathcal{M}_s = \mathcal{M},$$

where (i) $\mathcal{M}_0 = D^d = D^0 \times D^d$; and (ii) each \mathcal{M}_i is obtained from \mathcal{M}_{i-1} by attaching a k -handle. A d -dimensional **k -handle** $H^{k,d}$ is a product of balls $D^k \times D^{(d-k)}$; a k -handle $H^{k,d}$ attaches to \mathcal{M}_{i-1} along the $S^{(k-1)} \times D^{(d-k)}$ part of its boundary onto the boundary of \mathcal{M}_{i-1} .

$$\mathcal{M}_i = \mathcal{M}_{i-1} \cup_{f_i} H^{k,d}, \text{ where } f_i : \partial H^{k,d} \rightarrow \partial \mathcal{M}_{i-1} \text{ is an embedding.}$$

The sequence always starts with $\mathcal{M}_0 = H^{0,d}$, a d -dimensional 0-handle. Note that it is important to distinguish between the dimension of the handlebody d and the index of the handle k . When the dimension of the handlebody is understood, we may write $H^k = H^{k,d}$. A d -sphere always has at least one 0-handle and one d -handle.

The Akbulut–Kirby spheres form an infinite family of homotopy 4-spheres that have all been shown to be diffeomorphic to the standard sphere [5, 31]. They are described by a handlebody description indexed by a family of nontrivial balanced presentations of the trivial group: $G_{AK(r)} = \langle x, y \mid xyx = xyx, x^r = y^{r-1} \rangle$ and can be obtained as follows.

- Start with a 4-ball, i.e., a 4-dimensional 0-handle.
- Attach two 4-dimensional 1-handles. Label them x and y and give them some orientation.
- Attach two 4-dimensional 2-handles according to the relators of $G_{AK(r)}$.
- Attach a 4-dimensional 4-handle.

¹Kirby defines handlebodies for smooth manifolds. A manifold that admits a smooth handlebody decomposition will admit a compatible PL handlebody decomposition [9].




Figure 3.1: The red curve in the interior of H^3 in dimension 3 ends up on the boundary of $H^3 \times \mathbb{I}$.

We want to triangulate the Akbulut–Kirby spheres for reasons outlined in Chapter 1. To do this we follow a similar recipe as the above handlebody decomposition. However, to have better control over the way the 2-handles are attached, we start in dimension 3 and build a 3-ball with two 1-handles. In fact, a 4-dimensional 2-handle is a product $D^2 \times D^2$ which is attached via $S^1 \times D^2$ along its boundary. This $S^1 \times D^2$ is a solid 3-torus that can be interpreted as a regular neighborhood of the closed circle S^1 running through the interior of the 3-ball with the two 1-handles. We then bring the handlebody up to a higher dimension. In doing so we are able to “push” the circle out to the boundary so that the 2-handles can be glued onto the thickened circle. Consider Figure 3.1 as an example. An arbitrary (closed, tame) red curve lies in the interior of a 3-dimensional handlebody H^3 . This red curve represents the S^1 which is thickened to $S^1 \times D^2$. After taking the product $H^3 \times \mathbb{I}$, we can find (a copy of) the red curve on the boundary of $H^3 \times \mathbb{I}$. So the entire red curve is pushed out to the boundary of $H^3 \times \mathbb{I}$ even though in dimension 3, the red curve was in the interior of H^3 .

There is much freedom for choosing a construction method for the triangulation. This is just one example of many design decisions we have made to produce an automated program that will output three different triangulation types for the Akbulut–Kirby spheres for any given r . We begin with a short history of the Akbulut–Kirby spheres in Section 3.1. Then we describe, in detail, our construction method for producing the three triangulation types of the Akbulut–Kirby spheres in Section 3.2. Finally, we discuss the experiments we ran on the triangulations in Section 3.3

3.1 Background

In the seventies, Cappell and Shaneson [19] introduced an infinite family of homotopy 4-spheres as exotic candidates, that is, they presented these spheres as potential counterexamples to the smooth Poincaré conjecture in dimension 4. They indexed this family of spheres by a conjugacy

class of Cappell–Shaneson matrices $A \in SL(3, \mathbb{Z})$ with $\det(A - I) = \pm 1$ and a choice of two framings (trivial=untwisted or nontrivial=Gluck construction). The spheres with the trivial framings were all shown to be standard by Aitchinson and Rubinstein [4].

Soon after, Akbulut and Kirby [6] simplified one of the twisted cases with

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

to reveal a handlebody whose fundamental group can be read off from a diagram of the attaching maps of its handles, see Figure 29 in [6], as a nontrivial, balanced group presentation of the trivial group:

$$\langle x, y \mid xyx = yxy, x^5 = y^4 \rangle.$$

This associated group presentation can be generalized over a single parameter to index an infinite family of homotopy 4-spheres called the Akbulut–Kirby spheres. Each sphere in this family has two 1-handles, two 2-handles, and no 3-handles (in addition to the one 0-handle and one 4-handle). The attaching maps for the 2-handles are described by the presentation

$$G_{AK(r)} = \langle x, y \mid xyx = yxy, x^r = y^{r-1} \rangle \text{ over } r \in \mathbb{Z}. \quad (3.1)$$

Akbulut and Kirby [6] showed that these manifolds are homeomorphic to S^4 and that they naturally occurred as double covers of homotopy $\mathbb{R}P^4$'s. Since these so-called fake $\mathbb{R}P^4$'s were shown to be exotic by Cappell and Shaneson [19] some years earlier, these twisted examples became especially interesting to investigate as potential counterexamples to the smooth Poincaré conjecture.

Akbulut and Kirby wondered whether the group presentation associated with these spheres are counterexamples to another conjecture—the Andrews–Curtis conjecture². They also noted that these spheres could be counterexamples to yet another conjecture—the Schoenflies conjecture—which states that a smoothly embedded 3-sphere in a 4-sphere bounds a 3-ball³.

What started as one interesting example from a large family of Cappell–Shaneson 4-spheres had grown to become a mystical dragon that if defeated could disprove three conjectures. Over three decades after the spheres

²The Andrews–Curtis conjecture [7] states that any finite balanced presentation of the trivial group can be shown to be trivial by a finite number of Nielsen moves (three moves that changes the relators by: changing one of them with its inverse, swapping one with another, or combining two of them) plus one other move (changing one with its conjugate). These algebraic Andrews–Curtis moves correspond to geometric handle moves. If the presentation is found to be Andrews–Curtis trivial, that sequence of Andrews–Curtis moves would describe the diffeomorphism from the handlebody to S^4 .

³In particular, Akbulut and Kirby showed in [6] that $B = AK(5)$ minus the final 4-handle can be smoothly embedded in S^4 and, furthermore, that $\partial B = S^3$.

were first introduced, the Akbulut–Kirby spheres—and every one of the Cappell–Shaneson spheres—were shown to be diffeomorphic to S^4 [5, 31] disqualifying the Akbulut–Kirby spheres as exotic candidates, and therefore also disqualifying them as counterexamples to the Schoenflies conjecture⁴.

The (relative) simplicity of the group presentation and the difficulty smooth topologists had to prove them to be diffeomorphic to S^4 suggested that explicit triangulations of these spheres may produce some interesting objects to study.

3.2 Construction

In this section we describe our construction for explicit triangulations of the Akbulut–Kirby spheres. We begin with the general idea in Section 3.2.1. Section 3.2.2 explains the three construction types topologically. Section 3.2.3 gives a step-by-step summary of the construction procedure described in Section 3.2.2. Section 3.2.4 gives the details of the construction. And, finally, in Section 3.2.5, we discuss some important considerations with respect to our implementation of the automated program that produces triangulations of the Akbulut–Kirby sphere for any given r .

3.2.1 General Idea

Any finitely presented group G can be the fundamental group of a compact 4-manifold⁵, see e.g. [50, Ch.IV, Notes] or [69, 9.4]. Start with a ball; one can think of this as a base point for the fundamental group. Attach as many 1-handles to the ball as there are generators. Label and orient these handles. Then for each relator, glue in a 2-handle whose attaching map is described by its associated relator; here we pay attention to the orientation of the 1-handles. This produces a manifold with boundary whose fundamental group is isomorphic to the group G .

For the Akbulut–Kirby spheres, the group we start with is the above presentation (3.1) $G_{AK(r)}$ of the trivial group. Akbulut and Kirby [6] showed that for any r , the 4-manifold we obtain is (diffeomorphic to) a 4-ball $\mathcal{B}_{AK(r)}^4$. Glue on a 4-handle and we would get a (homotopy) 4-sphere.

⁴The presentation $G_{AK(3)}$ has been shown to be Andrews–Curtis trivial by Casson, but $r > 3$ still remain as candidates for counterexamples to the Andrews–Curtis conjecture; $r < 3$ is trivial. We want to further remark that since Akbulut [5] showed that the entire family of spheres is standard, by Pachner [58] we know that the existence of a sequence of bistellar moves between any (combinatorial) triangulation of the Akbulut–Kirby spheres to the boundary of a 5-simplex is guaranteed. How long that sequence may be or how one can find it is, however, unknown for $r > 3$. With the Andrews–Curtis conjecture still open, bistellar flips offers an alternative “sequence of moves” one can attempt to recognize spheres.

⁵This was a problem that was asked by Poincaré [61] and answered by Dehn [24]. It was used by Markov [49] for his proof of the unsolvability of the d -manifold recognition problem for $d \geq 4$.

LEMMA 3.1

Let B be a topological (or PL) d -ball. Then taking the double of B , i.e., taking a copy of B and identifying them along their common boundaries $B \times \{0, 1\} / \sim$, where $(x, 0) \sim (x, 1)$ for all $x \in \partial B$, is homeomorphic (or PL-homeomorphic) to attaching a 4-handle to B .

LEMMA 3.2

Let M be a d -dimensional topological manifold with boundary. Then $\partial(M \times \mathbb{I})$ is homeomorphic to taking the double of M .

The preceding lemmas merely follow by definition, but will help to clarify the next statement. Since $\mathcal{B}_{AK(r)}^4$ are proper 4-balls, we can take $\partial(\mathcal{B}_{AK(r)}^4 \times \mathbb{I})$, which a priori are homotopy spheres. Gompf [31] (for $r = 5$) and Akbulut [5] (for all r) showed that they are also diffeomorphic to the standard 4-sphere.

Theoretical Procedure

From a (nontrivial) presentation of the trivial group

$$G = \langle x_1, \dots, x_n \mid \varphi_1, \dots, \varphi_m \rangle,$$

construct a topological d -sphere ($d \geq 4$) that respects the handlebody decomposition described by the group G .

Step T1: Start with a d -ball $H^{0,d}$.

Step T2: Attach n 1-handles $H_{x_1}^1, \dots, H_{x_n}^1$ corresponding to the generators x_1, \dots, x_n to H^0 .

$$H = H^0 \bigcup_{i=1}^n H_{x_i}^1.$$

Step T3: Place disjoint curves $\varrho_1, \dots, \varrho_m$ on ∂H according to the relators $\varphi_1, \dots, \varphi_m$.

Step T4: Thicken the curves to $d - 1$ dimensions:

$$\tau_1 = \varrho_1 \times D^{d-2}, \dots, \tau_m = \varrho_m \times D^{d-2}.$$

Step T5: Glue in 2-handles $H_{\varphi_1}^2, \dots, H_{\varphi_m}^2$ along τ_1, \dots, τ_m :

$$\mathcal{B} := H \bigcup_{\varphi_1} H_{\varphi_1}^2 \bigcup \dots \bigcup_{\varphi_m} H_{\varphi_m}^2.$$

Step T6: Take the cross product $\mathcal{B} \times \mathbb{I}$.

Step T7: Take the d -dimensional boundary $\partial(\mathcal{B} \times \mathbb{I})$.

3.2.2 Topological Description of Three Construction Types

We now explain the topological theory behind our three construction types for triangulating the Akbulut–Kirby spheres. We begin with the finitely presented group

$$G_{AK(r)} = \langle x, y \mid \varphi_{blue} : xyx = yxy, \varphi_{red} : x^r = y^{r-1} \rangle.$$

To actualize the plan described above, we choose to change the order slightly. Rather than starting with a d -dimensional space H and finding (thickened) curves representing the gluing maps for the 2–handles on its boundary ∂H , we instead lay the curves first in 3-dimensional space, thicken them, and fill up the space around the thickened curves with tetrahedra to obtain H^3 , making sure to put some buffers between the curves to avoid unwanted identifications. Then take the cross product $H^3 \times \mathbb{I}^{d-3}$ to a higher dimensional space to glue in two d -dimensional 2–handles.

This gluing of the two 2–handles is the crucial step that motivated our idea to build H^3 in dimension 3. Recall Figure 3.1. The 3-dimensional thickened curves τ_{blue}, τ_{red} (which correspond to the two relators $\varphi_{blue}, \varphi_{red}$, respectively) will be further thickened to dimension 4 to $\tau_{blue} \times \mathbb{I}$ and $\tau_{red} \times \mathbb{I}$. After this process, two copies of τ_{blue} (and two copies of τ_{red}) can be found on the boundary $\partial(H^3 \times \mathbb{I}) = (\partial H^3 \times \mathbb{I}) \cup (H^3 \times \partial \mathbb{I})$; we can think of them as $\tau_{blue} \times \{0\}$ and $\tau_{blue} \times \{1\}$ (or $\tau_{red} \times \{0\}$ and $\tau_{red} \times \{1\}$, respectively). We can then choose one of these, say $\tau_{blue} \times \{0\}$, as the attaching map for the 4-dimensional 2–handle $H_{blue}^{2,4}$ or, similarly, $\tau_{red} \times \{0\}$ for $H_{red}^{2,4}$.

NOTATION 3.3

Let M be a PL $(d-1)$ -manifold with boundary with $\tau \subset M$. We will write

$$(M \times \mathbb{I}) \bigcup_{\tau} H^{2,d} := (M \times \mathbb{I}) \bigcup_f H^{2,d},$$

where $f : \partial(M \times \mathbb{I}) \rightarrow \partial H^{2,d}$ is an embedding and $f|_{\tau \times \{0\}}$ is the identity map. This means that the 2–handle $H^{2,d}$ is a d -ball ($D^2 \times D^{d-2}$) such that an $S^1 \times D^{d-2}$ part of this 2–handle’s boundary, specifically $\tau \times \{0\}$, is identified with a part of the boundary $\partial(M \times \mathbb{I})$. That is, $\tau \times \{0\} \subset \partial(M \times \mathbb{I})$ and $\tau \times \{0\} \subset \partial H^{2,d}$ and $\tau \cong S^1 \times D^{d-2}$.

