On minimal ideal triangulations of cusped hyperbolic 3–manifolds

William Jaco, Hyam Rubinstein, Jonathan Spreer and Stephan Tillmann

Abstract   Previous work of the authors studies minimal triangulations of closed 3–manifolds using a characterisation of low degree edges, embedded layered solid torus subcomplexes and 1–dimensional $\mathbb{Z}_2$–cohomology. The underlying blueprint is now used in the study of minimal ideal triangulations. As an application, it is shown that the monodromy ideal triangulations of the hyperbolic once-punctured torus bundles are minimal.

1   Introduction

Thurston observed that ideal triangulations are a useful tool to study the geometry and topology of cusped hyperbolic 3–manifolds. Geometric triangulations are useful to study geometric properties of a manifold. Minimal triangulations, i.e. topological ideal triangulations using the least number of ideal 3–simplices, are used in census enumeration, and as a platform to study the topology of the manifold using normal surface theory. In this paper we establish basic facts about minimal ideal triangulations of cusped hyperbolic 3–manifolds in analogy with our previous work for closed 3–manifolds [20, 21, 22].

Throughout this paper, by a *cusped hyperbolic 3–manifold* we mean an orientable non-compact 3–manifold M that admits a complete hyperbolic structure of finite volume. Such a manifold is known to be the interior of a compact, irreducible, ∂–irreducible, atoroidal, anannular 3–manifold $\overline{M}$ with non-empty boundary a finite union of tori. If $\overline{M}$ has exactly n boundary components, we say that M is n–cusped.

Let $c(M)$ be the minimal number of ideal tetrahedra in a (topological) ideal triangulation of the cusped hyperbolic 3–manifold M. In this paper, we use the *topology* of M to study the anatomy of a minimal ideal triangulation, and with this derive lower bounds on the complexity, and apply it to describe some infinite families of minimal ideal triangulations. Before stating our results, we describe previous work on minimal ideal triangulations using the *geometry* of the manifold.

1.1   Minimal ideal triangulations via geometry

The canonical cell decompositions of cusped hyperbolic 3–manifolds due to Epstein and Penner [7] are generically ideal triangulations. It follows from work of Guéritaud [11, Theorem 2.1.7 in Section 2.1.4] that the canonical cell decompositions of hyperbolic 2–bridge link complements are ideal triangulations. However, these are not minimal in general. We thank Jessica Purcell and the anonymous referee for pointing us to these examples.

If four ideal tetrahedra in a geometric ideal triangulation form a non-convex ideal octahedron, then there are 4–4 moves on this triangulation that result in non-geometric ideal triangulations. Neil Hoffman verified that in the current censuses of cusped hyperbolic 3–manifolds up to 6 tetrahedra, there is always at least one minimal

triangulation that is geometric, but that non-geometric minimal ideal triangulations appear through the presence of 4–4 moves as described above. In particular, it is currently not known whether there is always at least one minimal ideal triangulation that is geometric.

Let v_3 denote the volume of a regular ideal hyperbolic tetrahedron; this is approximately 1.0149. An argument due to Thurston [34] shows that

$$c(M) \geq \frac{\text{Vol}(M)}{v_3}. \tag{1.1}$$

In particular, any geometric triangulation only involving regular ideal hyperbolic tetrahedra is minimal. The figure eight knot complement has such a triangulation. By taking finite sheeted coverings this leads to infinitely many minimal triangulations that are geometric and for which the above inequality is an equality. A census of such manifolds decomposed into at most 25 ideal regular ideal hyperbolic tetrahedra was given by Fominykh, Garoufalidis, Goerner, Tarkaev and Vesnin [9].

In general the gap in inequality (1.1) can be arbitrarily large. To see this neither requires explicit minimal triangulations nor computation of volumes: Thurston's theory of hyperbolic Dehn surgery implies that for any $n \geq 1$ there are sequences of n–cusped hyperbolic 3–manifolds with bounded volume and whose complexities are an unbounded sequence.

An equivalent approach to complexity is via Matveev's theory of special spines. From this point of view, Petronio and Vesnin [29] give a lower bound on complexity using a volume estimate. Ishikawa and Nemoto [15] combined this with an upper bound on the complexity of 2–bridge links and determined the complexity of an infinite family of 2–bridge links. These ideal triangulations were described by Sakuma and Weeks [32] and proved to be canonical by Guéritaud [11, Theorem 2.1.7 in Section 2.1.4]. Independently, Akiyoshi, Sakuma, Wada and Yamashita have announced a proof of this statement [3]. A proof of this fact for the fibered 2–bridge links was given by Sakata [31].

1.2   Minimal ideal triangulations via topology

In this paper, we are first concerned with the anatomy of a minimal ideal triangulation. In particular, in §2 we characterise edges of low degree, and analyse subcomplexes that we call *maximal layered ∂–punctured solid tori*. These are solid tori with a point missing on the boundary, and imbued with a special ideal triangulation. These subcomplexes were exhibited by Guéritaud and Schleimer [13] in canonical decompositions of cusped hyperbolic 3–manifolds satisfying certain genericity hypotheses, and are thus a natural part of the structure of an ideal triangulation to analyse from a geometrical viewpoint. Our methods show that they also have an important topological role to play. In §3 we describe normal surfaces representing $\mathbb{Z}_2$–homology classes in ideal triangulations. We use this in §4 to prove our main result Theorem 1, which is analogous to the main result of [22]. Before we can state it, we need to introduce some extra definitions.

Let M be a cusped hyperbolic 3–manifold, and let S be an embedded closed surface representing a given $c \in H_2(M; \mathbb{Z}_2)$. We wish to emphasise that we do not work with $H_2(\overline{M}, \partial \overline{M}; \mathbb{Z}_2)$. An analogue of Thurston's norm [35] is defined in [22] as follows. If S is connected, let $\chi_-(S) = \max\{0, -\chi(S)\}$, and otherwise let

$$\chi_-(S) = \sum_{S_i \subset S} \max\{0, -\chi(S_i)\},$$

where the sum is taken over all connected components of S. Note that S_i is not necessarily orientable. Define:

$$\| c \| = \min\{\chi_-(S) \mid [S] = c\}.$$

The surface S representing a class $c \in H_2(M; \mathbb{Z}_2)$ is said to be $\mathbb{Z}_2$–*taut* if no component of S is a sphere or torus and $\chi(S) = -\| c \|$. Note that M contains no projective plane or Klein bottle. As in [35], one observes that every component of a $\mathbb{Z}_2$–taut surface is non-separating and geometrically incompressible.

Theorem 1 *Let $\mathscr{T}$ be an ideal triangulation of the cusped hyperbolic 3–manifold M. If $H \le H_2(M, \mathbb{Z}_2)$ is a rank 2 subgroup, then*

$$|\mathscr{T}| \ge \sum_{0 \neq c \in H} \| c \|.$$

Moreover, in the case of equality the triangulation is minimal, each canonical normal representative of a non-zero element in $H \le H_2(M, \mathbb{Z}_2)$ is taut and meets each tetrahedron in a quadrilateral disc, and the number of tetrahedra in the triangulation is even.

As an application, we obtain the following statement.

Theorem 2 *Monodromy ideal triangulations of once-punctured torus bundles are minimal.*

Monodromy ideal triangulations of hyperbolic once-punctured torus bundles are described by Floyd and Hatcher [8] and attributed to Thurston. These triangulations are canonically determined by the monodromy of the bundle acting on the Farey graph; they are an early special case of Agol's veering triangulations. The fact that the monodromy ideal triangulation of each hyperbolic once-punctured torus bundle is canonical, and hence also geometric, is due to Akiyoshi [2], Lackenby [25] and Guéritaud [11].

Such a triangulation may contain edges of arbitrarily high degree, hence arbitrarily many ideal hyperbolic tetrahedra of arbitrarily small volume, a fact which is also highlighted by explicit upper bounds on the volume of torus bundles in [12, Theorem B.1] and in [1, Corollary 2.4]. Thus, our family contains explicit minimal triangulations where the gap between complexity and volume can be arbitrarily large.

The monodromy ideal triangulation $\mathscr{T}$ of a hyperbolic once-punctured torus bundle M with the property that $H_2(M, \mathbb{Z})$ contains a Klein 4–group H satisfies the equality

$$|\mathscr{T}| = \sum_{0 \neq c \in H} \| c \|.$$

There are other once-cusped hyperbolic 3–manifolds, for which this equality is achieved. These are not fibred, and are described in §5.4. It remains an open problem to *classify all triangulations* for which the lower bound in Theorem 1 is achieved. Such a classification is given in [20] for our lower bound on the complexity for closed atoroidal 3–manifolds.

1.3   Other classes of hyperbolic 3–manifolds

We remark that for closed hyperbolic 3–manifolds, currently no infinite families of minimal triangulations are known. In the case of compact manifolds with non-empty boundary, lower bounds on complexity are known through work of [4] and [16], where it is independently shown that they are attained by $S_g \times [0, 1]$, where S_g is a closed orientable surface of genus $g \ge 1$.

There are infinite families of minimal triangulations of manifolds with a totally geodesic boundary component of genus $g \ge 2$ and some cusps due to Frigerio, Martelli and Petronio [10]. These triangulations can be thought of as subdivisions into partially truncated tetrahedra; or alternatively as ideal triangulations with one ideal vertex having link a surface of genus g and all others having link a torus. We remark that our methods and results do not generalise to this setting of higher genus vertex links. The main reason is that there is less control over the anatomy of these triangulations (cf. Remark 7).

2 The anatomy of a minimal ideal triangulation

2.1 Triangulations and pseudo-manifolds

Let $\widetilde{\Delta}$ be a finite union of pairwise disjoint Euclidean 3–simplices with the standard simplicial structure. Every k–simplex τ in $\widetilde{\Delta}$ is contained in a unique 3–simplex σ_τ. A 2–simplex in $\widetilde{\Delta}$ is termed a *face*.

Let Φ be a family of affine isomorphisms pairing the faces in $\widetilde{\Delta}$, with the properties that $\varphi \in \Phi$ if and only if $\varphi^{-1} \in \Phi$, and every face is the domain of a unique element of Φ. The elements of Φ are termed *face pairings*.

Consider the quotient space $\widehat{M} = \widetilde{\Delta}/\Phi$ with the quotient topology, and denote the quotient map $p\colon \widetilde{\Delta} \to \widehat{M}$. We make the additional assumption on Φ that for every k–simplex τ in $\widetilde{\Delta}$ the restriction of p to the interior of τ is injective. Then the set of non-manifold points of $\widehat{M}$ is contained in the 0–skeleton, and in this case $\widehat{M}$ is called a *closed 3–dimensional pseudo-manifold.* (See Seifert-Threfall [33].) We also always assume that $\widehat{M}$ is connected. The triple $\mathcal{T} = (\widetilde{\Delta}, \Phi, p)$ is called a *(singular) triangulation* of $\widehat{M}$.

Let $\widehat{M}^{(k)} = \{ p(\tau^k) \mid \tau^k \subseteq \widetilde{\Delta} \}$ denote the set of images of the k–simplices of $\widetilde{\Delta}$ under the projection map. Then the elements of $\widehat{M}^{(k)}$ are precisely the equivalence classes of k–simplices in $\widetilde{\Delta}$. The elements of $\widehat{M}^{(0)}$ are termed the *vertices* of $\widehat{M}$ and the elements of $\widehat{M}^{(1)}$ are the *edges*. The triangulation is a *k–vertex triangulation* if $\widehat{M}^{(0)}$ has size k.

2.2 Edge paths and edge loops

To each edge e in $\widehat{M}^{(1)}$ we associate a continuous path $\gamma_e\colon [0,1] \to \widehat{M}$ with endpoints in $\widehat{M}^{(0)}$ via the composition of maps $[0,1] \to \widetilde{\Delta} \to \widehat{M}$, where the first map is an affine map onto a 1–simplex in the equivalence class and the second is the quotient map. We call γ_e an *edge path*. An edge path is termed an *edge loop* if its ends coincide. For instance, in a 1–vertex triangulation every edge path is an edge loop.

We often abbreviate "edge path" or "edge loop" to "edge", and we denote the edge path γ_e with reversed orientation by $-\gamma_e$. The notions of interest in this paper are independent of the parametrisation.

2.3 Ideal triangulations

Let $M = \widehat{M} \setminus \widehat{M}^{(0)}$. Then M is a topologically finite, non-compact 3–manifold, and the pseudo-manifold $\widehat{M}$ is referred to as the *end-compactification* of M, the elements of $\widehat{M}^{(0)}$ as the *ideal vertices* of M and the elements of $\widehat{M}^{(1)}$ as the *ideal edges* of M. We also refer to the triangulation $\mathcal{T} = (\widetilde{\Delta}, \Phi, p)$ of $\widehat{M}$ as an *ideal (singular) triangulation* of M.

The adjective *singular* is usually omitted: we do not need to distinguish between simplicial and singular triangulations. If $\widehat{M}$ is a closed 3–manifold, we may write $M = \widehat{M}$, and hope this does not cause any confusion.

We suppress the notation $\mathcal{T} = (\widetilde{\Delta}, \Phi, p)$ and simply say that $\widehat{M}$ is triangulated or $M = \widehat{M} \setminus \widehat{M}^{(0)}$ is ideally triangulated.

The following result is implicit in Matveev [27] and shows that every topologically finite 3–manifold arises in this way from a closed pseudo-manifold. A stronger version for compact orientable irreducible an-annular 3–manifolds with ∂M a non-empty collection of incompressible tori, avoiding the use of spines can be found in [24, Theorem 1].

Proposition 3 (Topologically finite manifolds have ideal triangulations) *If M is the interior of a compact 3–manifold $\overline{M}$ with non-empty boundary, then M admits an ideal triangulation. The ideal vertices of the ideal triangulation are in one-to-one correspondence with the boundary components of $\overline{M}$.*

Proof The statement is implied by the following results in [27]. Theorem 1.1.13 due to Casler asserts that $\overline{M}$ possesses a special spine Σ; it has the property that $\overline{M}$ is homeomorphic to a regular neighbourhood of Σ in M. Theorem 1.1.26 implies that Σ is dual to an ideal triangulation of M with the property that the ideal vertices are in one-to-one correspondence with the boundary components of $\overline{M}$. $\qquad\square$

2.4   Layered solid tori

A *layered solid torus* LST is a triangulation of a solid torus iteratively constructed in the following way: Start with a one-triangle Möbius strip as shown in Figure 2(c). Then attach additional tetrahedra without a twist along pairs of (unpaired) faces of the existing complex sharing an edge e. This operation is called a *layering* on e, where the first layering onto the Möbius strip must always go on the interior edge of the Möbius strip, see Figure 1. Layered solid tori are close cousins of the monodromy ideal triangulations of once-punctured torus bundles defined in Section 5.1, since both types of triangulations are obtained from layering tetrahedra onto an existing complex. See [18] for a comprehensive introduction into layered solid tori and, more generally, layered triangulations. Here, we summarise material of [21, Section 2.3].
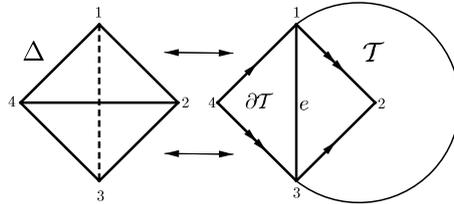
Figure 1: Layering tetrahedron Δ to triangulation $\mathcal{T}$ on edge e. Identifications of the triangles are indicated by the vertex labels.

By construction, the boundary of a layered solid torus containing at least one tetrahedron is always the standard triangulation of a torus with one vertex, two triangles, and three edges. Layered solid tori are distinguished by how often their meridional disc intersects their three boundary edges geometrically and we write $\mathrm{LST}(a,b,c)$ for a layered solid torus with meridional disc intersecting the boundary edges a, b and c times respectively. Assuming $a \le b \le c$ it is straightforward to see that $a + b = c$ and that a layered solid torus $\mathrm{LST}(a,b,c)$ can be turned into a layered solid torus of type $\mathrm{LST}(a,b,b-a)$, $\mathrm{LST}(a,c,a+c)$ or $\mathrm{LST}(b,c,b+c)$ by layering an additional tetrahedron on the edge with label c, b, or a respectively.

There exists only one layered solid torus with one tetrahedron, $\mathrm{LST}(1,2,3)$. Hence, by construction, every layered solid torus with at least one tetrahedron contains this layered solid torus, termed its *core tetrahedron*.

Let LST be a layered solid torus occurring as a subcomplex of some triangulation $\mathcal{T}$ of a 3-manifold. We say that LST is *maximal with respect to $\mathcal{T}$* if it is not strictly contained in any other layered solid torus subcomplex

in $\mathcal{T}$. By a *layered ∂–punctured solid torus*, we mean a *layered solid torus* with its vertex removed. Whenever we talk about the punctured version of a layered solid torus LST(a,b,c) we denote it by LST$^\star(a,b,c)$ to emphasise the difference. The notion of a layered ∂–punctured solid torus is very convenient when discussing the combinatorial structure of an ideal triangulation. For instance, they play a role in [13].

2.5 Properties of minimal triangulations

An ideal triangulation of a topologically finite 3–manifold M is *minimal* if it uses the smallest number of simplices in an ideal triangulation of M. The number of ideal simplices in a minimal ideal triangulation of M is denoted $c(M)$ and termed the *complexity* of M.

We assume that the reader is familiar with normal surface theory. We refer to [36] for a thorough introduction, but remark that we only need to study *closed* normal surfaces in ideal triangulations.

