

Finding Large Counterexamples by Selectively Exploring the Pachner Graph

Benjamin A. Burton  

The University of Queensland, Brisbane, Australia

Alexander He   

The University of Queensland, Brisbane, Australia

Abstract

We often rely on censuses of triangulations to guide our intuition in 3-manifold topology. However, this can lead to misplaced faith in conjectures if the smallest counterexamples are too large to appear in our census. Since the number of triangulations increases super-exponentially with size, there is no way to expand a census beyond relatively small triangulations – the current census only goes up to 10 tetrahedra. Here, we show that it is feasible to search for large and hard-to-find counterexamples by using heuristics to *selectively* (rather than exhaustively) enumerate triangulations. We use this idea to find counterexamples to three conjectures which ask, for certain 3-manifolds, whether one-vertex triangulations always have a “distinctive” edge that would allow us to recognise the 3-manifold.

2012 ACM Subject Classification Mathematics of computing → Topology

Keywords and phrases Computational topology, 3-manifolds, Triangulations, Counterexamples, Heuristics, Implementation, Pachner moves, Bistellar flips

Digital Object Identifier 10.4230/LIPIcs.SoCG.2023.21

Related Version *Full Version:* <https://arxiv.org/abs/2303.06321> [9]

Supplementary Material *Software (Source Code):*
<https://github.com/AlexHe98/triang-counterex>

Funding *Alexander He:* Supported by an Australian Government Research Training Program Scholarship.

Acknowledgements We thank the referees for their helpful comments.

1 Introduction

Many conjectures in computational geometry and topology are true in small cases, but turn out to have a relatively large counterexample. Perhaps the most notable example of this is the *Hirsch conjecture*, which posited that in a d -dimensional polytope with n facets, any two vertices can be connected by a path of at most $n - d$ edges. This is true when the dimension is small (specifically, when $d \leq 3$ [20]), as well as when the number of facets is small (specifically, when $n \leq d + 6$ [2]). However, this is not indicative of the general behaviour: Santos showed that there is a 43-dimensional counterexample with 86 facets [31].

For an example from 3-manifold topology, consider the Seifert fibre spaces (see Hatcher’s notes [15] for a definition). There are 302 Seifert fibre spaces that can be triangulated with no more than 7 tetrahedra; for all of these, at least one **minimal triangulation** (triangulation with the smallest possible number of tetrahedra) uses a standard prism-and-layering construction. This pattern does not persist: there is a Seifert fibre space whose unique (8-tetrahedron) minimal triangulation is given by the isomorphism signature `iLLLPQcbcgffghhhtsmhgosof` (the software *Regina* [8] can convert this back into a triangulation), and this is instead



© Benjamin A. Burton and Alexander He;

licensed under Creative Commons License CC-BY 4.0

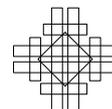
39th International Symposium on Computational Geometry (SoCG 2023).

Editors: Erin W. Chambers and Joachim Gudmundsson; Article No. 21; pp. 21:1–21:16

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



constructed from a gadget (called the *brick* B_5) that was first identified in Martelli and Petronio’s census of minimal triangulations up to 9 tetrahedra [22]; although the first author [3] and Matveev [24] have since extended the census, no other gadgets like this have been found.

We present a technique for finding large counterexamples in a similar setting; we study arbitrary one-vertex triangulations, not just minimal ones. The obvious source of counterexamples is a census of all triangulations up to a given number of tetrahedra. However, this only captures small counterexamples because the number of triangulations grows super-exponentially as we increase the number of tetrahedra; to date, the census of all closed 3-manifold triangulations only goes up to 10 tetrahedra, and this already includes over 2 billion triangulations, constituting over 63 GB of data [3].

How can we find a counterexample that is too large to appear in our census? We showcase a method of *selectively* (rather than exhaustively) enumerating triangulations, which yields large counterexamples to three conjectures posed by Saul Schleimer at the 2022 Dagstuhl workshop on *Computation and Reconfiguration in Low-Dimensional Topological Spaces*.

1.1 The conjectures

Each conjecture concerns edges of one-vertex triangulations. Since such edges realise embedded closed curves, we can ask whether they are embedded in an interesting way.

For instance, in a lens space, consider an edge e that forms a **core curve**, meaning that the complement of a regular neighbourhood of the curve is a solid torus; in this case, we call e a **core edge**. Finding such an edge certifies that the 3-manifold is a lens space.

► **Conjecture 1.** *Every one-vertex triangulation of a lens space has a core edge.*

Since solid torus recognition can be solved efficiently in practice [11, 5], proving Conjecture 1 would have provided a relatively efficient method for recognising lens spaces. Lens space recognition can also be used to determine whether a 3-manifold is **elliptic** (i.e., admits spherical geometry). Indeed, Lackenby and Schleimer [21] recently used the following variant of Conjecture 1 to show that recognising elliptic manifolds is in NP:

► **Theorem** ([21], Theorem 9.4). *Let \mathcal{M} be a lens space that is neither $\mathbb{R}P^3$ nor a prism manifold, and let \mathcal{T} be any triangulation of \mathcal{M} . Then the 86th iterated barycentric subdivision of \mathcal{T} contains a sequence of edges that forms a core curve.*

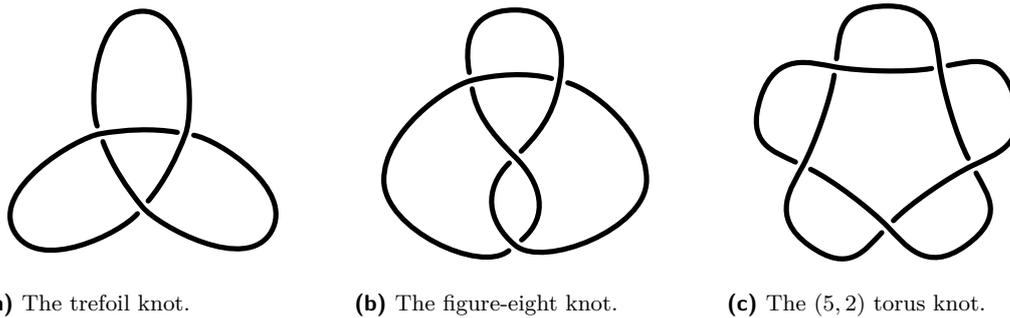
The other two conjectures have similar motivations: special edges can help us recognise certain 3-manifolds, so we would like to know whether such edges always exist.

