

AG
T

*Algebraic & Geometric
Topology*

Volume 25 (2025)

Crushing surfaces of positive genus

BENJAMIN A BURTON

THIAGO DE PAIVA

ALEXANDER HE

CONNIE ON YU HUI



Crushing surfaces of positive genus

BENJAMIN A BURTON

THIAGO DE PAIVA

ALEXANDER HE

CONNIE ON YU HUI

The operation of crushing a normal surface has proven to be a powerful tool in computational 3-manifold topology, with applications both to triangulation complexity and to algorithms. The main difficulty with crushing is that it can drastically change the topology of a triangulation, so applications to date have been limited to relatively simple surfaces: 2-spheres, discs, annuli, and closed boundary-parallel surfaces. We give the first detailed analysis of the topological effects of crushing closed essential surfaces of positive genus. To showcase the utility of this new analysis, we use it to prove some results about how triangulation complexity interacts with JSJ decompositions and satellite knots; although similar applications can also be obtained using techniques of Matveev, our approach has the advantage that it avoids the machinery of almost simple spines and handle decompositions.

57K30, 57Q15

1. Introduction	4949
2. Preliminaries	4951
3. Atomic moves on cell decompositions with ideal vertices	4967
4. Crushing surfaces of positive genus	4994
5. Triangulation complexity of 3-dimensional submanifolds	5004
References	5010

1 Introduction

The idea of crushing a normal surface was first developed by Jaco and Rubinstein [18] as part of a broader program of giving a theory of “efficient” 3-manifold triangulations. This led to new insights on minimal triangulations [18], and has also been the key to developing “efficient” (in various senses of the word, depending on the particular application) algorithms to solve a number of fundamental problems in low-dimensional topology [2; 3; 5; 6; 7; 8; 9; 13].

The key obstacle in developing new applications of crushing is that this operation can drastically alter the topology of a triangulation. This difficulty was initially compounded by the complicated formulation of

crushing that was originally given by Jaco and Rubinstein; although they were able to give a number of applications, these required intricate arguments about the topological effects of crushing 2-spheres, discs and closed boundary-parallel surfaces [18]. More recent applications rely on simpler formulations of crushing that are easier to understand and use:

- Following unpublished ideas of Casson, Fowler [9] used the language of special spines to understand the effect of crushing 2-spheres.
- Burton introduced a way to break crushing down into a sequence of simple atomic moves, and used this atomic approach to describe the topological effects of crushing 2-spheres and discs [3]; this has proven to be extremely useful for turning crushing into an accessible algorithmic tool for working with 3-manifolds [2; 3; 5; 6; 7; 8]. This atomic approach has also recently been applied to crushing certain types of properly embedded annuli [13].

We emphasise that although it is, in principle, possible to crush any normal surface, the applications to date have only involved 2-spheres, discs, annuli and closed boundary-parallel surfaces. Probably the main reason for this is that as the surfaces get more complicated, the topological effects of crushing also appear to get more complicated. Nevertheless, we demonstrate in this paper that it is possible to push through this challenge by building upon the atomic approach to crushing from [3].

To be precise, we use the atomic approach to understand the topological effects of crushing closed normal surfaces of positive genus; in particular, we are able to crush essential surfaces, not just boundary-parallel ones. This work is distributed across two sections of this paper. First, in Section 3, we carefully work through the necessary details to extend the atomic approach to crushing. Then, in Section 4, we apply the work from Section 3 to actually understand the effect of crushing a surface of positive genus.

To state the main theorem from Section 4, we require some notation and terminology which we now outline (see Section 4 for the precise definitions). Given a normal surface S in a 3-manifold triangulation \mathcal{T} , our goal is to triangulate a submanifold X of \mathcal{T} that is “cut out” by the surface S . More precisely, we fix a component of $\mathcal{T} - S$, which we call the *chosen region*, and then take X to be the closure of the chosen region. After crushing S to obtain a new triangulation \mathcal{T}' , each component of $\mathcal{T} - S$ “falls apart” to yield some subset of the components of \mathcal{T}' . In particular, the chosen region yields some subset \mathcal{T}^* of \mathcal{T}' , and our hope is that (a component of) \mathcal{T}^* actually gives a triangulation of X .

For our purposes, it turns out to be important to “push” or “expand” S as far into the chosen region as possible, to obtain what we call a *maximal* surface. We show in Lemma 14 that we can always, without loss of generality, assume that S is maximal. With this groundwork, together with our analysis from Section 3, we are able to prove the following theorem in Section 4:

Theorem 1 *Suppose that X is irreducible, ∂ -irreducible and anannular, and that it contains no two-sided properly embedded Möbius bands. Also, suppose S is maximal. Then \mathcal{T}^* is a valid triangulation such that:*

- *One of its components is an ideal triangulation of X .*
- *Every other component is a triangulation of the 3-sphere.*

In Section 5, we apply Theorem 1 to study the *triangulation complexity* $\Delta(\mathcal{M})$: the minimum number of tetrahedra required to triangulate some particular 3-manifold \mathcal{M} . In particular, we obtain the following general result as a relatively straightforward consequence of Theorem 1:

Theorem 2 *Let \mathcal{M} be a compact 3-manifold with no 2-sphere boundary components. Suppose \mathcal{M} contains a (possibly disconnected) closed incompressible surface S with no 2-sphere components, no projective plane components, and no boundary-parallel components. Let \mathcal{R} be a component obtained after cutting \mathcal{M} along S . If \mathcal{R} is irreducible, ∂ -irreducible, annular, and does not contain any two-sided properly embedded Möbius bands, then $\Delta(\mathcal{R}) < \Delta(\mathcal{M})$.*

We continue in Section 5 by specialising Theorem 2 to the particularly interesting setting where S is a collection of essential tori. This gives various nice results about how triangulation complexity interacts with JSJ decompositions and satellite knots.

The applications that we obtain in Section 5 are not entirely new, since they can also be obtained by combining various pieces of machinery from Matveev's book [28] (we discuss this in a little more detail in Section 5). Nevertheless, our applications demonstrate that crushing normal surfaces of positive genus has nontrivial consequences for objects that are of independent interest. This provides hope that future refinements of our techniques could lead to further applications, such as new algorithms involving decompositions along surfaces of positive genus.

It is worth noting that whilst Matveev's techniques use almost simple spines and handle decompositions, our work does not require such machinery; instead, our analysis of crushing only uses triangulations and cell decompositions. Some readers might therefore find our approach more accessible than that of Matveev. Moreover, in contrast to handle decompositions, crushing has the advantage that it is well established in software such as Regina [2; 4]; thus, our approach is probably more amenable for practical algorithmic applications.

Acknowledgements

This project is the result of a conversation that began at the *PhD Student Symposium: Graduate Talks in Geometry and Topology Get-Together*, or *(GT)³*, hosted by MATRIX in July 2022; we would therefore like to thank the MATRIX Institute and the organisers of the symposium.

Alexander He was supported by an Australian Federal Government Research Training Program Scholarship.

We would like to thank the reviewer for their many suggestions, which greatly helped to improve the readability of this paper.

2 Preliminaries

The main purpose of this section is to review all the definitions that we will require for our analysis of crushing.




Figure 1: Two tetrahedra glued together along a single pair of triangular faces.

As a convention that we will use throughout this paper, except where we explicitly state otherwise, all 3-manifolds will be compact. We will call a (compact) 3-manifold *closed* if it has empty boundary, and *bounded* if it has nonempty boundary.

Also, whenever we are working with an object X (such as a knot or a surface) embedded in a 3-manifold \mathcal{M} , we will often refer to ambient isotopies of X in \mathcal{M} simply as isotopies of X . For example, when we speak of isotoping a knot K (embedded in the 3-sphere S^3), we really mean that we are applying an *ambient* isotopy to the embedding of K in S^3 .

2.1 Triangulations and cell decompositions

A (*generalised*) *triangulation* \mathcal{T} consists of finitely many (abstract) tetrahedra with some or all of their triangular faces *glued* together in pairs via affine identifications (Figure 1 illustrates a single such gluing); denote the number of tetrahedra in \mathcal{T} by $|\mathcal{T}|$. We allow faces from the same tetrahedron to be glued together, which means that \mathcal{T} need not be a simplicial complex; indeed, generalised triangulations can usually be made much smaller than topologically equivalent simplicial complexes, which is often important for computational purposes.

In this paper, we also work with cell decompositions, which generalise the triangulations that we just defined. We build gradually towards a definition of cell decompositions, starting with an explanation of how we generalise tetrahedra to obtain a larger class of “building blocks”.

Topologically, we can think of a tetrahedron as a 3-ball whose boundary 2-sphere is divided into triangles by an embedding of the complete graph on four vertices. To generalise this, consider a topological 3-ball Δ with a multigraph Γ embedded in $\partial\Delta$. We call Δ an (*abstract*) *3-cell* if:

- Γ has no degree one vertices.
- The closure of each component of $(\partial\Delta) - \Gamma$ forms an embedded disc, which we call a *face* of Δ , whose boundary circle contains two or more vertices of Γ .

Assuming that these conditions are indeed satisfied, we refer to the vertices and edges of Γ as *vertices* and *edges*, respectively, of the 3-cell Δ . Intuitively, each face of an abstract 3-cell forms a curvilinear




Figure 2: Some examples of the nontetrahedron cells that we will encounter.

polygon with two or more edges; indeed, depending on the number of edges, we will often describe 3-cell faces as bigons, triangles, quadrilaterals, and so on.

There are infinitely many types of 3-cells. However, for our purposes, we will only need to deal with a finite number of these; some examples are shown in Figure 2. For details on precisely which types of 3-cells we need, see Definitions 3 and Section 2.5.

We now explain how we glue 3-cells together to obtain a cell decomposition. Endow every edge e of a 3-cell with an affine structure—a homeomorphism from e to the interval $[0, 1]$. We *glue* two distinct faces of two (not necessarily distinct) 3-cells via a homeomorphism that

- maps vertices to vertices;
- maps edges to edges; and
- restricts to an affine map on each edge.

A *cell decomposition* is a collection of finitely many 3-cells with some or all of their faces glued together in pairs; we emphasise again that we allow faces from the same 3-cell to be glued together. Since triangulations are a special case of cell decompositions, all of the subsequent definitions for cell decompositions apply to triangulations too.

Let \mathcal{D} denote a cell decomposition. The gluings that define \mathcal{D} give an equivalence relation on the faces of the 3-cells of \mathcal{D} ; call each equivalence class a *face* or *2-cell* of \mathcal{D} . More explicitly, a face of \mathcal{D} is either:

- A pair of 3-cell faces that have been glued together, in which case we say that the face is *internal*.
- A single 3-cell face that has been left unglued, in which case we say that the face is *boundary*.

The *boundary* of \mathcal{D} is the (possibly empty) union of all its boundary faces.

The gluings that define \mathcal{D} also merge vertices and edges of the 3-cells into equivalence classes; call each such vertex class a *vertex* or *0-cell* of \mathcal{D} , and call each such edge class an *edge* or *1-cell* of \mathcal{D} . For each $k \in \{0, 1, 2\}$, define the *k-skeleton* of \mathcal{D} , denoted by $\mathcal{D}^{(k)}$, to be the union of all n -cells of \mathcal{D} , where n runs over all dimensions up to and including k .

In general, if we consider the quotient topology arising from the face gluings that define a cell decomposition \mathcal{D} , the resulting topological space might fail to be a 3-manifold. Specifically, although nothing goes wrong in the interiors of 3-cells and the interiors of faces, we need to be careful with vertices and with midpoints of edges.

We begin by considering the midpoint p of an edge e . If e lies entirely in the boundary of \mathcal{D} , then the frontier of a small regular neighbourhood of p is a disc; in this case, nothing goes wrong, and we say that e is *boundary*. However, if e does not lie in the boundary, then we have two possibilities:

- If e is identified with itself in reverse, then the frontier of a small regular neighbourhood of p is a projective plane; this cannot occur in a 3-manifold. In this case, we say that e is *invalid*.
- Otherwise, the frontier of a small regular neighbourhood of p is a 2-sphere. In this case, nothing goes wrong and we say that e is *internal*.

We also say that e is *valid* if it is either boundary or internal.

For a vertex v , consider the surface given by the frontier of a small regular neighbourhood of v ; we call this surface the *link* of v . When v lies in the boundary of \mathcal{D} , its link is a surface with boundary. If the link is a disc, then nothing goes wrong and we say that v is *boundary*; otherwise, if the link is any other surface with boundary, then we say that the vertex is *invalid*.

On the other hand, when v does not lie in the boundary, its link is a closed surface. If the link is a 2-sphere, then nothing goes wrong and we say that v is *internal*; otherwise, if the link is any other closed surface, then we say that v is *ideal*.

A cell decomposition is *valid* if it has no invalid edges or vertices, and *invalid* otherwise. Given a (possibly invalid) cell decomposition \mathcal{D} , we often find it useful to *truncate* a vertex v by deleting a small open regular neighbourhood of v . In particular, by truncating each ideal or invalid vertex in \mathcal{D} , we obtain a pseudomanifold \mathcal{P} that we call the *truncated pseudomanifold* of \mathcal{D} ; the reason \mathcal{P} is a pseudomanifold (and not necessarily a manifold) is that midpoints of invalid edges in \mathcal{D} would give nonmanifold points in \mathcal{P} .

Observe that if \mathcal{D} has no invalid edges, then the truncated pseudomanifold \mathcal{P} is actually a (compact) 3-manifold. In this case, we will often refer to \mathcal{P} as the *truncated 3-manifold* of \mathcal{D} , and we will say that \mathcal{D} *represents* the 3-manifold \mathcal{P} ; when \mathcal{D} happens to be a triangulation, we will also often say that \mathcal{D} *triangulates* \mathcal{P} . Moreover, in the case where \mathcal{D} is valid and has no ideal vertices, since we do not need to truncate any vertices to obtain the truncated 3-manifold \mathcal{P} , we will sometimes find it more natural to refer to \mathcal{P} as the *underlying 3-manifold* of \mathcal{D} .

If we assume that \mathcal{D} is actually valid (so it has neither invalid edges nor invalid vertices), then the boundary components of the truncated 3-manifold \mathcal{P} come in two possible types, namely

- *ideal* boundary components, which are the boundary components that arise from truncating the ideal vertices; and
- *real* boundary components, which are built from boundary faces of \mathcal{D} .

In this case, it will be convenient to distinguish the following special types of cell decompositions:

- A valid cell decomposition is *closed* if every vertex is internal. For a closed cell decomposition, the truncated 3-manifold is a closed 3-manifold.
- A valid cell decomposition is *bounded* if it has at least one boundary vertex, and has no ideal vertices. For a bounded cell decomposition, the truncated 3-manifold is a bounded 3-manifold whose boundary components are all real.
- A valid cell decomposition is *ideal* if it has at least one ideal vertex, and has no boundary vertices. For an ideal cell decomposition, the truncated 3-manifold is again a bounded 3-manifold, but this time the boundary components are all ideal.

Remark When we have an ideal cell decomposition \mathcal{D} , we use the notion of the truncated 3-manifold to turn \mathcal{D} into a *compact* 3-manifold \mathcal{M} . A very common alternative (which we do not use in this paper) is to turn \mathcal{D} into a *noncompact* 3-manifold \mathcal{M}' by simply deleting (rather than truncating) each ideal vertex. Observe that \mathcal{M}' is homeomorphic to the interior of \mathcal{M} , so this distinction is not too important.

Remark Suppose \mathcal{T} is either a closed or ideal triangulation, and let \mathcal{M} denote the truncated 3-manifold of \mathcal{T} . Since we do not truncate the internal vertices of \mathcal{T} , observe that \mathcal{M} is a 3-manifold with no 2-sphere boundary components. For this reason, we will often find it convenient to make the mild assumption that a 3-manifold has no 2-sphere boundary components.

2.2 Decomposing along curves and surfaces

The goal in this section is to introduce some terminology that will streamline our descriptions of the topological effects of crushing. The idea is that crushing often changes the truncated 3-manifold or pseudomanifold by “decomposing along” a properly embedded surface; we will build gradually towards defining precisely what we mean by this. We start by going one dimension down, and defining what we mean by decomposing a surface along an embedded curve; this is useful in its own right, since it will help us describe how crushing changes the links of vertices.

Consider an embedded closed curve γ in a compact surface S . Let S^\dagger denote the surface obtained from S by *cutting along* γ — that is, removing a small open regular neighbourhood of γ from S . If γ is a two-sided curve in S , then we have two new copies of γ in ∂S^\dagger ; on the other hand, if γ is one-sided, then we have a single new curve in ∂S^\dagger . Call each of these new curves in ∂S^\dagger a *remnant* of γ ; see Figure 3. Consider the surface S' given by *filling* each remnant of γ with a disc; we say that S' is obtained from S by *decomposing along* γ .

We now aim to define similar terminology for truncated pseudomanifolds. Consider a (possibly disconnected) properly embedded surface S in a truncated pseudomanifold \mathcal{P} . Let \mathcal{P}^\dagger denote the pseudomanifold obtained from \mathcal{P} by *cutting along* S — similar to before, this means that we obtain \mathcal{P}^\dagger by removing a small open regular neighbourhood of S from \mathcal{P} . For each two-sided component E of S , we have two




Figure 3: Cutting along an embedded closed curve in a surface. Left: cutting along a two-sided curve yields a pair of remnants. Right: cutting along a one-sided curve yields a single remnant.

new copies of E in $\partial\mathcal{P}^\dagger$; on the other hand, for each one-sided component E of S , we have a single new double cover of E in $\partial\mathcal{P}^\dagger$. Call each of these new pieces in $\partial\mathcal{P}^\dagger$ a *remnant* of S .

For our purposes, it will be useful to have a notion of “decomposing along” S when S is one of the following seven types of (properly embedded) surface:

- A 2-sphere — which means that cutting along S yields a pair of 2-sphere remnants.
- A two-sided annulus — which means that cutting along S yields a pair of annulus remnants.
- A one-sided annulus — which means that cutting along S yields a single annulus remnant.
- A two-sided projective plane — which means that cutting along S yields a pair of projective plane remnants.
- A one-sided projective plane — which means that cutting along S yields a single 2-sphere remnant.
- A two-sided Möbius band — which means that cutting along S yields a pair of Möbius band remnants.
- A one-sided Möbius band — which means that cutting along S yields a single annulus remnant.

Notice that for these types of surface, the remnants are always either 2-spheres, annuli, projective planes or Möbius bands.

Similar to what we did with curves on surfaces, we construct the result of “decomposing along” S by “filling” the remnants of S . To do this for projective plane and Möbius band remnants, we use the following terminology: define an *invalid cone* to be a pseudomanifold given by taking a cone over a projective plane. With this in mind, let S^\dagger denote a remnant of S in \mathcal{P}^\dagger , and suppose S^\dagger is either a 2-sphere, annulus, projective plane or Möbius band. We define the operation of *filling* S^\dagger as follows:

- If S^\dagger is a 2-sphere, then filling means attaching a 3-ball B by identifying S^\dagger with the 2-sphere boundary of B .
- If S^\dagger is an annulus, then filling means attaching a thickened disc $D \times [0, 1]$ by identifying S^\dagger with the annulus $(\partial D) \times [0, 1]$.
- If S^\dagger is a projective plane, then filling means attaching an invalid cone \mathcal{C} by identifying S^\dagger with the projective plane boundary of \mathcal{C} .
- If S^\dagger is a Möbius band, then filling means attaching an invalid cone \mathcal{C} by choosing a small open disc D in $\partial\mathcal{C}$, and identifying S^\dagger with the Möbius band given by $(\partial\mathcal{C}) - D$.

Putting everything together, suppose S is one of the seven types of surface listed above, and let \mathcal{P}' denote the pseudomanifold obtained from \mathcal{P}^\dagger by filling each remnant of S . We say that \mathcal{P}' is obtained from \mathcal{P} by *decomposing along S* .

2.3 Normal surfaces

A *normal surface* in a triangulation \mathcal{T} is a (possibly disconnected) properly embedded surface that

- is disjoint from the vertices of \mathcal{T} ;
- meets the edges and faces of \mathcal{T} transversely; and
- intersects each tetrahedron Δ of \mathcal{T} in a (possibly empty) disjoint union of finitely many discs, called *elementary discs*, where each such disc forms a curvilinear triangle or quadrilateral whose vertices lie on different edges of Δ .

Two normal surfaces are *normally isotopic* if they are related by a *normal isotopy*—that is, an ambient isotopy that preserves each vertex, edge, face and tetrahedron of the triangulation. Up to normal isotopy, the elementary discs in each tetrahedron Δ come in seven possible types,

- four *triangle types*, each of which separates one vertex of Δ from the other three, as shown in Figure 4 (left image); and
- three *quadrilateral types*, each of which separates a pair of opposite edges of Δ , as shown in Figure 4 (middle three images).

Observe that if a tetrahedron contains two elementary quadrilaterals of different types, then these two quadrilaterals will always intersect each other; since normal surfaces are embedded, this means that if a tetrahedron contains quadrilaterals, then these quadrilaterals must all be of the same type.

We call a normal surface *nontrivial* if it includes at least one elementary quadrilateral, and *trivial* otherwise. It is easy to see that trivial normal surfaces always exist, and that every component of such a surface is just a vertex link. The existence of nontrivial normal surfaces is less obvious. In fact, it is possible to prove that many “interesting” embedded surfaces appear as (nontrivial) normal surfaces; we will get a glimpse of why this is the case when we discuss the theory of barriers and normalisation in Section 2.4.