An ideal triangulation of M is *0–efficient* if no normal surface is a 2–sphere. An ideal triangulation of M is *∂–efficient* if it has the property that any closed normal surface that is isotopic to a vertex linking surface is *normally isotopic* to that vertex-linking surface. A normal isotopy is an isotopy of M that preserves each simplex of each dimension.

Theorem 4 *Suppose M is the interior of a compact, irreducible, ∂–irreducible anannular 3–manifold. Then each minimal ideal triangulation of M is 0–efficient and ∂–efficient.*

Proof It is shown in [17, Corollary 7.3] that each minimal triangulation of M is 0–efficient. In [19, Theorem 4.7], the existence of a ∂–efficient triangulation is proven by a crushing method. Since crushing reduces the number of tetrahedra, this means that a triangulation that fails to be ∂–efficient is not minimal. $\square$

We now give a definition of what we mean by an ideal edge to be homotopic or isotopic *into the boundary*. The key is that intermediate paths in the homotopy are not allowed to pass through the ideal vertices, except for their endpoints. The reader is reminded than an ideal edge of M is an edge of $\widehat{M}$. An expanded discussion of the following material can be found in [14].

A path homotopy (resp. isotopy)[1] $H\colon [0,1]\times[0,1]\to\widehat{M}$ between two edge paths is *admissible* if $H(\,(0,1)\times[0,1]\,)\subset M$, where the first factor parametrises the paths. The edge path γ is *admissibly null-homotopic (resp. null-isotopic)* if there is a path homotopy (resp. isotopy) $H\colon [0,1]\times[0,1]\to\widehat{M}$ with $H(x,0)=\gamma(x)$, $H(x,1)=\gamma(0)$ for all $x\in[0,1]$ and $H(\,(0,1)\times[0,1)\,)\subset M$.

We say that an ideal edge e of M is *homotopic (resp. isotopic) into the boundary* if it is admissibly null-homotopic (resp. null-isotopic) in $\widehat{M}$.

Corollary 5 *Suppose M is the interior of a compact, irreducible, ∂–irreducible anannular 3–manifold other than a 3–ball, and $\mathcal{T}$ is a minimal ideal triangulation of M. Then no ideal edge is isotopic into the boundary. In particular, no edge in $\widehat{M}$ bounds an embedded disc in $\widehat{M}$.*

Proof Suppose the edge e is isotopic into the boundary. Choose a compact core M^c of M with the property that each boundary component of M^c is a vertex linking surface. There exists an embedded disc D in M^c whose boundary consists of $e\cap M^c$ and a subarc on a boundary component B of M^c. The boundary of a regular neighbourhood of $B\cup D$ has two components; one is a vertex linking surface normally isotopic to B and the other is a surface S isotopic to B and disjoint from e. Since S is incompressible and not a 2–sphere and M is irreducible, it follows that there is a normal surface isotopic but not normally isotopic to B. Since the triangulation is ∂–efficient according to Theorem 4, it follows that M is homeomorphic with $S\times(0,1)$. However, this contradicts our topological hypotheses on M. $\square$

---

[1]i.e. a homotopy (resp. isotopy) keeping endpoints fixed

Note that the above result does not apply to the unique minimal ideal triangulation $\mathscr{T}$ of the 3–ball. This is the 1–tetrahedron, 1–vertex triangulation of S^3 with the vertex removed. The triangulation of S^3 has precisely two edges; one is a trefoil knot in S^3 and the other the unknot in S^3. The former gives an ideal edge in $\mathscr{T}$ that is only homotopic into the boundary, whilst the latter is isotopic into the boundary. The above proofs breaks down for the 3–ball because the surface S in the proof shrinks to a sphere contained in a tetrahedron.

We now specialise to the class of manifolds of interest, namely cusped hyperbolic 3–manifolds, where more specific results can be shown about their minimal triangulations. In order to limit the cases to consider, we not only assume that all *vertex links are of Euler characteristic zero*, but also that the manifolds are *orientable*.

2.6 Faces in minimal ideal triangulations of cusped hyperbolic 3–manifolds

(a) triangle   (b) cone   (c) Möbius   (d) 3–fold   (e) dunce
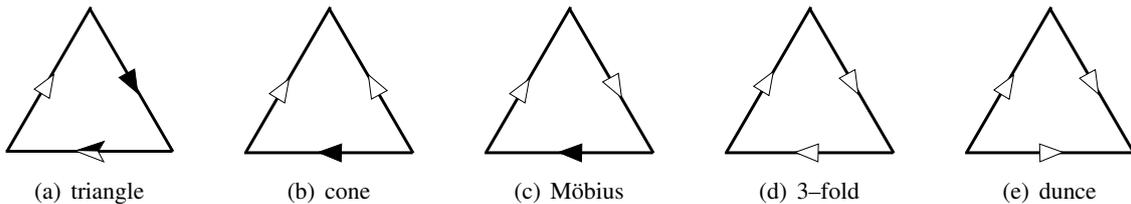
Figure 2: Types of faces. The Möbius, 3–fold and dunce faces only have one vertex due to the edge identifications. A cone face may have one or two vertices and a triangle face may have up to 3 vertices.

An ideal triangle in an ideal triangulation of a cusped hyperbolic 3–manifold M may have some of its edges identified. There are eight possible types of triangular faces in the triangulation of the end-compactification $\widehat{M}$, depending on how vertices or edges are identified (see Figure 1 in [17]). Ignoring possible identifications of vertices as in [26], this reduces to five types in M based on edge identifications only (see Figure 2). We name the faces in M according to their topology in $\widehat{M}$. A face with no edge-identifications is termed a *triangle face*; it is a 2–simplex, possibly with some or all of its vertices identified. If a pair of edges is identified, the face is either a cone (possibly with the tip identified with the vertex on the boundary) and called a *cone face* or it is a Möbius band and called a *Möbius face*. In a Möbius face, we distinguish the *boundary edge* and the *core edge*. If all three edges are identified, the face is either a 3–fold or a dunce hat, and called a *3–fold face* or *dunce face* respectively.

For minimal ideal triangulations, we can prove the following statement.

Theorem 6 *Let $\mathscr{T}$ be a minimal ideal triangulation of a cusped hyperbolic 3–manifold M. Then there is no 3–fold face and there is no dunce face.*

Proof Choose a compact core M^c of M with the property that each boundary component of M^c is a vertex linking surface. Let t be a triangle of $\mathscr{T}$.

Case 1: t is a 3–fold face. Refer to Figure 3. Let e be the edge of t, $N(e)$ a small regular neighbourhood of e in M and $N^c(e) = N(e) \cap M^c$. The frontier of $N^c(e)$ in M^c is an annulus; we call this the *exchange annulus around e*. Let $N(t)$ be a regular neighbourhood of t in M. Then the boundary of $N^c(t) = N(t) \cap M^c$ has a natural cell decomposition into two hexagons corresponding to the two sides of t and three quadrilaterals that are subsurfaces of the exchange annulus around e. The edges of the hexagons through the interior of M^c are termed the *long edges*. Seen with respect to $(t \cap M^c) \setminus N(e)$ the long edges of the hexagons have a natural labelling $1, 2, 3$ in cyclic order. See Figure 3 on the right for a drawing of $N^c(t)$, in particular the two hexagons and their long edges.

Give t a transverse orientation. This allows us to refer to one of the hexagons as the *top* hexagon and to the other as the *bottom* hexagon. The intersection $t \cap N^c(e)$ is a tripod times $[0,1]$. Figure 3 on the left depicts a cross-section of the tripod times $[0,1]$ at some $x \in (0,1)$. The transverse orientation of t is in the same direction at the legs of the tripod (either all clockwise or all anti-clockwise around the vertex of the tripod). It follows that the quadrilaterals join each long edge of the top hexagon to a long edge of the bottom hexagon. Moreover, on the labels this induces a 3–cycle. Thus, modulo orientation, the identification must be as shown in Figure 3 (both on the left and on the right).

After collapsing the sections of the exchange annulus to the long edges of the hexagons, we obtain a connected, orientable, bounded surface $\mathscr{S}$ decomposed into two hexagons, nine edges and six vertices – and thus of Euler characteristic -1. As can be deduced from the vertex labels, the surface has three boundary components and thus must be a 3–punctured sphere.
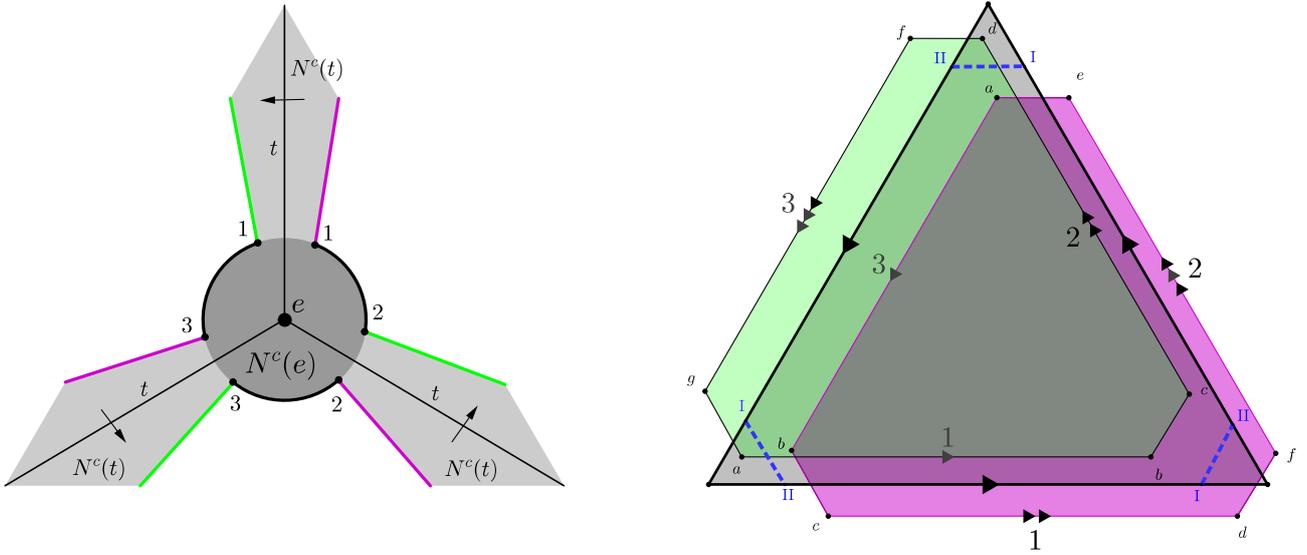
Figure 3: Left: Cross-section through the neighbourhood $N^c(t)$ of an ideal 3–fold t near its boundary edge e. The cross-section through the top and bottom hexagon of the right picture appear as green and purple lines respectively. Their endpoints mark the cross-section through the long edges of the respective hexagons. Labels coincide with the ones on the right. Right: Truncated boundary $N^c(t)$ of regular neighbourhood of the 3–fold face t (drawn in the centre in black). Its boundary surface is connected, of Euler characteristic -1 with three boundary components (each containing two edges, as indicated by the edge labels) and thus a 3–punctured sphere. The dotted blue lines with labels represent the intersection of t with the boundary torus T_0 – a triple arc between nodes 1 and 2, or the Θ–graph.

Because t is a 3–fold, the intersection of t with a boundary torus $T_0 \subset \partial M^c$ consists of two nodes and three (normal) arcs, all running between the two nodes, also known as the Θ–*graph*, see the dotted arcs in Figure 3. Since $\mathscr{S}$ has three boundary components meeting T_0 along the boundary of a neighbourhood of $t \cap T_0$, the complement of the embedding of the Θ–graph into T_0 must have three boundary components.

If no two arcs form a separating (i.e., inessential) curve on T_0, then the Θ–graph forms a spine of T_0. In particular, its complement is a disc with a single boundary component, a contradiction. Hence, two arcs must bound a disc in T_0. The remaining arc now either runs parallel to the first two, or completes each of the other two to an essential curve on T_0. In both cases the complement has three boundary components. Both embeddings are shown in Figure 4.

Assume that we are in the latter situation and the Θ–graph contains an essential curve on T_0. Since the embedding features two parallel arcs, one of the boundary curves of $\mathscr{S}$ bounds a disc in T_0. Pasting in this disc and

pushing it into the interior of the manifold yields a properly embedded annulus which, because M is hyperbolic, must be compressible or boundary parallel. If it is compressible, then T_0 is compressible. We conclude that $\mathscr{S}$ must be boundary parallel.

Hence assume that the Θ–graph consists of three parallel arcs and is thus contained in a disc of T_0. This time two of the boundary curves of $\mathscr{S}$ bound discs in T_0 and pasting them in and pushing them into the manifold yields a properly embedded disc, and we conclude again that $\mathscr{S}$ must be boundary parallel.

In both cases, uniting $\mathscr{S}$ with the complement of $N(t \cap T_0)$ in T_0 thus produces a torus

$$T_0 \setminus N(t \cap T_0) \cup \mathscr{S}$$

boundary parallel to T_0 and disjoint from e (which is itself incident to T_0). It follows that normalising this boundary parallel torus cannot result in a surface normally isotopic to T_0 and thus $\mathscr{T}$ cannot be boundary efficient. This is impossible since we assume that $\mathscr{T}$ is minimal and thus boundary efficient by Theorem 4.
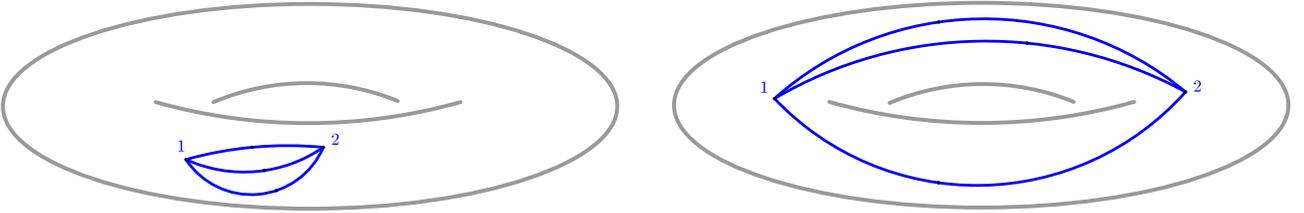

Figure 4: Two embeddings of the intersection of the 3–fold face t with torus boundary component T_0. Left: trivial embedding contained in a disc of T_0; in this case the boundary parallel torus is made up of a once-punctured torus and two discs that are subsurfaces of T_0 and the 3–punctured sphere $\mathscr{S}$. Right: two parallel essential loops; in this case the boundary parallel torus is made of a disc and an annulus that are subsurfaces of T_0 and the 3–punctured sphere $\mathscr{S}$.

Case 2: t **is a dunce face.** Refer to Figure 5. Let e be the edge of t. We use the notation and set-up as in the previous case. The intersection of t with $N^c(e)$ is again a tripod times an interval (see Figure 5 on the left for a cross-section). However, now the transverse orientations do not all agree: two are in one direction and the last is in the opposite direction. In particular, this implies that each of the hexagons has two of its long edges connected by a quadrilateral on the exchange annulus, and exactly one long edge of the top hexagon is joined to a long edge of the bottom hexagon (and these edges do not share the same label). It follows that, up to symmetry, there is only a single way to identify the long edges of the hexagons, shown in Figure 5 (both on the left and on the right).

The two hexagons again form a properly embedded 3–punctured sphere $\mathscr{S} \subset M^c$. The dunce hat t itself meets ∂M^c in two nodes and three normal arcs: two loops, one at each node, and one arc connecting the two nodes – also known as the *barbell graph*. Up to the action of the mapping class group, there are four possibilities of embedding the barbell graph into a torus boundary component $T_0 \subset \partial M^c$, producing three boundary components in a regular neighbourhood of the graph: One where both loops are essential in T_0 and thus necessarily parallel; one where one loop is inessential and the other one is an arbitrary essential curve in T_0; and two where both loops are inessential. See Figure 6 for a picture of all four cases. Using the same pasting argument as in Case 1 we can conclude that $\mathscr{S}$ is boundary parallel. Again, uniting $\mathscr{S}$ with the complement of $N(t \cap T_0)$ produces a torus boundary parallel to T_0 and disjoint from e. Thus, Theorem 4 produces a contradiction. □

Remark 7 The preconditions of Theorem 6 are necessary: the minimal triangulation of the Gieseking manifold has dunce faces, and there exists a minimal 2-tetrahedron triangulation of a 3–manifold with two Klein-bottle cusps, in which one of the four faces is a 3–fold face.
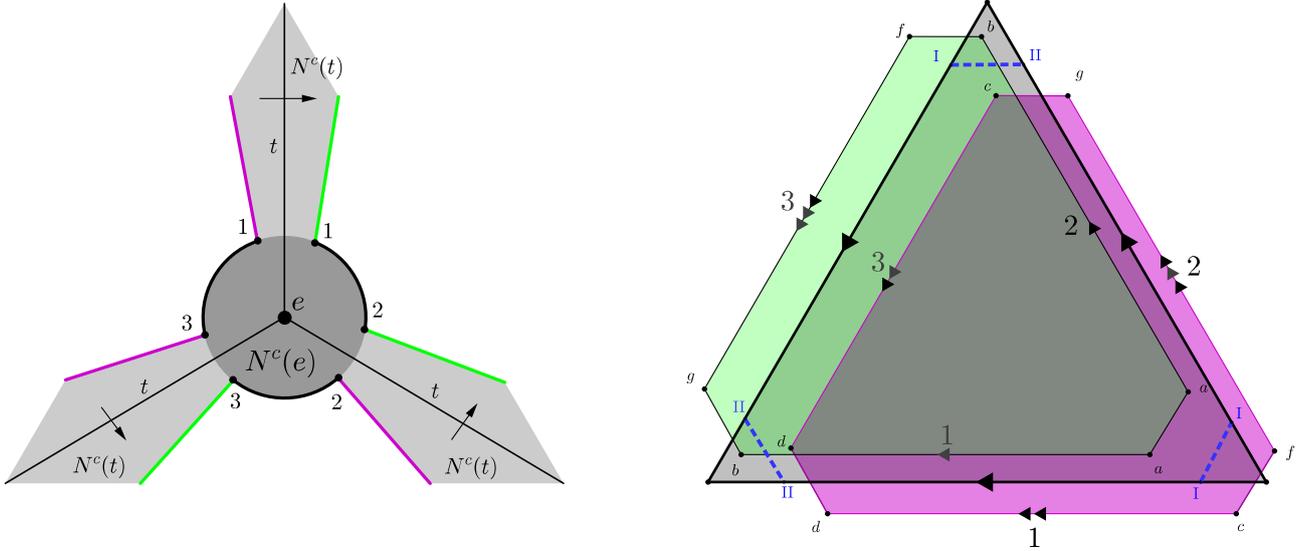
Figure 5: Left: Cross-section through the neighbourhood $N^c(t)$ of an ideal dunce hat t near its boundary edge e. Labels coincide with the ones on the right. Right: Truncated boundary $N^c(t)$ of regular neighbourhood of the dunce hat t (drawn in the centre in black). The boundary surface is connected, of Euler characteristic -1 with three boundary components (one containing four edges and two containing a single edge, as indicated by the edge labels) and thus a 3–punctured sphere. The dotted lines with vertex labels represent the intersection of t with the boundary torus T_0, two disjoint loops and a connecting arc.