The second conjecture concerns **tunnel number**: the smallest number of arcs that we need to add to a knot so that the complement becomes a handlebody. Let \mathcal{T} be an ideal triangulation of the complement of a knot K , so that each edge e in \mathcal{T} forms an arc α that meets K at its endpoints. We call e a **tunnel edge** if the complement of $K \cup \alpha$ is a genus-2 handlebody. If K is not the unknot, then the existence of a tunnel edge implies that K has tunnel number equal to one; Figure 1 shows three knots with tunnel number one.

► **Conjecture 2.** *Let K be a knot with tunnel number equal to one. Then every one-vertex ideal triangulation of the complement of K has a tunnel edge.*

The third conjecture concerns **Seifert fibre spaces**, which are fibred by circles in a particular way (again, see Hatcher’s notes [15]); we call each circle fibre a **Seifert fibre**. Seifert fibre spaces play a central role in torus decompositions for irreducible 3-manifolds:

- The JSJ decomposition (due to Jaco and Shalen [17], and Johannson [19]) cuts such 3-manifolds into pieces that are either atoroidal or Seifert fibred.



■ **Figure 1** Three knots with tunnel number equal to one.

■ The geometric decomposition from Thurston’s geometrisation conjecture (famously proved by Perelman [25, 27, 26]) cuts such 3-manifolds into pieces that are either hyperbolic or graph manifolds; the graph manifolds can be further decomposed into Seifert fibre spaces. The **small** Seifert fibre spaces do not contain any embedded two-sided incompressible surfaces, which makes them relatively difficult to work with.

► **Conjecture 3.** *Every one-vertex triangulation of a small Seifert fibre space has an edge isotopic to a Seifert fibre.*

1.2 Using a targeted search to find counterexamples

As mentioned earlier, we have found counterexamples to all three conjectures listed in section 1.1. The key ingredient for finding these examples is a heuristic for measuring how “far away” a triangulation is from having no “distinguished” edges (such as core edges).

To see why such a heuristic is necessary, we note that the census of triangulations up to 10 tetrahedra contains 422 533 279 one-vertex triangulations of the 3-sphere (the simplest possible lens space). All of these have at least one core edge; verifying this required 22 hours of wall time on 12 threads, but this does not include the time required to: (1) generate the census, and (2) identify the 3-spheres in this census (which had been done previously [3, 4]).

The upshot is that this exhaustive search was both expensive and unsuccessful. In contrast, our heuristic enabled a targeted search, which produced all of our counterexamples in just *minutes* of wall time. See section 5 for detailed computational results, as well as an explanation of how to turn our counterexamples into infinite families.

We introduce our heuristic and our targeted search algorithm in section 4. In section 3, we discuss some auxiliary algorithms, which are of independent interest; in particular, we present (and implement) an improved algorithm for handlebody recognition. We finish by mentioning some unanswered questions in section 6.

2 Preliminaries

2.1 Triangulations

A (**generalised**) **triangulation** \mathcal{T} is a finite collection of tetrahedra whose triangular faces may be affinely identified in pairs; each equivalence class of such faces is a **face** of \mathcal{T} . The **size** of \mathcal{T} , denoted $|\mathcal{T}|$, is the number of tetrahedra in \mathcal{T} . The **boundary** of \mathcal{T} consists of all the faces that are not identified with any other faces. We allow faces of the same tetrahedron to be identified, so our triangulations need not be simplicial complexes.

The face identifications also yield equivalence classes of vertices and edges of the tetrahedra; we call these classes the **vertices** and **edges** of \mathcal{T} , respectively. In this paper, we are particularly interested in **one-vertex triangulations**, which have exactly one vertex. The **degree** of an edge e of \mathcal{T} is the number of tetrahedra that meet e , counted with multiplicity. The **1-skeleton** of \mathcal{T} is the subcomplex consisting only of the vertices and edges of \mathcal{T} .

If a point lies on the boundary of a triangulation, then we say that it is **boundary**; otherwise, we say that it is **internal**. If a face, edge or vertex consists *entirely* of boundary points, then we say that it is **boundary**; otherwise, we say that it is **internal**.

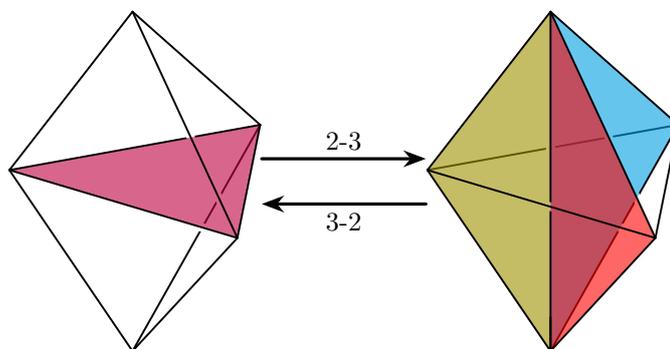
For a triangulation \mathcal{T} to be a 3-manifold, every point p in \mathcal{T} must be **non-singular** – i.e., have a small neighbourhood bounded by either a disc (if p is boundary) or a 2-sphere (if p is internal). This could fail for vertices and midpoints of edges. For this paper, we insist that no edge is identified with itself in reverse; this ensures that midpoints of edges are non-singular.

We call \mathcal{T} a **3-manifold triangulation** if all its vertices are non-singular, because in this case \mathcal{T} will genuinely be a 3-manifold. However, we also allow vertices to be **ideal**, meaning that a small neighbourhood of the vertex is bounded by a closed surface other than the 2-sphere. **Ideal triangulations** (which contain one or more ideal vertices) give a useful representation of the 3-manifold obtained by truncating the ideal vertices; this idea originated with Thurston’s ideal triangulation of the figure-eight knot complement [33, Example 1.4.8].

To concretely encode a triangulation, we give each tetrahedron both a label and an ordering of its four vertices. Two triangulations are **isomorphic** if they are identical up to relabelling tetrahedra and/or reordering tetrahedron vertices. We can uniquely identify any isomorphism class of triangulations using an efficiently-computable string called an **isomorphism signature**. There are many ways to formulate isomorphism signatures; we use the implementation in **Regina** [5, 8], which is described in [4].