Using this notation, we can now attach two 4-dimensional 2–handles onto $H^3 \times \mathbb{I}$ and precisely specify the attaching maps. We then obtain a topological ball

$$\mathcal{B}_I^4 := H^3 \times \mathbb{I} \bigcup_{\tau_{blue}} H_{blue}^{2,4} \bigcup_{\tau_{red}} H_{red}^{2,4}.$$

By Lemmas 3.1 and 3.2, we can now take $AK_I(r) = \partial(\mathcal{B}_I^4 \times \mathbb{I})$ to obtain a 4-sphere. This completes our first construction type.

We now describe two other construction types. For our Type 2 and Type 3 construction, we will glue in the two 2-handles in dimension 5. To understand the reason we choose to do this, we begin with the following Corollary, which result from definitions of products and quotients of topological spaces from basic topology, see, e.g., [56].

COROLLARY 3.4

Let M^4 be a PL 4-manifold with boundary and let $\tau \times \{0\} \subset \partial M^4$ where $\tau \cong S^1 \times D^2$. Let $H^{2,4}$ be a 4-dimensional 2-handle such that $\tau \times \{0\} \subset \partial H^{2,4}$. Then

$$\left(M^4 \bigcup_{\tau} H^{2,4}\right) \times \mathbb{I} \cong (M^4 \times \mathbb{I}) \bigcup_{\tau \times \mathbb{I}} H^{2,5},$$

where $H^{2,5}$ is a 5-ball such that $(\tau \times \mathbb{I}) \times \{0\} \subset \partial H^{2,5}$.

COROLLARY 3.5

$$AK_I(r) \cong \partial \left((H^3 \times \mathbb{I}^2) \bigcup_{\tau_{blue} \times \mathbb{I}} H_{blue}^{2,5} \bigcup_{\tau_{red} \times \mathbb{I}} H_{red}^{2,5} \right) =: AK_{II}(r)$$

Proof. By our construction of H^3 , we know that $\tau_{blue} \cong S^1 \times D^2$ and also $\tau_{blue} \subset H^3$. Then $\tau_{blue} \times \mathbb{I}^2 \subset H^3 \times \mathbb{I}^2$ and, in particular, we know that $(\tau_{blue} \times \mathbb{I}) \times \{0\} \subset \partial(H^3 \times \mathbb{I}^2)$. The same argument applies for τ_{red} . We can construct two 5-balls that have $\tau_{blue} \times \mathbb{I}$ and $\tau_{red} \times \mathbb{I}$ on their boundaries. Then apply Corollary 3.4. \square

This completes our Type 2 construction. For the final Type 3 construction, we start with one very simple idea: the 2-simplex $\Delta^2 \cong \mathbb{I}^2$.

$$\triangle \cong \square$$

This leads us to the following Corollary and Type 3 Construction.

COROLLARY 3.6

$$AK_{II}(r) \cong \partial \left((H^3 \times \Delta^2) \bigcup_{\tau_{blue} \times \mathbb{I}} H_{blue}^{2,5} \bigcup_{\tau_{red} \times \mathbb{I}} H_{red}^{2,5} \right) =: AK_{III}(r)$$

Putting together Corollaries 3.5, 3.6 we obtain:

COROLLARY 3.7

$$AK_I(r) \cong AK_{II}(r) \cong AK_{III}(r)$$

3.2.3 Procedure for Three Construction Types

Here are step-by-step instructions for constructing triangulated PL 4-spheres $AK_I(r), AK_{II}(r), AK_{III}(r)$, for $r \geq 3$, from the group

$$G_{AK(r)} = \langle x, y \mid \varphi_{blue} : xyx = yxy, \varphi_{red} : x^r = y^{r-1} \rangle.$$

Procedure for Constructing H^3

Step 1: Imagine an empty 3-dimensional space, which when filled will have a 3-ball H^0 with two 1-handles H_x^1, H_y^1 for the generators x, y . (The 1-handles are labeled and oriented.)

$$H^3 = H^0 \cup H_x^1 \cup H_y^1.$$

Step 2: Place disjoint (unlinked, unknotted) curves $\varrho_{blue}, \varrho_{red}$ in the empty H^3 according to the relators $\varphi_{blue}, \varphi_{red}$.

Step 3: Thicken $\varrho_{blue}, \varrho_{red}$ to solid tori τ_{blue}, τ_{red} .

- Use chains of triangular prisms to decompose the tori.
- Triangulate the triangular prisms.

Step 4: Fill up the space around the tori with tetrahedra to obtain the filled H^3 .

- First add rectangular buffers in the 1-handles H_x^1, H_y^1 .
- Triangulate them.
- Then fill H^0 .

Procedure for Type 1 Construction

Step I.1: Take $H^3 \times \mathbb{I}$.

Step I.2: Glue in 4-dimensional 2-handles $H_{blue}^{2,4}, H_{red}^{2,4}$. (See Notation 3.3.)

$$\mathcal{B}_I^4 := H^3 \times \mathbb{I} \cup_{\tau_{blue}} H_{blue}^{2,4} \cup_{\tau_{red}} H_{red}^{2,4}.$$

Step I.3: Take $\mathcal{B}_I^4 \times \mathbb{I}$.

Step I.4: Take the boundary to produce $AK_I(r) = \partial(\mathcal{B}_I^4 \times \mathbb{I})$.

Procedure for Type 2 Construction

Step II.1: Take $H^3 \times \mathbb{I}^2$.

Step II.2: Glue in 5-dimensional 2–handles $H_{blue}^{2,5}, H_{red}^{2,5}$.

$$\mathcal{B}_{II}^5 := H^3 \times \mathbb{I}^2 \bigcup_{\tau_{blue} \times \mathbb{I}} H_{blue}^{2,5} \bigcup_{\tau_{red} \times \mathbb{I}} H_{red}^{2,5}.$$

Step II.3: Take the boundary to produce $AK_{II}(r) = \partial \mathcal{B}_{II}^5$.

Procedure for Type 3 Construction

Step III.1: Take $H^3 \times \Delta^2$.

Step III.2: Glue in 5-dimensional 2–handles $H_{blue}^{2,5}, H_{red}^{2,5}$.

$$\mathcal{B}_{III}^5 := H^3 \times \Delta^2 \bigcup_{\tau_{blue} \times \mathbb{I}} H_{blue}^{2,5} \bigcup_{\tau_{red} \times \mathbb{I}} H_{red}^{2,5}.$$

Step III.3: Take the boundary to produce $AK_{III}(r) = \partial \mathcal{B}_{III}^5$.

3.2.4 Detailed Construction

Step 1: Setting up H^3

The space H^3 , see Figure 3.2, consists of a 3-ball H^0 (in the middle) and attached to it are two 3-dimensional 1–handles H_x^1 (on top) and H_y^1 (on the bottom) corresponding to the two generators x and y , respectively.

Next we want to choose two curves that represent the two relators inside this space.

Step 2: Planning the Curves

We have the freedom to choose any embedding for the curves $\varrho_{blue}, \varrho_{red}$ as long as they run over the two 1–handles H_x^1 and H_y^1 in a way that corresponds to the relators $\varphi_{blue}, \varphi_{red}$. Figure 3.3 is a sketch for our choice of the curves ϱ_{blue} and ϱ_{red} corresponding to the two relators for the Akbulut–Kirby sphere with $r = 5$. We choose this particular arrangement of the




Figure 3.2: The handlebody $H^3 := H^0 \bigcup_{H_x^1} \bigcup_{H_y^1}$.




Figure 3.3: The two curves ϱ_{blue} and ϱ_{red} corresponding to the relators $xyx = yxy$ (blue) and $x^5 = y^4$ (red).

curves to help us achieve certain steps in Step 4, which we explain in detail below. We call the curve corresponding to the first relator, $xyx = yxy$, the blue curve and the second curve corresponding to the relator, $x^5 = y^4$, the red curve.

To understand Figure 3.3, let us begin with the blue curve ϱ_{blue} corresponding to the relator $xyx = yxy$. We read this as $xyxy^{-1}x^{-1}y^{-1} = e$, that is, $\varphi_{blue} = xyxy^{-1}x^{-1}y^{-1}$. Remember that H^3 has two 1-handles: one of which corresponds to the generator x and the other to y . These handles are oriented as indicated by the black arrows. To see ϱ_{blue} , begin at the blue dot in Figure 3.3. Similarly for the relator φ_{red} , which is x^5y^{-4} , begin at the red dot to find ϱ_{red} .

The crossings indicate which sections of the curves run over or under. This will matter in Step 4.

Step 3: Thicken the Curves to Solid Tori

In this step, we finally start producing simplices. The simplices of the simplicial complex will be recorded as a collection of facets. The facets are now 3-dimensional so they are sets of four vertices as indicated by their vertex labels.

We thicken the $\varrho_{blue}, \varrho_{red}$ curves to solid tori τ_{blue}, τ_{red} built up of blocks of triangular prisms. We use triangular prisms for the following reasons.




Figure 3.4: The blue curve from Figure 3.3 is thickened to a chain of triangular prisms.

Firstly, a thickened curve needs to be a tubular neighborhood of S^1 . Secondly, a prism is $P \times \mathbb{I}$ for some polygon P . The cheapest polygon—in terms of memory, that means using the fewest number of vertices—is a triangle. Finally, a triangular prism is a $(2\text{-simplex}) \times (1\text{-simplex})$. There is a systematic way to breakdown such products into tetrahedra using a method called the product triangulation, see [14] and references therein. See Section 3.2.5, Figure 3.12 for our implementation of the product triangulation.

One convenient side effect of this method is that on any rectangular face (along the sides) of a triangular prism, the diagonal on that rectangular face will always include the lowest label number. This is an important feature that we will revisit in the following steps. From now on, whenever we see a rectangular face, we will always choose to take the diagonal that includes the lowest label number. If later we, for example, have neighboring triangular prisms that share a rectangular face, we want the triangulation of those prisms to predictably and consistently find the same diagonal on that shared face. In other words, we want to avoid having neighboring rectangular faces that have two (different) diagonals.

In Figure 3.4, we have the prisms for τ_{blue} for the Akbulut–Kirby sphere with $r = 5$. Notice that the vertex labels repeat at the ends as they are identified to form a solid torus. We will cut up the blue curve ϱ_{blue} in Figure 3.3 into 12 strands, exactly where the curves run into a different section of the handlebody H^3 (recall Figure 3.2). There are three strands on the handle H_x^1 , three on the handle H_y^1 , and six inside H^0 . One triangular prism is used for each of these strands and we let the pointy side always point upwards. We use three vertices every time ϱ_{blue} crosses a black horizontal line in Figure 3.3. So $12 \times 3 = 36$ vertices are used to build the blue solid torus τ_{blue} in Figure 3.4. We use $(2r + 2(r - 1)) \times 3 = 12r - 6$ vertices to build the red solid torus τ_{red} . The product triangulation method will triangulate a triangular prism into three tetrahedra. So we have three tetrahedra for each triangular prism. The vertex labeling scheme we choose here (the highest vertex label of each triangular face is on the pointy top of the face) comes in handy in the next step.

Step 4: Fill H^3

Step 4 will be broken up into the three sections of H^3 . The procedures for filling H_x^1 and H_y^1 are identical and will be described first.

Fill the 1–handles H_x^1 and H_y^1 . One reason we chose the particular arrangement for the curves ϱ_x, ϱ_y shown in Figure 3.3 is so that we have the curves running parallel to each other inside of the 1–handles. We want to add buffers between the tori to avoid having any unwanted identifications. By arranging the tori this way, we can fill the 1–handles just by filling the spaces between the triangular prisms with rectangular prisms that run alongside the triangular prisms. To be clear, we will be using exactly the vertices from the triangular prisms of Step 3 to form new rectangular prisms which will lay between certain strands of those triangular prisms.




Figure 3.5: The 1-handles can be filled using rectangular prisms between the triangular prisms to avoid unwanted identifications.

Figure 3.5 gives a front view of the handle H_x^1 as an example. The three blue triangles on the two ends are the parts of the blue solid 3-torus τ_{blue} that are in H_x^1 . The red triangles, respectively, correspond to the parts of the red torus τ_{red} in H_x^1 . The rectangular prisms between the triangular prisms are shown in white. We need $2 + 1 + (r - 1) = r + 2$ rectangular buffer prisms for the handle H_x^1 and $2 + 1 + (r - 2) = r + 1$ rectangular buffer prisms for the handle H_y^1 .

Now we must triangulate these rectangular prisms. Triangulating rectangular prisms is a two-step process.

Procedure

Triangulating rectangular prisms:

- Break down the rectangular prisms into two triangular prisms.
- Use product triangulation method to break down triangular prisms into tetrahedra.

Here we will find that the vertex labeling scheme we chose is quite practical. A nice way to break down a rectangular prism is

$$\begin{aligned} (\text{a rectangle}) \times \mathbb{I} &= (\text{two triangles}) \times \mathbb{I} \\ &= (\text{one triangle} \times \mathbb{I}) \cup (\text{the other triangle} \times \mathbb{I}). \end{aligned}$$

This simple decomposition from a rectangular prism into two triangular prisms is possible because the labeling scheme we chose for our triangular prisms fits well, see Figure 3.6. Recall the vertex labeling from Figure 3.4. The labeling on the triangles of the triangular prisms are so that the largest label is on the “pointy top” of the triangle and the two smaller ones are on the “flat bottom” of the triangle. Also, along the “long end” of the triangular prisms, the labeling goes up predictably by 3, e.g., after vertex 1 is 4, 7, \dots , 34 on τ_{blue} . The rectangles of the rectangular prisms take the vertices of these triangular prisms. Notice that when we split a rectangular prism into two triangular prisms, like in Figure 3.6, the rectangular faces that split into two triangles split consistently. Note that the rectangular face we are talking about now is not one that is along the side of a triangular prism. So we can split these faces into triangles in any way. We will choose

to be consistent with the product triangulation method used to triangulate the triangular prisms and select the diagonal on this face that includes the smallest vertex label. Our labeling scheme always places the larger vertex label (of the triangle of the triangular prism) on “top”. Moreover, the next triangle (on the long end of the triangular prism) is always just ± 3 . If we choose the diagonal that includes the smallest label (on these rectangular faces of the rectangular prisms), we will consistently match up on the other rectangular face of the rectangular prism.