Moreover, Frigerio, Martelli and Petronio describe classes of orientable hyperbolic manifolds $\mathscr{M}_{g,k}$ with k torus cusps and one end of type an orientable surface of genus g [10]. They show, that these manifolds admit minimal $(g+k)$–tetrahedra ideal triangulations if and only if $g > k$, or $g = k$ and g even, see [10, Proposition 1.4]. Such triangulations must necessarily have $3k$ cone faces, and $(2g - k)$ dunce- or 3–fold faces, see [10, Lemma 2.1]. Hence, for g large enough, there exist minimal ideal triangulations of orientable hyperbolic 3–manifolds with all but an arbitrarily small portion of faces being dunce- or 3–fold-faces.

2.7 Low degree edges in minimal ideal triangulations

We now show that there are no edges of degree one or two in a minimal triangulation of a cusped hyperbolic 3–manifold, and that edges of degree three must be contained in layered ∂–punctured solid tori.

Lemma 8 *If $\mathscr{T}$ is a minimal ideal triangulation of a cusped hyperbolic 3–manifold M, then $\mathscr{T}$ has no edge of degree one.*

Proof Suppose e_1 is an edge of $\mathscr{T}$ of degree 1. Let $\tilde{\Delta}_1$ be the single tetrahedron containing e_1. Then $\tilde{\Delta}_1 = e_1 *$ e is the join of e_1 and the edge e opposite e_1. The edge e of $\mathscr{T}$ bounds a disc. This contradicts Corollary 5. □

Lemma 9 *If $\mathscr{T}$ is a minimal ideal triangulation of a cusped hyperbolic 3–manifold M, then $\mathscr{T}$ has no edge of degree two.*

Proof The *figure-eight* knot complement m004[2] and its sister m003 are the only cusped hyperbolic 3–manifold having an ideal triangulation with two ideal tetrahedra. Neither of these triangulations has an edge of degree

---

[2]We include the names of manifolds from the cusped hyperbolic manifold census of SnapPea.
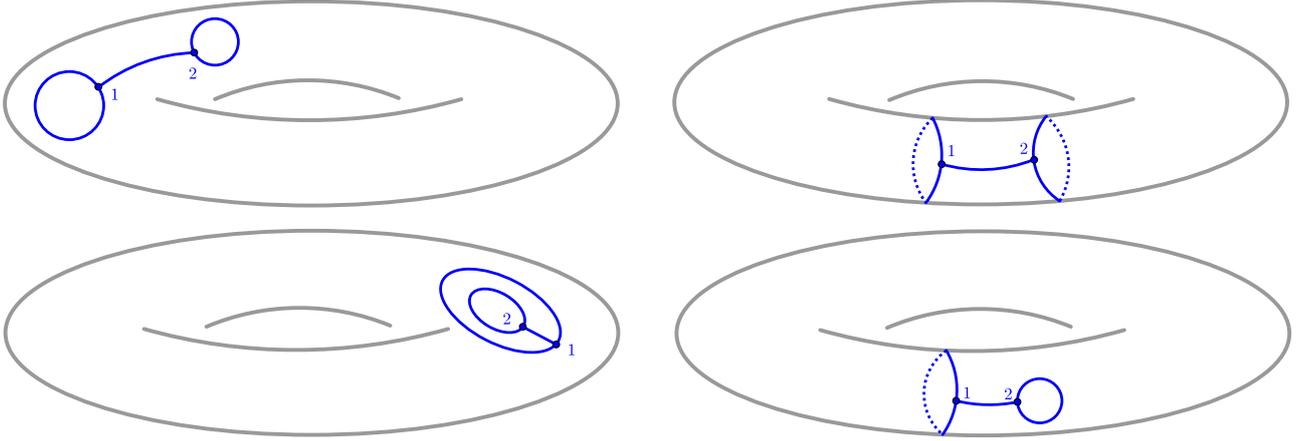
Figure 6: Four embeddings of the intersection of the dunce hat t with torus boundary component T_0. Top left: two inessential loops, not nested. Bottom left: two inessential loops, nested. Top right: two parallel essential loops. Bottom right: one essential loop and one inessential loop.

two. Hence, we may assume $\mathscr{T}$ has at least three tetrahedra and that the preimage of an edge e of degree two consists of two edges e' and e'' in two distinct tetrahedra, $\tilde{\Delta}'$ and $\tilde{\Delta}''$, respectively (note that an edge of degree two cannot be contained in a single tetrahedron).

We fix the following notation. The vertices of $\tilde{\Delta}'$ are denoted by A', B', C', and D' with $e' = A'D'$. The vertices of $\tilde{\Delta}''$ are denoted by A'', B'', C'', and D'' with $e'' = A''D''$. The face identifications are $(A'B'D') \leftrightarrow (A''B''D'')$ and $(A'C'D') \leftrightarrow (A''C''D'')$, with $(A'D') = (A''D'') = e$.
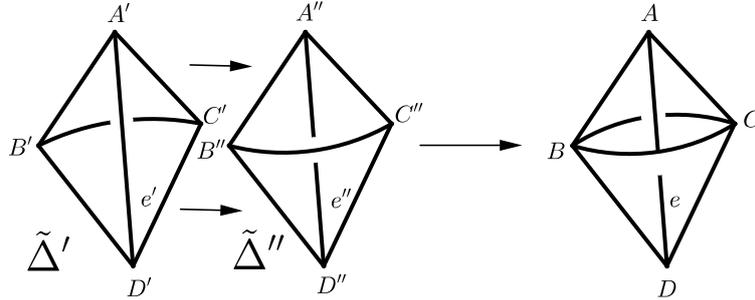
Figure 7: Edge of degree two.

Typically, in such a case of an edge of degree two, one can discard the two tetrahedra $\tilde{\Delta}'$ and $\tilde{\Delta}''$ and make new identifications $(A'B'C') \leftrightarrow (A''B''C'')$ and $(B'C'D') \leftrightarrow (B''C''D'')$.

If this is possible, then there is a triangulation of M having fewer tetrahedra than $\mathscr{T}$, contradicting that $\mathscr{T}$ is minimal. It follows that under the assumption of $\mathscr{T}$ minimal and e an edge of degree two, there must be an obstruction to such a re-identification, that is, some of the faces to be identified must already be identical. This can happen on the level of edges (Obstruction 1 below), or on the level of triangles (Obstruction 2 below). We consider both obstructions and show that each leads to a contradiction, establishing that there are no edges of degree two.

Obstruction 1. The edges $B'C'$ and $B''C''$ are identified, preventing the collapse, see Figure 8. In this situation we have two subcases:

(1) Edges are identified with opposite orientation (i.e., $B'C'$ is identified with $C''B''$). If this is the case, then in the truncated 3–manifold M^c of M with an open neighbourhood of the vertices removed, we have a

11

properly embedded Möbius band $\mathcal{N} \subset M^c$. If $\partial \mathcal{N}$ is an inessential curve on the vertex linking torus $T_0 \subset \partial M^c$, M admits an embedded projective plane and we have a connected summand of $\mathbb{R}P^3$ in M, contradiction to M being hyperbolic. Hence, assume that $\partial \mathcal{N}$ is an essential curve on T_0. From this situation we can construct a Seifert fibred space properly embedded in M which is impossible since our manifold is hyperbolic. To see this, note that the boundary of a small regular neighbourhood of $T_0 \cup \mathcal{N}$ is a torus. Since our manifold is hyperbolic this torus must be inessential, and thus it must bound a solid torus outside $T_0 \cup \mathcal{N}$. Hence the Seifert structure on $T_0 \cup \mathcal{N}$ extends to a Seifert fibering over this solid torus with at most two exceptional fibres, one at the centre line of the Möbius band $\mathcal{N}$.

(2) Edges are identified with coinciding orientation (i.e., $B'C'$ is identified with $B''C''$). In this case we have a properly embedded annulus $\mathscr{A} \subset M^c$. Note that, since M is hyperbolic, $\mathscr{A}$ must be either compressible or boundary-parallel. Hence, if one boundary component of $\mathscr{A}$ is an essential curve, the other one is, too, and both are parallel curves on the same torus boundary component $T_0 \subset \partial M^c$. In this case the annulus $\mathscr{A}$ runs parallel to T_0 with either AB and AC or BD and CD contained in the region between $\mathscr{A}$ and T_0. Now consider a torus running parallel to T_0 and bounding a collar neighbourhood of T_0 in M^c containing $\mathscr{A}$. Normalising this torus cannot result in a surface normally isotopic to T_0 because $\mathscr{A}$ acts as a barrier. Accordingly, the triangulation is not boundary efficient. This gives a contradiction to the triangulation being minimal due to Theorem 4.

If the boundary components are inessential, colour the bounding discs in the respective torus boundary components. More precisely, colour the disc near vertex B *purple*, and the one near vertex C *green*. W.l.o.g, edges AB and AC (truncated to run inside M^c) run inside the solid torus bounded by $\mathscr{A}$. By construction, the endpoint near B is purple and the endpoint near C is green. Moreover, the endpoints near A must either both be purple or both be green (note that $\mathscr{A}$ must be compressible). It follows that one of the edges AB and AC has endpoints with equal colours. It thus follows that this edge must be in the neighbourhood of the respective torus boundary components of M^c, see Figure 9. Again, we use Theorem 4 to produce a contradiction to the fact that the triangulation is minimal.

It follows that Obstruction 1. cannot occur.
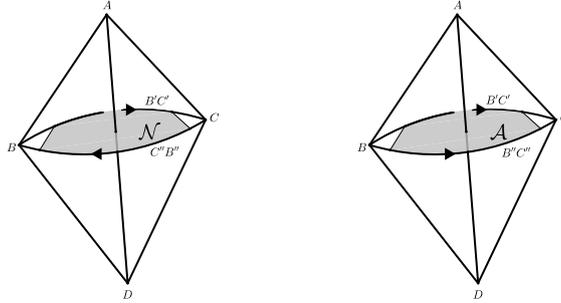
Figure 8: Obstruction 1 (1). Edge $B'C'$ identified with edge $C''B''$ forming a punctured $\mathbb{R}P^2$, or Möbius strip $\mathcal{N}$ (left). Obstruction 1 (2). Edge $B'C'$ identified with edge $B''C''$ forming an annulus $\mathscr{A}$ (right).

Obstruction 2. Faces from the front are identified with faces of the back (cf. Figure 7). As a result of such a face identification, the two edges $B'C'$ and $B''C''$ may become identified. If this the case we fall back onto the cases dealt with in Obstruction 1. Again, we have two subcases:

(1) Face $(A'B'C')$ is identified with the face $(B''C''D'')$ or $(A''B''C'')$ is identified with the face $(B'C'D')$. These situations are symmetric. As explained above we may assume that $B'C'$ and $B''C''$ are not identified. Suppose we flatten the two tetrahedra so that $(A'B'C')$ is identified with $(A''B''C'')$ and $(B'C'D')$ is identified with $(B''C''D'')$. Since the interiors of $B'C'$ and $B''C''$ were initially disjoint, the preimage of a point under this operation is either a point or an interval. In particular this flattening operation is cell-like
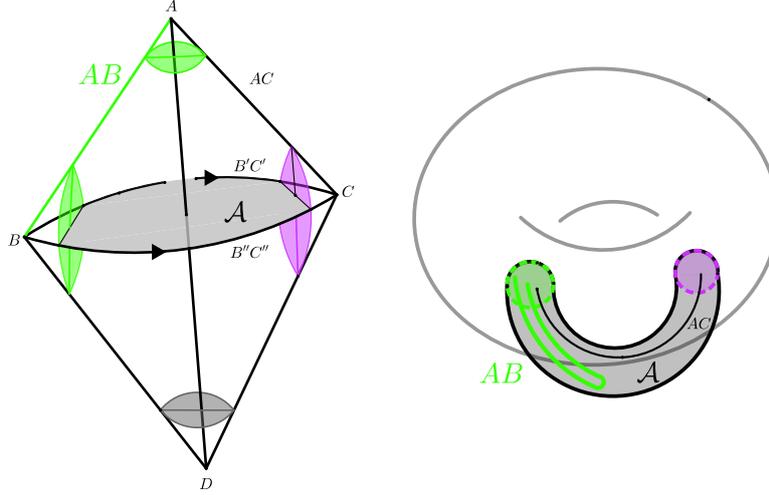
Figure 9: Obstruction 1 (2). Both boundary components of $\mathscr{A}$ are inessential ∂M^c.

and hence preserves the homeomorphism type of the manifold M. This gives a new triangulation of M with fewer tetrahedra which contradicts our assumption that the triangulation $\mathscr{T}$ is minimal.

(2) The face $(A'B'C')$ is identified with the face $(A''B''C'')$ or the face $(B'C'D')$ is identified with the face $(B''C''D'')$. Again, these situations are symmetric. However, if face $(A'B'C')$ is glued to face $(A''B''C'')$ without a twist, $B'C'$ is glued to $B''C''$ and we are done (alternatively, if $(A'B'C') \leftrightarrow (A''B''C'')$, then the vertex at $A' = A''$ is a manifold point and $\mathscr{T}$ is not an ideal triangulation). If the two faces are identified with a twist, all edges of the two triangles become identified. In particular $B'C'$ and $B''C''$ are identified and, again, we are done.

Since we have that $\mathscr{T}$ contains at least three tetrahedra, and that all possibilities of having an edge of degree two lead to a contradiction, there are no edges of degree two. $\qquad\square$

Lemma 10 *Let $\mathscr{T}$ be a minimal ideal triangulation of a cusped hyperbolic 3–manifold M. Suppose $\mathrm{LST}^{\star}$ is a layered ∂–punctured solid torus and there is a combinatorial map $\varphi\colon \mathrm{LST}^{\star} \to M$. Then φ is an embedding.*

Proof We need to show that no identifications of the boundary faces of $\mathrm{LST}^{\star}$ are possible in $\mathscr{T}$. If the two triangles of $\partial\,\mathrm{LST}^{\star}$ are identified, then M is a punctured lens space, contrary to the assumption that M is hyperbolic.

Hence, identifications can occur at most at the edges of $\partial\,\mathrm{LST}^{\star}$. But it follows from Theorem 6 that not all three edges of $\partial\,\mathrm{LST}^{\star}$ can be identified, and thus we only need to consider the case that two edges of $\partial\,\mathrm{LST}^{\star}$ are glued together.

A priori there are a total of six choices of how two boundary edges of $\partial\,\mathrm{LST}^{\star}$ can be identified (three pairs of edges and two possible orientations each). Ignoring the third edge, all of them result in a quadrilateral with all of its edges identified in the pattern shown in Figure 10 in the centre and the third boundary edge is one of he two diagonals of the quadrilateral. Denote the unique boundary edge of the quadrilateral by e.