Figure 4: The seven types of elementary disc. Left image: the four triangle types. Middle three images: the three quadrilateral types. Right image: a portion of a normal surface built entirely out of triangles.




Figure 5: A normal surface S can induce parallel cells (two types), wedge cells (three types) and central cells (five types). The faces that lie inside S are shaded red. First row: a parallel triangular cell (left) and a parallel quadrilateral cell (right). All nonshaded faces are bridge faces. Second row: a wedge cell with no bridge faces (left), a wedge cell with one bridge face (middle) and a wedge cell with two bridge faces (right). Third row: four types of nontetrahedron central cell. A tetrahedron is the fifth type of central cell.

A normal surface naturally splits a triangulation into a finer cell decomposition. To describe this idea more precisely, we introduce the following definitions, which are partly based on some terminology used by Jaco and Rubinstein [18, page 91]:

Definitions 3 Let S be a normal surface in a triangulation \mathcal{T} . The surface S divides each tetrahedron Δ of \mathcal{T} into a collection of *induced cells* of the following types:

- *Parallel cells* of two types (see Figure 5 (first row)):
 - *Parallel triangular cells* lie between two parallel triangles of S .
 - *Parallel quadrilateral cells* lie between two parallel quadrilaterals of S .
- *Nonparallel cells* of nine types:
 - *Corner cells* are tetrahedra that lie between a single triangle of S and a single vertex of Δ .

- *Wedge cells* of three types (see Figure 5 (second row)) only occur when S meets Δ in one or more quadrilaterals. In this case, if we ignore any parallel and corner cells in Δ , then the two cells left over are the wedge cells.
- *Central cells* of five types (see Figure 5 (third row)) only occur when S does not meet Δ in any quadrilaterals. In this case, if we ignore any parallel and corner cells in Δ , then the single cell left over is the central cell.

Amongst the faces of these induced cells, we will find it useful to distinguish the *bridge faces*, which are the quadrilateral faces that intersect S precisely in a pair of opposite edges. Note that bridge faces only appear in parallel and wedge cells (see Figure 5 (first and second rows)).

Let \mathcal{P} denote the truncated pseudomanifold of \mathcal{T} , and let \mathcal{P}^\dagger denote the pseudomanifold obtained from \mathcal{P} by cutting along S . The induced cells naturally yield a cell decomposition \mathcal{D} of \mathcal{P} , such that the surface S is given by a union of faces of \mathcal{D} . Moreover, ungluing the faces of \mathcal{D} that lie inside S yields a cell decomposition \mathcal{D}^\dagger of \mathcal{P}^\dagger . We say that the cell decompositions \mathcal{D} and \mathcal{D}^\dagger , and any cell decompositions given by components of \mathcal{D} and \mathcal{D}^\dagger , are *induced* by the normal surface S .

Since a tetrahedron can contain many parallel elementary discs, we could have arbitrarily many parallel cells. However, there are always at most six nonparallel cells per tetrahedron Δ :

- If Δ meets the normal surface in one or more quadrilaterals, then we have no central cells, exactly two wedge cells, and up to four corner cells.
- If Δ does not meet the normal surface in any quadrilaterals, then we have no wedge cells, exactly one central cell, and again up to four corner cells.

We will find this simple observation useful in Section 4.1.

2.4 Barriers and normalisation

We now review the theory of normalisation, which gives a procedure for transforming any properly embedded surface S into a normal surface (not necessarily isotopic to S). We also review the notion of a barrier surface, which gives a tool for “controlling” the result of the normalisation procedure. The material here is essentially an abridged and informal version of Section 3 of [18], focusing only on the details that are necessary for our purposes in this paper.

Throughout Section 2.4, let S , S' and B denote (possibly disconnected) surfaces that are properly embedded in a triangulation \mathcal{T} . Assume that these surfaces are disjoint from the vertices of \mathcal{T} , and transverse to the 2-skeleton of \mathcal{T} .

The idea of the normalisation procedure is to reduce the number of “anomalies” in a surface S until it becomes a normal surface. For instance, for S to be a normal surface, it cannot intersect any tetrahedron Δ in anomalous pieces such as

- a 2-sphere component that is *trivial* in the sense that it lies entirely inside Δ ; or
- a disc component that is *trivial* in the sense that its boundary curve lies entirely inside a single boundary face, and its interior lies entirely inside Δ .

To keep track of these and other anomalous features of S , we use the following measures of “complexity”:

- Define the *weight* $\text{wt}(S)$ to be the number of times S meets the 1-skeleton $\mathcal{T}^{(1)}$: $\text{wt}(S) = |S \cap \mathcal{T}^{(1)}|$. In general, S could meet a tetrahedron Δ of \mathcal{T} in a nonnormal piece that “doubles back” on itself to meet a single edge twice (for example, see Figure 6 (top right and bottom left)); the weight of S gives a proxy for counting the number of such anomalies.

- For each tetrahedron Δ of \mathcal{T} , let

$$x_{\Delta} = \sum_{c \neq S^2} (1 - \chi(c)),$$

where c runs over all components of $S \cap \Delta$ other than 2-spheres. Also define the *local Euler number*

$$\lambda(S) = \sum_{\Delta} x_{\Delta}.$$

Recall that a normal surface must, in particular, meet each tetrahedron of \mathcal{T} in a disjoint union of discs; apart from trivial 2-spheres (which we handle separately), the local Euler number detects any anomalies that violate this requirement.

- Let $\sigma(S)$ denote the number of closed curves in which S intersects the internal faces of \mathcal{T} . A normal surface cannot have any such anomalous curves.
- Let $\tau(S)$ denote the number of components of S that form trivial 2-spheres or trivial discs.

Define the *complexity* of S , denoted $C(S)$, to be the tuple

$$(\text{wt}(S), \lambda(S), \sigma(S), \tau(S)).$$

We will consider S to have smaller complexity than some other surface S' if $C(S)$ occurs before $C(S')$ in the lexicographical ordering. As suggested earlier, normalisation consists of a series of steps, each of which reduces the complexity.

Before we define the steps involved in normalisation, we introduce some useful terminology. Call a disc D an *edge-compression disc* for S if it is embedded so that

- the interior of D lies entirely in the interior of a tetrahedron Δ of \mathcal{T} ; and
- the boundary of D consists of two arcs α and γ that intersect each other only at their endpoints, such that $\alpha = D \cap S$ and γ is a subarc of an edge e of Δ .

Examples of edge-compression discs are shown in Figure 6 (top right and bottom left). Call an edge-compression disc *internal* if it meets an internal edge of \mathcal{T} , and *boundary* if it meets a boundary edge of \mathcal{T} ; notice that a boundary edge-compression disc is, in particular, a ∂ -compression disc for S .

With all the preceding setup in mind, the normalisation procedure proceeds by performing the following *normal moves* on a surface S :




Figure 6: Examples of normal moves. First row, left images: type (1), a compression of a surface (shaded red) along a compression disc (shaded blue) lying entirely inside a tetrahedron. First row, right images: type (2), an isotopy of a surface (shaded red) along an internal edge-compression disc (shaded blue). Second row, left images: type (3), a ∂ -compression of a surface (shaded red) along a boundary edge-compression disc (shaded blue). Second row, right images: type (4), a compression of a surface (shaded red) along a compression disc (shaded blue) lying entirely inside an internal face.

(1) **Compressions along discs that lie entirely in the interior of a tetrahedron** (see Figure 6 (top left)) Each such compression reduces the complexity $C(S)$ because it leaves the weight $\text{wt}(S)$ unchanged and reduces the local Euler number $\lambda(S)$. These compressions can be performed until S meets each tetrahedron of \mathcal{T} in a union of 2-spheres and discs, at which point $\lambda(S) = 0$. We assume for the rest of the normal moves that we have already reduced $\lambda(S)$ to 0 in this way.

(2) **Isotopies along internal edge-compression discs** (see Figure 6 (top right)) Each such isotopy reduces the complexity $C(S)$ because it reduces the weight $\text{wt}(S)$.

(3) **∂ -compressions along boundary edge-compression discs** (see Figure 6 (bottom left)) Like the isotopies in the previous step, each such ∂ -compression reduces the complexity $C(S)$ because it reduces the weight $\text{wt}(S)$. For the remaining two normal moves, we assume that we have performed all possible isotopies and ∂ -compressions along edge-compression discs, which ensures that S meets each tetrahedron Δ of \mathcal{T} in a union of

- elementary discs;
- trivial 2-spheres; and
- discs whose boundary curves lie entirely in the interior of some face of Δ .

(4) **Compressions along discs that lie entirely in the interior of an internal face** (see Figure 6 (bottom right)) Each such compression reduces the complexity $C(S)$ because it leaves $\text{wt}(S)$ and $\lambda(S)$ unchanged, and reduces $\sigma(S)$. After performing these compressions until no more such moves are possible, S meets each tetrahedron of \mathcal{T} in a union of elementary discs, trivial 2-spheres, and trivial discs; we assume for the final normal move that this has already been done.

(5) **Deletion of trivial 2-sphere and disc components** This final “clean-up” step reduces the complexity $C(S)$ because it leaves $\text{wt}(S)$, $\lambda(S)$ and $\sigma(S)$ unchanged, and reduces $\tau(S)$ to zero. At the end of this step, S is a normal surface.

For a complete explanation of why normalisation works as we have claimed, see Section 3.2 of [18]. In general, the normal surface that we obtain might not be isotopic to the original surface, because of the steps where we perform compressions and ∂ -compressions. However, if we assume that the original surface was incompressible and ∂ -incompressible, and also that the ambient 3-manifold is irreducible and ∂ -irreducible, then normalising must produce a normal surface with one component isotopic to the original surface.

We can get even more control over the result of normalisation using the notion of a barrier surface; we now review the aspects of barrier surfaces that we require for our purposes. Given a properly embedded surface B in \mathcal{T} , let \mathcal{N} denote a fixed but arbitrary component of $\mathcal{T} - B$. Call B a *barrier* for \mathcal{N} if any surface S that is properly embedded in \mathcal{N} can actually be normalised inside \mathcal{N} ; that is, the discs along which we compress, isotope or ∂ -compress always lie entirely inside \mathcal{N} , and at every stage the surface S remains properly embedded in \mathcal{N} .

In Theorem 3.2 from [18], Jaco and Rubinstein list a number of examples of barrier surfaces. For our purposes, we will need part (5) of this theorem, which we restate here:

Theorem 4 *Consider a (compact) 3-manifold \mathcal{M} with no 2-sphere boundary components. If \mathcal{M} is closed, let \mathcal{T} be a closed triangulation of \mathcal{M} ; otherwise, if \mathcal{M} is bounded, let \mathcal{T} be an ideal triangulation of \mathcal{M} . Let S be a normal surface in \mathcal{T} , and let A be a subcomplex of the cell decomposition of \mathcal{M} induced by S . The boundary B of a small regular neighbourhood of $S \cup A$ is a barrier surface for any component of $\mathcal{M} - B$ that does not meet $S \cup A$.*

2.5 Crushing via atomic moves

The main purpose of this section is to review the atomic formulation of crushing that was introduced by Burton [3]. We augment this with some new terminology, as this will be useful for our purposes in Section 4. To begin, we state a version of Definition 1 from [3]:

Definitions 5 (crushing procedure) Let S be a normal surface in a triangulation \mathcal{T} . Each of the following operations builds on the previous one:

(1) Cut along S , and let \mathcal{D} denote the resulting induced cell decomposition.




Figure 7: In addition to tetrahedra, a destructible cell decomposition \mathcal{D}^* can contain three other types of 3-cells. To recover a triangulation from \mathcal{D}^* , we need to flatten the nontetrahedron cells. Left: flattening a 3-sided football to an edge. Middle: flattening a 4-sided football to an edge. Right: flattening a triangular purse to a triangular face.

(2) Using the quotient topology, collapse each remnant of S to a point. This turns \mathcal{D} into a new cell decomposition \mathcal{D}' with 3-cells of the following four possible types (see Figure 7):

- 3-sided footballs, which are obtained from corner cells and parallel triangular cells;
- 4-sided footballs, which are obtained from parallel quadrilateral cells;
- triangular purses, which are obtained from wedge cells; and
- tetrahedra, which are obtained from central cells.

We say that \mathcal{D}' is obtained by *nondestructively crushing* S . Also, if a cell decomposition \mathcal{D}^* is built entirely from 3-cells of the four types listed above (even if it was not directly obtained by nondestructive crushing), then we call \mathcal{D}^* a *destructible* cell decomposition.

(3) To recover a triangulation from a destructible cell decomposition \mathcal{D}^* , we first build an intermediate cell complex \mathcal{C}^* by using the quotient topology to flatten

- all 3-sided and 4-sided footballs to edges; and
- all triangular purses to triangular faces.

This is illustrated in Figure 7. Since triangulations are defined only by face gluings between tetrahedra, there are two ways in which \mathcal{C}^* might fail to form a triangulation:

- (a) \mathcal{C}^* could contain vertices, edges and/or triangles that are *isolated*, meaning that they do not belong to any tetrahedra.
- (b) \mathcal{C}^* could contain vertices or edges that are *pinched*, meaning that they include identifications that are independent of any face gluings. In contrast, recall that every vertex of a triangulation is an equivalence class of vertices of tetrahedra, where all vertex identifications arise as consequences of face gluings. Similarly, every edge of a triangulation is given by edge identifications arising solely as consequences of face gluings.




Figure 8: Extracting a triangulation by deleting isolated edges and triangles (highlighted in blue), and separating pinched vertices and edges (highlighted in red).

Thus, as illustrated in Figure 8, we need to perform the following two operations to *extract* a triangulation \mathcal{T}^* from \mathcal{C}^* :

- (a) Delete all isolated vertices, edges and triangles.
- (b) Separate pieces of the cell complex that are only joined together along pinched vertices or edges (thereby ensuring that all vertex and edge identifications arise solely as consequences of face gluings).

We say that \mathcal{T}^* is obtained by *flattening* \mathcal{D}^* . Consider the triangulation \mathcal{T}' obtained by flattening the cell decomposition \mathcal{D}' that results from nondestructively crushing S ; we say that \mathcal{T}' is obtained by (*destructively*) *crushing* S .

It is not too difficult to see what happens if we crush a *trivial* normal surface S in a triangulation \mathcal{T} . Cutting along S yields one central cell per tetrahedron, together with some number of corner and parallel triangular cells. All the corner and parallel cells together form components that do not contain any central cells, so after the nondestructive crushing and flattening steps, these components become isolated edges that do not appear in the final triangulation. For the central cells, observe that nondestructive crushing turns these into tetrahedra that are glued together in the same way as the original triangulation. The upshot is that crushing a trivial normal surface always leaves the triangulation unchanged.

Suppose now that S is a *nontrivial* normal surface in a triangulation \mathcal{T} , and let \mathcal{T}' denote the triangulation obtained by crushing S . As before, each tetrahedron of \mathcal{T}' comes from a central cell in the cell decomposition \mathcal{D} given by cutting along S . However, this time, at least one tetrahedron of \mathcal{T} contains an elementary quadrilateral, which means that not every tetrahedron of \mathcal{T} gives rise to a central cell in \mathcal{D} . Thus, we see that crushing has the following useful feature:

Observation 6 *Let \mathcal{T} be a triangulation, and let \mathcal{T}' denote the triangulation obtained by crushing a nontrivial normal surface in \mathcal{T} . Then $|\mathcal{T}'| < |\mathcal{T}|$.*

The difficulty with crushing a nontrivial normal surface is that this operation could drastically change the topology of our triangulations. In particular, the triangulations before and after crushing could represent




Figure 9: The three atomic moves for flattening a destructible cell decomposition. Left: flattening a triangular pillow. Middle: flattening a bigon pillow. Right: flattening a bigon face.

different 3-manifolds, assuming they even represent 3-manifolds at all. In [18], Jaco and Rubinstein work through this difficulty using a complicated global analysis of their version of the crushing procedure.

In contrast, the formulation of crushing given in Definitions 5 is simpler to work with. This is because the process of flattening a destructible cell decomposition can always be realised by a sequence consisting of atomic moves of three types. The following lemma [3, Lemma 3] gives a precise statement of this idea:

Lemma 7 (crushing lemma) *Let \mathcal{T}^* be the triangulation given by flattening some destructible cell decomposition \mathcal{D}^* . Then \mathcal{T}^* can be obtained from \mathcal{D}^* by performing a sequence of zero or more of the following **atomic moves** (see Figure 9), one at a time, in some order:*

- flattening a **triangular pillow** to a triangular face;
- flattening a **bigon pillow** to a bigon face; and
- flattening a bigon face to an edge.

Since our cell decompositions are defined only by face gluings between 3-cells, after each atomic move we implicitly **extract** a cell decomposition by

- deleting all **isolated** vertices, edges, bigons and triangles that do not belong to any 3-cells; and
- separating pieces of the cell complex that are only joined together along **pinched** vertices or edges.

As part of the proof of the crushing lemma, Burton showed [3] that if we are careful about the order in which we perform the atomic moves, then we only ever encounter cell decompositions with 3-cells of the following seven types:

- 3-sided footballs;
- 4-sided footballs;
- triangular purses;
- tetrahedra;
- triangular pillows;
- bigon pillows; and
- *bigon pyramids* (see Figure 10).




Figure 10: A bigon pyramid.

The crushing lemma allows us to understand the topological effects of crushing by examining atomic moves one at a time. In particular, Burton proved the following result [3, Lemma 4] (which, among other things, paved the way for a practical algorithm for nonorientable prime decomposition [2; 3]):

Lemma 8 *Let \mathcal{D}_0 be a valid cell decomposition with no ideal vertices. If the underlying 3-manifold \mathcal{M}_0 contains no two-sided projective planes, then performing one of the atomic moves of Lemma 7 will yield a (valid) cell decomposition of a 3-manifold \mathcal{M}_1 such that one of the following holds:*

- $\mathcal{M}_0 = \mathcal{M}_1$.
- We flattened a triangular pillow, and \mathcal{M}_1 is obtained from \mathcal{M}_0 by deleting a single component \mathcal{C} , where \mathcal{C} is either a 3-ball, a 3-sphere or a copy of the lens space $L_{3,1}$.
- We flattened a bigon pillow, and \mathcal{M}_1 is obtained from \mathcal{M}_0 by deleting a single component \mathcal{C} , where \mathcal{C} is either a 3-ball, a 3-sphere, or a copy of real projective space $\mathbb{R}P^3$.
- We flattened a bigon face, and \mathcal{M}_1 is related to \mathcal{M}_0 in one of the following ways:
 - (i) \mathcal{M}_1 is obtained by cutting along a properly embedded disc in \mathcal{M}_0 .
 - (ii) \mathcal{M}_1 is obtained by filling a boundary 2-sphere of \mathcal{M}_0 with a 3-ball.
 - (iii) \mathcal{M}_1 is obtained by decomposing along an embedded 2-sphere in \mathcal{M}_0 .
 - (iv) $\mathcal{M}_0 = \mathcal{M}_1 \# \mathbb{R}P^3$ —that is, \mathcal{M}_1 removes a single $\mathbb{R}P^3$ summand from the connected sum decomposition of \mathcal{M}_0 .

One of our main goals in this paper is to extend Lemma 8 to cell decompositions that may be invalid and may have ideal vertices; this, in particular, allows us to study the topological effects of crushing a closed surface of positive genus, since nondestructively crushing such a surface produces a cell decomposition with ideal vertices. To do this, we will find it helpful to have “flattening maps” that keep track of how the points in a cell decomposition are affected by an atomic move. Although an atomic move “looks like” a quotient operation, the corresponding quotient map does not account for the implicit operation of extracting a cell decomposition, so a little care is required to define “flattening maps” appropriately:

Definitions 9 Let \mathcal{D}_1 be a cell decomposition obtained by performing a single atomic move on some cell decomposition \mathcal{D}_0 . In Lemma 7, each atomic move implicitly finishes with the operation of extracting a cell decomposition; consider the intermediate cell complex \mathcal{C} that we obtain by performing the atomic move *without* subsequently extracting a cell decomposition. Note that \mathcal{C} is obtained as a quotient of \mathcal{D}_0 , so we have a quotient map $q: \mathcal{D}_0 \rightarrow \mathcal{C}$.

We use q to construct a map $\hat{\varphi}_0: \mathcal{D}_0 \rightarrow 2^{\mathcal{D}_1}$ (here, 2^X denotes the *power set* of a set X) that acts on points p in \mathcal{D}_0 as follows:

- If $q(p)$ is part of an isolated vertex, edge, bigon or triangle — which means that $q(p)$ is deleted when we extract a cell decomposition — then take $\hat{\varphi}_0(p)$ to be the empty set.
- If $q(p)$ is part of a pinched edge or vertex — which means that $q(p)$ gets separated into multiple points when we extract a cell decomposition — then take $\hat{\varphi}_0(p)$ to be the set of points in \mathcal{D}_1 that originate from $q(p)$.
- Otherwise, $q(p)$ remains untouched when we extract a cell decomposition, in which case we take $\hat{\varphi}_0(p) = \{q(p)\}$ (here, by an abuse of notation, we are viewing $q(p)$ as a point in \mathcal{D}_1).

Intuitively, $\hat{\varphi}_0$ keeps track of how points in \mathcal{D}_0 are affected when we perform an atomic move.