Once the rectangular prisms have been broken up into triangular prisms, we use the product triangulation again to triangulate the two triangular prisms into tetrahedra. (For the implementation, there is one extra consideration, see Section 3.2.5.)

Now the two 1-handles H_x^1 and H_y^1 are filled; they are topological balls. We must do the same for the middle section H^0 to completely fill the handlebody H^3 .

Fill the 0-handle H^0 . Here we will first try to pair up as many parallel strands as we can, just like we did for the 1-handles. Let us zoom in to the middle section representing what will become H^0 , see Figure 3.7.

The triangular prisms in the middle section H^0 are not all parallel as they were along the 1-handles. Some cross each other. Remember that the vertices of the triangular prisms are all on the same level; let us call this the ground floor. The prisms themselves, however, may run over or under other prisms. The arrangement of the curves we chose places most of the prisms on the same floor—the ground floor `level:0`—and only three on other floors. We color the lines in Figure 3.7 to help distinguish the different floors of our ball. The ground floor `level:0` is colored gray; the basement floor `level:-1` is colored pink; and the two upper floors are colored green `level:+1` and orange `level:+2`. Notice that the gray lines on the ground floor do not cross any other gray line. This arrangement will make filling each floor easier.

We start on the ground floor, see also Figure 3.8, `level:0`. First glue in rectangular buffer prisms between the parallel strands of the ground floor. The number of buffers needed between the loops of τ_{red} will depend on r (in Figure 3.7 they are the gray half circles on the top and bottom). The two pairs of parallel strands from τ_{blue} only use one rectangular prism each

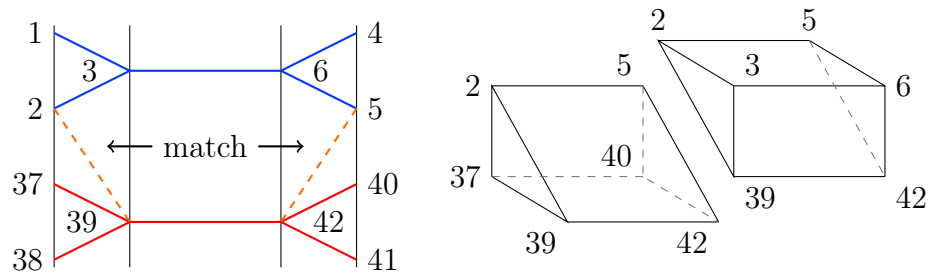


Figure 3.6: Rectangular prisms are placed as buffers between the triangular prisms. They can be nicely decomposed into two triangular prisms.

(in Figure 3.7 they are the two pairs of parallel gray lines on the left and right). In Figure 3.8, `level:0`, the buffers are colored gray. After gluing in the buffers, there will be five solid chunks in `level:0`. We then glue in 2-dimensional membranes (like installing a carpet) on the ground floor between the five solid chunks to close the floor.

For the remaining three floors `level:-1`, `level:+1`, `level:+2` we use a method that we call “filling a cupula” which is sketched in Figure 3.9. Say we have a green disc and a 2-dimensional blue strip whose ends are glued onto two sections on the boundary of the disc, like a picnic basket. We can now trace two circles along the boundary of this picnic basket. Make sure these two circles do not meet. Add a cone over each of the two circles to form a cupula (a pink dome over the green disc). Insert a vertex inside the void of this topological 2-sphere and take the cone over the triangles of that sphere to obtain a 3-ball. The blue strip is a part of the boundary of a single triangular prism and the tets of that prism lie outside of the dome.

The carpet on the ground floor forms parts of the green disc for the `level:-1`-ball and `level:+1`-ball. The cupula of `level:+1`—together with some other triangles from the faces on `level:0` and the `level:+1` triangular prism—is the green disc for `level:+2`. We make sure to trace the two circles without intersecting each other. In Figure 3.8 we draw the two circles on each level in different colors.

We now have the handlebody H^3 : a 3-ball with two solid 1-handles.

Glue in the 2-handles H_{blue}^2 and H_{red}^2

To understand a d -dimensional 2-handle, we describe it first in dimension 3. A 3-dimensional 2-handle is a topological 3-ball $H^{2,3} = D^2 \times D^1$ glued along the part of the boundary that is $S^1 \times D^1$, i.e., $H^{2,3}$ is a solid cylinder. Suppose we have a 3-dimensional donut formed by a chain of triangular prisms which represents the thickened curve along whose boundary the 2-handle will be glued on. Gluing in a 3-dimensional 2-handle onto the boundary of a 3-dimensional solid torus essentially means filling the hole of this donut. We will glue $H^{2,3}$ along one of the (three) flat sides of τ_3 . In Figure 3.10 the (rectangular) faces onto which the 2-handle will be glued are highlighted in green. In this example, we use the side of the triangular prisms that are positioned on the side that “sees” the center of the donut. Note that we could very well have chosen any of the sides as long as we are consistent. We use the inside faces in this demonstration because it is




Figure 3.7: The portion that will be filled to become the 3-ball H^0 . The gray lines are on `level:0`, the pink line is on `level:-1`, the green line is on `level:+1`, and the orange line is on `level:+2`.

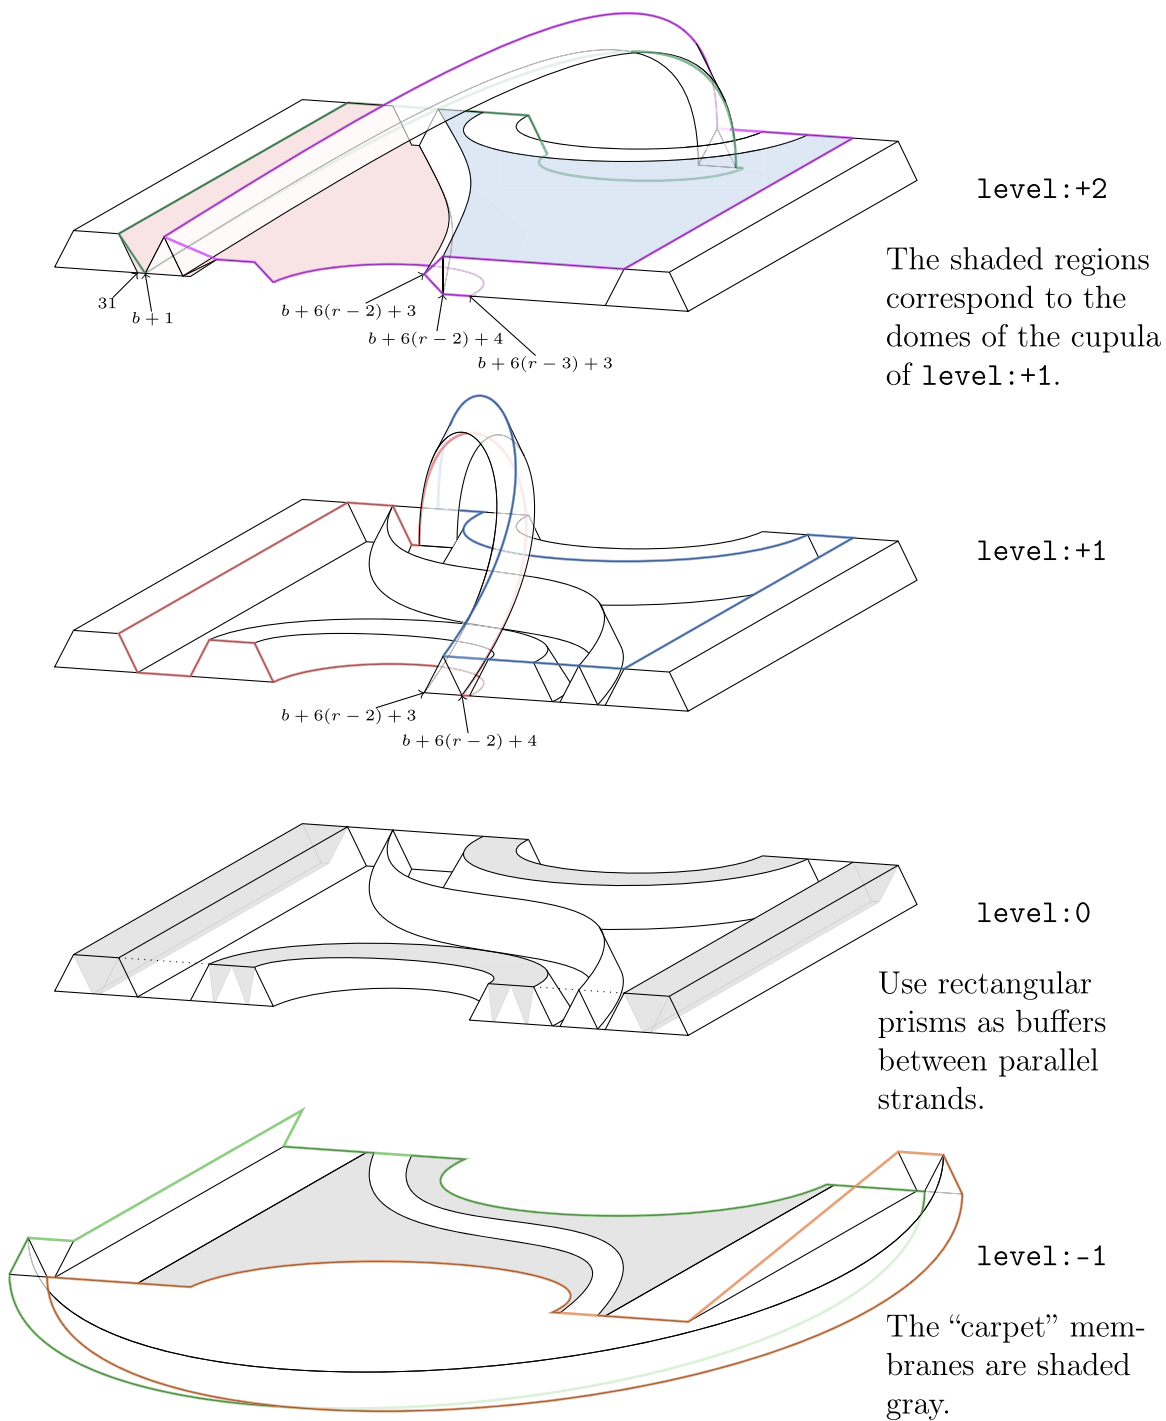


Figure 3.8: The two non-intersecting circles used to form parts of the domes for each of the cupula of level:-1, level:+1, and level:+2 are shown in different colors.




Figure 3.9: Fill each floor by forming a cupula and taking a cone over it.




Figure 3.10: Left: A thickened S^1 formed by a chain of triangular prisms. Right: the rectangular faces onto which we attach a 3-dimensional 2-handle are highlighted in green.

easier to imagine. The donut of triangular prisms will then be triangulated and the rectangular faces will become two triangles each. The solid cylinder representing the 2-handle is then constructed by successively adding cones onto those faces. Figure 3.11 demonstrates how we construct a triangulated 3-dimensional 2-handle to fill the hole in the green cylinder. We begin by setting T_2 which consists of the triangles of $S^1 \times D^1$ (that subdivide the rectangular faces of Figure 3.10), where the 2-handle will be glued on.

Procedure

Construct a triangulated 3-dimensional 2-handle H^2 whose gluing map is already known to be the set of triangles T_2

- Take a cone h_2 over T_2 with apex a_2 .
- Let $T_1 \subset \partial T_2$ be one of the $S^1(\times D^0)$ on the boundary of T_2 . Take a joint cone h_1 over $T_1 * a_2$ with apex a_1 .
- Then $H^2 = h_1 \cup h_2$.

Notice that the lower indices of T_2, T_1 are keeping track of the dimension of the faces contained in the set. In higher dimensions, the set T_{d-1} collects the shared $(d-1)$ -dimensional faces of the part of the boundary of the d -dimensional 2-handles (i.e., $S^1 \times D^{(d-2)}$) and the boundary of the d -dimensional solid tori depicting the thickened curve of the relators (e.g., τ_{blue}, τ_{red}).

The 3-dimensional 2-handle $H^{2,3}$ (glued along a part of the boundary of a 3-dimensional solid donut) is the collection of all the pink (h_2) and orange (h_1) tetrahedra $H^{2,3} = h_2 \cup h_1$ in Figure 3.11.




Figure 3.11: Constructing a 2-handle. Left: the green cylinder from Figure 3.10. Middle: take a cone over the green triangles to form the pink tetrahedra. Right: take the joint cone over the upper boundary of the green cylinder and the apex of the first cone to form the orange tetrahedra.

For the case of dimensions $d > 3$, the procedure is not much more difficult. In fact, the hardest part has already been done in the 3-dimensional example. We start with a 3-dimensional solid torus $\tau_3 \subset \mathbb{H}^3$. We go up to dimension d by successively taking products with the unit interval $H^3 \times \mathbb{I}^{(d-3)}$. So we will always find the (d -dimensional) solid torus $\tau_3 \times \mathbb{I}^{(d-3)} \subset \mathbb{H}^3 \times \mathbb{I}^{(d-3)}$. On the boundary of $H^3 \times \mathbb{I}^{(d-3)}$ are then two copies of the solid torus $\tau_3 \times \mathbb{I}^{(d-4)}$, see also Figure 3.1.

For a simplicial d -complex K , going up in dimension $K \times \mathbb{I}$ means that for each d -simplex σ in K , we make a copy of it and then use the product triangulation to triangulate the prism $\sigma \times \mathbb{I}$, see Section 3.2.5 for an example and details of our implementation. Making a copy means incrementing the vertex label by a constant k . That constant is always the largest vertex label used so far so that we do not introduce unintentional identifications.

Then one can easily differentiate between the two copies of $\tau_3 \times \mathbb{I}^{d-4}$ on the boundary of $H^3 \times \mathbb{I}^{d-3}$ because one is just the same constant k more than the other on every vertex label. We will always choose the lexicographically smaller copy, which we can think of as corresponding to $\tau_3 \times \mathbb{I}^{d-4} \times \{0\}$.