We now pass to a small regular neighbourhood $N = \mathrm{Nhd}(\mathrm{LST}^{\star} \cap M^c)$ of $\mathrm{LST}^{\star}$ in the compact core M^c of M. By construction, its boundary ∂N must be a surface with boundary properly embedded in M^c. We can think of ∂N as the surface consisting of the octagon lying in the interior of the quadrilateral of $\partial\,\mathrm{LST}^{\star}$, and two quadrilaterals from the annulus that is the regular neighbourhood of the ideal edge e in M^c, see Figure 10 on the right. The surface is obtained by gluing the dotted edges of the octagon to the dotted edges of the quadrilateral. The solid edges then denote the boundary of ∂N.
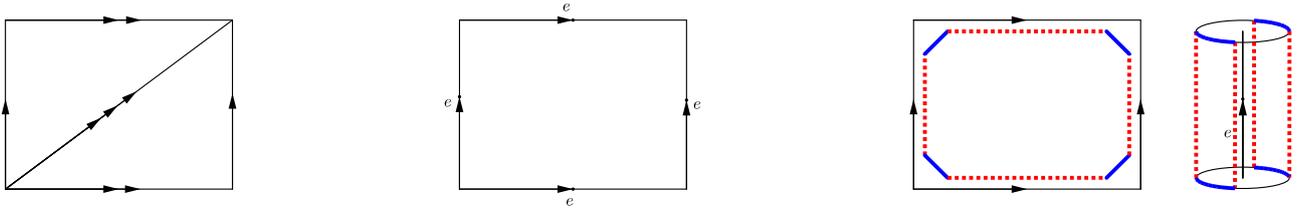
Figure 10: Left: Boundary of LST*. Centre: Boundary of LST* after disregarding the third boundary edge with two boundary edges identified. Right: Octagon and two quadrilaterals making up the boundary ∂N of a regular neighbourhood of LST* in M^c. Dotted edges are identified, solid edges denote the boundary of ∂N.

Up to symmetry, there are three ways two perform these gluings (notice that the red edges of the octagon can be paired with the red edges of the two quadrilaterals arbitrarily, but the orientation of the gluings is fixed). All of them are shown in Figure 11. The first set of gluings results in a once-punctured torus. This type of identification implies that the only identifications of e are the ones from inside LST*, a contradiction to the assumption that a pair of boundary edges of LST* are identified in $\mathscr{T}$. The second results in ∂N a once-punctured Klein bottle, and the third results in ∂N a thrice-punctured sphere.
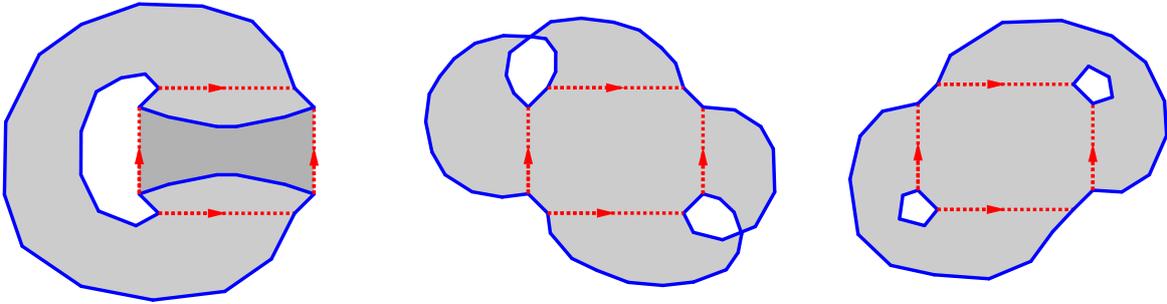
Figure 11: Three possible ways to obtain ∂N. Left: a once-punctured solid torus. Centre: a once-punctured Klein bottle. Right: a thrice-punctured sphere.

Since ∂N is separating it cannot be a once-punctured Klein bottle. Hence, it is a separating thrice-punctured sphere with all boundary components on the same boundary component T_0 of M^c. A parity argument shows that the number of essential boundary components of ∂N must be even. Moreover, note that, by construction, $T_0 \cap N$ is the neighbourhood of a pinched disc in T_0 and thus connected.

It follows that $T_0 \cap N$ must fall into one of two categories.

(1) Two boundary components of ∂N are essential and hence $T_0 \cap N$ must be equal to an annulus with a disc removed. Since N is separating, this disc must be outside N and thus disjoint from e. Hence, pasting this disc into $T_0 \cap N$ yields a necessarily boundary-parallel annulus.

(2) No boundary component of ∂N is essential and hence $T_0 \cap N$ must be equal to a subsurface of T_0 with two discs removed. Again, these discs must be outside N and thus disjoint from e. Pasting the discs into $T_0 \cap N$ results in a necessarily boundary-parallel disc.

Now, as in earlier arguments, consider a torus running parallel to T_0 and bounding a collar neighbourhood of T_0 in M^c containing the annulus or disc. Normalising this torus cannot result in a surface normally isotopic to T_0 because it must be disjoint from e. Accordingly, the triangulation is not boundary efficient. This gives a contradiction to the triangulation being minimal due to Theorem 4 and thus $\varphi\colon \text{LST}^\star \to M$ must be an embedding. $\qquad\square$

Lemma 11 *Let $\mathscr{T}$ be a minimal ideal triangulation of a cusped hyperbolic 3–manifold M and e be an ideal edge of degree three. Then e is contained in an embedded two tetrahedron subcomplex of $\mathscr{T}$ combinatorially equivalent to the layered ∂–punctured solid torus $\mathrm{LST}^\star(1,3,4)$. In particular, e is contained in a maximal layered ∂–punctured solid torus $\mathrm{LST}^\star$, twice in its core tetrahedron and once in the tetrahedron adjacent to the core, and no other edge in $\mathrm{LST}^\star$ is of degree three.*

Moreover, $\mathscr{T}$ contains at least three ideal tetrahedra and each ideal edge of degree three is contained in its own subcomplex of this type.

Proof We may and will work with the pseudo-manifold $\widehat{M}$. Let $e \in \mathscr{T}$ be an edge of degree 3. If e is incident with three distinct tetrahedra, then a 3–2–move produces a triangulation with fewer tetrahedra, contradicting minimality of the triangulation. If e is incident with exactly one tetrahedron, it cannot be of degree 3. Hence, e is incident with exactly two tetrahedra. A brute-force enumeration shows that there exists only one 2–tetrahedra complex with an edge of degree 3 in its interior – the layered solid torus $\mathrm{LST}(1,3,4)$ (see case $|\tilde{\Delta}_e| = 2$ in the proof of [21, Proposition 9] for details). Moreover, $\mathrm{LST}(1,3,4)$ has exactly one such degree 3 edge in its interior. This edge occurs twice in its core tetrahedron and once in the other tetrahedron (i.e., the tetrahedron adjacent to the core tetrahedron).

It follows from Lemma 10 that no identifications between the faces of this layered solid torus are possible, hence the subcomplex is embedded, contained in a maximal layered ∂–punctured solid torus and, in particular, the triangulation contains at least three tetrahedra. $\qquad\square$

We remark that an analysis of edges of degree four and five can be carried out following [21], but this is not needed for the main applications of this paper and hence omitted.

2.8 Layered punctured solid tori in ideal triangulations

The argument for closed 3–manifolds in [22] hinged on an understanding of the intersections of *maximal* layered solid tori. We now carry such an analysis out in the case of cusped hyperbolic 3–manifolds. This requires a different approach due to the fact that the combinatorics of minimal ideal triangulations is less restricted than the combinatorics of minimal and 0–efficient (material) triangulations.

As in [22] we focus on subcomplexes in triangulations that are maximal layered ∂–punctured solid tori (see Section 2.4 for a brief introduction).

Lemma 12 *Let $\mathscr{T}$ be a minimal ideal triangulation of a cusped hyperbolic 3–manifold M. Then the intersection of any two maximal layered ∂–punctured solid tori is either empty, an ideal vertex, or an ideal edge of $\mathscr{T}$.*

Proof Let $\mathrm{LST}_a, \mathrm{LST}_b \subset \mathscr{T}$ be two maximal ∂–layered solid tori. If their boundaries $\partial \mathrm{LST}_a$ and $\partial \mathrm{LST}_b$ meet in two faces, $\mathscr{T}$ must be homeomorphic to a punctured lens space. A contradiction. If LST_a and LST_b share a common tetrahedron, remove the tetrahedra in the intersection from one of them and observe, that they now meet in two faces. Again, a contradiction. Hence, we can assume that two maximal ∂–layered solid tori meet along a single face, three edges, or two edges.

(1) Assume that $\partial \mathrm{LST}_a$ and $\partial \mathrm{LST}_b$ meet along exactly two edges. It follows, that they must be identified along a spine of each of $\partial \mathrm{LST}_a$ and $\partial \mathrm{LST}_b$, referred to as their *common spine*. Let $\mathscr{D}_a$ and $\mathscr{D}_b$ be the discs obtained by taking the interior of $\partial \mathrm{LST}_a$ and $\partial \mathrm{LST}_b$ after truncating their vertex and removing their common spine. $\mathscr{D}_a$ and $\mathscr{D}_b$ can be glued together along sections of the exchange annuli in a

neighbourhood of the common spine yielding a 4–punctured sphere $\mathscr{S}$ properly embedded in M^c (the truncated compact core of M) and disjoint from the two edges of the common spine.

By construction, the four boundary components of $\mathscr{S}$ in one of the torus boundary components $T_0 \subset \partial M^c$ bound a regular neighbourhood of the intersection of the maximal layered solid tori LST_a and LST_b with T_0 – a regular neighbourhood of two discs pinched along four points (the four points being the intersection of the common spine with T_0). Since this regular neighbourhood is connected, of Euler characteristic -2, and with four boundary components, it must be a 4–punctured sphere contained in T_0. There are two cases: either all components of $\partial\mathscr{S}$ are trivial in T_0 (that is, $\mathscr{S}$ equals a disc of T_0 with three points removed), or two components of $\partial\mathscr{S}$ are non-trivial (that is, $\mathscr{S}$ equals an annulus of T_0 with two points removed).

First suppose $(\mathrm{LST}_a \cup \mathrm{LST}_b) \cap T_0$ is contained in a disc on T_0 and hence three of its boundary components bound discs in the complement of the intersection. Pasting these three discs of T_0 into $\mathscr{S}$ and pushing them slightly off T_0 yields a disc properly embedded in M^c. This disc can be extended to a torus parallel to boundary component T_0 of M^c. Moreover, this torus is disjoint from the two edges of $\mathscr{T}$ of the common spine, which both are incident to T_0. It follows that a normalised version of this torus cannot be equal to the vertex linking torus T_0 and thus $\mathscr{T}$ cannot be boundary efficient. A contradiction to Theorem 4.

Hence suppose $(\mathrm{LST}_a \cup \mathrm{LST}_b) \cap T_0$ is contained in an annulus on T_0 and hence two of its boundary components bound discs in the complement of the intersection. Pasting these two discs of T_0 into $\mathscr{S}$ and pushing them slightly off T_0 yields an annulus properly embedded in M^c and disjoint from the two edges of $\mathscr{T}$ of the common spine (which are both incident to T_0). This annulus together with the annulus in T_0 complementary to $(\mathrm{LST}_a \cup \mathrm{LST}_b) \cap T_0$ extends to a boundary parallel torus since otherwise M would contain a non-trivial Seifert fibred space. Again, this boundary parallel torus leads to a contradiction to the assumption that $\mathscr{T}$ is boundary efficient (cf. Theorem 4).

(2) Assume that $\partial\,\mathrm{LST}_a$ and $\partial\,\mathrm{LST}_b$ meet along exactly three edges. Following an analogous procedure as in the previous case, we arrive at two 3–punctured spheres $\mathscr{S}_1$ and $\mathscr{S}_2$ (instead of one 4–punctured sphere $\mathscr{S}$), both properly embedded in M^c, both meeting the same boundary component $T_0 \subset \partial M^c$. Again, it follows that a regular neighbourhood of $(\mathrm{LST}_a \cup \mathrm{LST}_b) \cap T_0$ is connected, of Euler characteristic -4 with six boundary components and thus a 6–punctured sphere.

Pasting in some of the boundary discs of this 6–punctured sphere, analogously to the procedure carried out in the previous case, yields a torus parallel to boundary component T_0 and disjoint from three edges of the triangulation which are all incident to T_0 – and we conclude as before by using Theorem 4.

(3) Assume that $\partial\,\mathrm{LST}_a$ and $\partial\,\mathrm{LST}_b$ meet along a single face. Again, we argue as in the previous cases. Here, we end up with one 3–punctured sphere $\mathscr{S}$ (the other one being filled in as a result of the face identification). A regular neighbourhood of $(\mathrm{LST}_a \cup \mathrm{LST}_b) \cap T_0$ now is connected, of Euler characteristic -1, has three boundary components, and thus is contained in a disc or an annulus on T_0. We again obtain a disc or an annulus that can be extended to a torus running parallel to T_0 which is disjoint of three edges of the triangulation incident to T_0 and we use Theorem 4 once again to finish the proof.

This completes the proof of the lemma. $\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\square$

3   Normal surfaces representing $\mathbb{Z}_2$–homology classes

Normal surfaces in 1–vertex triangulations of closed 3–manifolds that are dual to non-trivial first cohomology classes were used in [21, 22] in order to obtain lower bounds on the complexity of 3–manifolds. This section develops a similar theory for ideal triangulations. In this section, $\overline{M}$ is a compact, irreducible 3–manifold with non-empty boundary a finite union of tori, and M is its interior.

We consider surfaces representing elements of $H_2(M;\mathbb{Z}_2)$. We wish to emphasise that we do not work with $H_2(\overline{M}, \partial\overline{M};\mathbb{Z}_2)$. We recall the following definitions from the introduction. If S is connected, let $\chi_-(S) = \max\{0, -\chi(S)\}$, and otherwise let

$$\chi_-(S) = \sum_{S_i \subset S} \max\{0, -\chi(S_i)\},$$

where the sum is taken over all connected components of S. Note that S_i is not necessarily orientable. Define:

$$\| c \| = \min\{\chi_-(S) \mid [S] = c\}.$$

Then the surface S representing the class $c \in H_2(M;\mathbb{Z}_2)$ is $\mathbb{Z}_2$–*taut* if no component of S is a sphere, real projective plane, Klein bottle or torus and $\chi(S) = -\| c \|$. As in [35], one observes that every component of a $\mathbb{Z}_2$–taut surface is non-separating and geometrically incompressible.

3.1 Rank-1 colouring of edges and the canonical surface

We have $M = \widehat{M} \setminus \widehat{M}^{(0)}$, and thus Lefschetz duality gives a natural isomorphism

$$D_M \colon H_2(M;\mathbb{Z}_2) \cong H^1(\widehat{M}, \widehat{M}^{(0)};\mathbb{Z}_2).$$

An element $\varphi \in H^1(\widehat{M}, \widehat{M}^{(0)};\mathbb{Z}_2)$ can be viewed as an assignment of values 0 or 1 to each edge in the triangulation such that in each triangle the sum of all values at its edges equals zero modulo two. We say that the edge e is φ–*even* if $\varphi(e) = 0$ and φ–*odd* if $\varphi(e) = 1$. It follows that a tetrahedron falls into one of the following categories, which are illustrated in Figure 12:

Type Δ_q: A pair of opposite edges are φ–even, all others are φ–odd.

Type Δ_t: The three edges incident to a vertex are φ–odd, all others are φ–even.

Type $\Delta_\emptyset$: All edges are φ–even.
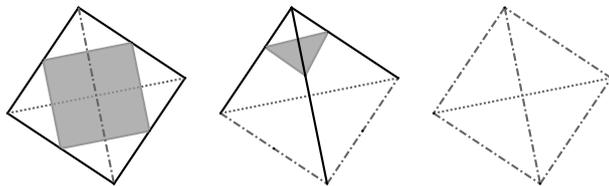
Figure 12: Tetrahedra of type Δ_q, Δ_t and $\Delta_\emptyset$. Grey and dashed edges are φ–even.

If φ is non-trivial, one obtains a unique normal surface, $S_\varphi = S_\varphi(\mathcal{T})$, with respect to $\mathcal{T}$ by introducing a single vertex on each φ–odd edge. This surface is disjoint from the tetrahedra of type $\Delta_\emptyset$; it meets each tetrahedron of type Δ_t in a single triangle meeting all φ–odd edges; and each tetrahedron of type Δ_q in a single quadrilateral dual to the φ–even edges. Moreover, S_φ is Lefschetz dual to φ and is termed the *canonical (normal) surface dual to φ*. Given $0 \neq c \in H_2(M;\mathbb{Z}_2)$ we also write $S_c = S_{D_M(c)}$ and call S_c the *canonical (normal) representative* of c.

Lemma 13 *Let M be a cusped hyperbolic 3–manifold with ideal triangulation $\mathcal{T}$ and $0 \neq c \in H_2(M;\mathbb{Z}_2)$. Then $\| c \| \leq \chi_-(S_c) = -\chi(S_c)$.*

Proof It follows from the construction and definitions that $[S_c] = c$ and we need to prove the equality in the statement. For this, it suffices to show that no component of S_φ is a sphere or a projective plane. Since M is irreducible and not $\mathbb{R}P^3$, no component can be a projective plane. Since each component of S_φ meets some edge, and it meets each edge in at most one point, no component can be a sphere. $\qquad\square$

3.2 Rank-2 colouring of edges

Given the subgroup $H = \langle \varphi_1, \varphi_2 \rangle \cong \mathbb{Z}_2 \oplus \mathbb{Z}_2$ of $H^1(\widehat{M}, \widehat{M}^{(0)}; \mathbb{Z}_2)$, we now introduce a refinement of the above colouring. Since $\varphi_1 + \varphi_2 = \varphi_3$, there are four types of edges:

edge e is called *H–even* or *0–even* if $\varphi_i[e] = 0$ for each $i \in \{1,2,3\}$; and

edge e is called *i–even* if $\varphi_i[e] = 0$ for a unique $i \in \{1,2,3\}$.