Observe that the nonempty sets in the image of $\hat{\varphi}_0$ give a partition of the points in \mathcal{D}_1 . Thus, we can construct a map $\hat{\varphi}_1: \mathcal{D}_1 \rightarrow 2^{\mathcal{D}_0}$ as follows: for each point p in \mathcal{D}_1 , let U be the (unique) set in the image of $\hat{\varphi}_0$ that contains p , and define $\hat{\varphi}_1(p)$ to be the set $\hat{\varphi}_0^{-1}(U)$. Intuitively, $\hat{\varphi}_1$ keeps track of how points in \mathcal{D}_1 would be affected if we perform an atomic move in reverse.

For each $i \in \{0, 1\}$, define a map $\varphi_i: 2^{\mathcal{D}_i} \rightarrow 2^{\mathcal{D}_{1-i}}$ that sends any subset S of \mathcal{D}_i (when we actually use the ideas defined here, S will usually be a vertex, edge, face or 3-cell of \mathcal{D}_i) to the set

$$\bigcup_{p \in S} \hat{\varphi}_i(p) \subseteq \mathcal{D}_{1-i}.$$

We call φ_0 the *flattening map* associated to the atomic move, and φ_1 the *inverse flattening map* (although, strictly speaking, these maps are not actually inverses of each other).

3 Atomic moves on cell decompositions with ideal vertices

Let S be a normal surface in a triangulation \mathcal{T} . When S is either a 2-sphere or a disc, nondestructively crushing S creates new vertices whose links are either 2-spheres or discs. Thus, if the vertices of \mathcal{T} are all either internal or boundary, then the topological effect of destructively crushing S only depends on how atomic moves affect cell decompositions whose vertices are all either internal or boundary; this was the motivation for Lemma 8 in [3].

Our main goal in this section is to extend this atomic approach to crushing beyond the case where S is a 2-sphere or disc. This requires us to study atomic moves on cell decompositions that are allowed to have ideal or invalid vertices. A similarly general understanding of atomic moves is necessary to

understand crushing if we allow the initial triangulation \mathcal{T} to have ideal or invalid vertices. Moreover, when \mathcal{T} triangulates a nonorientable 3-manifold, it turns out to be possible for an atomic move to create an invalid *edge*. The upshot is that, for a completely general analysis of atomic moves, we should not restrict the links of the vertices involved, and we should not exclude the possibility of invalid edges.

Of the three atomic moves, flattening a triangular pillow and flattening a bigon pillow are relatively straightforward to understand. We study these two atomic moves in full generality in Section 3.1.

We then devote Section 3.2 to understanding the topological effects of flattening a bigon face. In contrast to the other two atomic moves, there would be a tediously large number of cases to consider if we wanted to give a complete analysis. Thus, for the sake of brevity and clarity, we will focus mainly on flattening bigon faces in valid cell decompositions whose vertices are all either internal or ideal. This is sufficient to understand the effects of crushing S if the following conditions are satisfied:

- S is a closed surface.
- \mathcal{T} is valid and has no boundary vertices.
- The truncated 3-manifold of \mathcal{T} contains no two-sided properly embedded projective planes or Möbius bands (which is true, in particular, for all orientable 3-manifolds).

We leave to whomever may require them in future work the details of the cases we did not cover.

3.1 Flattening triangular and bigon pillows

Lemma 10 (flattening triangular pillows) *Let \mathcal{D}_0 be a (possibly invalid) cell decomposition, and let \mathcal{D}_1 be the cell decomposition obtained by flattening a triangular pillow F in \mathcal{D}_0 . One of the following holds:*

- (a) *The two triangular faces of F are not identified and not both boundary, in which case the truncated pseudomanifolds of \mathcal{D}_0 and \mathcal{D}_1 are homeomorphic.*
- (b) *F forms a (bounded) cell decomposition of a 3-ball, in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this 3-ball component.*
- (c) *F forms a (closed) cell decomposition of either S^3 (the 3-sphere) or $L_{3,1}$ (a lens space), in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this closed component.*
- (d) *F forms a two-vertex component \mathcal{C} of \mathcal{D}_0 with exactly one invalid edge e ; one of the vertices is incident to e and has 2-sphere link, while the other vertex is not incident to e and has projective plane link. In this case, \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this invalid component \mathcal{C} .*

Proof Throughout this proof, let t and t' denote the triangular faces that bound the triangular pillow F . We have several cases to consider, depending on how t and t' are glued to other faces of \mathcal{D}_0 (if at all).

First, suppose t and t' are not glued to each other. In this case, the triangular pillow F forms a 3-ball. If t and t' are not both boundary, then this ball lives inside some larger component of \mathcal{D}_0 , and flattening F does not change the truncated pseudomanifold; this corresponds to case (a). On the other hand, if t and t'




Figure 11: Cases where the two faces of a triangular pillow are glued to each other. Left: a triangular pillow that forms a (closed) cell decomposition of S^3 . Middle: a triangular pillow that forms a (closed) cell decomposition of $L_{3,1}$. Right: a triangular pillow that forms a component with an invalid edge (highlighted red).

are both boundary, then F forms the entirety of a 3-ball component of \mathcal{D}_0 , and flattening F deletes this 3-ball component; this corresponds to case (b).

With that out of the way, suppose t and t' are glued to each other. Up to symmetry, there are two possibilities for an orientation-reversing gluing:

- If t and t' are glued without a twist, then F forms a cell decomposition of S^3 (see Figure 11 (left)).
- If t and t' are glued with a twist, then F forms a cell decomposition of $L_{3,1}$ (see Figure 11 (middle)).

In either case, we see that F forms a closed component of \mathcal{D}_0 . Moreover, flattening F has the effect of deleting this closed component. This corresponds to case (c).

For an orientation-preserving gluing of t and t' , there is only one possibility up to symmetry. With this gluing, F forms a two-vertex component \mathcal{C} of \mathcal{D}_0 with exactly one invalid edge e (see Figure 11 (right)). One of the vertices of \mathcal{C} is given by identifying the two endpoints of e , and has 2-sphere link. The other vertex of \mathcal{C} is given by the vertex of F disjoint from e , and has projective plane link. This corresponds to case (d). \square

Lemma 11 (flattening bigon pillows) *Let \mathcal{D}_0 be a (possibly invalid) cell decomposition, and let \mathcal{D}_1 be the cell decomposition obtained by flattening a bigon pillow F in \mathcal{D}_0 . One of the following holds:*

- The two bigon faces of F are not identified and not both boundary, in which case the truncated pseudomanifolds of \mathcal{D}_0 and \mathcal{D}_1 are homeomorphic.*
- F forms a (bounded) cell decomposition of a 3-ball, in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this 3-ball component.*
- F forms a (closed) cell decomposition of either S^3 (the 3-sphere) or $\mathbb{R}P^3$ (real projective space), in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this closed component.*
- F forms an ideal cell decomposition of $\mathbb{R}P^2 \times [0, 1]$, in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this ideal component.*
- F forms a one-vertex component of \mathcal{D}_0 with exactly two invalid edges, in which case \mathcal{D}_1 is obtained from \mathcal{D}_0 by deleting this invalid component.*




Figure 12: Cases where the two faces of a bigon pillow are glued to each other. First image: a bigon pillow that forms a (closed) cell decomposition of S^3 . Second image: a bigon pillow that forms a (closed) cell decomposition of $\mathbb{R}P^3$. Third image: a bigon pillow that forms an ideal cell decomposition of $\mathbb{R}P^2 \times [0, 1]$. Fourth image: a bigon pillow that forms a component with two invalid edges.

Proof Throughout this proof, let b and b' denote the bigon faces that bound the bigon pillow F . We have several cases to consider, depending on how b and b' are glued to other faces of \mathcal{D}_0 (if at all).

First, suppose b and b' are not glued to each other. In this case, the bigon pillow F forms a 3-ball. If b and b' are not both boundary, then this ball lives inside some larger component of \mathcal{D}_0 , and flattening F does not change the truncated pseudomanifold; this corresponds to case (a). On the other hand, if b and b' are both boundary, then F forms the entirety of a 3-ball component of \mathcal{D}_0 , and flattening F deletes this 3-ball component; this corresponds to case (b).

With that out of the way, suppose b and b' are glued to each other. There are two possibilities for an orientation-reversing gluing:

- If b and b' are glued without a twist, then F forms a cell decomposition of S^3 (see Figure 12 (first image)).
- If b and b' are glued with a twist, then F forms a cell decomposition of $\mathbb{R}P^3$ (see Figure 12 (second image)).

In either case, we see that F forms a closed component of \mathcal{D}_0 . Moreover, flattening F has the effect of deleting this closed component. This corresponds to case (c).

Finally, for an orientation-preserving gluing, we again have two possibilities:

- One of these gluings causes the two edges of F to be identified together, and does not create invalid edges. In this case, F forms an ideal cell decomposition of $\mathbb{R}P^2 \times [0, 1]$ (see Figure 12 (third image)), and flattening F has the effect of deleting this ideal component. This corresponds to case (d).
- The other orientation-preserving gluing causes each edge of F to be identified with itself in reverse, so that F forms a one-vertex component of \mathcal{D}_0 with exactly two invalid edges (see Figure 12 (fourth image)). Flattening F has the effect of deleting this invalid component. This corresponds to case (e). □

3.2 Flattening bigon faces

We now study the effect of flattening a bigon face F . As mentioned earlier, our main goal is to give a detailed analysis in the case where F belongs to a valid cell decomposition whose vertices are all either internal or ideal. Our arguments only rely on the following properties:

- (a) F is an internal face.
- (b) Each edge incident to F is internal.
- (c) Each vertex incident to F is either internal or ideal.

Provided these properties hold, our analysis will apply even if F belongs to an invalid cell decomposition.

With this in mind, we assume throughout Section 3.2 that F is an internal bigon face. However, for the sake of generality, we do *not* assume that conditions (b) and (c) are satisfied; instead, we carefully enumerate the cases where these conditions hold, and for each such case we subsequently give a detailed description of the effect of flattening F .

We present our analysis in four parts. First, in Section 3.2.1, we give a brief user guide for the reader seeking to apply our results. Then, in Section 3.2.2, we make some preliminary observations by examining how flattening F interacts with the vertices incident to F . Finally, we partition the main analysis into two broad cases that we handle separately in Sections 3.2.3 and 3.2.4.

Before we dive into the details, we make some general comments about our proof strategy, and we introduce some notation and terminology to support this. One of the key ideas throughout our analysis is that, under our assumption that F is internal, flattening F has the side-effect that we lose the face-gluing along F . This means that flattening F has the same result as the following two-step procedure:

- (1) Undo the gluing along F , which yields two new boundary bigons F_0^\dagger and F_1^\dagger .
- (2) Flatten F_0^\dagger and F_1^\dagger ; since these are boundary faces, flattening these faces has no side-effects (unlike the original face F).

We will see that step (1) often corresponds to cutting along a properly embedded surface S , and that step (2) often corresponds to filling the remnants of S , so that the overall topological effect of flattening F is often to decompose along S (as defined in Section 2.2). With this in mind, we introduce the following notation (also see Figure 13), which we will use throughout the rest of this section:

Notation A As above, let F be an internal bigon face in a (possibly invalid) cell decomposition \mathcal{D}_0 . Let \mathcal{D}_1 be the cell decomposition obtained by flattening F , and let φ denote the associated flattening map. For each $i \in \{0, 1\}$, let V_i denote the set of ideal and invalid vertices in \mathcal{D}_i , and let \mathcal{P}_i denote the truncated pseudomanifold of \mathcal{D}_i ; recall that \mathcal{P}_i is obtained from \mathcal{D}_i by truncating the vertices in V_i .

As in step (1) above, let F_0^\dagger and F_1^\dagger denote the two new boundary bigons that we obtain after undoing the face-gluing along F , and let \mathcal{D}^\dagger denote the cell decomposition that we obtain after undoing this gluing.




Figure 13: The protagonists introduced in Notation A.

Let $g : \mathcal{D}^\dagger \rightarrow \mathcal{D}_0$ be the quotient map associated to the operation of regluing F_0^\dagger and F_1^\dagger to recover the original bigon face F .

We also introduce the following terminology, which will be useful not only for flattening bigon faces, but also for proving our main theorem in Section 4:

Definitions 12 Let \mathcal{D} be a (possibly invalid) cell decomposition, and let \mathcal{P} be the truncated pseudomanifold of \mathcal{D} . Since \mathcal{P} is obtained from \mathcal{D} by truncating the ideal and invalid vertices of \mathcal{D} , we can view \mathcal{P} as a subset of \mathcal{D} ; using this viewpoint, the *truncated bigon* associated to a bigon face B in \mathcal{D} is given by $B \cap \mathcal{P}$ (see Figure 14); we will see that in many cases, the truncated bigon forms a properly embedded surface in \mathcal{P} .

For some positive integer n , consider an embedded curve γ in \mathcal{D} that

- starts at the midpoint of an edge e_0 ;
- ends at the midpoint of an edge e_n (possibly equal to e_0 , to allow for the possibility that γ is a closed curve); and
- passes through the midpoints of a sequence e_0, \dots, e_n of edges, such that for each $i \in \{0, \dots, n-1\}$, the edges e_i and e_{i+1} together bound a single bigon face B_i that is bisected by γ .




Figure 14: There are three possibilities for the truncated bigon associated to a bigon face B . The portion of B that lives outside the truncated bigon is indicated by dashed edges and faint shading. Left: the case where both vertices of B are internal or boundary (either forming two distinct vertices, or identified to form a single such vertex). Middle: the case where one vertex of B is ideal or invalid, while the other is internal or boundary. Right: the case where both vertices of B are ideal or invalid (either forming two distinct vertices, or identified to form a single such vertex).




Figure 15: Three bigon faces bisected by a curve (drawn as a dashed red line) passing through midpoints of edges. These bigon faces together form a bigon path of length 3.

This is illustrated in Figure 15. We call the union $\mathcal{U} := B_0 \cup \dots \cup B_{n-1}$ a *bigon path* of length n in \mathcal{D}_0 , and we call the edges e_0 and e_n the *ends* of \mathcal{U} . If the bigon faces B_0, \dots, B_{n-1} are all boundary, then we say that \mathcal{U} is *boundary*; similarly, if B_0, \dots, B_{n-1} are all internal, then we say that \mathcal{U} is *internal*.

For each $i \in \{0, \dots, n-1\}$, let S_i denote the truncated bigon associated to B_i . We call the union $S_0 \cup \dots \cup S_{n-1}$ the *truncated bigon path* associated to \mathcal{U} ; similar to individual truncated bigons, truncated bigon paths often form properly embedded surfaces in \mathcal{P} .

We mentioned earlier that we divide our analysis into two cases that we handle separately in Sections 3.2.3 and 3.2.4. We now have the terminology to describe these two cases. Specifically, after ungluing F , the two new boundary bigons F_0^\dagger and F_1^\dagger could either

- share at least one common edge, so that they together form a single boundary bigon path of length two; or
- have no common edges, in which case they form two separate boundary bigon paths of length one.

There is no technical reason for dividing our analysis according to these two cases; we make this choice simply to help organise our analysis into smaller, more manageable pieces.

3.2.1 User guide We split the effects of flattening F into several parts:

- The effect on the vertices incident to F is described in Claim B.
- The effect on the edges incident to F is described in Claims D and E.
- The effect on the truncated pseudomanifold \mathcal{P}_0 is described in Claims D.1, D.2 and D.4, and in Claims E.1 and E.2.

Claims D, D.1, D.2 and D.4 all deal with the case where F_0^\dagger and F_1^\dagger form a single boundary bigon path, so they can be found in Section 3.2.3; on the other hand, Claims E, E.1 and E.2 all deal with the case where F_0^\dagger and F_1^\dagger form two separate boundary bigon paths, so they can be found in Section 3.2.4. The intended way to use all these results is to begin by referring to Claims D and E, as these two overarching claims will indicate which of the other claims are relevant for any given application.

The only other result that we prove is Claim C. This is a useful tool for our proofs in Sections 3.2.3 and 3.2.4, but it is otherwise not a crucial part of our description of the effects of flattening F . Having




Figure 16: A simple example of a v -cone (orange) over a subset S of a vertex link L . Here, L is a 2-sphere, and S is an embedded closed curve (red) in L .

said this, Claim C might be useful for the reader seeking to extend our analysis of flattening F to the cases that we do not study in detail.

3.2.2 Interaction with vertices Let v denote a vertex incident to F , and consider a small regular neighbourhood N of v . To describe how flattening F interacts with v , we will find it useful to view N as a cone over the link L of v ; that is, we view N as a union of lines, with each point in L being joined to v by one such line, and with any two such lines intersecting only at the vertex v . Under this viewpoint, any subset S of L defines a subset C_S of N consisting of the lines joining S to v (for example, see Figure 16); we will call C_S the v -cone over S . We will use this notion of v -cones to prove two claims:

- In Claim B, we describe how flattening F changes the vertex v .
- In Claim C, we give conditions under which we can, in some sense, “push F away from v ”; we will give a more precise formulation of this later. Roughly, the purpose of this is that it gives us a unified method to deal with some of the more inconvenient ways in which flattening F interacts with v ; this will become clearer when we see Claim C in action in Sections 3.2.3 and 3.2.4.

Since we are interested in the effect of flattening F , we devote particular attention to the subset of L given by $F \cap L$. Assuming that each edge incident to F is internal, we have the following possibilities (see Figure 17):




Figure 17: The three possibilities for a component γ (dotted red) of $F \cap L$. In each case, the v -cone (orange) over γ forms part of $F \cap N$. Left: when F forms a disc, γ forms an arc. Middle: when F forms a 2-sphere, γ forms a closed curve. Right: when F forms a projective plane, γ forms a closed curve.

- Suppose the edges of F are not identified, so that F forms a disc. In this case, $F \cap L$ consists of either one or two arcs:
 - If the two vertices of F form two distinct vertices in \mathcal{D}_0 , then v is one of these two vertices, and $F \cap L$ consists of a single arc in L .
 - If the two vertices of F are identified to form a single vertex in \mathcal{D}_0 , then $F \cap L$ consists of two disjoint arcs in L .
- Suppose the edges of F are identified to form a single edge e , and suppose this identification causes F to form a 2-sphere. In this case, $F \cap L$ consists of either one or two closed curves:
 - If the two vertices of F form two distinct vertices in \mathcal{D}_0 , then v is one of these two vertices, and $F \cap L$ consists of a single closed curve in L .
 - If the two vertices of F are identified to form a single vertex in \mathcal{D}_0 , then $F \cap L$ consists of two disjoint closed curves in L .
- Suppose the edges of F are identified to form a single edge e , and suppose this identification causes F to form a projective plane. In this case, the two vertices of F are identified to form a single vertex in \mathcal{D}_0 , and $F \cap L$ consists of a single closed curve in L .

In each of the above cases, observe that the v -cone over $F \cap L$ coincides exactly with $F \cap N$. Intuitively, this means that flattening F changes N in a way that “respects the cone structure”. This idea allows us to give a fairly straightforward description of how flattening F affects the vertex v :

Claim B *Assume that each edge incident to F is internal. Let v be a vertex incident to F , and let L denote the link of v . We have the following possibilities:*

- (a) *If the edges of F are not identified, then $\varphi(v)$ consists of a single vertex whose link is topologically equivalent to L .*
- (b) *If the edges of F are identified, then $F \cap L$ consists of either one or two closed curves in L . Let L'_0, \dots, L'_{k-1} denote the components of the surface obtained by decomposing L along the curves in $F \cap L$; there could be up to three such components (that is, we have $1 \leq k \leq 3$). After flattening F , the image $\varphi(v)$ consists of k vertices v'_0, \dots, v'_{k-1} such that for each $i \in \{0, \dots, k-1\}$, the vertex v'_i has link L'_i .*

Proof As above, let N denote a small regular neighbourhood of v , and view N as the v -cone over the vertex link L . We first consider the case where the two edges of F are not identified, so that $F \cap L$ consists of either one or two arcs in L . For each such arc γ , flattening F has the effect of collapsing γ to a single point p_γ , which leaves L topologically unchanged; the corresponding effect on N is to flatten the v -cone over γ to a single line joining p_γ to v , which means that we can continue to view N as the v -cone over L . As a result, we see that $\varphi(v)$ consists of a single vertex whose link is topologically equivalent to L . This proves case (a).

In the case where the edges of F are identified, recall that $F \cap L$ consists of either one or two closed curves in L . This time, we study the effect of flattening F by first ungluing F , and then flattening F_0^\dagger and F_1^\dagger :

(1) Ungluing F changes L by cutting along the curves in $F \cap L$. Since $F \cap L$ could have up to two components, each of which could possibly form a *separating* curve in L , cutting along $F \cap L$ could split L into up to three components; let k be the number of such components, and denote these components by L'_0, \dots, L'_{k-1} . The corresponding change to N is to cut along the v -cone over $F \cap L$, which has the following effects:

- v gets split into k new vertices v'_0, \dots, v'_{k-1} .
- N gets split into k components N'_0, \dots, N'_{k-1} such that for each $i \in \{0, \dots, k-1\}$, N'_i forms the v'_i -cone over L'_i .

At this stage, the vertices v'_i and the surfaces L'_i are not yet the same as the corresponding objects that appear in the claim statement. However, this situation will soon be fixed when we flatten F_0^\dagger and F_1^\dagger , which will have the consequence of modifying the vertices v'_i and the surfaces L'_i ; to keep notation as simple as possible, we will continue to use the same notation for these objects even after they are modified. Likewise, the neighbourhoods N'_i will also be modified, but we will continue to denote the modified neighbourhoods by N'_i .