2.2 The 2-3 and 3-2 moves

Given any one-vertex triangulation \mathcal{T} with at least two tetrahedra, we can produce a new one-vertex triangulation of the same 3-manifold by performing a **2-3 move** about a (triangular) face that meets two distinct tetrahedra. This replaces these two tetrahedra with three tetrahedra attached around a new edge e ; see Figure 2. We call the inverse move a **3-2 move** about e ; it is possible to perform a 3-2 move about an edge if and only if this edge is an internal degree-3 edge that actually meets three distinct tetrahedra.



■ **Figure 2** The 2-3 and 3-2 moves.

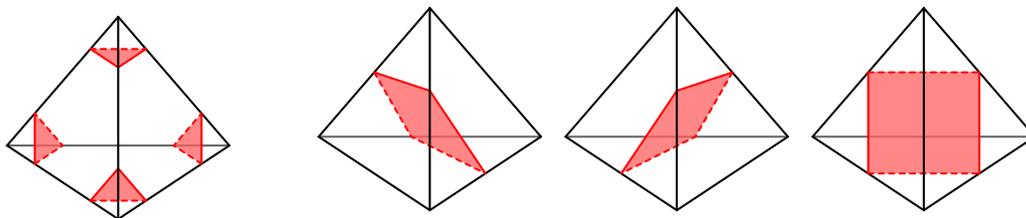
Consider a 3-manifold \mathcal{M} . We can think of the one-vertex triangulations of \mathcal{M} with at least two tetrahedra as nodes of an infinite graph, with two nodes connected by an undirected arc if and only if the corresponding triangulations are related by a 2-3 move. This graph, called the **Pachner graph** of \mathcal{M} , is known to be connected [1, 23, 28, 30].

2.3 Normal surfaces

We now outline the basics of normal surface theory; see [14] and [16] for more comprehensive discussion. A properly embedded surface S in a triangulation \mathcal{T} is a **normal surface** if:

- S meets each simplex (i.e., vertex, edge, triangle, or tetrahedron) of \mathcal{T} transversely; and
- S meets each tetrahedron Δ of \mathcal{T} in a finite (and possibly empty) collection of discs – known as **elementary discs** – where each such disc is a curvilinear triangle or quadrilateral whose vertices lie on different edges of Δ .

Up to **normal isotopy** – an ambient isotopy that preserves every simplex of \mathcal{T} – every elementary disc has one of the seven types shown in Figure 3.



(a) The four triangle types. (b) The three quadrilateral types.

■ **Figure 3** The seven elementary disc types in a tetrahedron.

The simplest example of a normal surface is a **vertex-linking surface**, which consists entirely of triangles. Such surfaces always exist, so finding these surfaces never gives us any new information about the underlying 3-manifold. For this reason, we consider a normal surface to be **non-trivial** when it includes at least one quadrilateral.

If \mathcal{T} has size n , then we can represent any normal surface in \mathcal{T} as a **normal coordinate vector** in \mathbb{R}^{7n} that counts the number of elementary discs of each type in each tetrahedron. It is often sufficient to focus on a class of surfaces called the **vertex normal surfaces**; roughly, these form a finite and algorithmically enumerable “basis” for the set of all normal surfaces.

3 Key tools

This section discusses some new algorithms and implementations that were crucial for our work. These algorithms may be of independent interest.

3.1 Handlebody recognition

To find counterexamples to Conjectures 1 and 2, we need algorithms to recognise solid tori and genus-2 handlebodies. **Regina** includes an implementation of solid torus recognition [13, 18] that is remarkably efficient in practice [8, 5, 11]. We generalise this to recognise handlebodies of arbitrary genus, improving the earlier algorithm of Jaco and Tollefson [18, Algorithm 9.3].

Handlebodies are **irreducible** and have **compressible** boundary (see Definitions 2.2.1 and 3.3.2 of [23]), so we can exploit the following consequence of Corollary 6.4 of [18]:

► **Theorem 4.** *Let \mathcal{T} be a triangulation of a compact irreducible 3-manifold with compressible boundary. Then \mathcal{T} contains a vertex normal essential compressing disc.*

The strategy of our algorithm is to repeatedly cut along vertex normal compressing discs; a triangulation \mathcal{T} is a handlebody if and only if \mathcal{T} eventually gets decomposed into a collection of 3-balls. For computational reasons, we modify this approach in two ways.

First, we expand our focus beyond discs, to include any vertex normal surface S whose Euler characteristic $\chi(S)$ is positive. The rationale is that $\chi(S)$ can be expressed as a linear function in normal coordinates, which allows us to exploit linear programming techniques. This leads to an algorithm for detecting non-trivial normal spheres and discs that is fast in practice (even though, in theory, this is actually the exponential-time bottleneck for both solid torus recognition and handlebody recognition); see [11] for details.

Second, because cutting along a normal surface can increase the number of tetrahedra exponentially, we use a technique called **crushing**. Crushing was introduced by Jaco and Rubinstein [16] following earlier unpublished work of Casson, and was later refined by the first author [6]. Crushing a non-trivial normal surface has the following crucial benefit: it produces a new triangulation with strictly fewer tetrahedra than before. The trade-off is that crushing could have topological side-effects, as detailed in the following theorem:

► **Theorem 5** ([6], Theorem 2). *Let \mathcal{T} be a triangulation of a compact orientable 3-manifold \mathcal{M} , and let S be a normal sphere or disc in \mathcal{T} . Crushing S yields a triangulation \mathcal{T}' whose underlying 3-manifold \mathcal{M}' is obtained from \mathcal{M} by zero or more of the following operations:*

- *undoing connected sums;*
- *cutting along properly embedded discs;*
- *filling boundary 2-spheres with 3-balls; or*
- *deleting 3-ball, 3-sphere, $\mathbb{R}P^3$, $L_{3,1}$ or $S^2 \times S^1$ components.*

For the rest of this section, call a connected orientable 3-manifold **interesting** if it has exactly one boundary component and its first homology is \mathbb{Z}^g , where g is the genus of the boundary surface. Our handlebody recognition algorithm relies on the following result:

► **Proposition 6.** *Let \mathcal{T} be a triangulation of an interesting 3-manifold, and suppose that \mathcal{T} contains a non-trivial normal sphere or disc S . Let \mathcal{T}' denote the triangulation obtained by crushing S . Then each component of \mathcal{T}' is either closed or interesting. Moreover, \mathcal{T} is a handlebody if and only if every component of \mathcal{T}' is either a 3-sphere or a handlebody.*

Proof Outline. To understand the effect of crushing the surface S , it suffices to consider the operations listed in Theorem 5 one at a time. The most intricate part of the proof is showing that if we perform one of these operations on an interesting 3-manifold, then the components of the resulting 3-manifold must all be either closed or interesting. This mostly involves some fairly routine homology arguments. See the full version [9] for details. ◀

The last thing we rely on is an algorithm to recognise the 3-ball; this is a well-known variant of 3-sphere recognition [16, 29, 32], and an implementation is available in **Regina** [5, 8].