Let $\{i,j,k\} = \{1,2,3\}$. The normal corners of the normal surface $S_{\varphi_i}(\mathcal{T})$ are precisely on the j–even and k–even edges. Edge e is φ_i–even if it is i–even or 0–even. It follows that a face of a tetrahedron either has all of its edges 0–even; or it has two 1–even (or 2-even or 3-even respectively) and one 0–even edge; or it has one 1–even, one 2–even and one 3–even edge. Whence a tetrahedron falls into one of the following categories:

Type Δ_{qtt}: One edge is 0–even, the opposite edge is i–even, and one vertex of the latter is incident with two j–even edges, and the other with two k–even edges, where $\{i,j,k\} = \{1,2,3\}$. (There are six distinct colourings of such a tetrahedron, in the following referred to as *sub-types*.)

Type Δ_{qq}: A pair of opposite edges are 0–even, all others are i–even for a unique $i \in \{1,2,3\}$. (There are hence three distinct sub-types.)

Type Δ_{tt}: The three edges incident to a vertex are i–even for a fixed $i \in \{1,2,3\}$, and all others are 0–even. (There are hence three distinct sub-types.)

Type $\Delta_{\emptyset}$: All edges are 0–even.

Type Δ_{qqq}: Each vertex is incident to an i–even edge for each $i \in \{1,2,3\}$. (In particular, no edge is 0–even, opposite edges are of the same type, and there are two distinct sub-types of tetrahedra.)

For each type, one sub-type is shown in Figure 13; the dashed grey edges correspond to 0–even edges and the normal discs in S_{φ_i} have the same colour as the i–even edges. The remaining sub-types are obtained by permuting the colours.
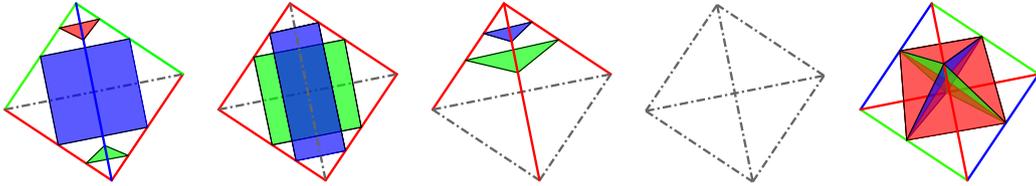
Figure 13: Tetrahedra of type Δ_{qtt}, Δ_{qq}, Δ_{tt}, $\Delta_{\emptyset}$ and Δ_{qqq}. Dashed grey edges are 0–even.

3.3 Combinatorial bounds for triangulations

The set-up and notation of the previous subsection is continued. Let

$$
\begin{aligned}
n_{\mathrm{qtt}}(\mathcal{T}) &= \text{ number of tetrahedra of type } \Delta_{\mathrm{qtt}}, \\
n_{\mathrm{qq}}(\mathcal{T}) &= \text{ number of tetrahedra of type } \Delta_{\mathrm{qq}}, \\
n_{\mathrm{tt}}(\mathcal{T}) &= \text{ number of tetrahedra of type } \Delta_{\mathrm{tt}}, \\
n_{\emptyset}(\mathcal{T}) &= \text{ number of tetrahedra of type } \Delta_{\emptyset}, \\
n_{\mathrm{qqq}}(\mathcal{T}) &= \text{ number of tetrahedra of type } \Delta_{\mathrm{qqq}}, \\
\mathfrak{e}(\mathcal{T}) &= \text{ number of 0–even edges}, \\
\tilde{\mathfrak{e}}(\mathcal{T}) &= \text{ number of preimages of 0–even edges in } \widetilde{\Delta}.
\end{aligned}
$$

The number of tetrahedra in $\mathscr{T}$ is

$$|\mathscr{T}| = n_{\mathrm{qtt}}(\mathscr{T}) + n_{\mathrm{qq}}(\mathscr{T}) + n_{\mathrm{tt}}(\mathscr{T}) + n_{\emptyset}(\mathscr{T}) + n_{\mathrm{qqq}}(\mathscr{T}).$$

For the remainder of this subsection, we write $n_{\mathrm{qtt}} = n_{\mathrm{qtt}}(\mathscr{T})$, etc.

Let $\hat{K}$ be the complex in $\mathscr{T}$ spanned by all 0–even edges (i.e., the complex of all edges, faces and tetrahedra of $\mathscr{T}$ containing only 0–even edges) and let K be the ideal complex obtained from $\hat{K}$ by removing its vertices.

Let N be a small regular neighbourhood of K. Then ∂N is a normal surface; it meets each tetrahedron in the same number and types of normal discs as $S_{\varphi_1} \cup S_{\varphi_2} \cup S_{\varphi_3}$ except for the tetrahedra of type Δ_{qqq}, which it meets in four distinct normal triangle types instead of three distinct normal quadrilateral types. The normal coordinate of ∂N is thus obtained by taking the sum of the normal coordinates of S_{φ_1}, S_{φ_2}, and S_{φ_3}, and adding to this the *tetrahedral solution* of each tetrahedron of type Δ_{qqq}. (The tetrahedral solution is obtained by adding all triangle coordinates of a tetrahedron and subtracting all quadrilateral coordinates; see, for instance, [23, Section], or [36, Section 2.8]. Each tetrahedral solution has both positive and negative coordinates.) Hence

$$\chi(S_{\varphi_1}) + \chi(S_{\varphi_2}) + \chi(S_{\varphi_3}) + n_{\mathrm{qqq}} = \chi(\partial N) = 2\chi(N) = 2\chi(K) = -2\mathfrak{e} + n_{\mathrm{tt}} + 2n_{\emptyset},$$

where the rightmost equality is a consequence of the following computation: By definition, the Euler characteristic of K equals minus the number of 0–even edges plus the number of triangles with only 0–even edges minus the number of tetrahedra of type $\Delta_{\emptyset}$ (note that K is ideal and hence has no vertices). Since every triangle occurs in exactly two tetrahedra, every tetrahedron of type Δ_{tt} accounts for one half of a triangle with 0–even edges and every tetrahedron of type $\Delta_{\emptyset}$ accounts for two triangles of 0–even edges. Hence, in total we have

$$\chi(K) = -\mathfrak{e} + \frac{1}{2}n_{\mathrm{tt}} + 2n_{\emptyset} - n_{\emptyset} = -\mathfrak{e} + \frac{1}{2}n_{\mathrm{tt}} + n_{\emptyset}.$$

In particular, $\chi(S_{\varphi_1}) + \chi(S_{\varphi_2}) + \chi(S_{\varphi_3}) + n_{\mathrm{qqq}}$ is even. Rearranging the above equality gives

$$n_{\mathrm{tt}} + 2n_{\emptyset} - n_{\mathrm{qqq}} = 2\mathfrak{e} + \chi(S_{\varphi_1}) + \chi(S_{\varphi_2}) + \chi(S_{\varphi_3}), \tag{3.1}$$

and hence:

$$\begin{aligned}
\tilde{\mathfrak{e}} &= n_{\mathrm{qtt}} + 2n_{\mathrm{qq}} + 3n_{\mathrm{tt}} + 6n_{\emptyset} \\
&= 2|\mathscr{T}| - n_{\mathrm{qtt}} + n_{\mathrm{tt}} + 4n_{\emptyset} - 2n_{\mathrm{qqq}} \\
&= 2|\mathscr{T}| - n_{\mathrm{qtt}} - n_{\mathrm{tt}} + 4\mathfrak{e} + 2(\chi(S_{\varphi_1}) + \chi(S_{\varphi_2}) + \chi(S_{\varphi_3})). \tag{3.2}
\end{aligned}$$

Lemma 14 *Let M be a cusped hyperbolic 3–manifold of finite volume, and suppose that $\varphi_1, \varphi_2 \in H^1(\widehat{M}, \widehat{M}^{(0)}; \mathbb{Z}_2)$ are non–trivial classes with $\varphi_1 + \varphi_2 = \varphi_3 \neq 0$. Let $\mathscr{T}$ be a minimal triangulation, S_{φ_i} be the canonical surface dual to φ_i, and let $\mathfrak{e}_d$ denote the number of 0–even edges of degree d. Then*

$$\mathfrak{e}_3 = n_{\mathrm{qtt}} + n_{\mathrm{tt}} - 2(|\mathscr{T}| + \chi(S_{\varphi_1}) + \chi(S_{\varphi_2}) + \chi(S_{\varphi_3})) + \sum_{d=5}^{\infty}(d-4)\mathfrak{e}_d. \tag{3.3}$$

Proof Lemmata 8 and 9 imply that the smallest degree of an edge in $\mathscr{T}$ is three. One has $\tilde{\mathfrak{e}} = \sum d\mathfrak{e}_d$ and $\mathfrak{e} = \sum \mathfrak{e}_d$. Inserting this into Equation (3.2) yields the desired equality. $\square$

4   Quadrilateral surfaces and minimal ideal triangulations

In this section we prove the following result, first stated in the introduction. The essence of the proof is a counting argument taking into account even edges, compression discs for the canonical normal representatives, and types of tetrahedra. This is modelled on the blueprint for closed 3–manifolds [20, 21, 22, 28].

Theorem 1 Let $\mathscr{T}$ be an ideal triangulation of a cusped hyperbolic 3–manifold M and let $H \leq H_2(M, \mathbb{Z}_2)$ be a rank 2 subgroup. Then

$$|\mathscr{T}| \geq \sum_{0 \neq c \in H} \| c \|.$$

Moreover, in the case of equality the triangulation is minimal, each canonical normal representative of a non-zero element in $H \leq H_2(M, \mathbb{Z}_2)$ is taut and meets each tetrahedron in a quadrilateral disc, and the number of tetrahedra in the triangulation is even.

Note that the conclusion that *each canonical normal representative meets each tetrahedron in a quadrilateral disc* is equivalent to saying that *all tetrahedra are of type* Δ_{qqq}.

Proof It suffices to work with a minimal ideal triangulation. For if we can show any such a triangulation satisfies the desired inequality, then clearly an arbitrary ideal one does also. Moreover, if any ideal triangulation admits equality, then it clearly must be minimal.

Let $\mathscr{T}'$ be a minimal ideal triangulation of a cusped hyperbolic 3–manifold M. Fix a subgroup $H < H_2(M, \mathbb{Z}_2)$ of rank 2, and let S_i, $i = 1, 2, 3$, be the canonical normal representatives of the non-zero classes $0 \neq \varphi_i \in H$. Let $\mathscr{T}$ be the minimal triangulation that has the smallest number of tetrahedra of type $\Delta_\emptyset$ amongst all minimal triangulations of M that can be obtained from $\mathscr{T}'$ by a sequence of 4-4-moves.

Following the notation used in Section 3, the union of the normal surfaces S_i, $i = 1, 2, 3$, intersects the tetrahedra of $\mathscr{T}$ in n_{qtt} tetrahedra of type Δ_{qtt}, n_{qq} tetrahedra of type Δ_{qq}, and so on.

Let d_i be the maximal number of pairwise disjoint compression discs for S_i, whose union does not separate any component of S_i. We have $\|c_i\| \leq -\chi(S_i)$ due to Lemma 13. Moreover, in the presence of compression discs, we obtain the stronger bound

$$\sum_{i=1}^{3} \| c_i \| \leq -\sum_{i=1}^{3} (\chi(S_i) + 2d_i)$$

since compressions do not change the homology class and do not introduce components that are spheres or projective planes. To simplify the expression, and with a view towards the second part of the theorem, we write

$$\sum_{i=1}^{3} \| c_i \| = -D - \sum \chi(S_i),$$

where $D \geq 0$. Here, we take care of the possibility that the canonical surfaces are incompressible but not norm minimising, and we note that each independent compression disc contributes 2 to D.

For the sake of contradiction, assume

$$|\mathscr{T}| < \sum_{i=1}^{3} \| c_i \|.$$

Lemma 14 now implies the strict inequality

$$n_{tt} + n_{qtt} + 2D + \sum_{d=5}^{\infty} (d - 4)\mathfrak{e}_d < \mathfrak{e}_3. \tag{4.1}$$

Hence, proving the first part of Theorem 1 is equivalent to showing that this is impossible. That is, we need to prove that the left hand side of inequality (4.1) must be at least as large as the right hand side.

The idea of the proof is to counter-balance each contribution to the right hand side with a contribution to the left hand side of (4.1) of equal or larger size. A contribution to the right hand side only arises from a degree three edge, and we have already shown that each such edge is contained in a layered ∂–punctured solid torus subcomplex. Any such complex contains at most one edge of degree three. The organising principles for the

counting argument are as follows. We identify sets of degree three edges that can be associated with an H–even edge f such that $d(f) - 4$ is as least as large as the size of the associated set of degree three edges (see Claim 2 and the paragraph directly after the proof of Claim 2). In doing so, we may double count some of the degree three edges, but we make sure not to double count any of edges f. For the remaining degree three edges, there are only a small number of special cases, in which we need to also take into account the remaining terms $n_{\mathrm{tt}} + n_{\mathrm{qtt}} + 2D$. Here we need to make sure that such contributions to the left hand side are only counted at most once. Moreover, in these special cases the contribution to the left hand side of (4.1) is strictly greater than the contribution to the right hand side.

The following terminology helps organise the counting argument. A *deficit in (4.1)* is when the contribution to the left is bigger than to the right, a *balance in (4.1)* is when they are equal, and otherwise we have a *gain in (4.1)*. Thus, for proving the first part of Theorem 1 it is sufficient to disregard all deficits and balances and show that it is impossible to construct a triangulation exhibiting a gain. For the second part of Theorem 1 we also need to keep track of balances.

Recall that the core tetrahedron of any layered ∂–punctured solid torus LST is a standard one-tetrahedron layered ∂–punctured solid torus $\mathrm{LST}^{\star}(1,2,3)$. A canonical normal representative of a class of $H_2(M, \mathbb{Z}_2)$ necessarily intersects $\mathrm{LST}^{\star}(1,2,3)$ in one particular quadrilateral type, or not at all. This can be seen by enumerating all $0/1$-colourings of edges in $\mathrm{LST}^{\star}(1,2,3)$ such that every face of $\mathrm{LST}^{\star}(1,2,3)$ has zero or two edges of colour 1. Accordingly, the intersection of $\mathrm{LST}^{\star}(1,2,3)$ with all three canonical normal representatives of non-trivial classes of H must result in a tetrahedron of type Δ_{qq} or $\Delta_{\emptyset}$.

It follows that either all tetrahedra in LST are of type Δ_{qq} or of type $\Delta_{\emptyset}$ (this follows directly from looking at the pattern of H–odd edges in the boundary). Accordingly we say that LST is of type Δ_{qq} or $\Delta_{\emptyset}$. A layered ∂–punctured solid torus LST of type Δ_{qq} has one H–even boundary edge and the other two boundary edges are H–odd.

Claim 0: All edges of degree 3 in $\mathscr{T}$ are H–even.

Proof of Claim 0: Due to Lemma 11, every edge $e \in \mathscr{T}$ of degree three can be associated to a maximal layered ∂–punctured solid torus $\mathrm{LST}_e \subset \mathscr{T}$ containing e in its interior. More precisely, e occurs twice in the core tetrahedron of LST_e and once in the tetrahedron adjacent to the core. If LST_e is of type $\Delta_{\emptyset}$, then naturally e must be H–even. Hence, let LST_e be of type Δ_{qq}. Since e occurs twice in the core tetrahedron and the other two edges in the core tetrahedron occur in it one and three times respectively, they only way the core tetrahedron of LST_e can be of type Δ_{qq} is for e to be H–even. □

A layered ∂–punctured solid torus LST has exactly one boundary edge of degree 1 with respect to LST. We call the other two boundary edges *non-unital*.

Claim 1: A maximal layered ∂–punctured solid torus LST of type $\Delta_{\emptyset}$ containing an H–even edge of degree three either has an interior H–even edge of degree ≥ 6, or non-unital boundary edges g and f with $d_{\mathrm{LST}}(g) = 3$ and $d_{\mathrm{LST}}(f) \geq 5$.

Proof of Claim 1: In a layered solid torus of type $\Delta_{\emptyset}$ every edge is necessarily H–even. We can thus ignore the edge-coulourings.

Every layered ∂–punctured solid torus LST in a minimal triangulation $\mathscr{T}$ and containing a degree three edge is obtained from a layered ∂–punctured solid torus of type $\mathrm{LST}^{\star}(1,3,4)$ by iteratively layering onto one of the non-unital boundary edges. In $\mathrm{LST}^{\star}(1,3,4)$, the non-unital boundary edges g and f have degrees $d_{\mathrm{LST}}(g) = 3$ and $d_{\mathrm{LST}}(f) = 5$. If we layer on f we create an interior edge of degree six and we are done. If we layer on g we create an interior edge of degree four, the formerly unital boundary edge is now of degree three, and the degree of g – still being in the boundary – is increased by two. Layering on the unital edge contradicts minimality, since it creates an edge of degree two. Iterating this procedure proves the claim. □

Claim 2: Let LST $\subset \mathscr{T}$ be a maximal layered ∂–punctured solid torus that contains an interior edge $e \in \mathscr{T}$ of degree three. Then LST is incident with an H–even edge f such that if m is the number of degree three edges of all maximal layered ∂–punctured solid tori incident with f, then $d(f) - 4 \geq m$ (i.e. contributing a deficit or balance in (4.1)), unless the following three conditions are satisfied for LST:

(1) all interior H–even edges distinct from e are of degree four,

(2) the solid torus is of type Δ_{qq}, and

(3) the unique H–even boundary edge f is of degree one with respect to LST.

We call a maximal layered ∂–punctured solid torus that contains an interior edge e of degree three and satisfies the conditions (1)–(3) *internally unbalanced*.