(2) For each $i \in \{0, \dots, k-1\}$, flattening $F^{\dagger 0}$ and $F^{\dagger 1}$ modifies L'_i by collapsing each remnant γ of $F \cap L$ to a single point p_γ , which is topologically equivalent to filling γ with a disc; the corresponding effect on N'_i is to flatten the v'_i -cone over γ to a single line joining p_γ to v'_i , which means that we can continue to view N'_i as the v'_i -cone over L'_i . The end result of all this is that $\varphi(v)$ consists precisely of the (now modified) vertices v'_0, \dots, v'_{k-1} , and that the links of these vertices are given by the (now modified) surfaces L'_0, \dots, L'_{k-1} , respectively. We also note that, topologically, L'_0, \dots, L'_{k-1} form the components of the surface obtained by decomposing L along $F \cap L$.

This proves case (b). □

We now turn our attention to the idea of “pushing F away from v ”; we will build up to this idea in a slightly roundabout way. Assume that the two edges of F are identified to form a single internal edge. As observed earlier, this means that $F \cap L$ consists of either one or two closed curves in L . Suppose a component γ of $F \cap L$ forms a separating curve that bounds a disc in L . Under these conditions, rather than beginning the process of flattening F by ungluing the entirety of F all at once, we will find it useful to follow a more fine-grained procedure for flattening F (see Figure 18):

(i) Cut along the subset of F given by the v -cone C_γ over γ . Since γ is a separating curve in L , this has the following effects:

- v gets split into two new vertices v'_0 and v'_1 .
- C_γ gets split into two remnants C_0^\dagger and C_1^\dagger such that for each $i \in \{0, 1\}$, C_i^\dagger forms the v'_i -cone over γ .




Figure 18: Flattening C_γ (orange) by first cutting along it, and then flattening the two remnants (orange and pink). This collapses γ (dotted red) to a point p_γ that we temporarily view as a vertex; it also causes $F - C_\gamma$ (blue) to become a new bigon F' . (These illustrations do not accurately reflect how everything is embedded in \mathcal{D}_0 .) Left images: the case where F forms a 2-sphere. In this case, the new bigon F' also forms a 2-sphere, and one of its two vertices is given by p_γ . Right images: the case where F forms a projective plane. In this case, the new bigon F' also forms a projective plane, and its two vertices are identified together to form p_γ .

We postpone ungluing or cutting along $F - C_\gamma$ (ie the rest of F) until later; as a result, the curve γ does not yet fall apart into two pieces because it is still “held together” by $F - C_\gamma$.

(ii) Flatten γ to a single point p_γ , and for each $i \in \{0, 1\}$ flatten the remnant C_i^\dagger to a single line α_i joining p_γ to v'_i . Intuitively, the lines α_0 and α_1 will eventually form segments of the edges in $\varphi(F)$. Viewing p_γ as a temporary vertex, this step causes $F - C_\gamma$ to become a new bigon F' .

(iii) Treating F' as if it were an internal bigon face in a cell decomposition, flatten F' by first cutting along it, and then flattening each of its remnants. Since γ was originally a separating curve in L , this step splits the temporary vertex p_γ into a pair of new points. We no longer view these new points as vertices (which is why we called p_γ a “temporary” vertex); instead, each of these new points will occur in the interior of an edge in $\varphi(F)$.

Together, we refer to steps (i) and (ii) as the operation of *flattening* C_γ . We emphasise that after performing this operation, the intermediate object that we obtain from \mathcal{D}_0 might not be a cell decomposition anymore; however, this problem is only temporary, since we will recover a cell decomposition once we complete step (iii).

In describing the operation of flattening C_γ , we used the assumption that γ is a separating curve, but made no mention of the assumption that γ bounds a disc in L . The purpose of the latter assumption is that it allows us to show that flattening C_γ leaves the topology of \mathcal{D}_0 unchanged, which means that flattening F is topologically equivalent to the operation of flattening the new bigon F' . Moreover, we can view F' as the bigon that results from “pushing F away from v ”:

Claim C Assume that each edge incident to F is internal. Let v be a vertex incident to F , and let L denote the link of v . Suppose a component γ of $F \cap L$ forms a separating curve that bounds a disc E in L , and suppose the interior of E is disjoint from F . Consider the subset of F given by the v -cone C_γ

over γ . Flattening C_γ creates a new internal vertex without changing the topology of \mathcal{D}_0 , and reduces the operation of flattening F to the operation of flattening a new bigon F' that is

- topologically equivalent to a bigon obtained from F by an isotopy that takes C_γ to E ; and
- incident to a temporary **internal** vertex given by the point p_γ that results from flattening γ .

Proof To see how flattening C_γ affects \mathcal{D}_0 topologically, we first claim that since γ bounds a disc E in L , we can find a 3-ball \mathcal{B} such that $C_\gamma - v$ lies in the interior of \mathcal{B} . If v were internal or boundary, we could simply take \mathcal{B} to be a small regular neighbourhood of the v -cone C_E over E . However, to account for the possibility that v is ideal or invalid, we instead construct \mathcal{B} as follows:

- Consider a regular neighbourhood N^* of v that is “large enough” so that C_E lies entirely in the interior of N^* .
- Slightly isotope the disc E so that it lies in the frontier of N^* , and then enlarge this disc slightly so that the v -cone over this disc forms the desired 3-ball \mathcal{B} (see Figure 19 (left images)).

The operation of flattening C_γ leaves everything outside of \mathcal{B} untouched, so we just need to understand how this operation affects \mathcal{B} topologically. We follow steps (i) and (ii) from above:

- As illustrated in Figure 19 (middle images), cutting along C_γ has the following effects:
 - The vertex v gets split into two new vertices v'_0 and v'_1 . One of the new vertices, say v'_0 , has link given by L minus a disc. The other new vertex v'_1 has link given by the disc E , as illustrated in Figure 19 (bottom middle).
 - The v -cone C_γ gets split into two remnants C_0^\dagger and C_1^\dagger such that for each $i \in \{0, 1\}$, C_i^\dagger forms the v'_i -cone over γ . These two remnants bound a newly created void inside our 3-ball \mathcal{B} .
- As illustrated in Figure 19 (right images) flattening $C^{\dagger 0}$ and $C^{\dagger 1}$ has the following effects:
 - The link of v'_0 gets “closed up” so that it becomes topologically equivalent to L . Thus, we can equate v'_0 with the original vertex v .
 - The link of v'_1 gets “closed up” to become the 2-sphere corresponding to $E/\partial E$, as illustrated in Figure 19 (bottom right). Thus, we can view v'_1 as a newly created internal vertex.
 - The void that we created in the previous step gets filled in, so that we once again have a 3-ball \mathcal{B}' .
 - The curve γ gets flattened to a single point p_γ that we temporarily view as a vertex. (Recall that p_γ is only a “temporary” vertex because after performing step (iii), p_γ gets split into two new points that we no longer view as vertices.) Since p_γ lies in the interior of the 3-ball \mathcal{B}' , we can think of p_γ as an *internal* vertex.

Topologically, all we have done is replaced the 3-ball \mathcal{B} with another 3-ball \mathcal{B}' , so we have not changed the topology of \mathcal{D}_0 .

To finish this proof, consider $F - C_\gamma$; this is the part of F that is being left “unflattened”. Recall that after flattening C_γ , this unflattened part of F becomes a new bigon F' , and that the operation of flattening F is




Figure 19: Since γ (red, top images) bounds a disc in L , flattening C_γ has no topological effect on \mathcal{D}_0 . The intersection of the 3-ball \mathcal{B} (grey) with the “unflattened” part of F is shaded blue. Top left: the 3-ball \mathcal{B} (grey) contains the entirety of the v -cone C_γ (orange), as well as a portion of $F - C_\gamma$ (blue). Top middle: cutting along the v -cone C_γ yields two remnants (orange and pink), and creates a void. Top right: flattening the remnants of C_γ fills the void back in, so we recover a 3-ball. Bottom row: schematic cross-sections of the 3-dimensional pictures in the top row. Here we also include the disc E (purple), which is not shown in the 3-dimensional pictures; the vertex v'_1 is repositioned slightly to accommodate this inclusion.

reduced to the operation of flattening F' . Topologically, observe that F' is equivalent to a bigon obtained from F by an isotopy that replaces C_γ with the disc E ; this can be seen by equating the vertices v and v'_0 , and then comparing how F intersects the grey 3-ball in Figure 19 (top left) with how F' intersects the grey 3-ball in Figure 19 (top right). □

3.2.3 The case where ungluing gives a single boundary bigon path We are now ready to present the main analysis of the effect of flattening F . We first consider the case where F_0^\dagger and F_1^\dagger together form a single boundary bigon path \mathcal{F}^\dagger . Depending on how the ends of \mathcal{F}^\dagger are identified (if at all), \mathcal{F}^\dagger could form a 2-sphere, projective plane or disc in the boundary of \mathcal{D}^\dagger . We refine this list of cases as follows:

Claim D *If F_0^\dagger and F_1^\dagger together form a single boundary bigon path \mathcal{F}^\dagger , then one of the following holds:*

- *The boundary bigon path \mathcal{F}^\dagger forms a 2-sphere, in which case the result $\varphi(F)$ of flattening F is a single internal edge in \mathcal{D}_1 . Letting e_0^\dagger and e_1^\dagger denote the edges incident to \mathcal{F}^\dagger , we have four cases depending on the behaviour of the gluing map g (see Figure 20):*




Figure 20: The four ways to glue F_0^\dagger and F_1^\dagger together when \mathcal{F}^\dagger forms a 2-sphere. They are gluings that result in F forming a disc (first image), F forming a projective plane (second image), F being incident to two invalid edges (third image) and F forming a 2-sphere (fourth image).

- (1) The map g realises an orientation-reversing gluing of F_0^\dagger and F_1^\dagger such that for each $i \in \{0, 1\}$, the edge e_i^\dagger gets identified with itself to form an internal edge of \mathcal{D}_0 . In this case, F forms a **disc** in \mathcal{D}_0 . (See Claim D.1 for details about the effect of flattening F in this case.)
 - (2) The map g realises an orientation-reversing gluing of F_0^\dagger and F_1^\dagger that causes e_0^\dagger and e_1^\dagger to be identified together to form a single internal edge of \mathcal{D}_0 . In this case, F forms a **projective plane** in \mathcal{D}_0 . (See Claim D.2 for details about the effect of flattening F in this case.)
 - (3) The map g realises an orientation-preserving gluing of F_0^\dagger and F_1^\dagger such that for each $i \in \{0, 1\}$, the edge e_i^\dagger gets identified with itself in reverse to form an invalid edge of \mathcal{D}_0 .
 - (4) The map g realises an orientation-preserving gluing of F_0^\dagger and F_1^\dagger that causes e_0^\dagger and e_1^\dagger to be identified together to form a single internal edge of \mathcal{D}_0 . In this case, F forms a **2-sphere** in \mathcal{D}_0 . (See Claim D.4 for details about the effect of flattening F in this case.)
- The boundary bigon path \mathcal{F}^\dagger forms a projective plane, in which case F is incident to an invalid edge in \mathcal{D}_0 .
 - The boundary bigon path \mathcal{F}^\dagger forms a disc, in which case F is incident to a boundary edge in \mathcal{D}_0 .

Proof We first consider the case where the ends of \mathcal{F}^\dagger are identified in such a way that \mathcal{F}^\dagger forms a 2-sphere in the boundary of \mathcal{D}^\dagger . In this case, observe that after flattening F_0^\dagger and F_1^\dagger , the image $\varphi(F)$ is a single internal edge in \mathcal{D}_1 . Let e_0^\dagger and e_1^\dagger denote the two edges of \mathcal{D}^\dagger that are incident to \mathcal{F}^\dagger . As illustrated in Figure 20, there are four ways to glue F_0^\dagger and F_1^\dagger together to recover the bigon face F , depending on whether the gluing is orientation-reversing, and on whether the gluing causes e_0^\dagger and e_1^\dagger to be identified together:

- The two orientation-reversing gluings are shown in Figure 20 (first and second images); these correspond to cases (1) and (2), respectively.
- The two orientation-preserving gluings are shown in Figure 20 (third and fourth images); these correspond to cases (3) and (4), respectively.

Suppose now that the ends of \mathcal{F}^\dagger are identified in such a way that \mathcal{F}^\dagger forms a projective plane in the boundary of \mathcal{D}^\dagger . Let e_0^\dagger and e_1^\dagger denote the two edges of \mathcal{D}^\dagger that are incident to \mathcal{F}^\dagger . Up to symmetry, there are two ways to glue F_0^\dagger and F_1^\dagger back together, depending on whether the gluing causes e_0^\dagger and e_1^\dagger to be identified together. As illustrated in Figure 21, F is incident to an invalid edge in both cases.




Figure 21: The two ways to glue F_0^\dagger and F_1^\dagger together when \mathcal{F}^\dagger forms a projective plane. They are gluings that result in F being incident to one internal edge and one invalid edge (left) and F being incident to an invalid edge (right).

Finally, suppose the ends of \mathcal{F}^\dagger are not identified, so that \mathcal{F}^\dagger forms a disc in the boundary of \mathcal{D}^\dagger . Observe that each end of \mathcal{F}^\dagger must be incident to a boundary face of \mathcal{D}^\dagger that is not part of \mathcal{F}^\dagger . This means that regardless of how we glue F_0^\dagger and F_1^\dagger back together, the bigon face F will always be incident to an edge lying in the boundary of \mathcal{D}_0 . \square

As mentioned earlier, we only give a detailed analysis of the effect of flattening F in the cases where F is not incident to any boundary or invalid edges. This corresponds to cases (1), (2) and (4) of Claim D.

Claim D.1 (disc) *In case (1) of Claim D, the truncated pseudomanifolds \mathcal{P}_0 and \mathcal{P}_1 are homeomorphic.*

Proof Recall that in case (1) of Claim D, the edges of F are not identified, so that F forms a disc. In this case, we can flatten F without changing the topology of \mathcal{D}_0 ; in particular, as we saw in Claim B, the links of the vertices incident to F remain unchanged after flattening F . Thus, we see that \mathcal{P}_0 and \mathcal{P}_1 are homeomorphic. \square

Claim D.2 (projective plane) *In case (2) of Claim D, the two vertices of F are identified to form a single vertex v , and one of the following holds:*

- (a) *The vertex v is internal, in which case F forms a one-sided properly embedded projective plane in \mathcal{P}_0 , and \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along this projective plane.*
- (b) *The vertex v is ideal, in which case the truncated bigon associated to F forms a one-sided properly embedded Möbius band S in \mathcal{P}_0 ; the boundary curve γ of S forms a two-sided curve in $\partial\mathcal{P}_0$. In this case, flattening F has one of the following effects:*
 - *If γ bounds a disc E in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a one-sided projective plane given by isotoping $S \cup E$ slightly off the boundary of \mathcal{P}_0 .*
 - *If γ does not bound a disc in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first decomposing along the Möbius band S , and then filling any new 2-sphere boundary components with 3-balls.*
- (c) *The vertex v is boundary or invalid.*




Figure 22: When \mathcal{F}^\dagger forms a 2-sphere, flattening \mathcal{F}^\dagger is equivalent to filling it with a 3-ball.

Proof Recall that in case (2) of Claim D, the edges of F are identified so that F forms a projective plane. In particular, this means that the two vertices of F are identified to form a single vertex v in \mathcal{D}_0 . We start by getting the easy cases out of the way:

- If v is boundary or invalid, then we are in case (c).
- If v is internal, then F forms an embedded projective plane that lies entirely in the interior of \mathcal{P}_0 . In this case, observe that ungluing F corresponds topologically to cutting along this projective plane; this yields a single 2-sphere remnant, corresponding precisely to the boundary bigon path \mathcal{F}^\dagger . This means, in particular, that F forms a *one-sided* projective plane in \mathcal{P}_0 . Moreover, as illustrated in Figure 22, flattening F_0^\dagger and F_1^\dagger is topologically equivalent to filling the 2-sphere \mathcal{F}^\dagger with a 3-ball. Altogether, we see that flattening F is topologically equivalent to decomposing along F . This proves case (a).

The rest of this proof is devoted to the case where v is ideal; we need to prove all the conclusions stated in case (b). For this, we first observe that the truncated bigon associated to F forms a properly embedded Möbius band S in \mathcal{P}_0 . Consider the pseudomanifold \mathcal{P}^\dagger obtained from \mathcal{D}^\dagger by truncating the vertices in $g^{-1}(V_0)$; viewing \mathcal{P}^\dagger as a subset of \mathcal{D}^\dagger , let S^\dagger denote the annulus in $\partial\mathcal{P}^\dagger$ given by $\mathcal{F}^\dagger \cap \mathcal{P}^\dagger$. Topologically, observe that \mathcal{P}^\dagger is obtained from \mathcal{P}_0 by cutting along the Möbius band S ; as shown in Figure 23, this yields a single remnant—namely, the annulus S^\dagger —so S must be a *one-sided* Möbius band in \mathcal{P}_0 .

Consider the ideal boundary component L of \mathcal{P}_0 given by truncating the vertex v , and let γ denote the boundary curve of S . To see that γ forms a two-sided curve in L , observe that cutting along S splits γ into the two disjoint curves that bound the annulus S^\dagger .




Figure 23: The truncated bigon associated to F forms a one-sided Möbius band S . Cutting along S yields a single annulus remnant.




Figure 24: Flattening \mathcal{F}^\dagger has the effect of filling the annulus remnant S^\dagger with a thickened disc.

All that remains is to understand the overall topological effect of flattening F . We begin with the case where γ bounds a disc E in L . In this case, we use Claim C to flatten the v -cone over γ . This reduces the operation of flattening F to the operation of flattening a new bigon F' given by pushing F slightly away from v ; topologically, F' is equivalent to a projective plane given by isotoping $S \cup E$ slightly off the boundary of \mathcal{P}_0 . Since the vertices of F' are identified to form the temporary internal vertex that results from flattening the curve γ , flattening F' has the same topological effect as flattening F in the case where v is internal (case (a)). That is, F' forms a *one-sided* projective plane in \mathcal{P}_0 , and \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along this projective plane. This completes the case where γ bounds a disc in L .

For the case where γ does *not* bound a disc in L , we flatten F by first ungluing F , and then flattening F_0^\dagger and F_1^\dagger . Earlier, we observed that ungluing F has the effect of cutting \mathcal{P}_0 along S , which yields a single annulus remnant S^\dagger in a new pseudomanifold \mathcal{P}^\dagger . With this in mind, consider the pseudomanifold \mathcal{P}^* obtained from \mathcal{D}_1 by truncating the vertices in $\varphi(V_0)$. Topologically, observe that \mathcal{P}^* is obtained from \mathcal{P}^\dagger by filling the annulus S^\dagger with a thickened disc; see Figure 24. In other words, \mathcal{P}^* is obtained from \mathcal{P}_0 by decomposing along the Möbius band S .

To see how \mathcal{P}^* is related to \mathcal{P}_1 , we need to compare the truncated vertex sets $\varphi(V_0)$ and V_1 . The only way these vertex sets can differ is if $\varphi(v)$ contains a vertex that is neither ideal nor invalid—such a vertex would be in $\varphi(V_0)$ but not in V_1 . We can use Claim B to determine the composition of $\varphi(v)$, and hence determine the relationship between \mathcal{P}^* and \mathcal{P}_1 :

- If γ is a nonseparating curve in L , then decomposing L along γ gives a single new closed surface L^* , and $\varphi(v)$ consists of a single vertex whose link is given by L^* . If L^* is not a 2-sphere, then $\varphi(v)$ is an ideal vertex, so \mathcal{P}^* is homeomorphic to \mathcal{P}_1 . However, if L^* is a 2-sphere, then $\varphi(v)$ is an internal vertex; topologically, L^* corresponds to a 2-sphere boundary component of \mathcal{P}^* , and we need to fill this 2-sphere with a 3-ball to recover \mathcal{P}_1 from \mathcal{P}^* .
- If γ is a separating curve in L , then decomposing along γ gives two new closed surfaces, and $\varphi(v)$ consists of two vertices whose links are given by these two new surfaces. Since γ does not bound a disc in L , both vertices in $\varphi(v)$ are ideal, so \mathcal{P}^* is homeomorphic to \mathcal{P}_1 .

This completes the proof of case (b). □




Figure 25: The truncated bigon associated to F forms a one-sided annulus S . Cutting along S yields a single annulus remnant.

Claim D.4 (2-sphere) *In case (4) of Claim D, each vertex incident to F is either ideal or invalid. Moreover, the truncated bigon associated to F forms a one-sided properly embedded annulus S in \mathcal{P}_0 , and each boundary curve of S forms a one-sided curve in $\partial\mathcal{P}_0$.*

Suppose F is only incident to ideal vertices (the two vertices of F could either be identified to form a single ideal vertex, or they could form two distinct ideal vertices). In this case, \mathcal{P}_1 is obtained from \mathcal{P}_0 by first decomposing along the annulus S , and then filling any new 2-sphere boundary components with 3-balls.

Proof Recall that in case (4) of Claim D, the edges of F are identified so that F forms a 2-sphere. The two vertices of F could either be identified to form a single vertex of \mathcal{D}_0 , or they could form two distinct vertices of \mathcal{D}_0 . Let L denote the union of the links of the vertices incident to F , and consider the two curves γ_0 and γ_1 in which F meets L . For each $i \in \{0, 1\}$, observe that ungluing F causes the curve γ_i to “unravel” to form a single new curve; thus, γ_i forms a one-sided curve in L . In particular, the vertices incident to F must have nonorientable vertex links, which implies that these vertices must be either ideal or invalid. This means that the truncated bigon associated to F forms a properly embedded annulus S in \mathcal{P}_0 .