Procedure

Construct a triangulated d -dimensional 2-handle $H^{2,d}$ to be glued onto $H^3 \times \mathbb{I}^{(d-3)}$ along a known thickened S^1 on both of their boundaries.

- Set $T_{d-1} = \tau_3 \times \mathbb{I}^{d-4} \times \{0\}$.
- Take a cone h_{d-1} over T_{d-1} with apex a_{d-1} .
- For $3 \leq i < d-1$, set $T_i = \tau_3 \times \mathbb{I}^{(i-3)} \times \{0\}^{(d-i)} \subset \partial T_{i+1}$.
- Take a joint cone h_i over $T_i * a_{d-1} * \dots * a_{i+1}$ with apex a_i .
- For T_2, T_1 and h_2, h_1 , see preceding Procedure.
- $H^{2,d} = \cup_{i=1}^{d-1} h_i$.

Note that the (5-dimensional) 2-handles for both the Type 2 construction and Type 3 construction are the same (up to shifting the apex vertex labels).

Take the boundary

Finally, we take the boundary to obtain our 4-spheres. The boundary of a d -ball can be found by looking for all of its $(d - 1)$ -faces that are contained in only one facet. All the other $(d - 1)$ -faces must be contained in exactly 2 facets.

Extra Step

There are a few more tetrahedra we pick up in our construction of H^3 that are not mentioned in this section. Those tets are topologically insignificant, but technically necessary. We include the reasons and procedure for producing these last few tetrahedra in the next section as it requires some details about our actual implementation to understand. The description of our construction serves as proof that if the construction is carried out as we described, the triangulations we end up with are indeed the triangulations of the Akbulut–Kirby spheres. For our computer implementation, we built tools to ensure we indeed produce what we intend to build, for example, automated checks which test for typos.

3.2.5 Implementation

Our motto, if you will, for designing the construction was to always keep things simple—choose curves that made filling H^0 easy, use rectangular prisms between parallel triangular prisms, try to use as few building blocks as possible. As a result, in a few areas, we were met with challenges we may not have encountered had we chosen a construction that, for example, used more vertices or subdivisions. The prescient warnings in the procedure described in Section 3.2.3 are a result of having labored over dealing with many of those challenges. Below we describe some custom programs we wrote to assist in our construction of the triangulation which can help to identify problem areas. We describe each tool and give examples of the possible user errors they can help to identify.

Product triangulation

The standard method for finding the tetrahedra of the product triangulation $(\text{triangle}) \times \mathbb{I}$ is shown in Figure 3.12. First we write the label numbers of the triangle $T = [1\ 2\ 3]$ along the bottom of a (3×2) -lattice. The copy of T is labeled by incrementing each of the vertex labels of T by the largest vertex in T , which is 3; they are written in the lattice above its corresponding vertex label, in this case $4 = 1 + 3$ is above 1, $5 = 2 + 3$ is above 2 and $6 = 3 + 3$ is above 3.

In our program, we input the vertex labels as an ordered list $[1\ 2\ 3\ 4\ 5\ 6]$ and the program will convert this into the lattice we just described. The tetrahedra triangulating this rectangular prism can then be found by walking along the lattice from the bottom-left corner (1) to the top-right corner (6) taking only right or left (monotone) moves. Finally we have three tetrahedra such that each rectangular face of the prism has a diagonal that

contains the smallest vertex label. For example, on the rectangular face $[1\ 2\ 4\ 5]$, the diagonal is $[1\ 5]$, which includes the smallest vertex 1.

The product triangulation method works well when taking a product between some simplicial complex K and, say, a single p -simplex. The vertex labels, in such cases, are given anew to the p copies of K , usually by incrementing every vertex label by some constant greater than or equal to the maximum label number used in K to avoid producing accidental identifications. This is exactly what we do to triangulate the triangular prisms τ_{blue}, τ_{red} , and whenever we go up in dimension by taking the product of a complex and \mathbb{I} . However, the rectangular prisms that we break up into triangular prisms (e.g., to fill H_x^1, H_y^1) are a different matter. The triangular prisms we want to triangulate (the two we obtain per rectangular prism) start with a prism whose vertex labels are already chosen—which may or may not be ordered in a predictable way—and we ask the product triangulation method to deliver a triangulation that complies with our chosen rule that for any rectangular face its diagonal contain the smallest vertex label.

Compare the example in Figure 3.13 with the earlier example in Figure 3.12. This new triangular prism has vertex labels that are already given as $[1\ 5\ 6]$ and $[2\ 3\ 4]$ on the triangular faces. Recall that the product triangulation is implemented as a program whose input is an ordered list of the vertices of the prism.

Notice in our example in Figure 3.13 that the order in which the vertex labels are inserted into the lattice of the product triangulation method produces different tetrahedra. In particular, the top input produces the wrong diagonals; but, the vertex labels of the triangles are in lexicographic order, which is perhaps the most ordinary order in which one would input this type of information. Again, the most ordinary input produces the *wrong* triangulation. It takes the diagonal $[4\ 5]$ in the rectangular face $[3\ 4\ 5\ 6]$; it should be $[3\ 6]$. To make sure we end up with the correct triangulation, we have to check that every (2×2) sub-lattice contains the minimum vertex label of that sub-lattice on its diagonal. We search for such a configuration over each permutation of the vertex labels, see Figure 3.14.

In every case where we need to take this extra consideration (i.e., any $K \times \mathbb{I}$), this process will find a working solution. In higher-dimensions, this method does not work, e.g., $K \times (2\text{-simplex})$ whose vertex labels are jumbled.

is_Pure

The complex is always recorded as a collection of its facets. Keeping track of all of the faces is memory heavy and cumbersome. Checking to see if every facet in the complex is of the same dimension is incredibly fast, but still useful. This test will fail if, for example, one of the facets repeats a vertex, $[1\ 2\ 3\ 3]$. Finding where in the code that bad facet was produced can help to remove this error. We run this check automatically before running any other check so that we can catch and deal with such simple errors early.




Figure 3.12: Use product triangulation method to break down triangular prisms that guarantees that the diagonal on rectangular faces includes the vertex having the smallest index.




Figure 3.13: Applying the product triangulation method directly is not enough. We take an extra step to ensure we find the correct diagonals on every rectangular face.




Figure 3.14: Test various permutations of the vertex labels until we find one such that the diagonals for every (2×2) sub-lattice contains the minimum vertex label of that sub-lattice.

The Euler characteristic

The Euler characteristic is also fast and easy to compute. The many parts of the construction are saved as separate complexes which we union at the end to obtain the final product. For example,

$$H^3 = \bigcup \{ \tau_{blue}, \tau_{red}, H_x^1\text{-filler}, H_y^1\text{-filler}, \\ \text{level:-1-ball}, \text{level:0-ball}, \text{level:+1-ball}, \text{level:+2-ball}, \\ \text{extra tets}^* \}.$$

We can determine in which subcomplex the problem may have occurred by checking the Euler characteristic of any subset of these 9 subcomplexes. (* We explain the “extra tets” below.) Figuring out what the correct Euler characteristic should be for any subset of these parts is a fun exercise.

If the Euler characteristic is not what you expected, it is a sure indicator that something has gone wrong. The correct Euler characteristic, however, does not guarantee that everything is fine. Further narrowing down the source of the problem requires the use of more checks described below.

The ridge check

Another easy check is to see whether ridges, or co-dim-1-faces, are contained in at most 2 facets. If the check fails, we output the ridge on which the check failed and the list of the facets it is a face of. This check is usually an effective way to fix the cases when the Euler characteristic check fails. Computationally, the ridge check requires many more steps and will take more time to compute than the Euler characteristic so we usually use the Euler characteristic test to narrow down the possible source of the problem before running the ridge check.

For spheres this condition is strengthened to having ridges contained in *exactly* 2 facets. In fact, this strengthened check can even be performed before the final step. The boundary of any ball (like a vertex star) should also have exactly 2 facets per ridge.

The manifold check

One very important check is to verify the “manifoldness” of the complex. We want all the vertex links to be PL spheres or balls. For a simplicial complex to be a closed PL manifold, the vertex links should be PL spheres. There are different ways to test for this and, in particular, a combination of heuristic algorithms can be used to solve this decision problem as explained in Section 2. But we are running these checks to identify the areas of the construction that produced some sort of badness. Just a “yes” or “no” does not suffice. Moreover, in our construction we are in dimension 3 most of the time which makes certain tests easier; for example, recognizing the 1-sphere (a circle) and the 2-sphere is easy. And while we are in the 3-dimensional part of the construction, we always (should) have a manifold with boundary. In fact every vertex in our construction is exactly on the boundary so vertex links are not spheres, they should be 2-balls. We also checked whether the boundaries of vertex stars are 2-spheres.




Figure 3.15: The dashed lines represent the diagonals chosen to split the rectangular prisms into two triangular prisms. We add one tetrahedra $[31 \ 33 \ b+1 \ b+2]$ so that we can use one *wrong* diagonal shown in green.

Manifold check

Run for every vertex v in a triangulated 3-manifold K with boundary.

- Ridge check on $\text{link}(v, K)$.
- Check Euler characteristic: $\chi(\text{link}(v, K)) = 1$.
- Check that boundary $\partial\text{link}(v, K)$ is a single circle.
- Ridge check on $\partial\text{star}(v, K)$.
- Check Euler characteristic: $\chi(\partial\text{star}(v, K)) = 2$.
- Check that for all vertices $w \in \partial\text{star}(v, K)$, $\text{link}(w, \partial\text{star}(v, K))$ is a single circle .

The “extra tets”

The tools we have described so far were helpful for catching small errors, like typos. But they were also instrumental in helping us making certain design choices. The following tets were added in addition to the simplices described in the construction in Section 3.2.3.

Recall that for producing the cupula, we need to first find two non-intersecting circles which will be coned over to form the dome for the cupula. To produce the cupula for `level:+1`, we wanted to use a “forbidden” diagonal for one of the circles, see Figure 3.15. Notice the rectangular




Figure 3.16: To remedy the pinched point at $b + 1$ we take a cone over the gray triangles.

face $[31 \ 33 \ b + 1 \ b + 2]$. Here we write b because its value depends on r . Specifically it is $b = 37 + 3 \cdot 2r + 1 = 6r + 38$.

The rule we set early on for our design is to take the diagonals on rectangular faces which include the smallest vertex label number. Here, this would mean the diagonal $[31 \ b + 2]$, as indicated by the gray dashed line (on all rectangular faces). For the circle to form one of the dome pieces, for the cupula of `level : +1` we wanted to use the edge $[33 \ b + 1]$. To be able to use this edge, we add one more facet. Exactly the facet $[31 \ 33 \ b + 1 \ b + 2]$. We highlight the problem diagonal in Figure 3.15 in green.

At the very end of the construction of H^3 , we noticed that the manifold check failed despite all other parts of the construction having been implemented correctly. It was caused by a pinched point. A pinched point is a vertex whose link has more than one connected component. Coincidentally, the problem occurred exactly in the same area as the extra tet we just added. Running the manifold check revealed that the vertex $[b + 1]$, where we had place the “wrong” diagonal was our culprit.

In particular, the $\text{link}(b + 1, \partial \text{star}(b + 1, H^3))$ turned out to be two circles, not one; the $\text{star}(b + 1, H^3)$ was a wedge of two balls. To fix the pinched point, we want to “get rid” of one of the extra circles. We do this by taking a cone over the triangles $[b + 1] * (\text{one circle})$. While the combinatorial description of how to repair pinched points is simple, the geometric picture of what we did is also quite nice.

The pinched point is aptly named as you can see in Figure 3.16. Of all the rectangular faces between the triangles in Figure 3.15, only the triangle $[33 \ b + 1 \ b + 2]$ is not shared by two tets. All other rectangular faces are shared on one side by the tets from the triangulation of the rectangular prism (on the side of the 1-handles) or by some cupula (on the side of the H^0). This final cone over $b + 1$ will essentially slap on some play-doh to fill the ugly gap around the pinched point. These extra tets complete the construction of H^3 .

3.3 Results & Experiments

THEOREM 3.8

The Akbulut–Kirby spheres—4-spheres with handlebody decompositions described by the group $G_{AK}(r)$ —have triangulations for $r \geq 3$ with face vectors:

$$\begin{aligned}
 f(AK_I(r)) &= (48r + 180, 744r + 2766, 2496r + 9284, 3000r + 11160, 1200r + 4464) \\
 f(AK_{II}(r)) &= (48r + 176, 720r + 2694, 2400r + 9036, 2880r + 10860, 1152r + 4344) \\
 f(AK_{III}(r)) &= (36r + 134, 516r + 1944, 1704r + 6456, 2040r + 7740, 816r + 3096)
 \end{aligned}$$

Vertices

To build the H^0 part of H^3 , we need

- the $3 \cdot 12 = 36$ vertices of τ_{blue} ;
- the $3 \cdot (2r + 2 \cdot (r - 1)) = 12r - 6$ vertices of τ_{red} ;
- the 2 vertices of the `level:0` carpet;
- the 2 vertices to build the dome for the cupula of `level:-1` plus 1 vertex taking its cone;
- the 3 vertices for the `level:+1` cupula;
- the 3 vertices for the `level:+2` cupula; and
- the 1 vertex used to produce the cone of the “extra tets”;

which gives us a total of $12r + 42$ vertices. The 1-handles H_x^1 and H_y^1 contribute no vertices. An i -dimensional 2-handle requires $i - 1$ vertices and we build two 2-handles in our construction, i.e., H_{blue}^2 and H_{red}^2 . Whenever we go up in dimension, we make a copy of whatever complex we have at the time and that copy will use exactly as many vertices as its original. For the Type 3 construction we go up by two dimensions in one step by crossing with a triangle, so instead of producing a single copy we produce two copies (plus the original makes three).