The following handlebody recognition algorithm is now also available in **Regina** [8]:

► **Algorithm 7** (Handlebody recognition). *To test whether a triangulation \mathcal{T} is a handlebody:*

- (1) *Check that \mathcal{T} is connected, orientable, and has exactly one boundary component. If \mathcal{T} fails to satisfy any of these conditions, then terminate and return **false**.*
- (2) *Compute the genus g of the boundary of \mathcal{T} .*
 - *In the case that $g = 0$, check whether \mathcal{T} is a 3-ball. If it is, terminate and return **true**; otherwise, terminate and return **false**.*

4.1 The degree defect and the multiplicity defect

► **Definitions 9.** Let e be an edge in a one-vertex triangulation. Let $d(e)$ denote the degree of e , and let $n(e)$ denote the number of *distinct* tetrahedra that meet e .

- The **degree defect** of e , denoted $\delta(e)$, is given by $|d(e) - 3|$.
- The **multiplicity defect** of e , denoted $\mu(e)$, is given by $d(e) - n(e)$.

Observe that we can perform a 3-2 move about an edge e , and hence remove this edge e , if and only if both the degree and multiplicity defects of e are zero. Thus, a natural way to remove all “unwanted” edges (such as core edges) from a triangulation is to minimise these defects across all “unwanted” edges. More precisely, our goal will be to minimise the following complexity with respect to the *lexicographical ordering*:

► **Definition 10.** Let \mathcal{T} be a one-vertex triangulation with some “unwanted” edges e_1, \dots, e_k , where $k \geq 1$. The **complexity** of \mathcal{T} (with respect to the “unwanted” edges) is given by

$$\left(k, \max_{1 \leq i \leq k} \mu(e_i), \max_{1 \leq i \leq k} \delta(e_i), |\mathcal{T}| \right).$$

Depending on context, the “unwanted” edges may be the core edges, tunnel edges, or edges isotopic to Seifert fibres. Our primary objective is to reduce the number of these edges, which is why this quantity appears as the first entry of our complexity. The last entry is the size of the triangulation because, all else being equal, it would be nice for our counterexamples to be as small as possible.

We have already discussed the rationale for the other two entries: we would like to reduce the degree and multiplicity defects across all “unwanted” edges. The reason for placing a higher priority on reducing multiplicity defect is best illustrated by recounting our initial approach to this problem. At first, our complexity did not involve the multiplicity defect at all. With this naïve approach, the search tended to get stuck enumerating lots of triangulations with small degree defect, but without ever actually reducing this defect to 0. We eventually realised that the search was getting trapped in a region of the Pachner graph where the “unwanted” edges had multiplicity defect equal to 2. Placing a high priority on reducing the multiplicity defect gives the search some impetus to avoid such regions.

4.2 The targeted search algorithm

► **Algorithm 11.** In a one-vertex triangulation, let P be a property of edges that is invariant under ambient isotopy. Call an edge **bad** if it satisfies property P , and **good** otherwise. This algorithm takes the following inputs:

- A one-vertex triangulation \mathcal{T} with n bad edges, where $n \geq 1$.
- A non-negative integer x (the number of extra bad edges that we are allowed to use).

To search for a new triangulation with $n - 1$ bad edges:

- (1) For each bad edge in \mathcal{T} , check whether it is possible to perform a 3-2 move about this edge. If such a move is possible, terminate and return the triangulation that results from performing this move.
- (2) Create a set \mathcal{S} . Also create a priority queue \mathcal{Q} that stores triangulations in order of increasing complexity (with respect to the bad edges). Add \mathcal{T} to both \mathcal{S} and \mathcal{Q} .
- (3) While \mathcal{Q} is non-empty:
 - (i) Remove the first triangulation \mathcal{F} (i.e., a triangulation with smallest complexity) from \mathcal{Q} , and let m denote the number of bad edges in \mathcal{F} .

- (ii) Call a 3-2 move about an edge e **eligible** if e is a good edge; for a 2-3 move, let e denote the new edge that is created by this move, and call this move **eligible** if either e is a good edge, or e is a bad edge but $m < n + x$. For each eligible move on \mathcal{F} , check whether (up to isomorphism) the set \mathcal{S} already contains the triangulation \mathcal{G} that we obtain after performing this move. If not:
- Add \mathcal{G} to both \mathcal{S} and \mathcal{Q} .
 - Perform all possible sequences of 3-2 moves about bad edges in \mathcal{G} . For each triangulation \mathcal{T}' that we obtain from such moves, if \mathcal{T}' does not already appear in the set \mathcal{S} (up to isomorphism), then add \mathcal{T}' to both \mathcal{S} and \mathcal{Q} .
 - If we can find such a sequence consisting of $m - n + 1$ 3-2 moves, then the final triangulation \mathcal{T}^* in this sequence has $n - 1$ bad edges. Terminate and return \mathcal{T}^* .

Visit <https://github.com/AlexHe98/triang-counterex> to see an implementation of this algorithm. We discuss the major details of this implementation in sections 4.3 and 4.4.

4.3 Troublesome regions, concurrent computation, and instability

In section 4.1, we mentioned regions of the Pachner graph where the “unwanted” edges have low degree defect, but multiplicity defect equal to 2 (such edges naturally occur, for instance, at the hearts of layered solid tori). Although placing a high priority on reducing multiplicity defect does help the search avoid such regions, if the search nevertheless gets trapped in such a region then it can be difficult to escape because the only way out is to perform moves that *increase* the degree defect of the “unwanted” edges.