To see that S is one-sided, consider the pseudomanifold \mathcal{P}^\dagger obtained from \mathcal{D}^\dagger by truncating the vertices in $g^{-1}(V_0)$; viewing \mathcal{P}^\dagger as a subset of \mathcal{D}^\dagger , let S^\dagger denote the annulus in $\partial\mathcal{P}^\dagger$ given by $\mathcal{F}^\dagger \cap \mathcal{P}^\dagger$. Topologically, \mathcal{P}^\dagger is obtained from \mathcal{P}_0 by cutting along the annulus S ; as shown in Figure 25, this yields a single remnant — namely, the annulus S^\dagger — which tells us that S is a one-sided annulus in \mathcal{P}_0 .

Suppose now that the vertices incident to F are all ideal. Consider the pseudomanifold \mathcal{P}^* obtained from \mathcal{D}_1 by truncating the vertices in $\varphi(V_0)$. Topologically, \mathcal{P}^* is obtained from \mathcal{P}^\dagger by filling the annulus S^\dagger with a thickened disc; see Figure 24. In other words, \mathcal{P}^* is obtained from \mathcal{P}_0 by decomposing along the annulus S .

To see how \mathcal{P}^* and \mathcal{P}_1 are related, we need to compare the truncated vertex sets $\varphi(V_0)$ and V_1 . For this, we first note that since F is only incident to ideal vertices, each component of L must be a closed surface other than a 2-sphere. Let L^* denote the (possibly disconnected) surface obtained by decomposing L along γ_0 and γ_1 ; each component of L^* must be a closed surface, but could possibly be a 2-sphere.




Figure 26: Gluing F_0^\dagger and F_1^\dagger together when each forms a separate 2-sphere or projective plane in the boundary of \mathcal{D}^\dagger . Left: gluing two separate 2-spheres results in F forming a 2-sphere. Middle: gluing two separate projective planes results in F forming a projective plane. Right: gluing a 2-sphere and a projective plane results in F being incident to an invalid edge.

By Claim B, the components of L^* correspond precisely to the boundary components of \mathcal{P}^* given by truncating the vertices in $\varphi(v)$. The only way \mathcal{P}^* can differ from \mathcal{P}_1 is if L^* has 2-sphere components; we need to fill each such 2-sphere with a 3-ball to recover \mathcal{P}_1 from \mathcal{P}^* . \square

3.2.4 The case where ungluing gives two separate boundary bigon paths We now consider the case where F_0^\dagger and F_1^\dagger form two separate boundary bigon paths. We have the following cases:

Claim E *If F_0^\dagger and F_1^\dagger form two separate boundary bigon paths, then one of the following holds (see Figure 26):*

- (1) *For each $i \in \{0, 1\}$, the boundary bigon path F_i^\dagger forms a 2-sphere. In this case, $\varphi(F)$ consists of two distinct internal edges in \mathcal{D}_1 . Moreover, the edges of F are identified to form a single internal edge in \mathcal{D}_0 , and F itself forms a **2-sphere** in \mathcal{D}_0 . (See Claim E.1 for details about the effect of flattening F in this case.)*
- (2) *For each $i \in \{0, 1\}$, the boundary bigon path F_i^\dagger forms a projective plane. In this case, $\varphi(F)$ consists of two distinct invalid edges in \mathcal{D}_1 . Moreover, the edges of F are identified to form a single internal edge in \mathcal{D}_0 , and F itself forms a **projective plane** in \mathcal{D}_0 . (See Claim E.2 for details about the effect of flattening F in this case.)*
- (3) *For some $i \in \{0, 1\}$, the boundary bigon path F_i^\dagger forms a 2-sphere, but the boundary bigon path F_{1-i}^\dagger forms a projective plane. In this case, the edges of F are identified to form a single invalid edge in \mathcal{D}_0 .*
- (4) *For some $i \in \{0, 1\}$, the boundary bigon path F_i^\dagger forms a disc, in which case (regardless of the behaviour of F_{1-i}^\dagger) F is incident to a boundary edge in \mathcal{D}_0 .*

Proof For each $i \in \{0, 1\}$, depending on how the ends of F_i^\dagger are identified (if at all), F_i^\dagger could form a 2-sphere, projective plane or disc in the boundary of \mathcal{D}^\dagger . Observe that:

- If F_i^\dagger forms a 2-sphere, then flattening F_i^\dagger yields a single internal edge.
- If F_i^\dagger forms a projective plane, then flattening F_i^\dagger yields a single invalid edge.

With this in mind, we have the following four cases, which correspond precisely to the cases stated in the claim:

- (1) Suppose that F_0^\dagger and F_1^\dagger both form 2-spheres. After the gluing these bigon faces back together, the edges of F are identified so that F forms a 2-sphere, as illustrated in Figure 26 (left).
- (2) Suppose that F_0^\dagger and F_1^\dagger both form projective planes. After gluing these bigon faces back together, the edges of F are identified so that F forms a projective plane, as illustrated in Figure 26 (middle).
- (3) Suppose that one of F_0^\dagger or F_1^\dagger forms a 2-sphere, whilst the other forms a projective plane. After gluing these bigon faces back together, the edges of F are identified to form a single invalid edge, as illustrated in Figure 26 (right).
- (4) Suppose that for some $i \in \{0, 1\}$, F_i^\dagger forms a disc in the boundary of \mathcal{D}^\dagger ; each end of F_i^\dagger must therefore be incident to a boundary face other than F_0^\dagger or F_1^\dagger . Thus, regardless of how we glue F_0^\dagger and F_1^\dagger back together, the bigon face F will always be incident to at least one boundary edge of \mathcal{D}_0 . \square

As before, we only give a detailed analysis of the effect of flattening F in the cases where F is not incident to any boundary or invalid edges. This corresponds to cases (1) and (2) of Claim E.

Claim E.1 (2-sphere) *In case (1) of Claim E, one of the following holds:*

- (a) *The bigon face F is only incident to internal vertices (the two vertices of F could either be identified to form a single internal vertex, or they could form two distinct internal vertices). In this case, \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere given by pushing F slightly away from its incident vertices.*
- (b) *The bigon face F is incident to one internal vertex and one ideal vertex. In this case, the truncated bigon associated to F forms a properly embedded disc S in \mathcal{P}_0 ; the boundary curve γ of S is a two-sided curve in $\partial\mathcal{P}_0$. Flattening F has one of the following effects:*
 - *If γ bounds a disc E in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere given by isotoping $S \cup E$ slightly off the boundary.*
 - *If γ does not bound a disc in \mathcal{P}_0 , then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first cutting along the disc S , and then filling any new 2-sphere boundary components with 3-balls.*
- (c) *The bigon face F is only incident to ideal vertices (the two vertices of F could either be identified to form a single ideal vertex, or they could form two distinct ideal vertices). In this case, the truncated bigon associated to F forms a two-sided properly embedded annulus S in \mathcal{P}_0 ; the boundary curves γ_0 and γ_1 of S form two-sided curves in $\partial\mathcal{P}_0$. Flattening F has one of the following effects:*
 - *If γ_0 and γ_1 respectively bound discs E_0 and E_1 in $\partial\mathcal{P}_0$, then either these discs are disjoint or one of these discs lies entirely in the interior of the other; choose $i \in \{0, 1\}$ so that E_i either lies entirely inside or entirely outside E_{1-i} . In this case, \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere S^* constructed as follows:*




Figure 27: When F' forms a properly embedded 2-sphere, flattening F' is equivalent to decomposing along this 2-sphere.

- (i) Isotope $S \cup E_i$ slightly off the boundary to obtain a properly embedded disc S' in \mathcal{P}_0 .
- (ii) Isotope $S' \cup E_{1-i}$ slightly off the boundary to obtain the desired 2-sphere S^* .

- If for some $i \in \{0, 1\}$, the curve γ_i bounds a disc E_i in $\partial\mathcal{P}_0$, but γ_{1-i} does not bound a disc in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first cutting along a disc given by isotoping $S \cup E_i$ slightly away from E_i , and then filling any new 2-sphere boundary components with 3-balls.
- If neither γ_0 nor γ_1 bounds a disc in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first decomposing along the annulus S , and then filling any new 2-sphere boundary components with 3-balls.

(d) There is a boundary or invalid vertex incident to F .

Proof Recall that in case (1) of Claim E, the edges of F are identified so that F forms a 2-sphere. The two vertices of F could either be identified to form a single vertex of \mathcal{D}_0 , or they could form two distinct vertices of \mathcal{D}_0 . Let L denote the union of the links of the vertices incident to F , and let γ_0 and γ_1 denote the two curves in which F meets L . For each $i \in \{0, 1\}$, let v_i denote the vertex of F that is cut off by the curve γ_i , and let L_i denote the link of v_i ; if v_i is ideal, then we also think of L_i as the ideal boundary component of \mathcal{P}_0 given by truncating v_i . If v_0 and v_1 are identified, then $L = L_0 = L_1$; otherwise, L is the disjoint union of L_0 and L_1 . With all this setup in mind, we start by getting the easy cases out of the way:

- If there is a boundary or invalid vertex incident to F , then we are in case (d).
- If the vertices incident to F are all internal, then F forms a 2-sphere in \mathcal{P}_0 . This 2-sphere might not be embedded, since the two vertices of F could be identified to form a single internal vertex. Thus, to ensure that we have a properly embedded 2-sphere, we use Claim C to flatten the v_0 -cone over γ_0 and the v_1 -cone over γ_1 , one at a time. This reduces the operation of flattening F to the operation of flattening a new bigon F' given by pushing F slightly away from its incident vertices. Since the vertices of F' form two distinct temporary internal vertices, F' forms the desired properly embedded 2-sphere in \mathcal{P}_0 . As shown in Figure 27, flattening F' is topologically equivalent to decomposing along this 2-sphere. This proves case (a).

We now consider the case where F is incident to one internal vertex and one ideal vertex; without loss of generality, suppose that v_0 is the internal vertex and v_1 is the ideal vertex. We need to prove




Figure 28: The truncated bigon associated to F forms a properly embedded disc S . Cutting along S yields two disc remnants.

the conclusions stated in case (b). Observe that the truncated bigon associated to F forms a properly embedded disc S in \mathcal{P}_0 . Consider the pseudomanifold \mathcal{P}^\dagger obtained from \mathcal{D}^\dagger by truncating the vertices in $g^{-1}(V_0)$; viewing \mathcal{P}^\dagger as a subset of \mathcal{D}^\dagger , for each $i \in \{0, 1\}$ let S_i^\dagger denote the disc in $\partial\mathcal{P}^\dagger$ given by $\mathcal{F}_i^\dagger \cap \mathcal{P}^\dagger$. Topologically, \mathcal{P}^\dagger is obtained from \mathcal{P}_0 by cutting along the disc S ; the two discs S_0^\dagger and S_1^\dagger form the remnants of cutting along S . This is illustrated in Figure 28. We also note that the boundary of the disc S is given by the curve γ_1 ; since cutting along S splits γ_1 into two remnants, one bounding each of the discs S_0^\dagger and S_1^\dagger , we see that γ_1 is a two-sided curve in L_1 . It remains to describe how flattening F changes \mathcal{P}_0 . This depends on whether γ_1 bounds a disc in L_1 :

- Suppose that γ_1 bounds a disc E in L_1 . We can use Claim C to flatten the v_1 -cone over γ_1 . This reduces the operation of flattening F to the operation of flattening a new bigon F' such that
 - one of the vertices of F' is the temporary internal vertex that results from flattening γ_1 ; and
 - the other vertex of F' is the internal vertex v_0 .

Topologically, F' is equivalent to a properly embedded 2-sphere given by isotoping $S \cup E$ slightly off the boundary of \mathcal{P}_0 . Moreover, by analogy with case (a), we see that \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along this 2-sphere F' .

- Suppose that γ_1 does not bound a disc in L_1 . Consider the pseudomanifold \mathcal{P}^* obtained from \mathcal{D}_1 by truncating the vertices in $\varphi(V_0)$. As shown in Figure 29, \mathcal{P}^* is obtained from \mathcal{P}^\dagger by collapsing S_0^\dagger and S_1^\dagger to arcs, which has no topological effect; in other words, \mathcal{P}^* is homeomorphic to \mathcal{P}^\dagger . To see how \mathcal{P}^* is related to \mathcal{P}_1 , consider the surface L^* obtained from L_1 by decomposing along γ_1 . Using Claim B, we




Figure 29: Collapsing the discs S_0^\dagger and S_1^\dagger to arcs has no topological effect.




Figure 30: The truncated bigon associated to F forms a two-sided annulus S . Cutting along S yields two annulus remnants.

see that each component of L^* corresponds to a boundary component of \mathcal{P}^* given by truncating one of the vertices in $\varphi(v_1)$. Thus, if $\varphi(v_1)$ contains any internal vertices (since γ is two-sided and does not bound a disc in L_1 , this is only possible if L_1 is a torus or a Klein bottle), then we need to fill each corresponding 2-sphere boundary component of \mathcal{P}^* with a 3-ball to recover \mathcal{P}_1 .

This proves case (b).

All that remains is to consider the case where the vertices incident to F are all ideal; we need to prove the conclusions stated in case (c). This time, the truncated bigon associated to F forms a properly embedded annulus S in \mathcal{P}_0 . Similar to before, consider the pseudomanifold \mathcal{P}^\dagger obtained from \mathcal{D}^\dagger by truncating the vertices in $g^{-1}(V_0)$; viewing \mathcal{P}^\dagger as a subset of \mathcal{D}^\dagger , for each $i \in \{0, 1\}$ let S_i^\dagger denote the annulus in $\partial\mathcal{P}^\dagger$ given by $\mathcal{F}_i^\dagger \cap \mathcal{P}^\dagger$. Topologically, \mathcal{P}^\dagger is obtained from \mathcal{P}_0 by cutting along the annulus S ; as shown in Figure 30, this yields a pair of remnants — namely, the annuli S_0^\dagger and S_1^\dagger — which means that S is a *two-sided* annulus in \mathcal{P}_0 . We also note that for each $i \in \{0, 1\}$, the curve γ_i forms a two-sided curve in L_i , since cutting along S causes this curve to split into two remnants, one meeting S_0^\dagger and the other meeting S_1^\dagger .

Depending on whether γ_0 and γ_1 bound discs in $\partial\mathcal{P}_0$, we have the following possibilities for how flattening F changes \mathcal{P}_0 :

- Suppose that for some $i \in \{0, 1\}$, γ_i bounds a disc E_i in L_i . We can assume without loss of generality that the interior of E_i is disjoint from F . To see why, note that the only way this assumption can fail is if γ_{1-i} lies in the interior of E_i ; in this case, γ_{1-i} bounds a “smaller” disc, so we can simply exchange the roles of γ_i and γ_{1-i} . This allows us to use Claim C to flatten the v_i -cone over γ_i . This reduces the operation of flattening F to the operation of flattening a new bigon F' such that

- one of the vertices of F' is given by the temporary internal vertex that results from flattening γ_i ; and
- the other vertex of F' is given by the ideal vertex v_{1-i} .

Thus, flattening F' has the same topological effect as flattening F in case (b). In more detail, after truncating the ideal vertex v_{1-i} , we see that F' becomes a properly embedded disc S' in \mathcal{P}_0 ; we can view γ_{1-i} as the boundary curve of this disc S' . Topologically, S' is obtained by isotoping $S \cup E_i$ slightly off the boundary of \mathcal{P}_0 , and the effect of flattening F' depends on whether γ_{1-i} bounds a disc in L_{1-i} :




Figure 31: For each $i \in \{0, 1\}$, flattening F_i^\dagger has the effect of filling the annulus remnant S_i^\dagger with a thickened disc.

- If γ_{1-i} bounds a disc E_{1-i} in L_{1-i} , then \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere given by isotoping $S' \cup E_{1-i}$ slightly off the boundary of \mathcal{P}_0 .
- If γ_{1-i} does not bound a disc in L_{1-i} , then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first cutting along the disc S' , and then filling any new 2-sphere boundary components with 3-balls.
- Suppose that neither γ_0 nor γ_1 bounds a disc in $\partial\mathcal{P}_0$. Consider the pseudomanifold \mathcal{P}^* obtained from \mathcal{D}_1 by truncating the vertices in $\varphi(V_0)$. Topologically, \mathcal{P}^* is obtained from \mathcal{P}^\dagger by filling the annuli S_0^\dagger and S_1^\dagger with thickened discs; see Figure 31. In other words, \mathcal{P}^* is obtained from \mathcal{P}_0 by decomposing along the annulus S . To see how \mathcal{P}^* is related to \mathcal{P}_1 , consider the surface L^* obtained from $L_0 \cup L_1$ by decomposing along γ_0 and γ_1 . Using Claim B, we see that each component of L^* corresponds to a boundary component of \mathcal{P}^* given by truncating one of the vertices in $\varphi(v_0)$ or $\varphi(v_1)$. Thus, if there are any internal vertices in $\varphi(v_0)$ or $\varphi(v_1)$, then we need to fill each corresponding 2-sphere boundary component of \mathcal{P}^* with a 3-ball to recover \mathcal{P}_1 .

This proves case (c). □

Claim E.2 (projective plane) *In case (2) of Claim E, the two vertices of F are identified to form a single vertex v , and one of the following holds:*

- (a) *The vertex v is internal, in which case F forms a two-sided properly embedded projective plane in \mathcal{P}_0 , and \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along this projective plane.*
- (b) *The vertex v is ideal, in which case the truncated bigon associated to F forms a two-sided properly embedded Möbius band S in \mathcal{P}_0 ; the boundary curve γ of S forms a two-sided curve in $\partial\mathcal{P}_0$. In this case, flattening F has one of the following effects:*
 - *If γ bounds a disc E in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a two-sided projective plane given by isotoping $S \cup E$ slightly off the boundary of \mathcal{P}_0 .*
 - *If γ does not bound a disc in $\partial\mathcal{P}_0$, then \mathcal{P}_1 is obtained from \mathcal{P}_0 by first decomposing along the Möbius band S , and then filling any new 2-sphere boundary components with 3-balls.*
- (c) *The vertex v is boundary or invalid.*

Before we prove Claim E.2, we need to understand the topological effect of flattening a boundary bigon face whose corresponding truncated bigon forms either a projective plane or a Möbius band. This is the purpose of the following lemma:

Lemma 13 *Let B be a boundary bigon face, with edges identified so that B forms a projective plane. Thus, the associated truncated bigon B' forms either a projective plane or a Möbius band. Topologically, flattening B has the effect of filling B' with an invalid cone, as described in Section 2.2.*

Proof Let \mathcal{P} denote the truncated pseudomanifold corresponding to the cell decomposition containing B . We first consider the case where the truncated bigon B' forms a projective plane boundary component of \mathcal{P} ; in other words, we are considering the case where the vertex incident to B does not get truncated, and hence $B' = B$. We produce a copy B_0 of B by an isotopy of B that

- fixes the vertex; and
- pushes the rest of B slightly into the interior of \mathcal{P} .

Let \mathcal{R} denote the region that is swept out by this isotopy, and let \mathcal{C} denote the quotient of \mathcal{R} obtained by flattening B .

By construction, deleting $\mathcal{R} - B_0$ preserves the ambient space \mathcal{P} up to homeomorphism, and replaces the boundary component B with its copy B_0 . Thus, we see that flattening B has the same topological effect as attaching a copy of \mathcal{C} to B . It therefore suffices to show that \mathcal{C} is actually an invalid cone. Our strategy for this is to express \mathcal{R} as a union of lines such that after flattening B , these lines exhibit the structure of a cone over the projective plane B_0 .

To this end, begin with $\mathcal{Q} := [0, 1]^3$. Let Λ denote the set of lines of the form $\{x\} \times \{y\} \times [0, 1]$, so that \mathcal{Q} is a (disjoint) union of the lines in Λ . Take \sim to be the minimal equivalence relation satisfying the following:

- $(x, 0, 0) \sim (x', 0, 0)$ for all $x, x' \in [0, 1]$.
- $(x, 1, 0) \sim (x', 1, 0)$ for all $x, x' \in [0, 1]$.
- $(x, y, 1) \sim (x, y', 1)$ for all $x, y, y' \in [0, 1]$.
- $(0, y, z) \sim (1, 1 - y, z)$ for all $y, z \in [0, 1]$.

In the quotient \mathcal{Q}/\sim , the $z = 0$ rectangle becomes a projective plane P_0 , and the $y = 0$ and $y = 1$ rectangles become a pair of triangles that together form a projective plane P_1 ; see Figure 32. Indeed, we can identify this quotient with \mathcal{R} in such a way that P_0 is identified with B_0 and P_1 is identified with B .

With this identification, we can express \mathcal{R} as a union of the lines given by Λ/\sim . Each of these lines joins a point in B_0 to a point on the curve $\{(x, *, 1)\}$. Moreover, we observe the following:

- For each point of the form $(x, *, 1)$, we have exactly two lines that join this point to the vertex of B_0 . The union of these lines is precisely the projective plane B .
- For every point in B_0 other than the vertex, there is a unique line that joins this point to the curve $\{(x, *, 1)\}$.




Figure 32: The region \mathcal{R} can be constructed as a quotient of the box $\mathcal{Q} := [0, 1]^3$. The $z = 0$ rectangle becomes the projective plane B_0 (shaded red). The $y = 0$ and $y = 1$ rectangles together become the projective plane B (shaded blue).