The vertices for the three construction types can be counted as follows:

$$AK_I(r): \quad \begin{array}{c} \nearrow \\ \text{go up to } d=5 \end{array} 2 \cdot \left(\begin{array}{c} \nearrow \\ \text{go up to } d=4 \end{array} 2 \cdot (12r + 42) + 2 \cdot 3 \right) = 48r + 180$$

vertices of H^0 4-dim 2-handles

$$AK_{II}(r): \quad \begin{array}{c} \nearrow \\ \text{go up to } d=5 \end{array} 2 \cdot \left(\begin{array}{c} \nearrow \\ \text{go up to } d=4 \end{array} 2 \cdot (12r + 42) \right) + 2 \cdot 4 = 48r + 176$$

vertices of H^0 5-dim 2-handles

$$AK_{III}(r): \quad \begin{array}{c} \nearrow \\ \text{cross by } \Delta \end{array} 3 \cdot (12r + 42) + 2 \cdot 4 = 36r + 134$$

vertices of H^0 5-dim 2-handles

We built $AK_I(r), AK_{II}(r), AK_{III}(r)$ for $r = 3, \dots, 20$ for all three construction types. We could have easily built more spheres for many more r 's but the computation time for the experiments increase as the size of the complex go up. Already at $r = 10$, the f -vector for the Type 1 construction is $f(AK_I(10)) = (660, 10206, 34244, 41160, 16464)$, which is a total of 102734 faces and therefore at least as many nodes in its Hasse diagram—not to mention its 370104 edges (corresponding to the incidence relations of the co-dim-1 faces)⁶. The bigger the Hasse diagram, the longer the computations discussed in Chapter 2 will take.

3.3.1 Homology


For completeness, we include here our results from the homology computation using `polymake`. The homology was found to be spherical $H_* = (\mathbb{Z}, 0, 0, 0, \mathbb{Z})$ for all three types for all $r = 3, \dots, 20$.

Our triangulations have also been used as test examples for a heuristic algorithm that uses an iterative discrete Morse reduction/co-reduction scheme as a preprocessor for homology computation. Our spheres were the first known examples in which the algorithm required more than one iteration of co-/reduction to reach the optimum. This work with Konstantin Mischaikow, Vidit Nanda, and Frank H. Lutz is still ongoing [48].


3.3.2 Random_Discrete_Morse

Recall that there are triangulations of contractible manifolds that are noncollapsible, see 2.2. A consequence of this fact is that a PL sphere


⁶In `polymake`, the Hasse diagram also includes a node for the entire complex and a node for the empty set. The actual numbers for the nodes and edges in the Hasse diagram are 102736 and 387228, respectively.



Using strategy random-random.



Using strategy random-lex-first.



Using strategy random-lex-last.

Figure 3.17

minus a facet is contractible, but may not be collapsible. That is, there are triangulations of PL spheres that do not admit a discrete Morse vector of $(1, 0, 0, 0, 1)$. In every experiment we attempted using `Random_Discrete_Morse`, we have not found a spherical Morse vector. We do not know whether this is due to the limitations of the algorithm or just a property of our triangulations. In Figure 3.17, we compare the average number of critical cells found in 10000 rounds using all three strategies.

In Section 2.2 we discussed that the `Random_Discrete_Morse` client in `polymake` may encounter contractible but noncollapsible complexes when run on complexes having many facets, even in dimension 3. This may be the reason the Type 3 construction finds fewer critical cells on average compared to the other two, and Type 2 finds fewer critical cells than Type 1. However, the data here suggests that this phenomenon may not be the only reason the algorithm does not find perfect discrete Morse vectors since, for example, the number of facets in $AK_I(5)$ is 10464 and in $AK_{III}(9)$ is 10440 (or the total face counts are 65294 and 65378, respectively). But the average number of critical cells that $AK_{III}(9)$ finds is greater than that of $AK_I(5)$ as Table 3.1 shows.

	$AK_I(5)$	$AK_{III}(9)$
# of facets	10464	10440
random-random	36.42	44.21
random-lex-first	20.73	24.17
random-lex-last	19.64	21.57

Table 3.1: The average number of critical cells found on $AK_{III}(9)$ is greater than that of $AK_I(5)$ even though it has about the same number of facets.

3.3.3 bistellar_simplification

The `bistellar_simplification` client implemented in `polymake` uses a simulated annealing technique that prefers moves that lower the f -vector, which is a reasonable strategy, but not necessarily the best. To be more precise, the algorithm attempts to reduce the g -vector the best it can. The g -vector for us is defined as

$$g_1 = f_0 - 6 \quad \text{and} \quad g_2 = (5 \cdot g_1) - 15 - f_1.$$

Since the g -vector and f -vector are linearly correlated, minimizing one will minimize the other, see [14] and references therein.

Attempting to minimize the g -vector is a reasonable strategy for several reasons:

- bistellar i – $(4-i)$ -moves change g_{i+1} by ± 1 so this is a cheap and easy value to track,
- stacked polytopes have zero g -vectors and stacked polytopes are PL-homeomorphic to spheres [71],
- only the face numbers f_0, f_1 are needed to determine the remaining face numbers of the f -vector of a triangulated 4-manifold without boundary by linear relations [43]; by the definition of the g -vector, we find that also the entire f -vector is determined by the g -vector.

The biggest problem with this strategy is that for any given g -vector, there may be an enormous number of different triangulations. Already at dimension 2, there are over 12 billion non-isomorphic triangulations of a 2-manifold with only 12 vertices [70]. Moreover, at any given triangulation, there may be thousands or even hundreds of thousands of available moves to choose from.

It can happen (quite often!) that the triangulation we have is a local minimum in terms of the g -vector, but not a global minimum. Say we have a triangulation $T_{\mathbf{m}}$ with g -vector $(\mathbf{m}_1, \mathbf{m}_2)$ and our goal is to show that $T_{\mathbf{m}} \stackrel{PL}{\simeq} T_{\mathbf{n}}$ where $T_{\mathbf{n}}$ has g -vector $(\mathbf{n}_1, \mathbf{n}_2)$ with $\mathbf{n}_1 \leq \mathbf{m}_1$ and $\mathbf{n}_2 \leq \mathbf{m}_2$, see Figure 3.18. To jiggle out of the local minimum at $T_{\mathbf{m}}$, we will have to make moves that increase the g -vector until we reach some triangulation that can find a nice downward path to $T_{\mathbf{n}}$.




Figure 3.18

However, the g -vector is so cheap to measure and track, we can try millions of moves in just a few minutes. If in fact $T_m \stackrel{PL}{\simeq} T_n$, then there must exist some sequence of bistellar moves between T_m and T_n . And, as we mentioned earlier, the `bistellar_simplification` algorithm has done a reasonably good job at finding those downward paths. That is, until we tested the algorithm on our triangulations of the Akbulut–Kirby spheres. To be fair, the algorithm was able to show that the triangulations for $r = 3$ of all three construction types are a PL sphere. We ran all of our “random” tests on several seed numbers, for a discussion on pseudorandom numbers, see [76]. For $r = 3$, `bistellar_simplification` flipped down our triangulations all the way 16 times out of 20 different seed numbers for Type 1, 17/20 for Type 2, and 18/20 for Type 3, which gives a computer-assisted verification of that these examples are indeed standard.

OBSERVATION 3.9

The triangulations $AK_I(3), AK_{II}(3), AK_{III}(3)$ of the Akbulut–Kirby spheres for $r = 3$ are bistellarly equivalent to the boundary of a 5-simplex, and therefore are indeed PL 4-spheres.

Figure 3.19 and Figure 3.20 displays our results from the bistellar test for $r = 4, \dots, 10$ for all three types. The `bistellar_simplification` client returned the smallest triangulation in its search. The graphs in Figure 3.19 and Figure 3.20 show the data obtained from the g -vectors of the smallest triangulations found over 20 different seed numbers.

In Figure 3.19, we see that there is very little difference between the three construction types. Looking closer at the smallest of and the average over the 20 g -vectors obtained from the smallest triangulation found by `bistellar_simplification` in Figure 3.20, we see that the Type 3 construction seems to behave similarly to the other two types. In the bottom graph in Figure 3.20, we also include the average number of moves the algorithm performed to obtain the data shown in Figure 3.19 and Figure 3.20. We display the average number of moves as opposed to the runtime since the tests were run on multiple machines and the runtime will vary depending on the speed of the processor on that machine. The

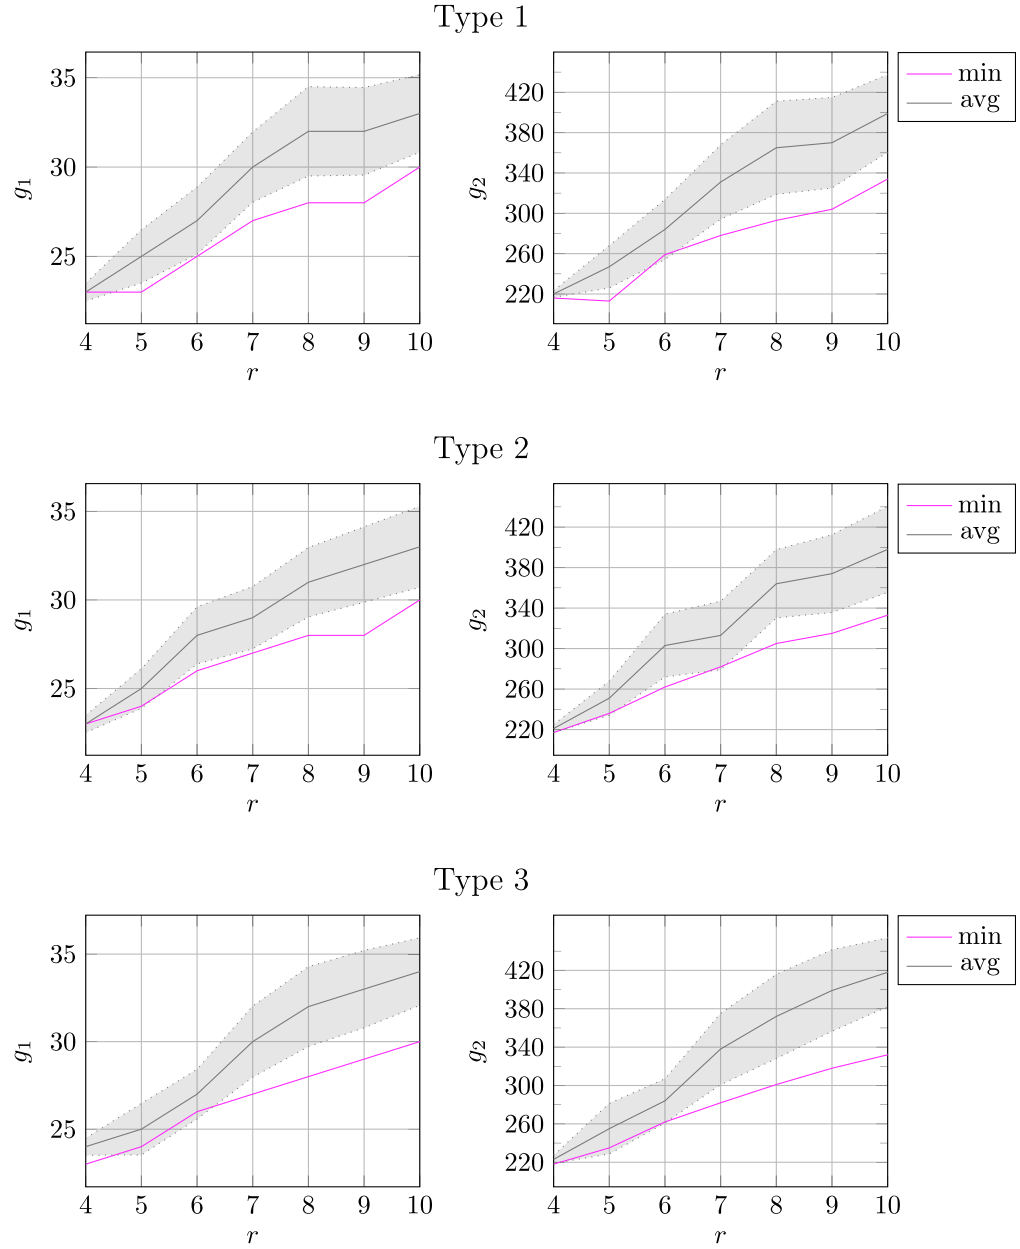


Figure 3.19: A summary of the g -vectors of the smallest triangulations found using 20 different seed numbers for all three types for $r = 4, \dots, 10$. The gray area spans the standard deviation from the average. Table 3.2 lists all the minimum (pink) values found over the 20 runs.

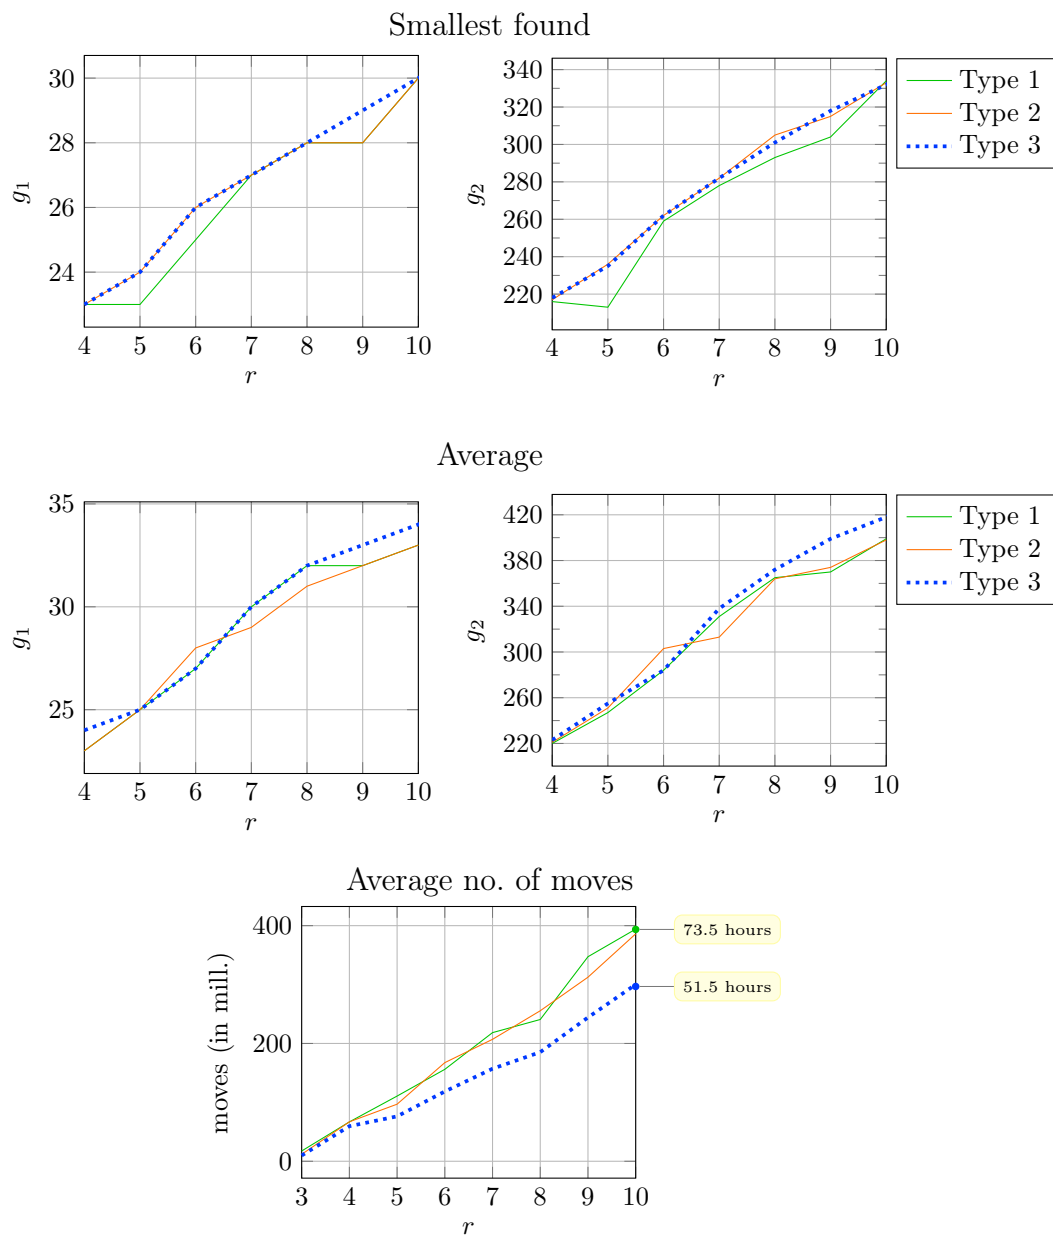


Figure 3.20: Comparing the different three different construction types.