This phenomenon occurs more generally: a locally optimal move can send the search into a “troublesome” region (which could, a priori, be *infinite*) of the Pachner graph where none of the subsequently available moves decrease the complexity. In the case above, we get stuck enumerating lots of triangulations with similar complexity; in section 5.1, we describe a case where we start enumerating lots of triangulations with *rapidly increasing* complexity.

Algorithm 11 is especially vulnerable to falling into such troublesome regions if we deal with triangulations one at a time in step 3. However, when we instead use multiple processes to deal with several triangulations concurrently, the search is sometimes able to either avoid or escape these troublesome regions. There are probably two drivers for this: (1) using multiple processes causes the search to explore with more “breadth” than a purely greedy approach, and (2) the search gains some randomness because the order in which triangulations are inserted into the priority queue \mathcal{Q} could vary each time we run the search. There may be more direct methods to achieve similar behaviour; our method was good enough, and had a low cost (both in human effort and in computational complexity).

One drawback is that it is impossible to know in advance how many concurrent processes we should use. We will see in section 5 that there is often a “sweet spot” where the search successfully terminates, but if we significantly increase or decrease the number of processes, then the search tends to get trapped in a troublesome region.

Our solution is to implement Algorithm 11 so that it periodically prints an update on the complexity of the triangulations that it is currently dealing with. If we manually observe that the search is not making satisfactory progress, then we simply restart the search, possibly with a change to the number of processes. Although it is not ideal that some human intervention is required, this is not unheard of for difficult problems in computational topology that only need to be solved once (such as finding a counterexample). Another example of this is the work done in [7] to extend the census of prime knots up to 19 crossings.

The last thing that we mention in this section is that our implementation is quite unstable. Indeed, we have already noted that the search can produce different results when we change the number of processes, or when we simply rerun the search. We have encountered one other manifestation of this instability: refactoring the source code can drastically change the behaviour of the search. We have not found a way to make the algorithm more stable.

4.4 Other implementation details

We now mention some other important details of our implementation of Algorithm 11.

First, since Algorithm 11 always considers triangulations up to isomorphism, we actually store isomorphism signatures in the set \mathcal{S} and the priority queue \mathcal{Q} . As described in [4], isomorphism signatures are an indispensable tool when exploring the Pachner graph.

Second, when the priority queue \mathcal{Q} contains multiple triangulations with the same complexity, we choose to prioritise such triangulations by insertion order.

Finally, we usually set the input variable x to be 0, in which case Algorithm 11 discards all triangulations with more bad edges than the input triangulation \mathcal{T} . However, there are occasions where taking $x > 0$ may be beneficial; we discuss one such case in section 5.3.

5 The counterexamples

We present our experimental results in sections 5.1, 5.2 and 5.3. We then show how to turn our counterexamples into infinite families in section 5.4.

Our computations were run on a laptop with an Intel Core i5-7200U processor, which has just two physical cores divided into four logical processors. It is therefore remarkable that we obtained all of our counterexamples in just a few *minutes* of wall time. This is mostly because our targeted search was able to home in on an extremely small portion of the search space: in total, we only needed to enumerate on the order of thousands of triangulations.

5.1 Triangulations with no core edges

The 20-tetrahedron triangulation \mathcal{T} with isomorphism signature

```
uLLvQQvLAPvPAQccdfeghhgmklnorsqssttthsaaggggaaaaaanaagb
```

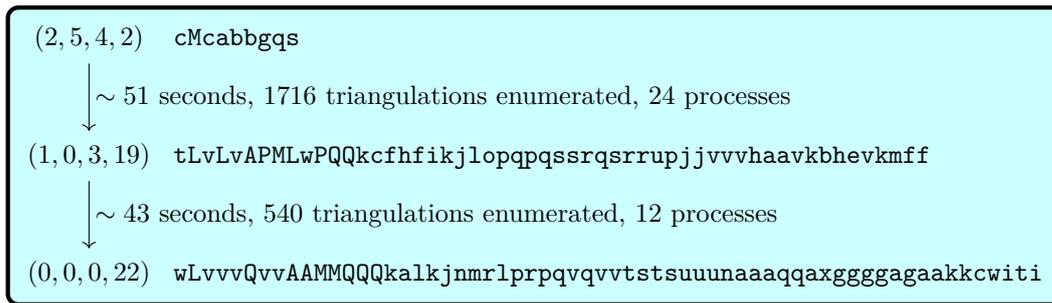
is a one-vertex 3-sphere with no core edges (for the 3-sphere, this means that every edge is non-trivially knotted). To find this, we arbitrarily selected a one-vertex 3-sphere \mathcal{T}_2 , which has isomorphism signature `cMcabbgqs` and complexity $(2, 5, 4, 2)$. After running Algorithm 11 twice, we obtained a 22-tetrahedron one-vertex triangulation \mathcal{T}_0 of the 3-sphere with no core edges; the results are summarised in Figure 4.

To turn \mathcal{T}_0 into the smaller triangulation \mathcal{T} , we ran a breadth-first search through the Pachner graph, but with the restriction that we ignored any 2-3 move that introduces a core edge. This took approximately 367 seconds of wall time.

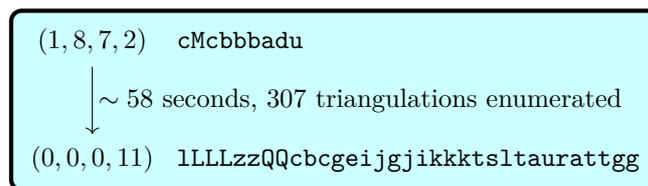
It is worth noting that to remove one core edge from \mathcal{T}_2 , we tried running Algorithm 11 with different numbers of processes, but using 24 processes produces the best results. In particular, we initially tried 12 processes, but this produced the isomorphism signature

```
sLvAAvLAzMMQQcdcefllkmjqonprqprhrvqnkkksqeeekocksf,
```

which has complexity $(1, 2, 1, 18)$. The core edge in this triangulation has multiplicity defect 2 but degree defect 1; this is exactly one of the troublesome cases that we mentioned in section 4.3. We needed to increase the number of processes to 24 to avoid this.