After flattening B , the curve $\{(x, *, 1)\}$ gets collapsed to a single “apex” point a , and each point in the projective plane B_0 is now joined to a by a unique line. This exhibits precisely the required cone structure on \mathcal{C} , and hence we conclude that flattening B is topologically equivalent to filling B' with this invalid cone \mathcal{C} .

What remains is the case where B' forms a Möbius band in $\partial\mathcal{P}$. In this case, flattening B is still equivalent to attaching an invalid cone, except we need to account for the fact that the vertex incident to B is truncated. This truncation only removes a small neighbourhood of the vertex, and this removed neighbourhood is disjoint from the apex of the cone. Thus, we still have an invalid cone \mathcal{C} , but the truncation means that there is a small disc $D \subset \partial\mathcal{C}$ that lies inside $\partial\mathcal{P}$, and \mathcal{C} is attached by identifying the Möbius bands B' and $\partial\mathcal{C} - D$. In other words, we again conclude that flattening B is topologically equivalent to filling B' with the invalid cone \mathcal{C} . □

Proof of Claim E.2 Recall that in case (2) of Claim E, the edges of F are identified so that F forms a projective plane. We will see that the proof is almost identical to the proof of Claim D.2; the main difference is that here, we end up working with embedded surfaces that are two-sided rather than one-sided.

Since F forms a projective plane, the two vertices of F are identified to form a single vertex v in \mathcal{D}_0 . When v is not ideal, the claim is easy to prove:

- If v is boundary or invalid, then we are in case (c).
- If v is internal, then F forms an embedded projective plane in the interior of \mathcal{P}_0 . Ungluing F yields two projective plane remnants corresponding to the boundary bigon paths F_0^\dagger and F_1^\dagger , which tells us that F forms a *two-sided* projective plane in \mathcal{P}_0 . Moreover, by Lemma 13, flattening F_0^\dagger and F_1^\dagger corresponds to filling these two projective plane boundary components with invalid cones. Altogether, we see that flattening F is topologically equivalent to decomposing along F . This proves case (a).




Figure 33: The truncated bigon associated to F forms a two-sided Möbius band S . Cutting along S yields two Möbius band remnants.

With that out of the way, suppose for the rest of this proof that v is ideal; we need to prove all the conclusions stated in case (b). Observe that the truncated bigon associated to F forms a properly embedded Möbius band S in \mathcal{P}_0 . Consider the pseudomanifold \mathcal{P}^\dagger obtained from \mathcal{D}^\dagger by truncating the vertices in $g^{-1}(V_0)$; viewing \mathcal{P}^\dagger as a subset of \mathcal{D}^\dagger , for each $i \in \{0, 1\}$ let S_i^\dagger denote the Möbius band in $\partial\mathcal{P}^\dagger$ given by $F_i^\dagger \cap \mathcal{P}^\dagger$. Topologically, \mathcal{P}^\dagger is obtained from \mathcal{P}_0 by cutting along the Möbius band S ; as shown in Figure 33, this yields two remnants — namely, the Möbius bands S_0^\dagger and S_1^\dagger — so S must be a *two-sided* Möbius band in \mathcal{P}_0 .

Consider the ideal boundary component L of \mathcal{P}_0 given by truncating the vertex v , and let γ denote the boundary curve of S ; we need to show that γ forms a two-sided curve in L . For this, it suffices to observe that cutting along S has the effect of splitting γ into two curves, one bounding the Möbius band S_0^\dagger and the other bounding the Möbius band S_1^\dagger .

All that remains is to understand the overall effect of flattening F . We begin with the case where γ bounds a disc E in L . In this case, we use Claim C to flatten the v -cone over γ . This reduces the operation of flattening F to the operation of flattening a new bigon F' given by pushing F slightly away from v ; topologically, F' is equivalent a properly embedded projective plane given by isotoping $S \cup E$ slightly off the boundary of \mathcal{P}_0 . Since the vertices of F' are identified to form a single temporary internal vertex, flattening F' has the same topological effect as flattening F in the case where v is internal (case (a)). In other words, F' forms a *two-sided* projective plane in \mathcal{P}_0 , and \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along this projective plane. This completes the case where γ bounds a disc in L .

For the case where γ does *not* bound a disc in L , consider the pseudomanifold \mathcal{P}^* obtained from \mathcal{D}_1 by truncating the vertices in $\varphi(V_0)$. Topologically, by Lemma 13, \mathcal{P}^* is obtained from \mathcal{P}^\dagger by filling the Möbius bands S_0^\dagger and S_1^\dagger with invalid cones; in other words, \mathcal{P}^* is obtained from \mathcal{P}_0 by decomposing along the Möbius band S . To see how \mathcal{P}^* is related to \mathcal{P}_1 , we compare the truncated vertex sets $\varphi(V_0)$ and V_1 . For this, let L^* denote the surface obtained by decomposing L along γ . Claim B tells us that the components of L^* correspond to boundary components of \mathcal{P}^* given by truncating the vertices in $\varphi(v)$. The only way \mathcal{P}^* can differ from \mathcal{P}_1 is if L^* has 2-sphere components; we need to fill each such 2-sphere with a 3-ball to recover \mathcal{P}_1 from \mathcal{P}^* . This completes the proof of case (b). \square

4 Crushing surfaces of positive genus

Consider a normal surface S in a 3-manifold \mathcal{M} . Roughly, our goal in this section is to give sufficient conditions under which crushing S gives an ideal triangulation of a component of $\mathcal{M} - S$. For this, we fix the following notation throughout this section:

- Let \mathcal{M} be a (compact) 3-manifold with no 2-sphere boundary components. If \mathcal{M} is closed, let \mathcal{T} be a closed triangulation of \mathcal{M} ; otherwise, if \mathcal{M} is bounded, let \mathcal{T} be an ideal triangulation of \mathcal{M} .
- Let S be a (possibly disconnected) separating normal surface in \mathcal{T} . (Since \mathcal{T} has no real boundary components, note that S must be a closed surface.) Assume that every component of S is two-sided, and that none of these components are 2-spheres.
- Fix any particular component of $\mathcal{M} - S$ that meets each component of S on exactly one side, and call it the *chosen region* for S ; also, for any (not necessarily normal) surface E isotopic to S , call the corresponding component of $\mathcal{M} - E$ the chosen region for E . Let X denote the compact 3-manifold given by the closure of the chosen region for S .

The assumptions that we have made on S and on the chosen region are not as restrictive as they might appear at first glance. If S has a component C that is either nonseparating or one-sided (or both), then we can always “repair” S as follows: build a new surface Σ by replacing C with the frontier of a regular neighbourhood of C . Up to homeomorphism, each component of $\mathcal{M} - S$ appears as a component of $\mathcal{M} - \Sigma$, so we lose nothing by repairing S in this way. A similar trick allows us to deal with components of $\mathcal{M} - S$ that meet a component of S on both sides.

With this in mind, let \mathcal{T}' denote the triangulation obtained by (destructively) crushing S . As we hinted earlier, our goal is to give sufficient conditions for \mathcal{T}' to include an ideal triangulation of X as one of its components (the precise statement is given in Theorem 1).

To this end, consider the cell decomposition \mathcal{D}' given by *nondestructively* crushing S . One of the components \mathcal{D}^* of \mathcal{D}' gives a destructible ideal cell decomposition of X ; see Figure 34. Call a 3-cell in \mathcal{D}' *benign* if and only if it belongs to the component \mathcal{D}^* .

From the formulation of crushing given in Section 2.5 (in particular, recall Definitions 5 and Lemma 7), \mathcal{T}' is obtained from \mathcal{D}' by a finite sequence of atomic moves. Since the atomic moves only either destroy or




Figure 34: Schematic illustration of nondestructively crushing S , in the case where S is connected.




Figure 35: Some notation that we use throughout Section 4.

modify the existing 3-cells in \mathcal{D}' , we can naturally speak about benign 3-cells in all of the intermediate cell decompositions that we encounter as we perform the atomic moves.

We will also call a component of a cell decomposition *benign* if this component is built entirely from benign 3-cells. Whilst \mathcal{D}' initially has exactly one benign component, namely \mathcal{D}^* , this number can change as we perform atomic moves.

With this in mind, let \mathcal{T}^* denote the triangulation consisting only of the benign components of \mathcal{T}' ; see Figure 35 for a visual summary of all the notation we have just introduced. We can now give a more precise statement of our main goal in this section: we want to show that one of the components of \mathcal{T}^* gives an ideal triangulation of X .

The proof boils down to checking that we do not make “drastic” topological changes when performing atomic moves on an ideal cell decomposition of X . In all but one of the cases, it is enough to require that X satisfies the following conditions:

- It is irreducible and ∂ -irreducible.
- It contains no essential annuli and no two-sided properly embedded Möbius bands.

The one difficult case is when we flatten a bigon whose corresponding truncated bigon forms a boundary-parallel annulus in X ; we will discuss how we circumvent this difficulty in Section 4.1. We then put everything together to prove Theorem 1 in Section 4.2.

4.1 Avoiding bad bigon paths

Throughout the rest of Section 4, call a bigon path (recall Definitions 12) *bad* if it is internal and its corresponding truncated bigon path forms a boundary-parallel annulus. Using this terminology, the difficult case that we mentioned earlier is when we flatten a bad bigon path of length one. In Section 4.2, we will see that the only way to have a bad bigon path of length one is if the cell decomposition \mathcal{D}^* initially contained a bad bigon path (of some arbitrary length). With this in mind, our goal is to cut this problem off at the source: we will give conditions on the surface S that will ensure that \mathcal{D}^* does not contain any bad bigon paths.

Roughly, the idea is that if \mathcal{D}^* contains a bad bigon path, then we can “push” or “expand” S further into the chosen region; thus, we would like to ensure that S cannot be “expanded” in this way. To make this




Figure 36: Let E and E' be two disjoint normal surfaces in the same isotopy class. If E' cannot be normally isotoped to lie outside the chosen region for E , then it is an expansion of E .

idea precise, we introduce the following terminology for any normal surface E in the isotopy class of S :

- Let E' be any normal surface that is, up to normal isotopy, disjoint from E . We call E' an *expansion* of E if it is isotopic to E , but cannot be normally isotoped to lie entirely outside the chosen region for E ; see Figure 36.
- Call E *maximal* if it does not admit such an expansion.

In Lemma 15, we will show that to avoid bad bigon paths in \mathcal{D}^* , it is enough to assume that X is irreducible, and that S is incompressible and maximal. Before we do so, it is worth noting that the maximality assumption is not very restrictive. Specifically, the following result says that, up to isotopy, we can always choose S to be maximal:

Lemma 14 *There is a maximal normal surface in the isotopy class of S .*

One way to prove Lemma 14 would be to appeal to Kneser's finiteness theorem (as is done, for instance, in [18, pages 160–161]); however, we do not actually need the full strength of this theorem. The following simple proof distills precisely the part of Kneser's finiteness theorem that is necessary:

Proof of Lemma 14 If S is itself maximal, then there is nothing to prove; thus, assume for the rest of this proof that S is not maximal. Let $E_0 = S$, and consider any sequence E_0, \dots, E_n of normal surfaces such that for each $i \in \{1, \dots, n\}$, E_i is an expansion of E_{i-1} . The idea is to show that such a sequence cannot be extended indefinitely, which means that after extending as much as possible, the final entry of the sequence will be maximal.

To this end, consider the surface $E_0 \cup \dots \cup E_n$; this is a normal surface, since E_0, \dots, E_n are mutually disjoint, up to normal isotopy. Let \mathcal{C} denote the induced cell decomposition obtained by cutting along $E_0 \cup \dots \cup E_n$. For each $i \in \{1, \dots, n\}$, let B_i denote the component of \mathcal{C} given by the trivial I -bundle between E_i and E_{i-1} . Since E_i and E_{i-1} are not normally isotopic, B_i must contain at least one nonparallel cell. However, each tetrahedron of \mathcal{T} gives rise to at most six nonparallel cells, so we see that $n \leq 6|\mathcal{T}|$. This implies that there is a sequence E_0, \dots, E_k of expansions whose length k is maximum among all such sequences; the surface E_k must therefore be a maximal normal surface in the isotopy class of S , otherwise we would be able to extend the sequence by adding an expansion of E_k . \square




Figure 37: An example of an induced cell Δ that meets the chosen region for S_1 . The red faces are parts of S_1 that originated from S , and the blue bridge face is a part of S_1 that originated from A . The intersection of S_2 with Δ is shaded orange. In this example, the orange piece forms an elementary triangle, which means that this piece is preserved by the normalisation procedure; thus, the final surface S_4 will include this orange triangle among its elementary discs.

Lemma 15 *If X is irreducible, S is maximal, and the remnants of S are incompressible in X , then \mathcal{D}^* contains no bad bigon paths.*

Proof Suppose \mathcal{D}^* contains a bad bigon path \mathcal{F} . We will show that S cannot be maximal by using \mathcal{F} to construct an expansion of S .

To do this, let \mathcal{D} denote the cell decomposition of X induced by S , and let $q: \mathcal{D} \rightarrow \mathcal{D}^*$ be the quotient map given by nondestructively crushing S . Observe that $q^{-1}(\mathcal{F})$ realises an annulus A that

- consists of bridge faces (as defined in Definitions 3) in the 2-skeleton of \mathcal{D} ; and
- is parallel to an annulus A^\parallel lying entirely inside a component of S .

We now aim to build a sequence S_1, S_2, S_3, S_4 of surfaces isotopic to S , each of which is “expanded further” into the chosen region than the last. Our goal is for S_4 to be the required expansion of S . We achieve this as follows:

- (1) Let S_1 be the surface obtained from S by replacing A^\parallel with A .
- (2) Consider the boundary B of a regular neighbourhood of $S \cup A$, and let S_2 be the union of the components of B that lie inside the chosen region for S_1 ; see Figure 37. Since the chosen region for S_1 meets each component of S_1 on exactly one of its two sides, observe that S_2 is isotopic to S . By Theorem 4, B is a barrier for any component of $\mathcal{M} - B$ that does not meet $S \cup A$; observe that the chosen region \mathcal{N} for S_2 is one such component of $\mathcal{M} - B$.
- (3) Let S_3 be the surface given by isotoping S_2 slightly into \mathcal{N} .
- (4) Using the barrier B , normalise S_3 to obtain a normal surface E in \mathcal{N} . Since the remnants of S are incompressible in X , we know that S_3 must be incompressible in \mathcal{N} . This, together with the fact that X is irreducible, implies that after deleting any 2-sphere components of E , we must be left with a normal surface S_4 isotopic to S . By construction, S_4 is disjoint from S , and it cannot be normally isotoped to lie outside the chosen region for S , so it is an expansion of S . \square




Figure 38: Schematic illustration of using a tube to turn a two-sided projective plane into a two-sided Möbius band.

4.2 Crushing the benign components

In Theorem 1 below, we give sufficient conditions so that after crushing S , one of the components of the triangulation \mathcal{T}^* gives an ideal triangulation of X . Specifically, our proof relies on the following:

- We require that S is maximal. As discussed in Section 4.1, this is not a serious restriction, thanks to Lemma 14.
- We require that X is irreducible, ∂ -irreducible and *anannular* (ie X contains no essential annuli). These are quite common “niceness” conditions for 3-manifolds. It is worth noting that we do *not* need to assume that X is *atoroidal* (ie X contains no essential tori).
- We require that X contains no two-sided properly embedded Möbius bands. This condition holds for all orientable 3-manifolds, but sometimes fails for nonorientable 3-manifolds.
- We require that X contains no two-sided properly embedded projective planes. This follows from the previous condition, together with the fact that X has at least one boundary component (given by a remnant of S). To see why, suppose X contains a two-sided projective plane E . Consider a path γ that starts at a point p_0 in E and ends at a point p_1 in ∂X . Remove a small disc around p_0 from E , and replace it with a thin tube that “follows” the path γ and ends with a curve that bounds a small disc around p_1 in ∂X ; see Figure 38. This turns E into a two-sided properly embedded Möbius band in X .

Theorem 1 *Suppose that X is irreducible, ∂ -irreducible and anannular, and that it contains no two-sided properly embedded Möbius bands. Also, suppose S is maximal. Then \mathcal{T}^* is a valid triangulation such that:*

- *One of its components is an ideal triangulation of X .*
- *Every other component is a triangulation of the 3-sphere.*

Proof Throughout this proof, call a cell decomposition *acceptable* if:

- One of its components is an ideal cell decomposition of X .
- Every other component is a (closed) cell decomposition of the 3-sphere.




Figure 39: Notation for the inductive step in Theorem 1.

Recall from Lemma 7 that the procedure of flattening \mathcal{D}^* to get \mathcal{T}^* can be realised by a finite sequence of atomic moves. Thus, our strategy will be to inductively prove that each atomic move preserves the property of being acceptable.

For the proof to work, we actually need to prove slightly more than this. The problem is that flattening a bad bigon path (as defined in Section 4.1) has the topological effect of decomposing along a boundary-parallel annulus, which could potentially yield a cell decomposition that is no longer acceptable. To circumvent this problem, we will show by induction that performing any number of atomic moves on \mathcal{D}^* always yields an acceptable cell decomposition *that contains no bad bigon paths*.

For the base case, consider the initial cell decomposition \mathcal{D}^* (which is obtained by performing zero atomic moves). Recall that \mathcal{D}^* is an ideal cell decomposition of X (with no extra 3-sphere components), so it is acceptable. By Lemma 15, we also know that \mathcal{D}^* contains no bad bigon paths.

For the inductive step, assume that we have some acceptable cell decomposition \mathcal{D}_0 that contains no bad bigon paths. We need to show that performing any atomic move on \mathcal{D}_0 yields a new acceptable cell decomposition \mathcal{D}_1 that has no bad bigon paths. In particular, to show that \mathcal{D}_1 remains acceptable, we will show that an atomic move always either: has no topological effect on the truncated pseudomanifold, or changes the truncated pseudomanifold by adding or removing a 3-sphere component.

Throughout the rest of this proof, let F denote the triangular pillow, bigon pillow or bigon face in \mathcal{D}_0 that we flatten to obtain \mathcal{D}_1 . Let φ denote the flattening map associated to this atomic move, and let ψ denote the inverse flattening map. For each $i \in \{0, 1\}$, let \mathcal{P}_i denote the truncated pseudomanifold of \mathcal{D}_i . This notation is summarised in Figure 39.

We first consider the case where F is a triangular pillow. Recall from Lemma 10 that the effect of flattening F depends on whether the two triangular faces of F are identified:

- If the faces of F are *not* identified, then we are in case (a) of Lemma 10. Thus, the truncated pseudomanifolds of \mathcal{D}_0 and \mathcal{D}_1 are homeomorphic, which implies that \mathcal{D}_1 is acceptable. Suppose for the sake of contradiction that \mathcal{D}_1 contains a bad bigon path B_1 . Observe that B_1 meets the triangle $\varphi(F)$ in some (possibly empty) subset of the edges of this triangle, which implies that $\psi(B_1)$ is a bad bigon path in \mathcal{D}_0 ; this violates the inductive hypothesis, and hence shows that \mathcal{D}_1 cannot contain a bad bigon path.




Figure 40: If F is a bigon pillow such that the bigon face $\varphi(F)$ forms part of a bad bigon path B_1 , then $\psi(B_1)$ contains a (bad) bigon path that is topologically equivalent to B_1 .

- If the faces of F are identified, then we are in case (c) of Lemma 10: F forms a (closed) cell decomposition of either S^3 or $L_{3,1}$. Since the only closed components of \mathcal{D}_0 are 3-spheres, \mathcal{D}_1 must be obtained from \mathcal{D}_0 by deleting the 3-sphere component given by F ; thus, \mathcal{D}_1 is acceptable. Moreover, since the ideal component of \mathcal{D}_0 is left entirely untouched by the operation of flattening F , we see that \mathcal{D}_1 cannot contain any bad bigon paths.

This completes the inductive step for the case where F is a triangular pillow.

The case where F is a bigon pillow is similar, but slightly more involved. Recall from Lemma 11 that the effect of flattening F depends on whether the two bigon faces of F are identified:

- If the faces of F are *not* identified, then we are in case (a) of Lemma 11. Thus, the truncated pseudomanifolds of \mathcal{D}_0 and \mathcal{D}_1 are homeomorphic, which implies that \mathcal{D}_1 is acceptable. Suppose for the sake of contradiction that \mathcal{D}_1 contains a bad bigon path B_1 . If B_1 is disjoint from the interior of the bigon $\varphi(F)$, then observe that $\psi(B_1)$ is a bad bigon path in \mathcal{D}_0 . This would violate the inductive hypothesis, so we conclude that the bigon $\varphi(F)$ must be part of the bigon path B_1 ; this situation is illustrated in Figure 40. Consider the internal bigon path B_0 in \mathcal{D}_0 given by $\psi(B_1) - F$; the ends of B_0 are precisely the two edges incident to the bigon pillow F . Observe that augmenting B_0 with one of the two bigon faces of F gives a bad bigon path in \mathcal{D}_0 . This again violates the inductive hypothesis, so we conclude that \mathcal{D}_1 cannot contain a bad bigon path.
- If the faces of F are identified, then we are in either case (c) or case (d) of Lemma 11. Actually, F cannot form an ideal cell decomposition of $\mathbb{R}P^2 \times [0, 1]$ (case (d)) because such a component would contain a two-sided projective plane. Thus, we must be in case (c): F must form a (closed) cell decomposition of either S^3 or $\mathbb{R}P^3$. By assumption, the only closed components of \mathcal{D}_0 are 3-spheres, so \mathcal{D}_1 must be obtained from \mathcal{D}_0 by deleting the 3-sphere component given by F ; this shows that \mathcal{D}_1 is acceptable. We also see that flattening F leaves the ideal component of \mathcal{D}_0 entirely untouched, which implies that \mathcal{D}_1 contains no bad bigon paths.