	4	5	6	7	8	9	10	11	12	13	14	15	16	17
Type 1	23	23	25	27	28	28	30							
Type 2	23	24	26	27	28	28	30							
Type 3	23	24	26	27	28	29	30	31	31	32	33	33	34	32

Table 3.2: The number of vertices of the triangulation having the smallest g -vector found using `bistellar_simplification` with 20 different seed numbers. Type 1 and Type 2 took very long to run so we only ran Type 3 for $r > 10$.

number of moves is a more fair way to get a feel for how much difficulty the algorithm had in finding the smallest triangulation.

As the number of faces grows, so does the number of moves⁷ and, therefore, also the runtime. The amount of memory and runtime the Type 1 and Type 2 constructions use compared to Type 3 seems to be a heavy cost seeing that the smallest found by the Type 3 construction is about the same as the other two. Note that in the bottom graph in Figure 3.20, we also indicated the average runtime (wallclock time) for $AK_I(10)$ and $AK_{III}(10)$ for comparison. That is, on average, it took about 73.5 hours to complete one run of `bistellar_simplification` on $AK_I(10)$, a total of about 1470.5 hours (~ 60 days) to run all 20 seed numbers. For $AK_{III}(10)$, the average runtime was about 51.5 hours taking a total of about 1029.5 hours (~ 40 days) for 20 runs. It is worth mentioning, however, that the number of moves and the size of the triangulation obtained are not correlated.

The middle table in Figure 3.20 shows that the Type 3 construction is performing slightly worse than the other two construction types. This is probably due to the fact that because the Type 1 and Type 2 constructions have more faces, there is initially more room to (bistellarly) shuffle the faces around making it easier to find downward paths. However, the minimum is what we are most interesting in finding and the Type 3 construction seems to go down to around the same as the other two types in that respect. That is, since the Type 3 construction will run faster, we can just run more instances (seed numbers) of it and feel fairly confident that the best triangulation found over those runs will be as good as the best found over a fewer number of runs of the other two types. And that is exactly what we did. We stopped the tests on the Type 1 and Type 2 constructions after $r = 10$. And we continued to run tests on higher r 's of the Type 3 construction. The table below displays the minimum vertex count found over those runs. We can see that as the r increases, the harder it was for `bistellar_simplification` to find a small triangulation, tho it seems to taper off a bit.

⁷The algorithm `bistellar_simplification` will “give up” when sufficiently many heating rounds have been run without any improvement to the g -vector. That number of heating rounds when the algorithm gives up is determined by the size of the complex—10 times the number of facets, by default. So the algorithm takes longer to run on larger complexes also by design.

3.3.4 Fundamental Group

To compute the fundamental group, we use `polymake` to find the edge-path group [67] and simplify the resulting presentation using `GAP`, which uses Tietze transformations as described in [34].

OBSERVATION 3.10

The presentation of the fundamental group computed using `polymake` and `GAP` on triangulations for $AK_I(r)$, $AK_{II}(r)$, $AK_{III}(r)$ for $r = 3, \dots, 20$ are all nontrivial presentations of the trivial group.

For $r = 4$, we started bistellar flips on $AK_{III}(4)$ for different random seeds and then computed the fundamental group on the small triangulations we obtained. In 100 out of 450 runs a trivial presentation was found for the fundamental group, giving a certificate that $AK_{III}(4)$ is a topological 4-sphere (by Freedman’s classification), but failing to give a certificate for whether $AK_{III}(4)$ is a PL 4-sphere (which a priori is known by Akbulut’s proof [5] that all the examples $AK_I(r)$, $AK_{II}(r)$, $AK_{III}(r)$ are standard).

We tried similar tests on the (bistellarly) simplified triangulations of different types $AK_I(r)$, $AK_{II}(r)$, $AK_{III}(r)$ and $r = 4, \dots, 10$ that we obtained in the previous experiment using 20 different seed numbers. One notable observation is that we often (see Table 3.3) encountered presentations that “resembled” the original presentation $G_{AK(r)}$ up to rearranging some letters of each relator and/or exchanging one of the generators with its inverse. Here are a few examples of such presentations.

$$\begin{aligned} AK_I(6)(\text{seed} = 16) &: \langle x, y \mid x^{-1}yx^{-1}y^{-1}xy^{-1}, x^4y^5x^2 \rangle \\ AK_{II}(7)(\text{seed} = 21) &: \langle x, y \mid x^{-1}y^{-1}xyxy^{-1}, yx^{-7}y^5 \rangle \\ AK_{III}(8)(\text{seed} = 16) &: \langle x, y \mid yxy^{-1}x^{-1}y^{-1}x, y^{-7}x^8 \rangle \end{aligned}$$

We also give a few examples of presentations that do *not* resemble $G_{AK(r)}$ for comparison.

$$\begin{aligned} AK_I(4)(\text{seed} = 14) &: \langle x, y \mid x(yx^{-1})^2y^{-1}x, yxyx^{-3}y^{-1}x, x^{-2}y^{-1}(xy)^2x^{-1} \rangle \\ AK_{II}(5)(\text{seed} = 20) &: \langle x, y \mid x^{-1}yxy^{-1}xyx^{-1}, y(x^{-1}y^{-1})^2xyx^2 \rangle \\ AK_{II}(9)(\text{seed} = 14) &: \langle x, y \mid x^{-1}y^2x^{-1}y(xy^{-3})^3xy^{-2}(x^{-1}y^3)^3, \\ &\quad x^{-1}y^{-3}(x^{-1}y^3)^3x^{-1}y^{-1}xyx^{-1}y^{-5}(x^{-1}y^3)^3, \\ &\quad x(y^{-2}xy^{-1})^4y^{-4}(x^{-1}y^3)^3x^{-1}y^{-1}xy^5 \rangle \\ AK_{III}(10)(\text{seed} = 28) &: \langle x, y \mid xy^{-1}xyx^{-2}y, x^{-5}y^{-1}xyx^7yx^{-1}y^{-1}x^{-1} \rangle \end{aligned}$$

The number of these familiar-looking presentations increased as r increased, see Table 3.3. The number of trivial presentations (the empty presentation), however, decreased as r went up. We will also mention that the familiar presentations were found quite often for the triangulation of the spheres before any bistellar flips were performed. This is not so surprising given our method of construction. The vertex labels influence the edge-path

	trivial							original						
	4	5	6	7	8	9	10	4	5	6	7	8	9	10
$AK_I(r)$	5	3	0	2	1	0	1	0	0	2	3	1	0	4
$AK_{II}(r)$	5	3	3	1	0	2	0	1	0	5	2	4	2	4
$AK_{III}(r)$	4	2	2	1	0	0	1	0	1	1	4	8	6	7

Table 3.3: We computed presentations of the fundamental groups of the smallest triangulations obtained by `bistellar_simplification` using 20 seed numbers for all three types for $r = 4, \dots, 10$. The table shows the number of those presentations that `GAP` simplified to be the empty set (trivial on Left) or a presentation that resembled $G_{AK(r)}$ (original on Right).

group that is generated and our labeling scheme exactly follows the relators of $G_{AK(r)}$ by design. It *is* surprising to find that much of that original structure is preserved after some 300 million bistellar moves.

We also note that every nontrivial presentation we found always had two generators. The number of relators was (zero or) two or three, but the number of generators was always either none or two.

3.3.5 Conclusion

We conclude this chapter with the following titular claim of this dissertation.

OBSERVATION 3.11

The complexes $AK_I(r)$, $AK_{II}(r)$, $AK_{III}(r)$ for $r > 3$ are complicated triangulations of the PL 4-sphere.

Chapter 4

Mazur & Friends

To build the Akbulut–Kirby spheres, we start with a ball with two 1-handles and then attach two 2-handles. Constructing Mazur’s 4-manifold is similar in that it starts with a ball with one 1-handle and then attaching one 2-handle. Topologically, however, the Akbulut–Kirby sphere and Mazur’s 4-manifold are of two entirely different species.

With the Akbulut–Kirby spheres, the two 2-handles are attached to kill off the homology generators (and also the generators of the fundamental group), that is, they close up the two holes that were created when the two 1-handles were attached to the ball. The resulting manifold is then contractible, but, a priori, it is not obvious that it is a (smooth or PL) 4-ball. Akbulut and Kirby [6] had to prove that the 4-manifold they obtained was indeed diffeomorphic to a 4-ball. One example where this construction fails to produce a ball is Mazur’s 4-manifold. The 4-manifold one obtains after gluing in the 2-handle is contractible, but *not* a 4-ball.

The construction method we developed to build the triangulation of the Akbulut–Kirby spheres provided all the tools necessary to triangulate Mazur’s 4-manifold. And so we did. In Section 4.1, we give a brief overview of Mazur’s 4-manifold and a construction of interesting non-PL manifolds described in [9]. In Section 4.2, we describe our construction of the triangulation of Mazur’s 4-manifold and triangulations (with perfect discrete Morse vectors) of a non-PL 5-ball and a non-PL 5-sphere which can be built using the triangulation of Mazur’s 4-manifold. In Section 4.3, we discuss our results from experiments run on the triangulations.

4.1 Background

Mazur’s 4-manifold [52] is a contractible 4-manifold with a non-simply connected boundary, a homology 3-sphere that is not S^3 . The boundary is an example of a Brieskorn homology sphere $\Sigma(2, 5, 7)$, which is a homology sphere that bounds a contractible manifold [20].

Cannon[17] and Edwards [25] showed that one can obtain a topological sphere by taking the double suspension of any homology sphere. As an example, they took the double suspension of the boundary of Mazur’s 4-manifold and showed that it is indeed S^5 .

Whitehead [77] proved that any compact manifold with boundary that has a PL triangulation that is collapsible is a PL ball. A natural question to ask next is whether there are also non-PL balls that are collapsible. Benedetti [9] describes Adiprasito’s construction of a collapsible non-PL triangulation of a 5-ball. For this construction, they, again, use Mazur’s 4-manifold.

4.2 Construction

The construction techniques developed to build the Akbulut–Kirby spheres in Section 3.2.3 were repurposed for the construction of Mazur’s 4-manifold which was then used to construct a collapsible non-PL 5-ball, which in turn was used to build a non-PL 5-sphere (for which subdivisions can have perfect discrete Morse vectors). In Section 4.2.1, we describe the construction for the triangulation of Mazur’s 4-manifold. In Section 4.2.2, we describe the construction for the triangulations of the non-PL 5-ball and the non-PL 5-sphere.

4.2.1 Mazur’s 4-manifold

In this section we describe the construction of our triangulations M of Mazur’s 4-manifold. To build M we use a construction procedure similar to that of the Akbulut–Kirby sphere, see Section 3.2.3, and therefore we will use similar notation.




Figure 4.1: Left: The curve Γ used to build a Mazur manifold is taken from [52]. Right: The same curve using our notation. The blue areas indicate where we place the vertices for the chain of triangular prisms.

Step 1: Set up the space H_M

We, again, will start in 3-space. We build H_M , a 3-ball with one (3-dimensional) 1-handle.

Step 2: Setting up ϱ_Γ

Mazur [52] sketches a very specific curve Γ that corresponds to a generator of the first homology group $H_1(H_M)$, see Fig. 4.1 left. To help us with our construction, we redraw Γ in our notation ϱ_Γ , see Figure 4.1 right. We concentrate any non-parallel lines into one section which we call H_M^0 and keep all the parallel lines in the 1-handle just like we did for the Akbulut–Kirby spheres. However, here we will split the 1-handle into two parts $H_{M,1}^1, H_{M,2}^1$ so that we do not accidentally introduce two edges that share two vertices. We also split the two crossing loops inside of H_M^0 for the same reason (the blue dots in Figure 4.1 indicate the respective extra vertices).

Step 3: Building τ_Γ

Mazur then instructs us to thicken up Γ to a tubular $\Gamma \times \mathbb{I}^2$. Just as in Section 3.2, we thicken up ϱ_Γ to a chain of triangular prisms τ_Γ . The vertices of the triangular prisms are located exactly at the blue horizontal lines and points in Figure 4.1 right. We chop up ϱ_Γ into 5 strands in H_M^0 and three strands each in the two parts of the 1-handle. There are $(5 + 2 \cdot 3) \times 3 = 11 \times 3 = 33$ vertices in τ_Γ . We triangulate the prisms using the product triangulation method as before.