■ **Figure 4** Removing core edges from cMcabbgqs (3-sphere); $x = 0$.



■ **Figure 5** Removing tunnel edges from cPcbbbadu (trefoil knot); $x = 0$, 8 processes.

Surprisingly, increasing the number of processes significantly beyond 24 also has an adverse affect on the effectiveness of Algorithm 11. For example, with 36 processes, although the search is sometimes able to remove a core edge, it seems to do so less reliably. Instead, we find that the search has a tendency to get trapped in a different type of troublesome region:

- After about 10 seconds, the search reaches a triangulation with complexity (2, 4, 16, 15).
- After about 20 seconds, the search reaches a triangulation with complexity (2, 4, 23, 22).
- After about 30 seconds, the search reaches a triangulation with complexity (2, 4, 26, 25).

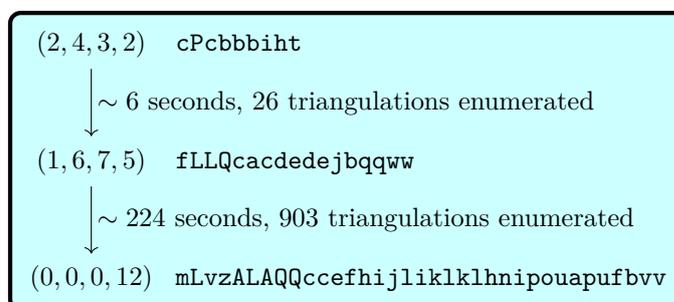
When this happens, the complexity only increases further if we allow the search to continue.

5.2 Triangulations with no tunnel edges

We found ideal triangulations with no tunnel edges for three knots with tunnel number one:

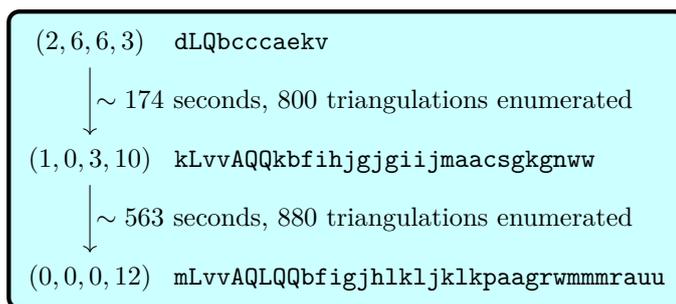
- the trefoil knot (results summarised in Figure 5);
- the figure-eight knot (results summarised in Figure 6); and
- the (5, 2) torus knot (results summarised in Figure 7).

Although the triangulations here are relatively small, and although we only enumerated relatively few triangulations, our running times are much longer than those in section 5.1.



■ **Figure 6** Removing tunnel edges from cPcbbbiht (figure-eight knot); $x = 0$, 24 processes.

21:12 Finding Large Counterexamples by Selectively Exploring the Pachner Graph



■ **Figure 7** Removing tunnel edges from `dLQbcccaekv` ((5, 2) torus knot); $x = 0$, 24 processes.

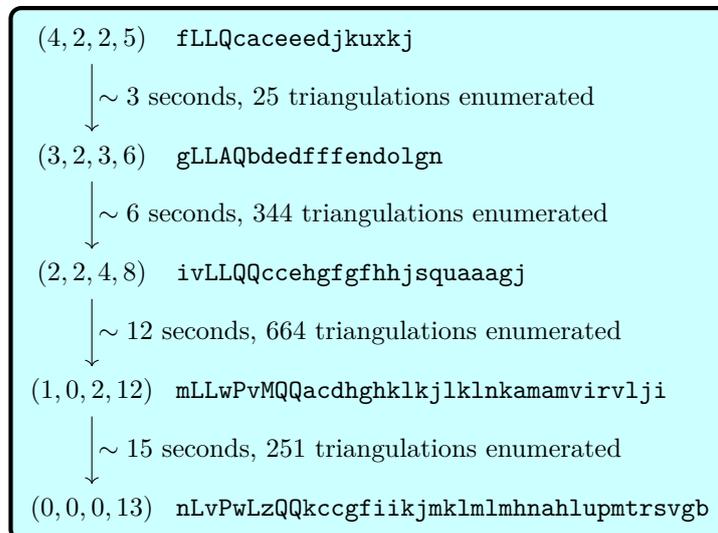
This is probably because genus-2 handlebodies are harder to recognise than solid tori.

5.3 Triangulations with no edges isotopic to Seifert fibres

We found the following 11-tetrahedron one-vertex triangulation \mathcal{T} of a small Seifert fibre space¹ with no edges isotopic to Seifert fibres:

`lLLLLPMQccddfjiihikkpkrwaaacttvc.`

To do this, we began with the isomorphism signature `fLLQcaceeedjkuxkj`. Running Algorithm 11 four times with $x = 1$ and 12 processes produced a 13-tetrahedron triangulation \mathcal{T}^* with no edges isotopic to Seifert fibres; the results are summarised in Figure 8. We then used a breadth-first search (similar to the one from section 5.1) to turn \mathcal{T}^* into \mathcal{T} .



■ **Figure 8** Removing edges isotopic to Seifert fibres from `fLLQcaceeedjkuxkj`; $x = 1$, 12 processes.

The reason for taking $x = 1$ instead of $x = 0$ comes from early versions of our algorithm for detecting good edges (in this context, edges not isotopic to Seifert fibres): we initially included fewer normal surfaces in the analysis, so the results were not particularly conclusive.

¹ Specifically, this is fibred over the 2-sphere with three exceptional fibres with slopes $\frac{1}{2}$, $\frac{2}{3}$ and $-\frac{1}{3}$.

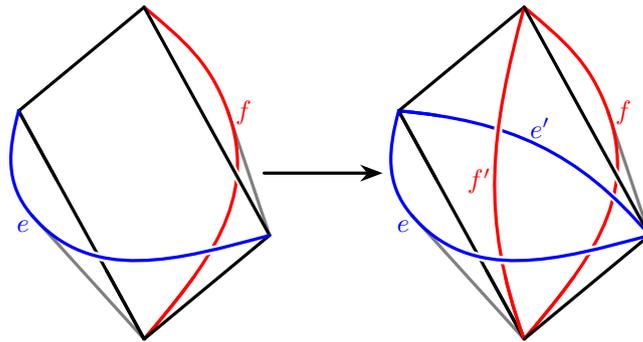
At the time, we encountered triangulations for which every 2-3 move would create an edge that we could not certify as good, so our short-term solution was to allow extra bad edges if required. In the current version of the code, taking $x = 1$ no longer seems strictly necessary.