Proof of Claim 2: No interior edge of LST distinct from e can be of degree less than four (see Lemma 11). If LST contains an interior H–even edge of degree > 4, then we choose f to be this edge. It follows that $m = 1$, and thus $d(f) - 4 \geq m$ holds. Hence assume that all interior H–even edges distinct from e are of degree four.

Suppose that LST is of type $\Delta_\emptyset$. All edges of LST are H–even and, by the above, all its interior edges apart from e are of degree four. But then, by Claim 1, there exists a boundary edge $f \in \partial$ LST satisfying $d_{\mathrm{LST}}(f) \geq 5$. Now $m - 1$ additional maximal layered solid tori containing degree three edges meet f in one of their H–even boundary edges. By Lemma 12, f is of degree at least $4 + 2m$. Thus $d(f) - 4 \geq 2m > m$ and we have a deficit in (4.1).

The last case is that LST is of type Δ_{qq}. Since LST is of type Δ_{qq} it has a unique H-even boundary edge $f \in \partial$ LST. If $d_{\mathrm{LST}}(f) > 1$, then $d_{\mathrm{LST}}(f) \geq 3$. Again $m - 1$ additional maximal layered solid tori containing degree three edges are incident with f. By Lemma 12, f is of degree at least $2 + 2m$. Thus $d(f) - 4 \geq 2m - 2$. So if $m \geq 2$, then $d(f) - 4 \geq m$. If $m = 1$, then we claim that the degree of f is at least five. Suppose it is not. Then the neighbourhood of f is modelled on a layered ∂–punctured solid torus having one more tetrahedron than LST but extra identifications between the boundary edges. This is not possible due to maximality and Lemma 10. Hence no identifications are possible amongst the edges and so $d(f) \geq 5$. $\qquad\square$

The set of internally unbalanced maximal layered ∂–punctured solid tori can now be canonically partitioned, by grouping together those which meet along their unique H–even boundary edge. Fix one such set of m internally unbalanced maximal layered ∂–punctured solid tori around a common H–even edge f. Its contribution to the right hand side of (4.1) is m. But f must be of degree at least $2m$ by Lemma 12, and its contribution to the left hand side of (4.1) is $\geq 2m - 4$. Hence, there is a deficit if $m > 4$. We therefore consider the cases where $m \leq 4$. If $m = 4$ and $d(f) > 8$, we also obtain a deficit.

Special cases: Since the degree $d = \deg(f)$ of f is bounded below by four, we obtain at most a balance if:

$$(m, d) \in \{(4, 8), (3, 7), (2, 6), (1, 5)\},$$

and we obtain at most a gain if:

$$(m, d) \in \{(3, 6), (2, 5), (2, 4), (1, 4)\}.$$

We call the edges in the above collections the *balancing edges* and the *gaining edges*, respectively.

The maximal potential gain is 2 in the case $(2, 4)$ and it is 1 in all other cases. Our strategy now is to visit each edge f that gives a potential gain, show that it either achieves a deficit, or (even stronger) that it achieves a deficit for itself and all H–even edges in its star.

Before taking a closer look at these cases, we first prove an auxiliary statement. For this and for the remainder of this proof, we denote the tetrahedra around f by $\Delta_1, \ldots, \Delta_d$ (labelled in cyclic order such that Δ_i shares a triangle with Δ_{i+1}; some tetrahedron may appear in this sequence multiple times). Moreover, without loss of

generality and following Lemma 12, the internally unbalanced maximal layered ∂–punctured solid torus LST_i incident to f can be assumed to contain Δ_{2i-1}, $1 \leq i \leq m$.

Claim 3: In all four cases $(m,d) \in \{(3,6),(2,5),(2,4),(1,4)\}$, the edge f is surrounded by d pairwise distinct tetrahedra.

Proof of Claim 3: Tetrahedra contained in distinct internally unbalanced maximal layered ∂–punctured solid tori must be pairwise distinct. Hence the tetrahedra Δ_{2i-1}, $1 \leq i \leq m$, are pairwise distinct since Δ_{2i-1} is contained in LST_i.

A tetrahedron of type Δ_{qtt} cannot occur twice around f because it has only one H–even edge and, of course, there can be no tetrahedron of type Δ_{qqq}. Moreover, only Δ_3 in case $(1,4)$ can be of type $\Delta_\emptyset$, and hence a tetrahedron of type $\Delta_\emptyset$ cannot occur more than once around f.

If f occurs more than once in a tetrahedron Δ of type Δ_{tt}, then we must be in $(2,5)$ or $(1,4)$ since f is contained in a face only having H–even edges. In case $(2,5)$ we have $\Delta = \Delta_4 = \Delta_5$; in case $(1,4)$ we either have $\Delta = \Delta_2 = \Delta_3$, or $\Delta = \Delta_3 = \Delta_4$, or $\Delta = \Delta_2 = \Delta_4$. Assume we have one of the three cases where $\Delta = \Delta_i = \Delta_{i+1}$. Then two faces of Δ are identified. However, there is only one way to identify two faces of a tetrahedron of type Δ_{tt} in an orientation preserving manner, and this creates an edge of degree one, which is not possible. Hence we are in case $(1,4)$ and $\Delta = \Delta_2 = \Delta_4$. Then Δ_3 is necessarily of type $\Delta_\emptyset$. But $\Delta_\emptyset$ cannot share two distinct faces with $\Delta_2 = \Delta_4$ because a tetrahedron of type Δ_{tt} only has one triangle of H–even edges.

The remaining tetrahedron type Δ_{qq} can occur in all cases and in all places around f. To start suppose Δ_2 is of type Δ_{qq}; one of the H–even edges of Δ_2 is f and we denote the other f'. If Δ_2 appears more than once around f, then $f = f'$. In particular the two faces of Δ_2 containing f' meet Δ_1 in at least two edges and hence cannot be identified with a tetrahedron Δ_{2i-1}, $2 \leq i \leq m$. This shows that in the cases $(3,6)$ and $(2,4)$ all tetrahedra must be pairwise distinct by the symmetry of the situation.

In the case $(2,5)$ this implies that Δ_2 cannot equal Δ_4. However, Δ_2 also meets a face of Δ_3 and by the same reasoning cannot equal Δ_5. The only remaining possibility in this case is that Δ_4 equals Δ_5. Here again, Δ_5 meets Δ_1 in a face, and hence if Δ_4 equals Δ_5, then Δ_1 and Δ_3 share two edges, which is not possible.

It remains to consider the case $(1,4)$. First suppose that Δ_2 and Δ_3 are the same tetrahedron. There are two possibilities that are combinatorially equivalent. Taking all the face pairings into account, we conclude that Δ_2 is a 1–tetrahedron $\mathrm{LST}^\star(1,2,3)$. But this meets Δ_1, and hence LST_1, in three edges, a contradiction. See Figure 14 for an illustration of this case.
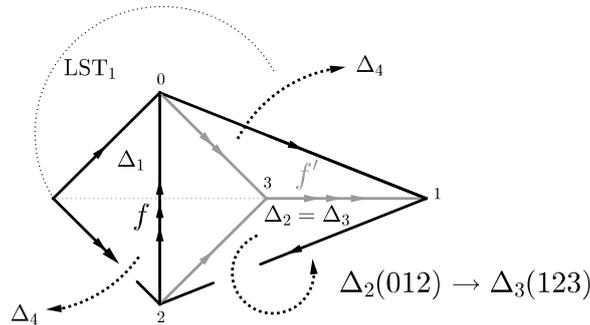
Figure 14: Identifying Δ_2 and Δ_3 in case $(1,4)$ yields a one-tetrahedron layered solid torus meeting LST_1 in three edges.

Hence suppose without loss of generality that Δ_2 and Δ_4 are the same tetrahedron. Analysing the two different gluings that arise, we see that Δ_2 is layered on one of the two H-odd boundary edges of Δ_1. But this either contradicts maximality of LST_1 or it contradicts Lemma 10, see Figure 15.
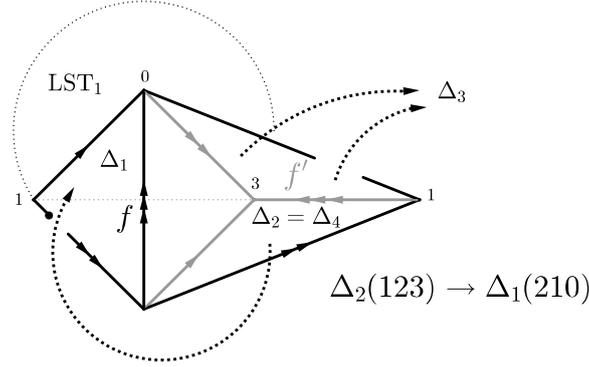
Figure 15: Identifying Δ_2 and Δ_4 in case $(1,4)$ yields a layering on the boundary of LST_1.

This completes the proof of the claim. □

The case $m = 2$ and $d = 4$: Since f is H–even of degree four, we have to find a contribution to the left hand side of (4.1) of 2 to obtain a balance and > 2 to obtain a deficit.

By Claim 3 and by the preliminary observations above, we can assume that all four tetrahedra surrounding f are distinct and tetrahedra Δ_1 and Δ_3 are of both of type Δ_{qq}. If the quadrilaterals in Δ_1 and Δ_3 are of distinct colours, Δ_2 and Δ_4 must both be of type Δ_{qtt}. Hence, they contribute 2 to the left hand side of (4.1). Moreover, since these two tetrahedra then do not have any H–even edges other than f, they cannot contribute to any other set of maximal layered solid tori united along an H–even edge, and their contribution to (4.1) is global. In particular, there is at most a balance and not a gain. Now one of the surfaces, S_i, links the edge f in an annulus. We can perform a compression of this annulus, using a disc transverse to f. This is indeed a compression disc since the resulting surface is still connected, due to its intersection with the H-odd boundary edges of the layered ∂–punctured solid tori. Whence in total, we have a deficit of 4. We also note that this compression disc is confined to a neighbourhood of f and that the tetrahedra Δ_2 and Δ_4 have a *unique H–even edge*. Thus, the quadrilateral discs of S_i affected by this compression allow no other compressions of this kind and we may associate the total deficit with f.

We can thus assume that Δ_1 and Δ_3 are coloured the same. This implies that Δ_2 and Δ_4 must both be of type Δ_{qq} with the same colours as well. In particular, the octahedron that is the boundary of the star of f has eight H-odd edges, all of degree two inside the star of f, and a 4–cycle of H–even edges, all of degree 1 inside the star of f. In this case we can perform compressions of two of the surfaces as above. Whence in total we have a contribution of 8 to the left hand side and a contribution of 2 to the right hand side. We now allocate a contribution of 3 to f so that in total we have a deficit of 1 associated with f. We then associate a contribution of 2.5 to each of the other two H–even edges of Δ_2 and Δ_4. Hence each of these (if it is in our collection of gaining or balancing edges) now has a deficit associated with it and is not visited again. In particular, any further compression discs we find in the star of an H–even edge is disjoint from the ones we just found.

This completes the case of gaining edges of type $(2,4)$ showing that they all contribute a deficit. We may thus assume that the maximal potential gain of an edge we have not yet considered is 1.

The case $m = 3$ and $d = 6$: Since f is H–even of degree six, it contributes 2 to the left hand side of (4.1). Hence, we have to find an additional contribution of 1 to obtain a balance and > 1 to obtain a deficit.

Since LST_1, LST_2 and LST_3 are internally unbalanced, Δ_1, Δ_3 and Δ_5 are all of type Δ_{qq}. If not all of them are coloured the same, at least two of Δ_2, Δ_4 and Δ_6 must be of type Δ_{qtt}, contributing 2 to the left hand side of (4.1) and we are done (again, note that the two tetrahedra of type Δ_{qtt} do not have any other H–even edges, and their contribution to the star of f is global).

Hence all tetrahedra Δ_i, $1 \leq i \leq 6$, are of type Δ_{qq}, all with the same colours. This implies that one can perform a compression of two of the surfaces S_i, $i \in \{1, 2, 3\}$, by cutting along the two tubes around the H–even edge f, and pasting two discs, each intersecting f once (see Figure 16 for an illustration of this process). The compressed surfaces are still embedded.

Moreover, note that three of the quadrilaterals in the tubes linking edge f are contained in layered solid tori LST_j, $1 \leq j \leq 3$. Hence, each such quadrilateral has its vertices identified in diagonally opposed pairs. In particular, the compressed surfaces still have the same number of connected components, and their components still have non-positive Euler characteristic and Lemma 13 applies. Altogether, the possibility of such compressions yields a contribution of 8 to the term $2D \geq 4\sum_{i=1}^{3} d_i$ of the left hand side of (4.1).
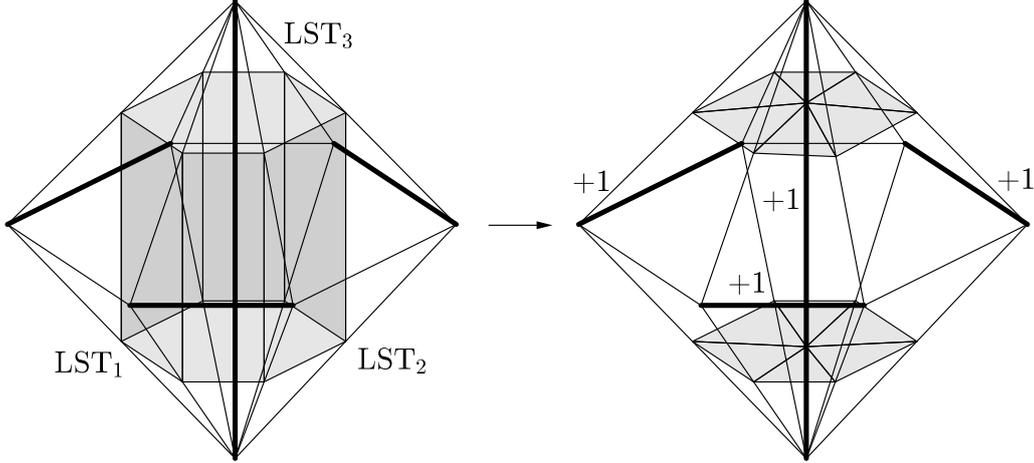
Figure 16: Compression along an H–even edge surrounded by tetrahedra of type Δ_{qq}. The picture only shows one of the two parallel surfaces.

We distribute this contribution onto f and the three remaining H–even edges of Δ_2, Δ_4 and Δ_6 equally. Thus, all of the four edges receive a sufficient contribution to give a deficit in (4.1), and are not visited again (except that they may receive additional deficit from other edges). In particular, any other compression discs we find do not intersect the tetrahedra in the star of f.

This completes the case of gaining edges of type $(3, 6)$ showing that they all contribute a deficit.

The case $m = 2$ and $d = 5$: Since f is H–even of degree five, it contributes 1 to the left hand side of (4.1). Hence, we have to find an additional contribution of 1 to achieve a balance and > 1 to achieve a deficit.

Again, Δ_1 and Δ_3 can be assumed to be of type Δ_{qq}. If Δ_1 and Δ_3 are of distinct colours, Δ_2 must be of type Δ_{qtt}. If there is at least one other tetrahedron of type Δ_{qtt} incident with f, then there is a contribution to the left hand side of 4.1 of 2 and we have a deficit. Otherwise there are two distinct tetrahedra of type Δ_{tt} also incident with f. We obtain a contribution of 2 to the left hand side of (4.1) from these tetrahedra which needs to be shared with the two other H–even edges, g and h of the triangle shared by Δ_4 and Δ_5. We can distribute a contribution of $\frac{2}{3}$ from these tetrahedra to the (possibly not pairwise distinct edges) f, g and h. This yields a total contribution of $\frac{5}{3} > 1$ to f and hence a deficit.

Hence, assume that Δ_1 and Δ_3 are equally coloured and thus Δ_2 must be of type Δ_{qq} with matching colours. If any of Δ_4 or Δ_5 is of type Δ_{qtt}, the other one must be as well, we obtain a contribution of 2, and we have a deficit.

If one of Δ_4 and Δ_5 is of type Δ_{qq} the other one must be as well. As before, this implies that one can perform a compression of two of the surfaces S_i, $i \in \{1, 2, 3\}$, by cutting along the two tubes around the H–even edge f.

Again, there is an additional contribution to the left hand side of 8, distributed equally over the four H–even edges of Δ_2, Δ_4 and Δ_5 (see Figure 16). In particular, each of these edges has an associated deficit.