Next we want to fill the space around τ_Γ to build a solid one-holed handlebody H_M .

Step 4: Filling H_M^0

The two parts of the 1-handle $H_{M,1}^1, H_{M,2}^1$ are easy to fill using rectangular buffer prisms between the triangular prisms of τ_Γ , just as we did in Section 3.2.3. The remaining piece, H_M^0 , will require more care.

Figure 4.2 is a zoomed in view of H_M^0 . We begin, as before, by dividing H_M^0 into different floors so we can use the fill-a-cupula method. As before, we use different colors to identify the strands on different floors: `level:0` in gray, `level:-1` in pink, `level:+1` in purple and `level:+2` in blue. The green triangular face of b is laying on the carpet of `level:0`.

We provide sketches that were used for the rest of this step in Figure 4.4.

Finish up M

With H_M^0 completed, we will then go up to dimension 4. And finally, glue in a 4-dimensional 2-handle H_Γ^2 to obtain M .

4.2.2 A perfect non-PL 5-ball and a perfect non-PL 5-sphere

The following is the procedure for building a non-PL 5-ball that admits a perfect discrete Morse vector $(1, 0, 0, 0, 0, 0)$ and a non-PL 5-sphere that admits a perfect discrete Morse vector $(1, 0, 0, 0, 0, 1)$ starting with Mazur’s 4-manifold M .




Figure 4.2: (Left) A zoomed in view of H_M^0 of Figure 4.1. (Right) A view of H_M^0 on level:0.







Figure 4.3



Forming the cupula (underneath) `level:-1`




Before cupula




After cupula

Forming the cupula in `level:+1`




Before cupula



After cupula

Forming the cupula in `level:+2`



Take a cone for the last extra tets.

Figure 4.4

Procedure

Construct a non-PL 5-ball \mathbf{B} starting with a triangulation of Mazur's 4-manifold M . We take some of the notation from [9, Proposition 3.20]

Step B1: Take the boundary $H = \partial M$.

Step B2: Let p be a vertex not in M and take the cone(H) = $p * H$.

Step B3: Let $C_0 = M \cup \text{cone}(H)$.

Step B4: Let a be a vertex not in C_0 and take the cone(C_0) = $a * C_0$.

Step B5: Make a copy cone(C'_0) of cone(C_0) by incrementing all vertex labels by the value of a .

Step B6: Build J (see below).

Step B7: Let $\mathbf{B} = \text{cone}(C_0) \cup \text{cone}(C'_0) \cup J$.

To build a perfect non-PL 5-sphere \mathbf{S} starting with \mathbf{B} .

Step S1: Let c be a vertex not in \mathbf{B} and take the cone($\partial\mathbf{B}$) = $c * \partial\mathbf{B}$.

Step S2: Let $\mathbf{S} = \mathbf{B} \cup \text{cone}(\partial\mathbf{B})$.

Step B6

The J in Step B6 is the collection of the triangulations of the prisms whose one face is a 4-simplex in cone(H) and the other face is the corresponding 4-simplex in the copy of cone(H) in C'_0 . That is, for every $[v_0 \ v_1 \ v_2 \ v_3 \ v_4] \in \text{cone}(H)$ we want to find the triangulation of the prism $[v_0 \ v_1 \ v_2 \ v_3 \ v_4 \ w_0 \ w_1 \ w_2 \ w_3 \ w_4]$ using the product triangulation method (with our extra check, see 3.2.5), where $w_i = v_i + a$ for $i = 0, \dots, 4$.

This step is equivalent to identifying cone(H) in C_0 to its copy cone(H') in C'_0 while avoiding unwanted identifications. That is, C_1 corresponds exactly to the suspension $C_1 = \Sigma C_0$ from [9].

4.3 Results & Experiments

4.3.1 Mazur’s 4-manifold

We refer to our triangulation of Mazur’s 4-manifold as M .

THEOREM 4.1

Mazur’s 4-manifold has a triangulation with face vector:

$$f(M) = (89, 850, 2175, 2151, 738).$$

All the vertices of M lie on its boundary which has face vector:

$$f(\partial M) = (89, 701, 1224, 612).$$

The boundary ∂M is a triangulation of the Brieskorn homology sphere $\Sigma(2, 5, 7)$.

The homology of M is $H_*(M) = (\mathbb{Z}, 0, 0, 0, 0, 0)$. It has trivial fundamental group. The best discrete Morse vector we obtained was $(1, 1, 1, 0, 0)$ which was found 8349 times out of 10000 runs of `Random_Discrete_Morse` using strategy `random-random`.

The boundary $H = \partial M$ has spherical homology. Its fundamental group was found with `GAP` to have as one of its presentations:

$$\langle x, y \mid y^{-1}xy^4xy^{-2}, y^2(xy^{-1})^3xy^2x^{-1} \rangle.$$

The best discrete Morse vector we found was $(1, 2, 2, 1)$ which was found 3458 times out of 10000 runs of `Random_Discrete_Morse` using strategy `random-random`.

Furthermore, we attempted `bistellar_simplification` on $H = \partial M$ using several seed numbers. A collection of small triangulations of homology 3-spheres can be found on the webpage [47]. In at least one case, the simplified complex was found to be isomorphic to the homology sphere `Sigma_2_5_7`. We allowed our program to stop running as soon as it found a match to any of the six spheres listed on that webpage. A match was found after only ten attempts.

Mazur’s 4-manifold is contractible, but M is not collapsible. By Whitehead’s theorem 2.1, this is because M as a PL manifold is not a PL-ball. There are, however, non-PL balls that admit a discrete Morse vector of $(1, 0, \dots, 0)$ and thus are collapsible. The non-PL 5-ball **B** is one such example.

One further remark is that this construction of Mazur’s 4-manifold produces a non-ball because of the choice of Mazur’s Γ , the generator for the first homology group. Notice that Γ is knotted¹ inside of the solid torus

¹Mazur [52] would perhaps describe Γ as “clucky”. “Presumably only particularly unlucky choices of Γ will have [the manifold] simply connected.” This could also be a typo for clunky. Or, perhaps, lucky.

in dimensions 3. If we require it to remain in the boundary in dimension 4, it will still be knotted. However, if we bring the whole torus (with the thickened Γ inside) up to dimension 5, we can find (a copy of) the thickened $\Gamma \times \mathbb{I}^3$, which in the 4-dimensional boundary will “unclasp”.

As a test, we constructed a 5-manifold gluing in a 5-dimensional 2-handle along a copy of a thickened Γ which we found on the boundary of a (5-dimensional) solid torus. This 5-manifold has f -vector $(133, 1743, 6319, 9814, 6984, 1878)$. And, as expected, we found the boundary of this new 5-manifold to be PL-homeomorphic to the boundary of the 5-simplex; it has f -vector $(133, 1743, 5662, 6750, 2700)$. In fact, `polymake`’s `bistellar_simplification` client recognized it to be a sphere very fast after only 34956 moves (using `seed=0`); the Akbulut–Kirby sphere $AK_{III}(3)$ —when successful—required over 7 million moves on average to be recognized as a 4-sphere.

4.3.2 A perfect non-PL 5-ball

We refer to our triangulation of the non-PL 5-ball as **B**.

THEOREM 4.2

The non-PL 5-ball **B** with face vector:

$$f(\mathbf{B}) = (182, 2938, 13060, 23785, 19242, 5760).$$

has perfect discrete Morse vector $(1, 0, 0, 0, 0, 0)$ and thus is collapsible.

The non-PL 5-ball **B** has homology $H_*(\mathbf{B}) = (\mathbb{Z}, 0, 0, 0, 0, 0)$. `GAP` found its fundamental group to be trivial. And out of 10000 runs of `Random_Discrete_Morse`, we found the discrete Morse vector $(1, 0, 0, 0, 0, 0)$ every time using any strategy.

Note that [1] describes a non-PL 5-manifold that admits a perfect discrete Morse vector, but that example is not a 5-ball since its boundary is not a 4-sphere.

4.3.3 A perfect non-PL 5-sphere

We refer to our triangulation of the non-PL 5-sphere as **S**.

THEOREM 4.3

The non-PL 5-sphere **S** has a triangulation with face vector:

$$f(\mathbf{S}) = (183, 3116, 15550, 31985, 29052, 9684)$$

which contains Mazur’s contractible 4-manifold M in its 4-skeleton.

The non-PL 5-sphere \mathbf{S} has spherical homology. Its fundamental group is also trivial as is its simplified presentation obtained via `polymake` and `GAP`. The best discrete Morse vector we found using `Random_Discrete_Morse` was $(1, 0, 1, 2, 1, 1)$ which was found 22 out of 100 runs.

According to Adiprasito's construction [9, Proposition 3.20], after sufficiently many barycentric subdivisions $\mathbf{B}' = sd^m(\mathbf{B}, m > 0)$ on \mathbf{B} we can obtain a non-PL 5-sphere $\mathbf{S}' = \mathbf{B}' \cup \text{cone}(\partial\mathbf{B}')$ which will admit a perfect discrete Morse vector. However, one barycentric subdivision of \mathbf{S} is a simplicial complex that has 89570 vertices (not counting its other dimensional faces). We cannot fit the Hasse diagram of such an enormous object into the memory of our 8GB machines. So we made random stellar subdivisions of faces of any dimension on the boundary of \mathbf{S} to see if any of these partially subdivided \mathbf{S}' would admit a perfect discrete Morse vector. Our experiments thus far in this direction have been inconclusive.

Alternatively, we can make a different type of local modification. We know that bistellar flips preserves the PL-type of (closed) PL manifolds. But we can also try bistellar flips on non-PL manifolds. A bistellar flip removes a ball (the star of some face) and replaces it with a topologically equivalent ball.

We tried bistellar flips (making only 1–4-moves or 2–3-moves so that we do not change the number of vertices) on \mathbf{S} to see if we can find a triangulation that admits a perfect discrete Morse vector. Out of over 1500 tries, where in each run we make 100000 bistellar moves, we found one triangulation T_{243} which admitted a triangulation with discrete Morse vector $(1, 0, 0, 1, 1, 1)$. We then produced another 100 triangulations which were obtained by an additional 10000 bistellar moves each on T_{243} . Of those 100 new triangulations, four of them admitted a perfect discrete Morse vector of $(1, 0, 0, 0, 0, 1)$, the smallest of which has f -vector $(183, 3281, 16004, 32522, 29424, 9808)$.

THEOREM 4.4

There is a collapsible non-PL triangulation of S^5 with face-vector

$$f = (183, 3281, 16004, 32522, 29424, 9808).$$

With enough time and attempts, we expect the stellar subdivision tests to also find a triangulation that admits a perfect discrete Morse vector. Each stellar subdivision adds an additional vertex increasing the size of the complex (and therefore also its Hasse diagram) and as a result, the stellar subdivision tests take much more time (and memory) compared to the bistellar tests. Moreover, the fact that we found only one instance out of 1500 runs of the bistellar test that found a better discrete Morse vector (which then required further flips) suggests that it may be difficult to find a random stellar subdivision of \mathbf{B} which can be used to construct some \mathbf{S}' that admits a perfect Morse vector.

Bibliography

- [1] K. Adiprasito, B. Benedetti, and F.H. Lutz. Extremal examples of collapsible complexes and random discrete Morse theory. [arXiv:1404.4239](#), 2014, 20 pages.
- [2] K. Adiprasito and I. Izmistiev. Derived subdivisions make every pl sphere polytopal. [arXiv:1311.2965](#), 2014, 7 pages.
- [3] S.I. Adyan. Algorithmic unsolvability of problems of recognition of certain properties of groups. (Russian). *Dokl. Akad. Nauk SSSR*, 103:533–535, 1955.
- [4] I.R. Aitchison and J.H. Rubinstein. Fibered knots and involutions on homotopy spheres. Four-manifold theory, Proc. AMS-IMS-SIAM Joint Summer Res. Conf., Durham/N.H. 1982, Contemp. Math. 35, 1-74, 1984.
- [5] S. Akbulut. Cappell–Shaneson homotopy spheres are standard. *Ann. Math.*, 171:2171–2175, 2010.
- [6] S. Akbulut and R. Kirby. A potential smooth counterexample in dimension 4 to the Poincaré conjecture, the Schoenflies conjecture, and the Andrews–Curtis conjecture. *Topology*, 24:375–390, 1985.
- [7] J.J. Andrews and M.L. Curtis. Free groups and handlebodies. *Proc. Am. Math. Soc.*, 16:192–195, 1965.
- [8] B. Bagchi and B. Datta. Combinatorial triangulations of homology spheres. *Discrete Math.*, 305:1–17, 2005.
- [9] B. Benedetti. Smoothing discrete Morse theory. [arXiv:1212.0885v2](#), 2014, 27 pages; *Annali della Scuola Normale Superiore di Pisa Classe di Scienze* (to appear).
- [10] B. Benedetti and F. H. Lutz. Random discrete Morse theory and a new library of triangulations. *Exp. Math.*, 23:66–94, 2014.
- [11] B. Benedetti and F.H. Lutz. The dunce hat and a minimal non-extendably collapsible 3-ball. Electronic Geometry Model, No. 2013.10.001, 2013. <http://www.eg-models.de/2013.10.001>.
- [12] B. Benedetti and F.H. Lutz. Knots in collapsible and non-collapsible balls. *Electron. J. Comb.*, 20(3):P31, 2013.