5.4 From one counterexample to infinitely many

The following result allows us to turn our counterexamples into infinite families:

► **Proposition 12.** *Let \mathcal{M} be either a closed 3-manifold, or a 3-manifold with a single boundary component of positive genus (and no other boundary components). In the closed case, let \mathcal{T} be a one-vertex triangulation of \mathcal{M} ; in the bounded case, let \mathcal{T} be a one-vertex ideal triangulation of \mathcal{M} . Let P be a property of edges of \mathcal{T} that is invariant under ambient isotopy. Call \mathcal{T} **interesting** if every edge in \mathcal{T} satisfies P . If \mathcal{M} has an interesting triangulation, then \mathcal{M} has infinitely many interesting triangulations.*

Proof. Let \mathcal{T} be an interesting triangulation of \mathcal{M} . Fix a tetrahedron Δ of \mathcal{T} , and let e and f denote a pair of opposite edges of Δ . We obtain a new triangulation \mathcal{T}' of \mathcal{M} by replacing Δ with a three-tetrahedron gadget, as shown in Figure 9. In terms of 1-skeletons, observe that all we have done is introduce two new edges e' and f' such that e' is isotopic to e and f' is isotopic to f . Thus, \mathcal{T}' is an interesting triangulation of \mathcal{M} . Repeating this procedure indefinitely gives the desired infinite family of interesting triangulations of \mathcal{M} . ◀



■ **Figure 9** Building a new interesting triangulation.

► **Corollary 13.** *The 3-sphere has infinitely many triangulations with no core edges.*

► **Corollary 14.** *Let K be the trefoil knot, figure-eight knot, or $(5, 2)$ torus knot. Then K has infinitely many ideal triangulations with no tunnel edges.*

► **Corollary 15.** *There exists a small Seifert fibre space that has infinitely many triangulations with no edges isotopic to Seifert fibres.*

6 Discussion

6.1 Unanswered questions

For knots with tunnel number one, we found triangulations with no tunnel edges for all three knots that we considered. There are many other such knots in the census of knots [7] – in particular, every torus knot has tunnel number one – so it would be interesting to see whether such counterexamples exist for all of these knots. A key obstacle to a systematic investigation is the requirement for human supervision when running Algorithm 11, as discussed in section 4.3. In any case, we tentatively propose the following conjecture:

► **Conjecture 16.** *Every knot with tunnel number one admits infinitely many ideal triangulations with no tunnel edges.*

In contrast, we have only been able to find triangulations with no core edges for the 3-sphere; we have tried two other lens spaces – $L_{3,1}$ and $L_{5,1}$ – but without success. Similarly, although we have considered several small Seifert fibre spaces, we currently only have counterexamples for one case. We nevertheless suspect that the answer to the following two questions is “yes”, but new insights may be required to verify this:²

► **Question 17.** *Is there a lens space other than the 3-sphere that admits infinitely many triangulations with no core edges?*

► **Question 18.** *Is there another small Seifert fibre space that admits infinitely many triangulations with no edges isotopic to Seifert fibres?*

6.2 Future applications

Elementary moves such as 2-3 and 3-2 moves have many computational applications beyond how we used them in this paper. The two most common are:

- (1) using moves to improve a triangulation, which is often critical for making exponential-time computations feasible; and
- (2) finding a sequence of moves that transforms a triangulation \mathcal{T} into another triangulation \mathcal{T}' , which gives a computational proof that \mathcal{T} and \mathcal{T}' are homeomorphic.

Finding the right sequence of moves can be difficult, so a targeted search could be useful. In particular, increasing the number of tetrahedra is sometimes unavoidable [4]:

- (1) There are many triangulations \mathcal{T} of the 3-sphere such that to simplify \mathcal{T} to a smallest-possible triangulation via 2-3 and 3-2 moves, we must visit at least one intermediate triangulation with two more tetrahedra than \mathcal{T} . There is also a triangulation of a graph manifold for which simplification requires *three* additional tetrahedra.
- (2) In the census of minimal triangulations up to 9 tetrahedra, there are many 3-manifolds for which two additional tetrahedra are required to connect all the minimal triangulations using 2-3 and 3-2 moves. There is also one 3-manifold – the lens space $L_{3,1}$ – for which *three* additional tetrahedra are required to connect all the minimal triangulations.

There are also situations where we might actually *want* to increase the number of tetrahedra, provided this improves the triangulation with respect to another quantity. For example, some algorithms in 3-manifold topology are known to be fixed-parameter tractable in a quantity called the **treewidth** [10, 12]; for such algorithms, minimising the treewidth is more important than minimising the number of tetrahedra.

As mentioned in section 1, increasing the number of tetrahedra is accompanied by a super-exponential increase in the number of triangulations; even if we fix a 3-manifold, this is still at least exponential. This means that exhaustive searches quickly run into problems with not only running time, but also (often more importantly) memory management. However, our work suggests that, with the right heuristics, we may be able to circumvent these problems; the challenge is to actually devise such heuristics in the settings mentioned above.

² The full version [9] presents new results that answer these questions (but also raise new questions).