Hence, we can assume that both Δ_4 and Δ_5 are of type Δ_{tt}, both with the same colouring. We obtain a contribution of 2 to the left hand side of (4.1) which needs to be shared with the two other H–even edges, g and h of the triangle shared by Δ_4 and Δ_5. However, for this contribution to be insufficient globally, both g and h must be surrounded by some internally unbalanced maximal layered solid tori, with the tetrahedra surrounding g and h and distinct from Δ_4 and Δ_5 all being of type Δ_{qq} and of matching colours. This implies that one can perform two compressions of two of the surfaces S_i, $i \in \{1, 2, 3\}$, by cutting along the three times two incomplete tubes around the edges f, g and h. This process replaces the two parallel triangles inside Δ_4 and Δ_5 by pairs of parallel normal triangles of all three other types. Again, quadrilaterals in the tubes linking edges which are contained in a layered solid torus have their vertices identified in diagonally opposed pairs. Hence this operation does not produce any new connected components. See Figure 17 for an illustration of this process in the case of one surface (rather than two parallel ones).
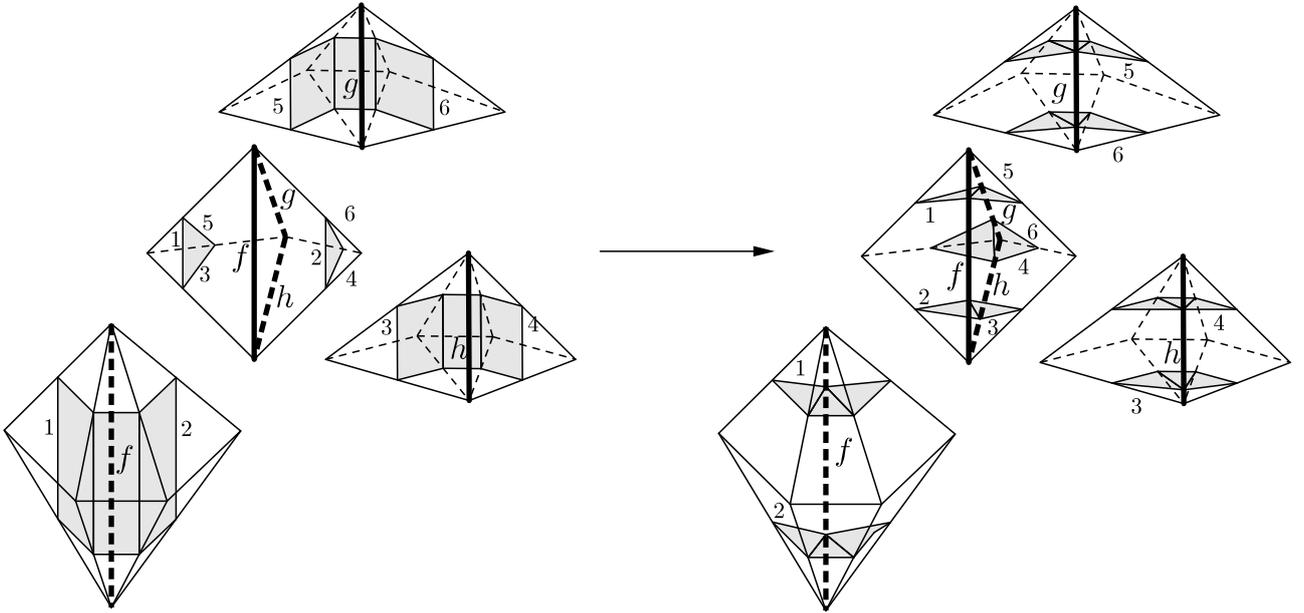
Figure 17: Double compression along a triangle of H–even edges. Only one of the two parallel copies of surfaces is drawn.

This double compression of two surfaces contributes a total of 16 to the left hand side of (4.1). We distribute two to each of f, g and h, as well as to the at most three other H–even edges in the neighbourhoods of f, g and h. The extra contribution of at least 4 is a surplus and not needed. We therefore again have a deficit.

This completes the case of gaining edges of type $(2, 5)$ showing that they all contribute a deficit.

The case $m = 1$ and $d = 4$: In this case, f does not contribute to the left hand side of (4.1) and we have to find a total contribution of 1 for a balance and > 1 for a deficit.

Again, Δ_1 can be assumed to be of type Δ_{qq}. If any of Δ_2, Δ_3 and Δ_4 is of type Δ_{qtt} then we again have two cases. Either at least two of these are of type Δ_{qtt} and we obtain a contribution of at least 2 to the left hand side of (4.1) and we are done, or one of them is of type Δ_{qtt} and two are of type Δ_{tt}. In this case we distribute the contribution of 2 to the tetrahedra of type Δ_{tt} evenly over their three H–even edges and obtain a total contribution > 1 for f. If all Δ_i, $1 \leq i \leq 4$ are of type Δ_{qq}, there exist compressions for two of the surfaces S_i, $i \in \{1, 2, 3\}$, with a total contribution of 8, which can be distributed equally over the four H–even edges contained in Δ_2, Δ_3 and Δ_4. This again gives a deficit for all H–even edges involved.

There are two cases remaining (up to symmetry). In the first case we have (without loss of generality) Δ_2 of type Δ_{qq} and Δ_3 and Δ_4 of type Δ_{tt}. Here, we fall back onto the double compression case along a triangle of H–even edges, as done in the case $m = 2$, $d = 5$ (see Figure 17). This again gives a deficit for all H–even edges involved.

In the second case we have Δ_2 and Δ_4 of type Δ_{tt} and Δ_3 of type $\Delta_\emptyset$. In this case, we obtain another minimal triangulation of M by performing a 4-4-move on f. This results in two tetrahedra of type Δ_{tt} and two of type Δ_{qq}. But this contradicts our hypothesis that $\mathscr{T}$ locally has the smallest number of tetrahedra of type $\Delta_\emptyset$. In particular, this case does not happen.

This completes the case of gaining edges, showing that they all contribute a deficit.

To sum up, it is not possible to achieve a gain, and hence we have proved the first part of the theorem, establishing the basic lower bound on the number of tetrahedra.

To prove the second part, we make some additional observations regarding the balancing edges.

If the balancing edge f received a contribution > 0 as a member of the star of a gaining edge, then it has an associated deficit. Hence assume that f has not received such a contribution. If f is incident with a tetrahedron of type Δ_{qtt} then there is a deficit. If f is incident with a tetrahedron of type Δ_{tt}, then this tetrahedron has not been counted yet since otherwise f would have received a contribution. We can therefore allocate $1/3$ to each edge incident with this tetrahedron and hence f has a deficit.

In particular, this shows that associated to the subcomplex that is the union of all internally unbalanced tori and all tetrahedra of type Δ_{qtt} or Δ_{tt} we have an overall deficit.

These observations regarding the balancing edges suffice to analyse the case of equality. Hence suppose $|\mathscr{T}| = \sum\limits_{0 \neq c \in H} \| c \|$ and so

$$n_{tt} + n_{qtt} + 2D + \sum_{d=5}^{\infty}(d-4)\mathfrak{e}_d = \mathfrak{e}_3.$$

The above analysis shows that this equality is only possible if there is no deficit since otherwise the right hand side is smaller than the left hand side. In particular, the subcomplex that is the union of all internally unbalanced tori and all tetrahedra of type Δ_{qtt} or Δ_{tt} must be empty. Whence $n_{tt} = n_{qtt} = 0$. From the possible face pairings between tetrahedra of the different types, this implies that either all tetrahedra in $\mathscr{T}$ are of type $\Delta_\emptyset$ (but then H has rank 0); or all tetrahedra are of type Δ_{qq} (but then H has rank 1); or all are of type Δ_{qqq}. Hence we have the last case and each canonical surface is a quadrilateral surface. Since all tetrahedra are of type Δ_{qqq}, there are no H–even edges. Whence $\mathfrak{e}_3 = 0$, which implies $D = 0$ and the three surfaces therefore are taut.

It remains to show that in *every* minimal triangulations of M the canonical normal representatives of non-trivial classes in H are taut and all tetrahedra are of type Δ_{qqq}. Recall that $\mathscr{T}$ was obtained from an arbitrary minimal triangulation $\mathscr{T}'$ by 4-4 moves.

We first note that every edge e of $\mathscr{T}$ is contained in at most $d(e) - 1$ distinct tetrahedra. Otherwise the unique canonical representative S disjoint from e can be altered by cutting out the tube around e and pasting in two discs orthogonal to e. Since S is taut, this operation cannot be a compression of S and thus it must produce two connected components – one of which must be a normal 2-sphere. This contradicts the fact that $\mathscr{T}$ is 0-efficient due to Theorem 4. But if no edge e of $\mathscr{T}$ is contained in $d(e)$ distinct tetrahedra, it is in particular true that $\mathscr{T}$ does not admit 4-4-moves. Hence $\mathscr{T}' = \mathscr{T}$. In particular, the minimality assumption on the number of tetrahedra of type $\Delta_\emptyset$ is true for all minimal triangulations of M.

We now show that if all tetrahedra are of type Δ_{qqq}, then the total number of tetrahedra in $\mathscr{T}$ is even. For this we use the fact that M is orientable and we assume that all tetrahedra are coherently oriented. In a tetrahedron of type Δ_{qqq}, every boundary triangle has one 1-even, one 2-even and one 3-even edge and the orientation on the triangle induced by the orientation of the tetrahedron determines a cyclic order of the labels 1, 2, and 3. This is the same cyclic order for all boundary triangles of the same oriented tetrahedron. Hence there are two types of oriented tetrahedra of type Δ_{qqq}. Tetrahedra meeting in a triangle must induce opposite cyclic orders on this triangle, see Figure 18. Hence, they are of different orientation types. It follows that the number of tetrahedra is even. $\qquad\square$
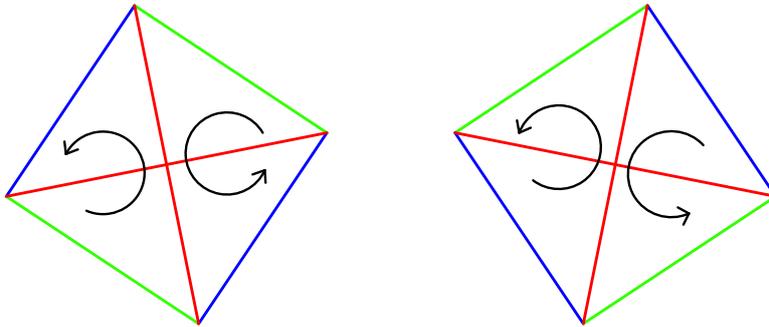
Figure 18: The two orientation types of a tetrahedron of type Δ_{qqq}.

5 The monodromy ideal triangulations of once-punctured torus bundles

In this section we apply Theorem 1 to a nice class of examples known as the (canonical) monodromy ideal triangulations of once-punctured torus bundles over the circle. In particular, we show that these are minimal. In the last subsection we describe an infinite family of triangulations of semi-bundles that we conjecture to be minimal.

5.1 Definition of the triangulation

Monodromy ideal triangulations were studied in Lackenby [25] and Gueritaud [12], and we refer the reader to these references for details.

The monodromy ideal triangulations all have the property that they are ideal triangulations for which all edges are of even degree and there is a single ideal vertex. As we are only interested in the case of hyperbolic 3-manifolds amongst once punctured torus bundles, we restrict to monodromy as a map $A \in SL(2,\mathbb{Z})$ having trace different from $0, \pm 1, \pm 2$.

Each tetrahedron in the monodromy ideal triangulation is *layered* on a once punctured triangulated torus with two ideal triangles. In particular, two sets of opposite edges of each tetrahedron are identified in pairs so that the top and bottom pairs of faces form such triangulated once punctured tori. Each tetrahedron induces a diagonal flip on the triangulation of the once punctured torus, and tracing the sequence of flips gives rise to a conjugate of $\pm A$. Indeed, each flip corresponds to one of the standard transvection matrices

$$R = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix} \quad \text{and} \quad L = \begin{pmatrix} 1 & 0 \\ 1 & 1 \end{pmatrix},$$

and hence one obtains a factorisation of a conjugate of $\pm A$ in terms of positive powers of L and R.

A key property of these triangulations is that the normal quadrilateral discs are of two types—called *horizontal* and *vertical*. A horizontal quadrilateral has vertices on the two pairs of opposite edges which are glued together. A vertical quadrilateral has two vertices on such a pair of edges and two vertices on the two edges which are not glued together.

A crucial fact about normal surfaces in these triangulations is that the Euler characteristic is exactly the negative of the total number of horizontal quadrilaterals, see Lemma 15. In particular no normal triangle and no vertical quadrilateral contributes to the Euler characteristic.

In Sections 5.2 and 5.3 we prove the following result, first stated in the introduction.

Theorem 2 Monodromy ideal triangulations of once-punctured torus bundles are minimal.

5.2 Reduction using coverings

Monodromy ideal triangulations of once punctured torus bundles M have the nice property that any cyclic covering induced by the bundle structure is again a once punctured torus bundle and the triangulation lifts to the monodromy ideal triangulation. Throughout this section, we assume that the monodromy has trace different from $0, \pm 1, \pm 2$. Hence there is a unique $\mathbb{Z}$–summand of $H_1(M, \mathbb{Z})$. The cyclic coverings we are using in this section come from the kernels of maps $\pi_1(M) \to H_1(M, \mathbb{Z}) \to \mathbb{Z} \to \mathbb{Z}_k$.

We claim that either $H_1(M, \mathbb{Z}_2)$ has rank 3 or there is a 2– or 3–fold cyclic covering $\widetilde{M}$ of M with the property that $H_1(\widetilde{M}, \mathbb{Z}_2)$ has rank 3. The natural presentation of the fundamental group of M is of the form

$$\langle\, t, a, b \mid t^{-1}at = \alpha(a), t^{-1}bt = \alpha(b) \,\rangle,$$

where α is the automorphism of the free group $\langle a, b \rangle$ induced by the monodromy A. It follows that the map induced by A on $\mathbb{Z}_2 \oplus \mathbb{Z}_2$ is precisely its image $\overline{A} \in SL(2, 2) \cong \mathrm{Sym}(3)$. Whence the rank of $H_1(M, \mathbb{Z}_2)$ is 3 if $\overline{A}$ is the identity. Let $k \in \{1, 2, 3\}$ be the order of $\overline{A}$. Since the monodromy of the k–fold cyclic covering of M corresponding to the kernel of the map $\pi_1(M) \to \mathbb{Z} \to \mathbb{Z}_k$ gives rise to the map $\overline{A}^k \in SL(2, 2)$, it follows that we obtain rank 3 after passing to a k–fold cyclic covering.

It now suffices to show that a monodromy ideal triangulation is minimal for a once punctured torus bundle M with $H_1(M, \mathbb{Z}_2)$ of rank 3. For if the rank is not three, we can pass to a 2– or 3–fold cyclic covering $\widetilde{M}$ with $H_1(\widetilde{M}, \mathbb{Z}_2)$ of rank 3 as in the previous paragraph. If we can prove that the lifted monodromy ideal triangulation of $\widetilde{M}$ is minimal then the monodromy ideal triangulation of M must also be minimal.

5.3 Closed one-sided incompressible surfaces

In what follows, the *weight* of a normal surface S in a triangulation $\mathcal{T}$ is defined to be the overall number of intersection points of S with the edges of $\mathcal{T}$.

Lemma 15 *Let M be a once-punctured hyperbolic torus bundle and let $\mathcal{T}$ be a monodromy ideal triangulation of M. Furthermore, let $S \subset \mathcal{T}$ be a normal surface of $\mathcal{T}$ which is a taut least weight representative of an even torsion class of $H_2(M, \mathbb{Z})$. Then S meets each tetrahedron of $\mathcal{T}$ in a single quadrilateral. Moreover, $\chi(S) = -(\# \text{ horizontal quads})$.*

Proof Let $S \subset \mathcal{T}$ be a normal surface of $\mathcal{T}$ which is a taut representative of a fixed even torsion class $c \in H_2(M, \mathbb{Z})$. We assume that S is of least weight amongst all taut representatives. The monodromy ideal triangulation $\mathcal{T}$ naturally endows M with a circular Morse type function: The layering defines a cyclic ordering of pairs of faces forming a once-punctured torus fibre. Let T be such a two-triangle fibre.

Consider the intersection $S \cap T$. This is a family of embedded disjoint normal simple closed curves in the two-triangle structure on T. A well-known elementary fact is that any such curve is either non separating or vertex linking. Next, note that $S \cap T$ cannot contain two parallel curves, since otherwise we can perform an annular compression using an innermost fibre 0 weight annulus A and find a new homologous normal surface of less weight or less genus or both.

To be more specific, consider $N(A)$, where $N(A)$ is a small regular neighbourhood of A. Then $S \cap N(A)$ contains two annuli and these are replaced by two annuli in $\partial N(A)$ to form a new surface with the same weight and Euler characteristic as S and representing the same class c. If this surface is disconnected, it is easy to see one component is one-sided and the other is two-sided, as c is a torsion class. Hence we can discard the two-sided component (it must bound in M) and the resulting surface still represents c. As this new surface is not normal, it can be isotoped to a normal surface of less weight, giving a contradiction.

Hence, from now on, we can assume that each fibre meets S in one or two curves. Note that in the latter case exactly one curve links the ideal vertex in the fibre and the other curve is non separating. Let γ denote the non separating curve in $S \cap T$.

The curve $\gamma \subset T$ is determined by how often it intersects the three edges of T. Since γ is connected, these three intersection numbers are pairwise coprime. This can be seen by the fact that, given two intersection numbers, the third one must either be the sum or the non-negative difference of the first two.

A key observation is that S propagates across the tetrahedron $\Delta \in \mathcal{T}$ layered on top of T in a very well-behaved way. In what follows, we denote the edge of Δ opposite T by e, and the new fibre opposite T (i.e., the two triangles of Δ containing e) by T'.

Assume that $S \cap T$ contains a vertex linking curve α. α intersects T in six normal arcs (one of each type). Four of them necessarily yield two normal triangles in Δ and two normal arcs in T'. For the remaining two arcs, three cases need to be considered.

1) They are in the boundary of a quadrilateral normal disc in Δ. In this case, the two normal triangles and the normal quadrilateral containing α intersect both triangles of T' in two parallel normal arcs. Moreover, all other normal arcs of $S \cap T'$ must be of the same type (otherwise S self-intersects inside Δ). In particular, $S \cap T'$ has parallel curves, which contradicts our assumption that S is least weight, using the annular compression argument from above. See Figure 19 on the left.