- [13] B. Benedetti and G. M. Ziegler. On locally constructible spheres and balls. *Acta Math.*, 206:205–243, 2011.
- [14] A. Björner and F.H. Lutz. Simplicial manifolds, bistellar flips and a 16-vertex triangulation of the Poincaré homology 3-sphere. *Exp. Math.*, 9(2):275–289, 2000.
- [15] W.W. Boone, W. Haken, and V. Poenaru. On recursively unsolvable problems in topology and their classification. *Contrib. Math. Logic, Proc. Logic Colloq., Hannover 1966*, 37-74 (1968), 1968.
- [16] Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for 3-manifold topology and normal surface theory. <http://regina.sourceforge.net/>, 1999–2014.
- [17] J.W. Cannon. Shrinking cell-like decompositions of manifolds. Codimension three. *Ann. Math.*, 110:83–112, 1979.
- [18] CAPD::RedHom. Redhom software library. <http://redhom.ii.uj.edu.pl>.
- [19] S.E. Cappell and J.L. Shaneson. Some new four-manifolds. *Ann. Math.*, 104(2):61–72, 1976.
- [20] A.J. Casson and J.L. Harer. Some homology lens spaces which bound rational homology balls. *Pac. J. Math.*, 96:23–36, 1981.
- [21] A.V. Chernavsky and V.P. Leksine. Unrecognizability of manifolds. *Ann. Pure Appl. Logic*, 141:325–335, 2006.
- [22] CHomP. Computational Homology Project. <http://chomp.rutgers.edu>.
- [23] 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.
- [24] M. Dehn. Über unendliche diskontinuierliche Gruppen. *Math. Ann.*, 71:116–144, 1912.
- [25] R.D. Edwards. The double suspension of a certain homology 3-sphere is S^5 . *Notices AMS*, 22:A–334, 1975.
- [26] R. Forman. Morse theory for cell complexes. *Adv. Math.*, 134:90–145, 1998.
- [27] R. Forman. A user’s guide to discrete Morse theory. *Sémin. Lothar. Comb.*, 48:B48c, 35 p., electronic only, 2002.
- [28] M.H. Freedman. The topology of four-dimensional manifolds. *J. Differ. Geom.*, 17(3):357–453, 1982.
- [29] The GAP Group. *GAP – Groups, Algorithms, and Programming, Version 4.7.6*, 2014.

- [30] E. Gawrilow and M. Joswig. polymake: a framework for analyzing convex polytopes. In *Polytopes—combinatorics and computation (Oberwolfach, 1997)*, volume 29 of *DMV Sem.*, pages 43–73. Birkhäuser, Basel, 2000.
- [31] R.E. Gompf. Killing the Akbulut–Kirby 4-sphere, with relevance to the Andrews–Curtis and Schoenflies problems. *Topology*, 30(1):97–115, 1991.
- [32] M. Gromov. Hyperbolic groups. Essays in group theory, Publ., Math. Sci. Res. Inst. 8, 75–263, 1987.
- [33] J. Hass and G. Kuperberg. The complexity of recognizing the 3-sphere. *Oberwolfach Reports*, 24:1425–1426, 2012.
- [34] G. Havas, P.E. Kenne, J.S. Richardson, and E.F. Robertson. A Tietze transformation program. Computational group theory, Proc. Symp., Durham/Engl. 1982, 69–73., 1984.
- [35] T. Homma, M. Ochiai, and M. Takahashi. An algorithm for recognizing S^3 in 3-manifolds with Heegaard splittings of genus two. *Osaka J. Math.*, 17:625–648, 1980.
- [36] J.F.P. Hudson. Piecewise Linear Topology. Mathematics Lecture Note Series. New York–Amsterdam: W.A. Benjamin, Inc. IX, 1969.
- [37] S.V. Ivanov. Recognizing the 3-sphere. *Ill. J. Math.*, 45(4):1073–1117, 2001.
- [38] M. Joswig, F. H. Lutz, and M. Tsuruga. Heuristics for sphere recognition. In H. Hong and C. Yap, editors, *Mathematical Software — ICMS 2014*, volume 8592 of *LNCIS*, pages 152–159. Springer, 2014. [arXiv:1405.3848](https://arxiv.org/abs/1405.3848), 2015, 30 pages, Extended version (Preprint).
- [39] M. Joswig and M.E. Pfetsch. Computing optimal Morse matchings. *SIAM J. Discrete Math.*, 20:11–25, 2006.
- [40] R. Kannan and A. Bachem. Polynomial algorithms for computing the Smith and Hermite normal forms of an integer matrix. *SIAM J. Comput.*, 8(4):499–507, 1979.
- [41] R.C. Kirby. *The Topology of 4-Manifolds*. Springer-Verlag, Berlin, 1989.
- [42] R.C. Kirby and L.C. Siebenmann. Foundational essays on topological manifolds, smoothings and triangulations, 1977.
- [43] V. Klee. A combinatorial analogue of Poincaré’s duality theorem. *Can. J. Math.*, 16:517–531, 1964.
- [44] D.E. Knuth. *The Art of Computer Programming. Vol. 2*. Addison-Wesley, Reading, MA, 1998. Seminumerical algorithms, Third edition [of MR0286318].

- [45] T. Lewiner, H. Lopes, and G. Tavares. Optimal discrete Morse functions for 2-manifolds. *Comput. Geom.*, 26:221–233, 2003.
- [46] W.B.R. Lickorish. Simplicial moves on complexes and manifolds. In *Proceedings of the Kirbyfest, Berkeley, CA, USA, June 22–26, 1998*, pages 299–320. University of Warwick, Institute of Mathematics, 1999.
- [47] F.H. Lutz. The manifold page, 2015. http://page.math.tu-berlin.de/~lutz/stellar/homology_3spheres.
- [48] F.H. Lutz, K. Mischaikow, V. Nanda, and M. Tsuruga. Hard instances for homology computations. (in preparation).
- [49] A. Markov. The insolubility of the problem of homeomorphy. *Dokl. Akad. Nauk SSSR*, 121:218–220, 1958.
- [50] W.S. Massey. *A Basic Course in Algebraic Topology*. Springer-Verlag, New York, NY, 1991.
- [51] S.V. Matveev. An algorithm for the recognition of 3-spheres (according to Thompson). *Sb. Math.*, 186(5):695–710 (1995); translation from *mat. sb.* 186, no. 5, 69–84, 1995.
- [52] B. Mazur. A note on some contractible 4-manifolds. *Ann. Math.*, 73(1):221–228, 1961.
- [53] C.F. Miller III. Turing machines to word problems. *Lecture Notes in Logic 42, Turing’s Legacy: Developments from Turing’s Ideas in Logic*, pages 329–385, 2014.
- [54] J.W. Milnor. On manifolds homeomorphic to the 7-sphere. *Ann. Math. (2)*, 64:399–405, 1956.
- [55] E.E. Moise. Affine structures in 3-manifolds. V: The triangulation theorem and hauptvermutung. *Ann. Math. (2)*, 56:96–114, 1952.
- [56] J.R. Munkres. *Elements of Algebraic Topology*. Addison-Wesley Publishing Company, Menlo Park, CA, 1984.
- [57] P.S. Novikov. *On the Algorithmic Unsolvability of the Word Problem in Group Theory*, volume 44 of *Trudy Matematicheskogo Instituta imeni V. A. Steklova*. Izdatel’sтво Akademii Nauk SSSR, Moscow, 1955.
- [58] U. Pachner. Konstruktionsmethoden und das kombinatorische Homöomorphieproblem für Triangulationen kompakter semilinearer Mannigfaltigkeiten. *Abh. Math. Sem. Univ. Hamburg*, 57:69–86, 1987.
- [59] G. Perelman. The entropy formula for the Ricci flow and its geometric applications. [arXiv:math/0211159](https://arxiv.org/abs/math/0211159), 2002, 39 pages.
- [60] Perseus. <http://www.math.rutgers.edu/~vidit/perseus.html>.
- [61] H. Poincaré. Analysis situs. *J. de l’Éc. Pol. (2)* I. 1-123, 1895.

- [62] M.O. Rabin. Recursive unsolvability of group theoretic problems. *Ann. Math. (2)*, 67:172–194, 1958.
- [63] T. Radó. Über den Begriff der Riemannschen Fläche. *Acta Litt. Sci. Szeged*, 2:101–121, 1925.
- [64] A.A. Ranicki. On the Hauptvermutung. In *The Hauptvermutung Book. A Collection of Papers of the Topology of Manifolds*, pages 3–31. Kluwer Academic Publishers, Dordrecht, 1996.
- [65] J.H. Rubinstein. An algorithm to recognize the 3-sphere. In S. D. Chatterji, editor, *Proc. Internat. Congr. of Mathematicians, ICM '94, Zürich*, Volume 1, pages 601–611, Basel, 1995. Birkhäuser Verlag.
- [66] S. Schleimer. Sphere recognition lies in NP. In *Low-dimensional and symplectic topology*, volume 82 of *Proc. Sympos. Pure Math.*, pages 183–213. Amer. Math. Soc., Providence, RI, 2011.
- [67] H. Seifert and W. Threlfall. *Lehrbuch der Topologie*. B.G. Teubner, Leipzig, 1934.
- [68] S. Smale. On the structure of manifolds. *Am. J. Math.*, 84:387–399, 1962.
- [69] J. Stillwell. *Classical Topology and Combinatorial Group Theory*. Springer-Verlag, New York, 2nd ed. edition, 1993.
- [70] T. Sulanke and F.H. Lutz. Isomorphism free lexicographic enumeration of triangulated surfaces and 3-manifolds. *Eur. J. Comb.*, 30:1965–1979, 2009.
- [71] E. Swartz. Topological finiteness for edge-vertex enumeration. *Adv. Math.*, 219(5):1722–1728, 2008.
- [72] A. Thompson. Thin position and the recognition problem for S^3 . *Math. Res. Lett.*, 1(5):613–630, 1994.
- [73] H. Tietze. Über die topologischen Invarianten mehrdimensionaler Mannigfaltigkeiten. *Monatsh. Math. Phys.*, 19:1–118, 1908.
- [74] I.A. Volodin and A.T. Fomenko. Manifolds, knots, and algorithms. *Sel. Math. Sov.*, 3:311–341, 1984.
- [75] I.A. Volodin, V.E. Kuznetsov, and A.T. Fomenko. The problem of discriminating algorithmically the standard three-dimensional sphere. *Russ. Math. Surv.*, 29(5):71–172, 1974.
- [76] J. von Neumann. Various techniques used in connection with random digits. *Proceedings of symposium on “Monte Carlo Method” 1949 Los Angeles. Nat’l. BuStandards, Applied Math. Series 12*, pages 36–38, 1951.

- [77] J.H.C. Whitehead. Simplicial spaces, nuclei and m -groups. *Proc. Lond. Math. Soc., II. Ser.*, 45:243–327, 1939.
- [78] Wolfram Research, Inc., Champaign, IL. *Mathematica, Version 7.0.0*, 2008.
- [79] E.C. Zeeman. Seminar on combinatorial topology. Institut des Hautes Etudes Scientifiques 1963. Chap. 1–6; Chap. 7, (1965); Chap. 8 (1966).
- [80] E.C. Zeeman. On the dunce hat. *Topology*, 2:341–358, 1963.
- [81] A.J. Zomorodian. *Topology for Computing*. Cambridge University Press, Cambridge, 2009.

Summary

The Akbulut–Kirby spheres are an infinite family of (smooth) 4-manifolds that are diffeomorphic to the standard sphere; they each have a handlebody decomposition that is described by a certain nontrivial presentation of the trivial group. These spheres were proposed as candidate counterexamples to the smooth Poincaré conjecture in dimension 4. The smooth Poincaré conjecture in dimension 4 is still open today, but the entire family of spheres are no longer exotic candidates. In dimension 4, the categories PL and DIFF coincide, so PL triangulations of the Akbulut–Kirby spheres are standard PL 4-spheres. We describe three different ways to construct explicit triangulations of the Akbulut–Kirby spheres.

However, to show that these PL triangulations are indeed PL 4-spheres is no simple task. The problem of PL sphere recognition in dimension 4 is also still open. In dimension 3, sphere recognition is known to be in the complexity class NP. In higher dimensions, PL spheres are unrecognizable. There are, however, heuristic algorithms that can recognize spheres effectively on many examples. We will discuss some of these algorithms and their limitations. Our experiments show that the heuristics—even when used in combination—were not able to recognize the majority of the Akbulut–Kirby triangulated PL 4-spheres.

We will also briefly discuss our construction of a triangulation of Mazur’s 4-manifold, which used many of the tools we designed to build the triangulated Akbulut–Kirby spheres. Mazur’s 4-manifold is a contractible 4-manifold that has a boundary that is a homology sphere that is not the standard 3-sphere. We use this triangulation to build a collapsible non-PL 5-ball and a non-PL 5-sphere.

Zusammenfassung

Die Akbulut–Kirby Sphären gibt eine unendliche Familie von (glatten) 4–Mannigfaltigkeiten, die diffeomorph zur Standardsphäre sind; sie haben jeweils eine handlebody Zersetzung, die durch eine bestimmte nicht-triviale Vorstellung des trivialen Gruppe beschrieben wird. Diese Sphären wurden als Kandidaten Gegenbeispiele zum reibungslosen Poincaré–Vermutung in Dimension 4. Die glatte Poincaré–Vermutung in Dimension 4 vorgeschlagen wird, ist noch offen heute, sondern die ganze Familie von Kugeln sind nicht mehr exotisch Kandidaten. In Dimension 4 die Kategorien PL und DIFF zusammenfallen, so PL Triangulierungen der Akbulut–Kirby Kugeln sind Standard PL 4–Sphären. Wir beschreiben drei verschiedene Möglichkeiten, explizite Triangulierungen der Akbulut–Kirby Sphären zu konstruieren.

Aber um zu zeigen, dass diese PL Triangulierungen sind in der Tat PL 4–Sphären ist keine einfache Aufgabe. Das Problem der PL Kugel Anerkennung in Dimension 4 ist auch noch offen. In Dimension 3 ist Sphäre Erkennung bekannt, in der Komplexitätsklasse NP sein. In höheren Dimensionen sind PL Sphären nicht wiederzuerkennen. Es gibt jedoch heuristische Algorithmen, die Kugeln effektiv auf vielen Beispielen erkennen kann. Wir werden einige dieser Algorithmen und ihre Grenzen zu diskutieren. Unsere Experimente zeigen, dass die Heuristik — selbst wenn sie in Kombination verwendet werden — nicht in der Lage, den Großteil der Akbulut–Kirby erkennen trianguliert PL 4–Sphären.

Wir werden auch kurz auf unsere Bau einer Triangulation von Mazur 4–Verteiler, die viele der Werkzeuge, die wir entwickelt, um die Dreiecks Akbulut–Kirby Kugeln bauen verwendet. Mazur 4–Verteilers ist eine zusammenziehbare 4–Verteiler, der eine Grenze, die eine Homologie Kugel, die nicht die Standard–3–Kugel hat. Wir verwenden diese Triangulation, eine zusammenklappbare nicht-PL 5–Ball und eine nicht-PL 5–Sphäre zu bauen.