References

- 1 Gennaro Amendola. A calculus for ideal triangulations of three-manifolds with embedded arcs. *Math. Nachr.*, 278(9):975–994, 2005. doi:10.1002/mana.200310285.
- 2 David Bremner and Lars Schewe. Edge-graph diameter bounds for convex polytopes with few facets. *Experiment. Math.*, 20(3):229–237, 2011. doi:10.1080/10586458.2011.564965.
- 3 Benjamin A. Burton. Detecting genus in vertex links for the fast enumeration of 3-manifold triangulations. In *Proceedings of the 36th International Symposium on Symbolic and Algebraic Computation*, pages 59–66. Association for Computing Machinery, 2011. doi:10.1145/1993886.1993901.
- 4 Benjamin A. Burton. Simplification paths in the Pachner graphs of closed orientable 3-manifold triangulations. [arXiv:1110.6080](https://arxiv.org/abs/1110.6080), 2011.
- 5 Benjamin A. Burton. Computational topology with Regina: Algorithms, heuristics and implementations. In *Geometry and Topology Down Under*, volume 597 of *Contemporary Mathematics*, pages 195–224. American Mathematical Society, 2013. doi:10.1090/conm/597.
- 6 Benjamin A. Burton. A new approach to crushing 3-manifold triangulations. *Discrete Comput. Geom.*, 52:116–139, 2014. doi:10.1007/s00454-014-9572-y.
- 7 Benjamin A. Burton. The Next 350 Million Knots. In Sergio Cabello and Danny Z. Chen, editors, *36th International Symposium on Computational Geometry (SoCG 2020)*, volume 164 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 25:1–25:17. Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2020. doi:10.4230/LIPIcs.SoCG.2020.25.
- 8 Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. <https://regina-normal.github.io>, 1999–2022.
- 9 Benjamin A. Burton and Alexander He. Finding large counterexamples by selectively exploring the Pachner graph. [arXiv:2303.06321](https://arxiv.org/abs/2303.06321), 2023.
- 10 Benjamin A. Burton, Clément Maria, and Jonathan Spreer. Algorithms and complexity for Turaev-Viro invariants. *Journal of Applied and Computational Topology*, 2:33–53, 2018. doi:10.1007/s41468-018-0016-2.
- 11 Benjamin A. Burton and Melih Özlen. A fast branching algorithm for unknot recognition with experimental polynomial-time behaviour. [arXiv:1211.1079](https://arxiv.org/abs/1211.1079), 2012. To appear in *Math. Program.*
- 12 Benjamin A. Burton and Jonathan Spreer. The complexity of detecting taut angle structures on triangulations. In *Proceedings of the 2013 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 168–183. Society for Industrial and Applied Mathematics, 2013. doi:10.1137/1.9781611973105.13.
- 13 Wolfgang Haken. Theorie der Normalflächen. *Acta Math.*, 105:245–375, 1961. doi:10.1007/BF02559591.
- 14 Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger. The computational complexity of knot and link problems. *J. ACM*, 46(2):185–211, 1999. doi:10.1145/301970.301971.
- 15 Allen Hatcher. Notes on Basic 3-Manifold Topology. <https://pi.math.cornell.edu/~hatcher/3M/3Mfds.pdf>, 2007.
- 16 William Jaco and J. Hyam Rubinstein. 0-efficient triangulations of 3-manifolds. *J. Differential Geom.*, 65:61–168, 2003. doi:10.4310/jdg/1090503053.
- 17 William Jaco and Peter B. Shalen. A new decomposition theorem for irreducible sufficiently large 3-manifolds. In R. James Milgram, editor, *Algebraic and Geometric Topology, Part 2*, volume XXXII of *Proceedings of Symposia in Pure Mathematics*, pages 71–84. American Mathematical Society, 1978. doi:10.1090/pspum/032.2.
- 18 William Jaco and Jeffrey L. Tolleson. Algorithms for the complete decomposition of a closed 3-manifold. *Illinois J. Math.*, 39(3):358–406, 1995. doi:10.1215/ijm/1255986385.
- 19 Klaus Johannson. *Homotopy Equivalences of 3-Manifolds with Boundaries*, volume 761 of *Lecture Notes in Mathematics*. Springer-Verlag, 1979. doi:10.1007/BFb0085406.
- 20 Victor Klee. Paths on polyhedra. II. *Pacific J. Math.*, 17(2):249–262, 1966. doi:10.2140/pjm.1966.17.249.

- 21 Marc Lackenby and Saul Schleimer. Recognising elliptic manifolds. [arXiv:2205.08802](https://arxiv.org/abs/2205.08802), 2022.
- 22 Bruno Martelli and Carlo Petronio. Three-manifolds having complexity at most 9. *Experiment. Math.*, 10(2):207–236, 2001. doi:10.1080/10586458.2001.10504444.
- 23 Sergei V. Matveev. *Algorithmic Topology and Classification of 3-Manifolds*, volume 9 of *Algorithms and Computation in Mathematics*. Springer-Verlag, second edition, 2007. doi:10.1007/978-3-540-45899-9.
- 24 Sergei V. Matveev. Atlas of 3-Manifolds. <http://www.matlas.math.csu.ru/>, accessed November 2022.
- 25 Grisha Perelman. The entropy formula for the Ricci flow and its geometric applications. [arXiv:math/0211159](https://arxiv.org/abs/math/0211159), 2002.
- 26 Grisha Perelman. Finite extinction time for the solutions to the Ricci flow on certain three-manifolds. [arXiv:math/0307245](https://arxiv.org/abs/math/0307245), 2003.
- 27 Grisha Perelman. Ricci flow with surgery on three-manifolds. [arXiv:math/0303109](https://arxiv.org/abs/math/0303109), 2003.
- 28 Riccardo Piergallini. Standard moves for standard polyhedra and spines. In *III Convegno Nazionale Di Topologia: Trieste, 9-12 Giugno 1986 : Atti*, number 18 in Rend. Circ. Mat. Palermo (2) Suppl., pages 391–414, 1988. doi:11581/244540.
- 29 J. Hyam Rubinstein. An algorithm to recognize the 3-sphere. In *Proceedings of the International Congress of Mathematicians*, pages 601–611. Birkhäuser, 1995. doi:10.1007/978-3-0348-9078-6_54.
- 30 J. Hyam Rubinstein, Henry Segerman, and Stephan Tillmann. Traversing three-manifold triangulations and spines. *Enseign. Math.*, 65(1/2):155–206, 2019. doi:10.4171/LEM/65-1/2-5.
- 31 Francisco Santos. A counterexample to the Hirsch Conjecture. *Ann. of Math.*, 176(1):383–412, 2012. doi:10.4007/annals.2012.176.1.7.
- 32 Abigail Thompson. Thin position and the recognition problem for S^3 . *Math. Res. Lett.*, 1(5):613–630, 1994. doi:10.4310/MRL.1994.v1.n5.a9.
- 33 William P. Thurston. *Three-Dimensional Geometry and Topology*. Princeton University Press, 1997. doi:10.1515/9781400865321.