This completes the inductive step for the case where F is a bigon pillow.

With the pillow cases out of the way, all that remains is to consider the case where F is a bigon face. As in Sections 3.2.3 and 3.2.4, we divide our study into cases depending on whether the two new boundary bigons given by ungluing F form a single boundary bigon path or two separate boundary bigon paths.




Figure 41: When the truncated bigon associated to F forms a one-sided Möbius band S (blue), the frontier of a small regular neighbourhood of S forms an annulus A (green). If A admits an essential compression disc E , then the boundary curve of E coincides with the boundary of the Möbius band B (orange).

First, suppose ungluing F yields a single boundary bigon path. Since \mathcal{D}_0 is valid and has no boundary edges, Claim D tells us that $\varphi(F)$ is a single internal edge in \mathcal{D}_1 . Validity of \mathcal{D}_0 also tells us that we are in either case (1), (2) or (4) of Claim D. Actually, cases (2) and (4) are both impossible:

- Consider case (2) of Claim D. In this case, F forms a *projective plane* in \mathcal{D}_0 , and Claim D.2 tells us that the two vertices of F are identified to form a single vertex v . Moreover, since \mathcal{D}_0 is valid and has no boundary vertices, we must be in either case (a) or case (b) of Claim D.2:

(a) If v is internal, then Claim D.2 tells us that F forms a one-sided properly embedded projective plane in \mathcal{P}_0 . In fact, F lies in X , since the 3-sphere components of \mathcal{P}_0 cannot contain an embedded projective plane. Observe that a small regular neighbourhood N of F is homeomorphic to $\mathbb{R}P^3$ minus a small open 3-ball, and that the frontier of N forms a properly embedded 2-sphere E in X . Notice that E does not bound a 3-ball in X : the region on the “inside” of E is N , and the region on the “outside” contains all the boundary components of X .

(b) If v is ideal, then by Claim D.2 the truncated bigon associated to F forms a one-sided properly embedded Möbius band S in \mathcal{P}_0 . In fact, S lies in X , since the 3-sphere components of \mathcal{P}_0 have empty boundary. Consider the annulus A given by the frontier of a small regular neighbourhood N of S . Since X is anannular, A must be either compressible or boundary-parallel. We claim that neither case is possible:

- If A is compressible, then consider an essential compression disc E for A . Up to isotopy, the boundary curve of E coincides with the boundary curve of a Möbius band B in N given by thickening the core curve of S ; see Figure 41. Observe that $E \cup B$ forms an embedded projective plane in X . By assumption, this projective plane cannot be two-sided. However, it also cannot be one-sided, since this would contradict the fact that X is irreducible, by the same argument as before. Thus, we conclude that A cannot be compressible.
- If A is boundary-parallel, then isotoping A into the boundary shows that the entire component X is homeomorphic to a regular neighbourhood of the Möbius band S . But this means that X is a solid torus, which contradicts the assumption that X is ∂ -irreducible.




Figure 42: The two cases when the bigon face F is a disc that forms part of $\psi(B_1)$. Left: if $\psi(B_1)$ is a bigon path, then it is topologically equivalent to B_1 . Right: if there is a bigon path B_0 such that $B_0 \cap F$ is a single edge and $B_0 \cup F = \psi(B_1)$, then B_0 is topologically equivalent to B_1 .

The upshot is that, under our assumptions on X , case (2) of Claim D can never occur.

- Consider case (4) of Claim D. In this case, F forms a 2-sphere in \mathcal{D}_0 , and Claim D.4 tells us that the truncated bigon associated to F forms a one-sided properly embedded annulus in \mathcal{P}_0 . But this gives an essential annulus in \mathcal{P}_0 , which is impossible since X is anannular and none of the 3-sphere components of \mathcal{P}_0 can contain properly embedded annuli.

We are left with case (1) of Claim D. In this case, F forms a disc in \mathcal{D}_0 , and Claim D.1 tells us that the truncated pseudomanifolds \mathcal{P}_0 and \mathcal{P}_1 are homeomorphic, and hence that \mathcal{D}_1 is acceptable. Suppose for the sake of contradiction that \mathcal{D}_1 contains a bad bigon path B_1 . We note that B_1 must contain the edge $\varphi(F)$; otherwise, flattening F would leave B_1 untouched, which would imply that $\psi(B_1)$ is a bad bigon path in \mathcal{D}_0 , contradicting the inductive hypothesis. Thus, $\psi(B_1)$ must contain the bigon face F , and we are left with the following two possibilities:

- If $\psi(B_1)$ itself forms an internal bigon path in \mathcal{D}_0 , then observe that flattening F reduces the length of this bigon path by one, but has no topological effect; see Figure 42 (left). Thus, $\psi(B_1)$ is a bad bigon path in \mathcal{D}_0 , contradicting the inductive hypothesis.
- Otherwise, there must be an internal bigon path B_0 in \mathcal{D}_0 such that $B_0 \cap F$ is a single edge and $B_0 \cup F = \psi(B_1)$; see Figure 42 (right). In this case, we have $\varphi(B_0) = B_1$, and flattening F essentially leaves B_0 untouched, which means that B_0 is a bad bigon path in \mathcal{D}_0 , again contradicting the inductive hypothesis.

Thus, we conclude that \mathcal{D}_1 contains no bad bigon paths. This completes the inductive step for the case where ungluing F yields a single boundary bigon path.

Suppose now that ungluing F yields two separate boundary bigon paths. Since \mathcal{D}_0 is valid and has no boundary edges, we are in either case (1) or case (2) of Claim E. Actually, case (2) is impossible, since all three possibilities in Claim E.2 contradict the assumption that \mathcal{D}_0 is acceptable:

- Possibility (a) requires \mathcal{P}_0 to contain a two-sided properly embedded projective plane.
- Possibility (b) requires \mathcal{P}_0 to contain a two-sided properly embedded Möbius band.
- Possibility (c) requires \mathcal{D}_0 to contain either a boundary vertex or an invalid vertex.

We are left with case (1) of Claim E. In this case, F forms a 2-sphere in \mathcal{D}_0 , and $\varphi(F)$ consists of two distinct internal edges in \mathcal{D}_1 . Since \mathcal{D}_0 is valid and has no boundary vertices, we are in either case (a), (b) or (c) of Claim E.1. In cases (a) and (b), it is relatively easy to see that \mathcal{D}_1 is acceptable and contains no bad bigon paths:

- Consider case (a) of Claim E.1. In this case, F is only incident to internal vertices, and \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere E . Since X is irreducible and every other component of \mathcal{P}_0 is a 3-sphere, we see that E bounds a 3-ball, which implies that flattening F only changes the truncated pseudomanifold by creating a new 3-sphere component. Thus, \mathcal{D}_1 remains acceptable. Moreover, observe that $\varphi(F)$ is not incident to any ideal vertices of \mathcal{D}_1 , which means that any bad bigon path B_1 in \mathcal{D}_1 must be disjoint from $\varphi(F)$; no such B_1 can exist, otherwise $\psi(B_1)$ would be a bad bigon path in \mathcal{D}_0 .
- Consider case (b) of Claim E.1. In this case, F is incident to one internal vertex and one ideal vertex, which means that the truncated bigon associated to F forms a properly embedded disc S in X . Since X is ∂ -irreducible, the boundary curve of S must bound a disc lying entirely in ∂X , in which case Claim E.1 tells us that \mathcal{P}_1 is obtained from \mathcal{P}_0 by decomposing along a properly embedded 2-sphere in X . As before, by irreducibility of X , we conclude that flattening F only changes the truncated pseudomanifold by creating a new 3-sphere component. Thus, \mathcal{D}_1 remains acceptable. To see that \mathcal{D}_1 contains no bad bigon paths, we first note that $\varphi(F)$ is incident to exactly one ideal vertex v . With this in mind, suppose for the sake of contradiction that \mathcal{D}_1 contains a bad bigon path B_1 . Observe that $B_1 \cap \varphi(F)$ is either empty or consists only of the ideal vertex v , which means that flattening F essentially leaves B_1 untouched. This implies that $\psi(B_1)$ is a bad bigon path in \mathcal{D}_0 , contradicting the inductive hypothesis.

All that remains is to consider case (c) of Claim E.1. In this case, F is only incident to ideal vertices, which means that the truncated bigon associated to F forms a properly embedded annulus S in X ; let γ_0 and γ_1 denote the two boundary curves of S . Observe that S cannot be boundary-parallel, because this would mean that F itself forms a bad bigon path in \mathcal{D}_0 . Combining this with the assumption that X is anannular, we see that S must be a compressible annulus.

Given what we know about the 3-manifold X and the annulus S , we claim that flattening F only changes the truncated pseudomanifold by creating a new 3-sphere component. To prove this, we start by compressing S along an essential compression disc, which yields two properly embedded discs E_0 and E_1 such that for each $i \in \{0, 1\}$, the boundary curve of E_i is γ_i . Since X is ∂ -irreducible, each curve γ_i must therefore bound a disc E'_i lying entirely in the boundary of X ; see Figure 43 (left). Thus, by case (c) of Claim E.1, flattening F corresponds to decomposing along a properly embedded 2-sphere in X . Since X is irreducible, this 2-sphere bounds a 3-ball, so the only topological effect is to create a new 3-sphere component. This shows that \mathcal{D}_1 remains acceptable.

To finish, we just need to verify that \mathcal{D}_1 contains no bad bigon paths. If a bad bigon path B_1 in \mathcal{D}_1 is disjoint from $\varphi(F)$ or meets $\varphi(F)$ only in (ideal) vertices, then observe that $\psi(B_1)$ would be a bad bigon




Figure 43: Left: for each $i \in \{0, 1\}$, the curve γ_i bounds a properly embedded disc E_i in X given by compressing the annulus S . The fact that X is ∂ -irreducible therefore implies that γ_i bounds a disc E'_i that lies entirely in ∂X . Right: if F is a bigon face that forms a 2-sphere, and if there is a bigon path B_0 such that $B_0 \cap F$ is a single edge and $B_0 \cup F = \psi(B_1)$, then B_0 is topologically equivalent to B_1 .

path in \mathcal{D}_0 , which is impossible. The only other possibility is that B_1 meets $\varphi(F)$ in an edge, in which case there would exist an internal bigon path B_0 in \mathcal{D}_0 such that $B_0 \cap F$ is a single edge and $B_0 \cup F = \psi(B_1)$; see Figure 43 (right). Observe that B_0 would be a bad bigon path in \mathcal{D}_0 , which is again impossible.

In summary, we have shown that in every possible case, performing an atomic move on \mathcal{D}_0 gives a new acceptable cell decomposition \mathcal{D}_1 (and also that \mathcal{D}_1 contains no bad bigon paths). By induction, this shows that after performing however many atomic moves we need to flatten \mathcal{D}^* , the triangulation \mathcal{T}^* that results from flattening will be acceptable. \square

5 Triangulation complexity of 3-dimensional submanifolds

We wish to showcase some applications of Theorem 1 to a notion — namely, *triangulation complexity* — whose significance is independent from crushing. There has been substantial effort devoted to finding upper and lower bounds on triangulation complexity for various families of 3-manifolds; for instance, see [1; 14; 15; 16; 17; 19; 20; 21; 22; 26; 27; 32; 33]. For this paper, the triangulation complexity for a closed 3-manifold \mathcal{M} will refer to the minimum number of tetrahedra in any (closed) triangulation of \mathcal{M} , and the triangulation complexity for a bounded 3-manifold \mathcal{M} will refer to the minimum number of tetrahedra in any *ideal* triangulation of \mathcal{M} ; in either case, we will denote this quantity by $\Delta(\mathcal{M})$.

Our main application of Theorem 1 is to prove that, under quite general conditions, the triangulation complexity of a 3-manifold \mathcal{M} is strictly bigger than the triangulation complexity of a 3-dimensional submanifold of \mathcal{M} bounded by surfaces of positive genus. The precise statement is given in Theorem 2 below.

To our knowledge, Theorem 2 has never previously been written down in the literature. However, as mentioned in Section 1, it is important to note that a similar result can also be obtained by combining, from Matveev's book [28],

- ideas about the duality of triangulations and special spines (from Section 1.1);
- ideas about the conversion of almost simple spines into special spines (from Section 2.1.1); and

- results about how the complexity of almost simple spines interacts with the operation of cutting along normal surfaces in handle decompositions (from Section 4.2).

In any case, a good reason to explicitly write down Theorem 2 is that its assumptions are relatively easy to check. This gives a way to streamline some applications by avoiding the need to directly use either our crushing machinery or Matveev's spine machinery. We demonstrate this in Sections 5.1 to 5.3 by giving some straightforward applications of Theorem 2 to JSJ decompositions and satellite knots.

Theorem 2 *Let \mathcal{M} be a compact 3-manifold with no 2-sphere boundary components. Suppose \mathcal{M} contains a (possibly disconnected) closed incompressible surface S with no 2-sphere components, no projective plane components, and no boundary-parallel components. Let \mathcal{R} be a component obtained after cutting \mathcal{M} along S . If \mathcal{R} is irreducible, ∂ -irreducible, anannular, and does not contain any two-sided properly embedded Möbius bands, then $\Delta(\mathcal{R}) < \Delta(\mathcal{M})$.*

Proof Let \mathcal{T} be a closed (if \mathcal{M} is closed) or ideal (if \mathcal{M} has boundary) triangulation of \mathcal{M} such that $|\mathcal{T}| = \Delta(\mathcal{M})$. Our goal is to find an ideal triangulation of \mathcal{R} with strictly fewer tetrahedra than \mathcal{T} . We do this by constructing a suitable normal surface S' in \mathcal{T} , and using Theorem 1 to ensure that crushing S' yields the desired triangulation of \mathcal{R} .

In detail, let $N(S)$ denote a closed tubular neighbourhood of S in \mathcal{M} . Viewing \mathcal{R} as the submanifold of \mathcal{M} given by deleting the interior of $N(S)$, let S' be the union of all the components of $\partial N(S)$ that meet \mathcal{R} . Since S has no 2-sphere or projective plane components, observe that S' cannot have any 2-sphere components. Also, each component of S' is two-sided, since it meets $N(S)$ on one side and \mathcal{R} on the other side; moreover, since \mathcal{R} meets each component of S' on exactly one side, we can take the interior of \mathcal{R} to be the *chosen region* (as defined at the beginning of Section 4) for S' . This already establishes most of what we require to apply Theorem 1; what remains is to show that S' is an essential surface in \mathcal{M} , which will allow us to assume that S' is a nontrivial normal surface with respect to \mathcal{T} .

Since S' is a closed (but possibly disconnected) surface, showing that S' is essential entails verifying that every component of S' is incompressible, and that at least one component of S' is not boundary-parallel; in fact, we will be able to show that every component of S' is not boundary-parallel, which is stronger than we require. To this end, consider any particular component C' of S' . Since C' lies in the boundary of a closed tubular neighbourhood $N(C)$ of some component C of S , we have the following two cases:

- If C is two-sided, then C' is isotopic to C , so the fact that C' is incompressible and not boundary-parallel follows from the assumption that these conditions are satisfied by every component of S .
- If C is one-sided, then $C' = \partial N(C)$. To see that C' is incompressible, consider a compression disc D for C' in \mathcal{M} ; we need to show that D cannot be an *essential* compression disc. We have the following cases:
 - If $D \subset \mathcal{R}$, then D is a compression disc for \mathcal{R} . Since \mathcal{R} is assumed to be ∂ -irreducible, D cannot be essential.

- If $D \subset N(C)$, then we can use a standard fundamental group argument. Note that $\pi_1(N(C)) \cong \pi_1(C)$, and that the double-covering map $p: \partial N(C) \rightarrow C$ induces an injective homomorphism $p_*: \pi_1(\partial N(C)) \rightarrow \pi_1(C)$. Since ∂D is homotopically trivial in $N(C)$, injectivity of p_* tells us that ∂D must also be homotopically trivial in $C' = \partial N(C)$. Thus, we again see that D cannot be essential.

The upshot is that C' does not admit an essential compression disc, so it is incompressible. To see that C' is not boundary-parallel, suppose instead that this is false. The isotopy of C' into $\partial \mathcal{M}$ defines a product region P in $\mathcal{M} - C'$. Note that C' meets two components of $\mathcal{M} - C'$: the interior of $N(C)$, and the interior of \mathcal{R} . The product region P must coincide with \mathcal{R} . However, this would contradict the assumption that \mathcal{R} is anannular, so we conclude that C' cannot be boundary-parallel.

As mentioned above, this suffices to show that S' is a closed essential surface in \mathcal{M} .

By incompressibility of S' , we can use the normalisation procedure (recall Section 2.4) to ensure that S' is normal with respect to \mathcal{T} . Moreover, since links of ideal vertices of \mathcal{T} correspond to boundary-parallel surfaces, the fact that S' is not boundary-parallel ensures that we have a *nontrivial* normal surface in \mathcal{T} . By Lemma 14, we may further assume that this normal surface S' is maximal.

To recap, we are now in the setting laid out at the beginning of Section 4: we have a suitable normal surface S' , together with a suitable chosen region given by the interior of \mathcal{R} . By assumption, we have that \mathcal{R} is irreducible, ∂ -irreducible and anannular, and also that \mathcal{R} contains no two-sided properly embedded Möbius bands. Thus, all the prerequisites for Theorem 1 are satisfied, and applying this theorem tells us that after crushing S' , one of the benign components (as defined in Section 4) forms an ideal triangulation \mathcal{T}^* of \mathcal{R} . Since S' is a *nontrivial* normal surface, we have $|\mathcal{T}^*| < |\mathcal{T}|$. Hence $\Delta(\mathcal{R}) \leq |\mathcal{T}^*| < |\mathcal{T}| = \Delta(\mathcal{M})$, as required. \square

5.1 Application: hyperbolic JSJ pieces

Let \mathcal{M} be an irreducible and ∂ -irreducible 3-manifold with no 2-sphere boundary components. Recall that by work of Jaco and Shalen [23; 24], and independently by Johannson [25], there is a canonical collection $\{S_i\}$ of finitely many disjoint essential tori in \mathcal{M} such that each piece resulting from cutting along $\bigcup_i S_i$ is either atoroidal or Seifert fibred; formal statements of this result can also be found in [10, Theorem 1.9; 31, Theorem 8.23]. This collection of tori is called the *JSJ decomposition* (or the *torus decomposition*) of \mathcal{M} ; it is closely related to (but not exactly the same as) the decomposition along tori described by the Thurston–Perelman geometrisation theorem. Theorem 2 almost immediately yields the following consequence for JSJ decompositions:

Theorem 16 *Let \mathcal{M} be an orientable 3-manifold with no 2-sphere boundary components. If \mathcal{M} is irreducible, ∂ -irreducible and has nonempty JSJ decomposition $\{S_i\}$, then any hyperbolic component \mathcal{H} that results from cutting \mathcal{M} along $\bigcup_i S_i$ satisfies $\Delta(\mathcal{H}) < \Delta(\mathcal{M})$.*

Proof Suppose that after cutting along the tori in the JSJ decomposition of \mathcal{M} , (at least) one of the resulting pieces \mathcal{H} is hyperbolic. By Thurston's hyperbolisation theorem, \mathcal{H} is irreducible, ∂ -irreducible and anannular. Moreover, orientability of \mathcal{M} implies that \mathcal{H} is orientable, and hence that \mathcal{H} contains no two-sided properly embedded Möbius bands. Thus, by Theorem 2, we have $\Delta(\mathcal{H}) < \Delta(\mathcal{M})$. \square

5.2 Application: satellite knots

Recall that the *exterior* of a knot or link L in S^3 is the 3-manifold obtained by deleting an open regular neighbourhood of L from S^3 . We take the *triangulation complexity* of a link L , denoted $\Delta(L)$, to mean the triangulation complexity of the exterior of L .

Our goal now is to present an easy consequence of Theorem 2 concerning the triangulation complexity of satellite knots; see Theorem 18 below. For this, we first review the definition of a satellite knot:

Definitions 17 Let V^* denote the solid torus given by the exterior of an unknot U^* , let C^* denote the core circle of V^* , and let $e: V^* \rightarrow S^3$ be an embedding such that the image of C^* under e is a nontrivial knot C . Consider a knot K^* in the interior of V^* such that

- K^* is not isotopic (inside V^*) to C^* ; and
- every meridional disc of V^* meets K^* at least once.

The image of K^* under e is a nontrivial knot K called a *satellite knot*. We call the knot C a *companion* of K , and we call the torus $e(\partial V^*)$ a *companion torus* of K . We also call the link $K^* \cup U^*$ a *pattern* of K .

Theorem 18 Let K be a satellite knot, and let L denote either a companion or a pattern of K . If the exterior of L is anannular, then $\Delta(L) < \Delta(K)$.

It is worth noting that every hyperbolic link is anannular. Thus, Theorem 18 applies to a very large class of satellite knots; Figure 44 shows one example of such a satellite knot.




Figure 44: An example of a satellite knot, together with a companion and a pattern; in this case, both the companion and the pattern are hyperbolic, and hence anannular. Left: a satellite knot K , the untwisted Whitehead double of the figure-eight knot. Middle: a companion of K , the figure-eight knot. Right: a pattern of K , the Whitehead link.