2) One of them yields a normal triangle t, and the other one is in the boundary of a normal quadrilateral in Δ together with a normal arc of γ. In this case, $S \cap T'$, again consists of one type of normal arc per triangle of T', with two exceptional normal arcs coming from t. Let one triangle of T' contain k and the other one contain ℓ copies of parallel normal arcs. Hence, one edge of T' intersects S once, the second one $k + 1 = \ell$ times and the third k times, hence, $\ell + 1 = k$, a contradiction. See Figure 19 on the right.

3) All of the six normal arcs of α are in the boundaries of normal triangles in Δ. Since 1) and 2) are not possible, this must be the case for all tetrahedra of $\mathcal{T}$ and S contains a boundary parallel torus. This contradicts the assumption that S is a least weight taut representative of c.

Hence, for the rest of the proof we can assume that $S \cap T$ is connected for each fibre T. Consequently, it follows that $S \cap T$ has only one or two normal arc types in each triangle of T, since otherwise $S \cap T$ contains a vertex linking component.

We claim that Δ can only contain vertical normal discs of S, unless $\Delta \cap S$ consists of a single horizontal quad and S intersects two edges of T exactly once, and is disjoint from e.

More precisely, there are two cases we have to consider:
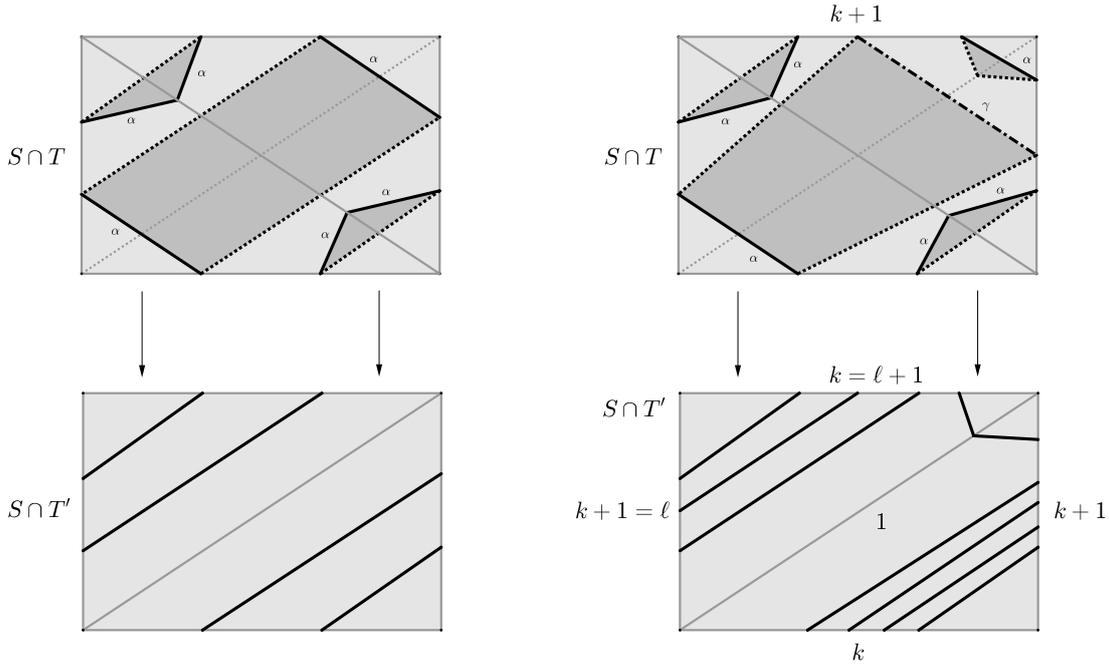
Figure 19: Left: The normal arcs of α are in two normal triangles and one normal quadrilateral. Right: The normal arcs of α are in three normal triangles and one normal quadrilateral.

(1) The intersection of S with T is disjoint from one of the edges of T. In this case the other two edges must intersect S exactly once. If we layer on the edge disjoint from S then there is a single horizontal quadrilateral in the tetrahedron, if we layer on one of the other edges, there is a single vertical quadrilateral. See Figure 20 on the right.

(2) The surface S intersects T in two normal arc types per triangle. In this case all normal discs of S in Δ must necessarily be vertical. To see this note that a horizontal normal quadrilateral can only have two normal arcs of $S \cap T$ in its boundary. Since by the assumption, at least one other normal disc must exist in Δ, this must be a vertical normal triangle. Arguing as in case 1) and 2) above we can then follow that all three normal arc types must be present per triangle, a contradiction. See Figure 20 on the left.

It is now easy to verify that the Euler characteristic of S is the negative of the number of horizontal quadrilateral discs, using the circular Morse type function induced by the layering. For each horizontal quadrilateral disc induces a simple saddle type singularity of Morse index one, whereas there are no critical points on vertical triangles or quadrilaterals.
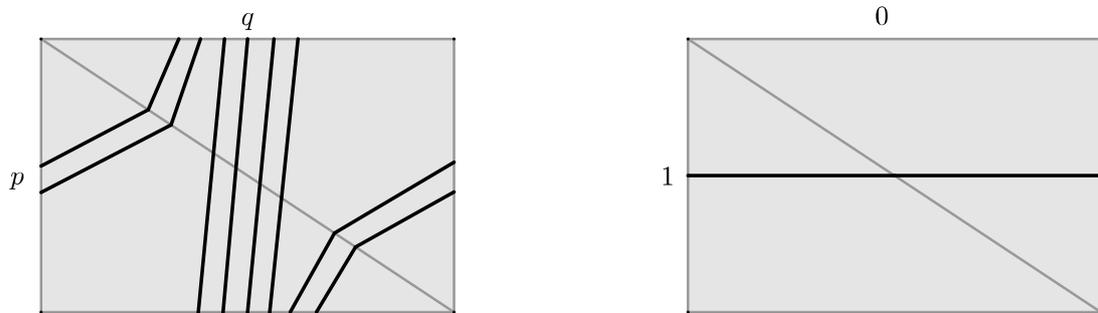
Figure 20: Left: $S \cap T$ has two normal arc types per triangle of T. Right: $S \cap T$ is disjoint from one edge of T.

Since S is taut, no component can be a torus or Klein bottle. Hence there must be at least one horizontal

quadrilateral. We can cut M open along a corresponding fibre T to form a product of the form $T \times I$ where T is a once punctured torus. This product admits an induced ideal triangulation which is minimal layered.

There are two steps to complete the proof. The first is to follow the argument in [30, Theorem 12] which shows that there is a unique isotopy class of incompressible surfaces in a lens space with even fundamental group. We want to apply a similar argument to an incompressible surface S^* in a product of a once punctured torus and an interval.

Consider a product $X = T \times I$ where T is a once punctured torus and there are two essential boundary curves C, C' in $T \times \{0\}, T \times \{1\}$ respectively and the intersection number of C, C' is even when projected to a copy of T. We claim there is a unique incompressible surface S^* up to isotopy with $\partial S^* = C \cup C'$.

Pick an annulus A in the product structure of X with boundary curve on $T \times \{1\}$ parallel to C'. We can assume without loss of generality that A is transverse to S^* and the intersection $A \cap S^*$ contains only arcs which are essential on S^* and have both ends on $T \times \{0\}$. An innermost such an arc on A cuts off a bigon which can be used to perform a boundary compression of S^* across $T \times \{0\}$. Exactly as in [30], this boundary compression must change the boundary curve C of ∂S^* to a new curve which has smaller intersection number with C'. It is now easy to show as in [30] that this process produces a unique family of bands attached to a collar about C forming S^* which is therefore unique up to isotopy.

The second step is to observe that the canonical quadrilateral surface in the induced ideal layered triangulation of the product is incompressible. The proof again follows [30]. The key idea is to use induction and the recursion formula for the genus of the incompressible surface. By removing the tetrahedra up to the first containing a horizontal quadrilateral, we obtain a new product with an induced minimal layered triangulation. The recursion formula for the change in boundary slope and the genus matches the formula in [30] and this shows that the canonical quadrilateral surface must be incompressible as it is a minimal genus cobordism surface connecting C and C'. This completes the proof. $\qquad\square$

Corollary 16 *Let M be a once-punctured hyperbolic torus bundle and let $\mathscr{T}$ be a monodromy ideal triangulation of M. Suppose that $H_2(M, \mathbb{Z})$ contains a Klein 4-group H. Then*

$$|\mathscr{T}| = \sum_{0 \neq c \in H} \| c \|,$$

and in particular $\mathscr{T}$ is a minimal triangulation.

Proof For each of the three non-trivial classes in H we have least weight taut normal representatives. We know from the first part of Lemma 15 that each of them meets each tetrahedron of $\mathscr{T}$ in exactly one normal quadrilateral. Moreover, since H is a Klein 4-group, the sum of all three representatives must have edge weight zero mod two. It follows that every tetrahedron contains all three quadrilateral types exactly once. In particular, in each tetrahedron exactly one of the three representatives has a single horizontal normal quadrilateral.

Hence, by the second part of Lemma 15 the negative of the sum of Euler characteristics of these three representatives must equal the number of tetrahedra of $\mathscr{T}$. It now follows from Theorem 1 that $\mathscr{T}$ is minimal. $\qquad\square$

Combining Corollary 16 with the covering argument provided in §5.2 proves Theorem 2.

5.4 Further examples

The orientable cusped hyperbolic census up to nine tetrahedra, described and verified in [5] and partially shipped with Regina [6], contains 162 182 minimal triangulations of 61 911 cusped hyperbolic 3-manifolds. There are exactly 26 triangulations for which the lower bound in Theorem 1 is attained. Of these, 22 are triangulations of once-punctured torus bundles, but the remaining four are not. Their isomorphism signatures are

- `gLLMQbeefffehhqxhqq` of manifold `s781`,

- `iLLLQPcbefgffhhhxxhaqxxqh` of manifold `t05624`,

- `iLLLQPcbefgffhhhhhqaxhhxq` of manifold `t06056`,

- `iLLwQPcbeefgehhhhhqhhqhqx` of manifold `t12546`.

These examples, as well as infinite families of minimal triangulations they are contained in, will be described in detail in a separate note.

References

[1] Ian Agol. Small 3-manifolds of large genus. *Geom. Dedicata*, 102:53–64, 2003.

[2] Hirotaka Akiyoshi. On the Ford domains of once-punctured torus groups. *Sūrikaisekikenkyūsho Kōkyūroku*, (1104):109–121, 1999. Hyperbolic spaces and related topics (Japanese) (Kyoto, 1998).

[3] Hirotaka Akiyoshi, Makoto Sakuma, Masaaki Wada, and Yasushi Yamashita. *Punctured torus groups and 2-bridge knot groups. I*, volume 1909 of *Lecture Notes in Mathematics*. Springer, Berlin, 2007.

[4] Michelle Bucher, Roberto Frigerio, and Cristina Pagliantini. The simplicial volume of 3-manifolds with boundary. *J. Topol.*, 8(2):457–475, 2015.

[5] Benjamin A. Burton. The cusped hyperbolic census is complete. `https://arxiv.org/abs/1405.2695/`, 2014. To appear in Trans. Amer. Math. Soc.

[6] Benjamin A. Burton, Ryan Budney, William Pettersson, et al. Regina: Software for low-dimensional topology. `https://regina-normal.github.io/`, 1999–2017.

[7] David B. A. Epstein and Robert C. Penner. Euclidean decompositions of noncompact hyperbolic manifolds. *J. Differential Geom.*, 27(1):67–80, 1988.

[8] William Floyd and Allen Hatcher. Incompressible surfaces in punctured-torus bundles. *Topology Appl.*, 13(3):263–282, 1982.

[9] Evgeny Fominykh, Stavros Garoufalidis, Matthias Goerner, Vladimir Tarkaev, and Andrei Vesnin. A census of tetrahedral hyperbolic manifolds. *Exp. Math.*, 25(4):466–481, 2016.

[10] Roberto Frigerio, Bruno Martelli, and Carlo Petronio. Dehn filling of cusped hyperbolic 3-manifolds with geodesic boundary. *J. Differential Geom.*, 64(3):425–455, 2003.

[11] François Guéritaud. Géométrie hyperbolique effective et triangulations idéales canoniques en dimension trois. `https://tel.archives-ouvertes.fr/tel-00119465/document`. PhD thesis, 154 pages, Mathématiques [math]. Université Paris Sud - Paris XI, 2006. Français.

[12] François Guéritaud. On canonical triangulations of once-punctured torus bundles and two-bridge link complements. *Geom. Topol.*, 10:1239–1284, 2006. With an appendix by David Futer.

[13] François Guéritaud and Saul Schleimer. Canonical triangulations of Dehn fillings. *Geom. Topol.*, 14(1):193–242, 2010.

[14] Craig D. Hodgson, J. Hyam Rubinstein, Henry Segerman, and Stephan Tillmann. Triangulations of 3-manifolds with essential edges. *Ann. Fac. Sci. Toulouse Math. (6)*, 24(5):1103–1145, 2015.

[15] Masaharu Ishikawa and Keisuke Nemoto. Construction of spines of two-bridge link complements and upper bounds of their Matveev complexities. *Hiroshima Math. J.*, 46(2):149–162, 2016.

[16] William Jaco, Jesse Johnson, Jonathan Spreer, and Stephan Tillmann. Bounds for the genus of a normal surface. *Geom. Topol.*, 20(3):1625–1671, 2016.

[17] William Jaco and J. Hyam Rubinstein. 0-efficient triangulations of 3-manifolds. *J. Differential Geom.*, 65(1):61–168, 2003.

[18] William Jaco and J. Hyam Rubinstein. Layered-triangulations of 3-manifolds. *ArXiv Mathematics e-prints*, March 2006.

[19] William Jaco and J. Hyam Rubinstein. Annular-Efficient Triangulations of 3-manifolds. *ArXiv e-prints*, August 2011.

[20] William Jaco, J. Hyam Rubinstein, Jonathan Spreer, and Stephan Tillmann. Z2-Thurston Norm and Complexity of 3-Manifolds, II. *ArXiv e-prints*, November 2017.

[21] William Jaco, J. Hyam Rubinstein, and Stephan Tillmann. Minimal triangulations for an infinite family of lens spaces. *J. Topol.*, 2(1):157–180, 2009.

[22] William Jaco, J. Hyam Rubinstein, and Stephan Tillmann. $\mathbb{Z}_2$-Thurston norm and complexity of 3-manifolds. *Math. Ann.*, 356(1):1–22, 2013.

[23] Ensil Kang and J. Hyam Rubinstein. Ideal triangulations of 3-manifolds. I. Spun normal surface theory. In *Proceedings of the Casson Fest*, volume 7 of *Geom. Topol. Monogr.*, pages 235–265. Geom. Topol. Publ., Coventry, 2004.

[24] Marc Lackenby. Taut ideal triangulations of 3-manifolds. *Geom. Topol.*, 4:369–395, 2000.

[25] Marc Lackenby. The canonical decomposition of once-punctured torus bundles. *Comment. Math. Helv.*, 78(2):363–384, 2003.

[26] Feng Luo and Stephan Tillmann. A new combinatorial class of 3-manifold triangulations. *Asian J. Math.*, 21(3):543–569, 2017.

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

[28] Kei Nakamura. The complexity of prime 3-manifolds and the first $\mathbb{Z}_{/2\mathbb{Z}}$-cohomology of small rank. *ArXiv e-prints*, December 2017.

[29] Carlo Petronio and Andrei Vesnin. Two-sided bounds for the complexity of cyclic branched coverings of two-bridge links. *Osaka J. Math.*, 46(4):1077–1095, 2009.

[30] J. Hyam Rubinstein. One-sided Heegaard splittings of 3-manifolds. *Pacific J. Math.*, 76(1):185–200, 1978.

[31] Naoki Sakata. Canonical decompositions of hyperbolic fibered two-bridge link complements. *Topology Appl.*, 196(part B):821–845, 2015.

[32] Makoto Sakuma and Jeffrey Weeks. Examples of canonical decompositions of hyperbolic link complements. *Japan. J. Math. (N.S.)*, 21(2):393–439, 1995.

[33] Herbert Seifert and William Threlfall. *Seifert and Threlfall: a textbook of topology*, volume 89 of *Pure and Applied Mathematics*. Academic Press, Inc., New York-London, 1980. Translated from the German edition of 1934 by Michael A. Goldman,.

[34] William P. Thurston. The geometry and topology of 3-manifolds. Lecture notes, Princeton University, 1978.

[35] William P. Thurston. A norm for the homology of 3-manifolds. *Mem. Amer. Math. Soc.*, 59(339):i–vi and 99–130, 1986.

[36] Stephan Tillmann. Normal surfaces in topologically finite 3-manifolds. *Enseign. Math. (2)*, 54(3-4):329–380, 2008.

William Jaco
Department of Mathematics, Oklahoma State University, Stillwater, OK 74078-1058, USA
jaco@math.okstate.edu
——

J. Hyam Rubinstein
School of Mathematics and Statistics, The University of Melbourne, VIC 3010, Australia
joachim@unimelb.edu.au
——

Jonathan Spreer
School of Mathematics and Statistics F07, The University of Sydney, NSW 2006 Australia
jonathan.spreer@sydney.edu.au
——

Stephan Tillmann
School of Mathematics and Statistics F07, The University of Sydney, NSW 2006 Australia
stephan.tillmann@sydney.edu.au