Proof of Theorem 18 Let \bar{K} and \bar{L} denote the exteriors of K and L , respectively. Since \bar{L} is one of the components given by cutting the exterior of \bar{K} along a companion torus (which is an essential torus), we can apply Theorem 2 provided that \bar{L}

- is irreducible, ∂ -irreducible and anannular; and
- contains no two-sided properly embedded Möbius bands.

We have assumed that \bar{L} is anannular, and the fact that \bar{L} is orientable implies that it contains no two-sided properly embedded Möbius bands. Moreover, in the case where L is a companion, irreducibility and ∂ -irreducibility follow from the fact that L must be a nontrivial knot; on the other hand, when L is a pattern, irreducibility and ∂ -irreducibility follow from the fact that L must be a nonsplit link. The upshot is that, by Theorem 2, we have $\Delta(L) = \Delta(\bar{L}) < \Delta(\bar{K}) = \Delta(K)$. \square

Corollary 19 Consider the connected sum $K \# K'$ of two nontrivial knots K and K' . If K is a hyperbolic knot, then $\Delta(K) < \Delta(K \# K')$.

Proof Recall that the connected sum $K \# K'$ can be viewed as a satellite knot with companion given by either of its summands. Thus, if the summand K is hyperbolic, in which case the exterior of K is anannular, then it follows immediately from Theorem 18 that $\Delta(K) < \Delta(K \# K')$. \square

5.3 Application: rod complements in the 3-torus

The crushing techniques developed in previous sections can also be used to study link exteriors in ambient spaces other than the 3-sphere, such as the 3-torus \mathbb{T}^3 . Hui and Purcell [12] initiated the use of 3-dimensional geometry and topology to study rod packing structures in crystallography. In crystallographic chemistry, a rod packing is a packing of uniform cylinders (also called rods) that represent linear or zigzag chains of particles. Readers may refer to [29; 30] for examples of rod packing structures.

Many rod packing structures exhibit translational symmetry in each dimension of 3-dimensional Euclidean space. Taking a quotient by this symmetry allows such rod packings to be encoded as geodesic links in the 3-torus, which we can then study using tools from 3-manifold geometry and topology. Each component of such a geodesic link is called a *rod-shaped circle*, or often simply a *rod*, in \mathbb{T}^3 ; the complement of such a link in \mathbb{T}^3 is a 3-manifold called a *rod complement*.

For \mathcal{M} belonging to a large family of rod complements, the JSJ decomposition gives a unique hyperbolic piece \mathcal{H} , and this hyperbolic piece is also a rod complement. Theorem 21 states this precisely, and then (by applying Theorem 16) relates the triangulation complexities of \mathcal{M} and \mathcal{H} .

To study the JSJ decompositions of rod complements, we rely on work of Hui [11], which completely classified the geometry of all rod complements using simple linear algebra conditions. The linear algebra

arises from the fact that a rod in \mathbb{T}^3 lifts to a straight line in the universal cover \mathbb{R}^3 . This linear structure leads to the natural concepts of

- *linear independence* of rods in \mathbb{T}^3 ; and
- *linear isotopy* of rods in \mathbb{T}^3 , which roughly means an isotopy along a planar annulus between two parallel rods.

These notions appear in the statements of Lemma 20 and Theorem 21 below; readers may refer to [11; 12] for definitions of these notions, as well as for explanations of other related terminology that we use in the proofs.

Lemma 20 *Let \mathcal{M} be a toroidal rod complement in the 3-torus with at least three linearly independent rods. Any essential torus in \mathcal{M} bounds a solid torus in \mathbb{T}^3 whose interior contains two or more linearly isotopic rods.*

Proof Since \mathcal{M} is toroidal, by Proposition 3.8 in [11], either \mathcal{M} is a rod complement with all rods spanning a plane torus, or there exist disjoint parallel rods that are linearly isotopic in the complement of the other rods (or possibly both). The assumption that \mathcal{M} is a rod complement with three linearly independent rods thus implies the existence of at least two linearly isotopic parallel rods.

Let T_e be an essential torus in \mathcal{M} . Note that T_e is not in the homotopy class of a plane torus because there exist three linearly independent rods for \mathcal{M} . Hence, the torus T_e , essential in \mathcal{M} , has at least one generator in $\pi_1(T_e) \cong \mathbb{Z} \times \mathbb{Z}$ represented by an essential loop that is homotopically trivial in \mathbb{T}^3 .

By Lemma 3.7 in [11], there exists a compression disc D for T_e in \mathbb{T}^3 . It then follows from the irreducibility of \mathbb{T}^3 that T_e is separating. As T_e is incompressible in \mathcal{M} , some rod must intersect D . Hence, the other generator of $\pi_1(T_e) \cong \mathbb{Z} \times \mathbb{Z}$ is represented by a loop that is homotopically nontrivial in \mathbb{T}^3 . By Lemma 3.9 in [11], T_e bounds a solid torus V_e in \mathbb{T}^3 . Since T_e is not boundary-parallel, the interior of V_e contains two or more parallel rods that are linearly isotopic in the complement of the other rods. \square

Theorem 21 *Let \mathcal{M} be a toroidal rod complement in the 3-torus with at least three linearly independent rods. The JSJ decomposition of \mathcal{M} gives a unique (up to homeomorphism) hyperbolic rod complement \mathcal{H} with $\Delta(\mathcal{H}) < \Delta(\mathcal{M})$.*

Proof Without loss of generality, assume that \mathcal{M} is the complement of an open neighbourhood of all the rods in the 3-torus.

We first show that \mathcal{M} satisfies all the assumptions in Theorem 16. Note that \mathcal{M} is an orientable 3-manifold with no 2-sphere boundary components. Since \mathcal{M} is the complement of finitely many rods in the 3-torus, Proposition 3.6 in [11] tells us that \mathcal{M} is irreducible and ∂ -irreducible. We also know that \mathcal{M} has nonempty JSJ decomposition because it is toroidal by assumption, and because it is not Seifert fibred by Theorem 4.1 in [11]. Thus, Theorem 16 applies to any hyperbolic component that results from cutting along the JSJ tori.

Let $\{T_i\}$ be the set of essential tori that gives the JSJ decomposition of \mathcal{M} ; recall that this set of tori is unique up to isotopy, which means that the JSJ pieces (the 3-manifold components obtained by cutting along these tori) are unique up to homeomorphism. By Lemma 20, each essential torus T_i bounds a solid torus V_i in \mathbb{T}^3 whose interior contains two or more linearly isotopic rods. Thus, one of the JSJ pieces is a 3-manifold $\mathcal{H} := \mathcal{M} - \bigcup_i \text{int}(V_i)$ whose interior is homeomorphic to a rod complement. We will show that \mathcal{H} is the unique hyperbolic JSJ piece for \mathcal{M} .

To do this, we repeatedly appeal to the classification given by Theorem 4.1 in [11]. Since \mathcal{M} has at least three linearly independent rods, observe that the same must be true for \mathcal{H} . Thus, by the classification, \mathcal{H} cannot be Seifert fibred. This means that \mathcal{H} must be atoroidal, since it is one of the JSJ pieces. Using the classification again, we therefore see that \mathcal{H} cannot have a pair of linearly isotopic rods. This has two implications:

- First, we must have exactly one essential torus T_i for each linear isotopy class containing at least two rods from \mathcal{M} , and the corresponding solid torus V_i in \mathbb{T}^3 must contain *all* of the rods in this class.
- Second, by the classification, the interior of \mathcal{H} admits a complete hyperbolic structure.

To see that \mathcal{H} is the only hyperbolic JSJ piece, observe that each of the other JSJ pieces is obtained from one of the solid tori V_i by deleting a small neighbourhood of all the (linearly isotopic) rods inside V_i ; in other words, all the other JSJ pieces are Seifert fibred, since they are homeomorphic to solid tori with at least two core curves removed.

Finally, since \mathcal{H} is a hyperbolic piece obtained after cutting along the JSJ tori for \mathcal{M} , it follows from Theorem 16 that $\Delta(\mathcal{H}) < \Delta(\mathcal{M})$. \square

References

- [1] **M Bucher, R Frigerio, C Pagliantini**, *The simplicial volume of 3-manifolds with boundary*, *J. Topol.* 8 (2015) 457–475 MR
- [2] **B A Burton**, *Computational topology with Regina: algorithms, heuristics and implementations*, from “Geometry and topology down under” (CD Hodgson, WH Jaco, MG Scharlemann, S Tillmann, editors), *Contemp. Math.* 597, Amer. Math. Soc., Providence, RI (2013) 195–224 MR
- [3] **B A Burton**, *A new approach to crushing 3-manifold triangulations*, *Discrete Comput. Geom.* 52 (2014) 116–139 MR
- [4] **B A Burton, R Budney, W Pettersson**, et al., *Regina: Software for low-dimensional topology* Available at <https://regina-normal.github.io>
- [5] **B A Burton, A Coward, S Tillmann**, *Computing closed essential surfaces in knot complements*, from “Computational geometry” (GD da Fonseca, T Lewiner, L Peñaranda, editors), ACM, New York (2013) 405–413 MR
- [6] **B A Burton, A He**, *Finding large counterexamples by selectively exploring the Pachner graph*, from “39th International Symposium on Computational Geometry” (E W Chambers, J Gudmundsson, editors), LIPIcs. Leibniz Int. Proc. Inform. 258, Schloss Dagstuhl. Leibniz-Zent. Inform., Wadern (2023) art. id. 21 MR

- [7] **B A Burton, M Ozlen**, *A fast branching algorithm for unknot recognition with experimental polynomial-time behaviour*, preprint (2014) arXiv 1211.1079v3
- [8] **B A Burton, S Tillmann**, *Computing closed essential surfaces in 3-manifolds*, *J. Appl. Comput. Topol.* 9 (2025) art. id. 18 MR
- [9] **J Fowler**, *Finding 0-efficient triangulations of 3-manifolds*, senior honors thesis, Harvard University (2003)
- [10] **A Hatcher**, *Notes on basic 3-manifold topology* Available at <https://pi.math.cornell.edu/~hatcher/3M/3Mdownloads.html>
- [11] **C O Y Hui**, *A geometric classification of rod complements in the 3-torus*, *Proc. Amer. Math. Soc.* 153 (2025) 381–394 MR
- [12] **C O Y Hui, J S Purcell**, *On the geometry of rod packings in the 3-torus*, *Bull. Lond. Math. Soc.* 56 (2024) 1291–1309 MR
- [13] **K Ichihara, Y Nishimura, S Tani**, *The computational complexity of classical knot recognition*, *J. Knot Theory Ramifications* 32 (2023) art. id. 2350069 MR
- [14] **A Jackson**, *Minimal triangulation size of Seifert fibered spaces with boundary*, preprint (2023) arXiv 2301.02085
- [15] **W Jaco, J Johnson, J Spreer, S Tillmann**, *Bounds for the genus of a normal surface*, *Geom. Topol.* 20 (2016) 1625–1671 MR
- [16] **W Jaco, H Rubinstein, J Spreer, S Tillmann**, *On minimal ideal triangulations of cusped hyperbolic 3-manifolds*, *J. Topol.* 13 (2020) 308–342 MR
- [17] **W Jaco, H Rubinstein, S Tillmann**, *Minimal triangulations for an infinite family of lens spaces*, *J. Topol.* 2 (2009) 157–180 MR
- [18] **W Jaco, J H Rubinstein**, *0-efficient triangulations of 3-manifolds*, *J. Differential Geom.* 65 (2003) 61–168 MR
- [19] **W Jaco, J H Rubinstein, J Spreer, S Tillmann**, *\mathbb{Z}_2 -Thurston norm and complexity of 3-manifolds, II*, *Algebr. Geom. Topol.* 20 (2020) 503–529 MR
- [20] **W Jaco, J H Rubinstein, J Spreer, S Tillmann**, *Complexity of 3-manifolds obtained by Dehn filling*, *Algebr. Geom. Topol.* 25 (2025) 301–327 MR
- [21] **W Jaco, J H Rubinstein, S Tillmann**, *Coverings and minimal triangulations of 3-manifolds*, *Algebr. Geom. Topol.* 11 (2011) 1257–1265 MR
- [22] **W Jaco, J H Rubinstein, S Tillmann**, *\mathbb{Z}_2 -Thurston norm and complexity of 3-manifolds*, *Math. Ann.* 356 (2013) 1–22 MR
- [23] **W Jaco, P B Shalen**, *A new decomposition theorem for irreducible sufficiently-large 3-manifolds*, from “Algebraic and geometric topology, part 2” (R J Milgram, editor), *Proc. Sympos. Pure Math.* 32, Amer. Math. Soc., Providence, RI (1978) 71–84 MR
- [24] **W H Jaco, P B Shalen**, *Seifert fibered spaces in 3-manifolds*, *Mem. Amer. Math. Soc.* 220 (1979) MR
- [25] **K Johannson**, *Homotopy equivalences of 3-manifolds with boundaries*, *Lecture Notes in Math.* 761, Springer (1979) MR
- [26] **M Lackenby, J S Purcell**, *The triangulation complexity of elliptic and sol 3-manifolds*, *Math. Ann.* 390 (2024) 1623–1667 MR

- [27] **M Lackenby, JS Purcell**, *The triangulation complexity of fibred 3-manifolds*, *Geom. Topol.* 28 (2024) 1727–1828 MR
- [28] **S Matveev**, *Algorithmic topology and classification of 3-manifolds*, 2nd edition, *Algorithms and Computation in Mathematics* 9, Springer (2007) MR
- [29] **M O’Keeffe, J Plévert, T Ogawa**, *Homogeneous cubic cylinder packings revisited*, *Acta Crystallogr. Sect. A* 58 (2002) 125–132 MR
- [30] **M O’Keeffe, J Plévert, Y Teshima, Y Watanabe, T Ogama**, *The invariant cubic rod (cylinder) packings: symmetries and coordinates*, *Acta Cryst. Sect. A* 57 (2001) 110–111 MR
- [31] **JS Purcell**, *Hyperbolic knot theory*, *Graduate Studies in Math.* 209, Amer. Math. Soc., Providence, RI (2020) MR
- [32] **JH Rubinstein, J Spreer, S Tillmann**, *A new family of minimal ideal triangulations of cusped hyperbolic 3-manifolds*, from “2021–2022 MATRIX annals” (DR Wood, J de Gier, CE Praeger, editors), *MATRIX Book Ser.* 5, Springer (2024) 5–28 MR
- [33] **A Y Vesnin, EA Fominykh**, *Exact values of complexity for Paoluzzi–Zimmermann manifolds*, *Dokl. Math.* (2011) 542–544 MR

BAB: *School of Mathematics and Physics, University of Queensland
Brisbane, QLD, Australia*

TdP, COYH: *School of Mathematics, Monash University
Melbourne, VIC, Australia*

AH: *Department of Mathematics, Oklahoma State University
Stillwater, OK, United States*

bab@maths.uq.edu.au, thiago.depaivasouza@monash.edu, alex.he@okstate.edu,
onyu.hui@monash.edu

<https://people.smp.uq.edu.au/BenjaminBurton/>,
<https://sites.google.com/view/thiago-de-paiva>,
<https://sites.google.com/view/alex-he>, <https://sites.google.com/view/oyhui>

Received: 10 April 2024 Revised: 9 October 2024

ALGEBRAIC & GEOMETRIC TOPOLOGY

msp.org/agt

EDITORS

PRINCIPAL ACADEMIC EDITORS

John Etnyre
etnyre@math.gatech.edu
Georgia Institute of Technology

Kathryn Hess
kathryn.hess@epfl.ch
École Polytechnique Fédérale de Lausanne

BOARD OF EDITORS

Julie Bergner	University of Virginia jeb2md@eservices.virginia.edu	Thomas Koberda	University of Virginia thomas.koberda@virginia.edu
Steven Boyer	Université du Québec à Montréal cohf@math.rochester.edu	Markus Land	LMU München markus.land@math.lmu.de
Tara E Brendle	University of Glasgow tara.brendle@glasgow.ac.uk	Christine Lescop	Université Joseph Fourier lescop@ujf-grenoble.fr
Indira Chatterji	CNRS & Univ. Côte d'Azur (Nice) indira.chatterji@math.cnrs.fr	Norihiko Minami	Yamato University minami.norihiko@yamato-u.ac.jp
Octav Cornea	Université de Montreal cornea@dms.umontreal.ca	Andrés Navas	Universidad de Santiago de Chile andres.navas@usach.cl
Alexander Dranishnikov	University of Florida dranish@math.ufl.edu	Jessica S Purcell	Monash University jessica.purcell@monash.edu
Tobias Ekholm	Uppsala University, Sweden tobias.ekholm@math.uu.se	Birgit Richter	Universität Hamburg birgit.richter@uni-hamburg.de
Mario Eudave-Muñoz	Univ. Nacional Autónoma de México mario@matem.unam.mx	Jérôme Scherer	École Polytech. Féd. de Lausanne jerome.scherer@epfl.ch
David Futер	Temple University dfuter@temple.edu	Vesna Stojanoska	Univ. of Illinois at Urbana-Champaign vesna@illinois.edu
John Greenlees	University of Warwick john.greenlees@warwick.ac.uk	Zoltán Szabó	Princeton University szabo@math.princeton.edu
Matthew Hedden	Michigan State University mhedden@math.msu.edu	Maggy Tomova	University of Iowa maggy-tomova@uiowa.edu
Kristen Hendricks	Rutgers University kristen.hendricks@rutgers.edu	Daniel T Wise	McGill University, Canada daniel.wise@mcgill.ca
Hans-Werner Henn	Université Louis Pasteur henn@math.u-strasbg.fr	Lior Yanovski	Hebrew University of Jerusalem lior.yanovski@gmail.com
Daniel Isaksen	Wayne State University isaksen@math.wayne.edu		


See inside back cover or msp.org/agt for submission instructions.

The subscription price for 2025 is US \$760/year for the electronic version, and \$1110/year (+\$75, if shipping outside the US) for print and electronic. Subscriptions, requests for back issues and changes of subscriber address should be sent to MSP. Algebraic & Geometric Topology is indexed by Mathematical Reviews, Zentralblatt MATH, Current Mathematical Publications and the Science Citation Index.

Algebraic & Geometric Topology (ISSN 1472-2747 printed, 1472-2739 electronic) is published 9 times per year and continuously online, by Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004. Periodical rate postage paid at Oakland, CA 94615-9651, and additional mailing offices. POSTMASTER: send address changes to Mathematical Sciences Publishers, 2000 Allston Way # 59, Berkeley, CA 94701-4004.

AGT peer review and production are managed by EditFlow® from MSP.

PUBLISHED BY

 **mathematical sciences publishers**
nonprofit scientific publishing

<https://msp.org/>

© 2025 Mathematical Sciences Publishers

ALGEBRAIC & GEOMETRIC TOPOLOGY

Volume 25 Issue 8 (pages 4437–5174) 2025

Hierarchies for relatively hyperbolic virtually special groups	4437
EDUARD EINSTEIN	
Intersection norms on surfaces and Birkhoff sections for geodesic flows	4499
MARCOS COSSARINI and PIERRE DEHORNOY	
Thin knots and the cabling conjecture	4547
ROBERT DEYESO III	
Linear linkless embeddings: proof of a conjecture by Sachs	4585
LYNN STANFIELD	
Constructing rational homology 3-spheres that bound rational homology 4-balls	4599
LISA LOKTEVA	
Homological stability for the ribbon Higman–Thompson groups	4633
RACHEL SKIPPER and XIAOLEI WU	
A group-theoretic framework for low-dimensional topology, or: how not to study low-dimensional topology?	4667
SARAH BLACKWELL, ROBION KIRBY, MICHAEL KLUG, VINCENT LONGO and BENJAMIN RUPPIK	
Classification of genus-two surfaces in S^3	4719
FILIPPO BARONI	
Meromorphic projective structures: signed spaces, grafting and monodromy	4787
SPANDAN GHOSH and SUBHOJOY GUPTA	
A new twist on modular links from an old perspective	4827
KHANH LE	
Flat fully augmented links are determined by their complements	4839
CHRISTIAN MILLICHAP and ROLLAND TRAPP	
BNSR-invariants of surface Houghton groups	4897
NOAH TORGERSON and JEREMY WEST	
Topological symmetry groups of the generalized Petersen graphs	4921
ANGELYNN ÁLVAREZ, ERICA FLAPAN, MARK HUNNELL, JOHN HUTCHENS, EMILLE LAWRENCE, PAUL LEWIS, CANDICE PRICE and RUTH VANDERPOOL	
Crushing surfaces of positive genus	4949
BENJAMIN A BURTON, THIAGO DE PAIVA, ALEXANDER HE and CONNIE ON YU HUI	
Annular links from Thompson’s group T	5013
LOUISA LILES	
Realizing pairs of multicurves as cylinders on translation surfaces	5031
JULIET AYGUN, JANET BARKDOLL, AARON CALDERON, JENAVIE LORMAN and THEODORE SANDSTROM	
Involutive Khovanov homology and equivariant knots	5059
TAKETO SANO	
A diagrammatic computation of abelian link invariants	5113
DAVID CIMASONI, LIVIO FERRETTI and JESSICA LIU	
Equivariant double-slice genus, stabilization, and equivariant stabilization	5137
MALCOLM GABBARD	
Coarse and bi-Lipschitz embeddability of subspaces of the Gromov–Hausdorff space into Hilbert spaces	5153
NICOLÒ ZAVA	