

ANNALES DE LA FACULTÉ DES SCIENCES DE TOULOUSE Mathématiques

CRAIG D. HODGSON, J. HYAM RUBINSTEIN,
HENRY SEGERMAN AND STEPHAN TILLMANN
Triangulations of 3–Manifolds with essential edges

Tome XXIV, n° 5 (2015), p. 1103-1145.

http://afst.cedram.org/item?id=AFST_2015_6_24_5_1103_0

© Université Paul Sabatier, Toulouse, 2015, tous droits réservés.

L'accès aux articles de la revue « Annales de la faculté des sciences de Toulouse Mathématiques » (<http://afst.cedram.org/>), implique l'accord avec les conditions générales d'utilisation (<http://afst.cedram.org/legal/>). Toute reproduction en tout ou partie de cet article sous quelque forme que ce soit pour tout usage autre que l'utilisation à fin strictement personnelle du copiste est constitutive d'une infraction pénale. Toute copie ou impression de ce fichier doit contenir la présente mention de copyright.

cedram

Article mis en ligne dans le cadre du
Centre de diffusion des revues académiques de mathématiques
<http://www.cedram.org/>

Triangulations of 3–Manifolds with essential edges

CRAIG D. HODGSON⁽¹⁾, J. HYAM RUBINSTEIN⁽²⁾,
HENRY SEGERMAN⁽³⁾ AND STEPHAN TILLMANN⁽⁴⁾

*Dedicated to Michel Boileau for his many contributions and leadership
in low dimensional topology.*

RÉSUMÉ. — Pour une variété de dimension 3 nous définissons les notions de triangulation essentielle et fortement essentielle. Sous certaines hypothèses sur la topologie ou la géométrie de la variété, nous donnons quatre constructions utilisant différentes méthodes (décompositions de Heegaard, hiérarchies des variétés de Haken, décompositions d’Epstein-Penner et cut-loci des variétés riemanniennes) pour obtenir des triangulations ayant ces propriétés.

Nous montrons aussi qu’une structure de semi-angle est une condition suffisante pour qu’une triangulation d’une 3-variété soit essentielle, et qu’une structure d’angle stricte est une condition suffisante pour qu’une triangulation soit fortement essentielle. De plus, des algorithmes pour tester si une triangulation d’une 3-variété est essentielle ou fortement essentielle sont donnés.

⁽¹⁾ Department of Mathematics and Statistics, The University of Melbourne, VIC 3010, Australia

craigdh@unimelb.edu.au

⁽²⁾ Department of Mathematics and Statistics, The University of Melbourne, VIC 3010, Australia

rubin@ms.unimelb.edu.au

⁽³⁾ Department of Mathematics, Oklahoma State University, Stillwater, Oklahoma, OK 74078, USA

segerman@math.okstate.edu

⁽⁴⁾ School of Mathematics and Statistics, The University of Sydney, NSW 2006, Australia

tillmann@maths.usyd.edu.au

This research is supported by the Australian Research Council under the Discovery Projects funding scheme (project DP130103694).

ABSTRACT. — We define essential and strongly essential triangulations of 3-manifolds, and give four constructions using different tools (Heegaard splittings, hierarchies of Haken 3-manifolds, Epstein-Penner decompositions, and cut loci of Riemannian manifolds) to obtain triangulations with these properties under various hypotheses on the topology or geometry of the manifold.

We also show that a semi-angle structure is a sufficient condition for a triangulation of a 3-manifold to be essential, and a strict angle structure is a sufficient condition for a triangulation to be strongly essential. Moreover, algorithms to test whether a triangulation of a 3-manifold is essential or strongly essential are given.

1. Introduction

Finding combinatorial and topological properties of geometric triangulations of hyperbolic 3-manifolds gives information which is useful for the reverse process—for example, solving Thurston’s gluing equations to construct an explicit hyperbolic metric from an ideal triangulation. Moreover, suitable properties of a triangulation can be used to deduce facts about the variety of representations of the fundamental group of the underlying 3-manifold into $\mathrm{PSL}(2, \mathbb{C})$. See for example [32].

In this paper, we focus on one-vertex triangulations of closed manifolds and ideal triangulations of the interiors of compact manifolds with boundary. Two important properties of these triangulations are introduced. For one-vertex triangulations in the closed case, the first property is that no edge loop is null-homotopic and the second is that no two edge loops are homotopic, keeping the vertex fixed. We call a triangulation with the first property *essential* and with both properties *strongly essential*. Similarly, for ideal triangulations of the interiors of compact manifolds, we say that the triangulation is *essential* if no ideal edge is homotopic into a vertex neighbourhood keeping its ends in the vertex neighbourhood, and is *strongly essential* if no two ideal edges are homotopic, keeping their ends in the respective vertex neighbourhoods. Precise definitions are given in §3.

Four different constructions are given of essential and strongly essential triangulations, each using a different condition on the underlying manifold. Each construction is illustrated with key examples.

The first construction applies to closed P^2 -irreducible 3-manifolds different from $\mathbb{R}P^3$ and with non-zero first homology with \mathbb{Z}_2 coefficients. For these manifolds, an essential triangulation is constructed using one-sided Heegaard splittings. A similar construction, using Heegaard splittings with

compression bodies each containing a torus boundary component, gives essential ideal triangulations of the interiors of compact orientable irreducible atoroidal 3-manifolds with boundary consisting of two incompressible tori.

The second construction gives strongly essential triangulations of closed Haken 3-manifolds using hierarchies. A variation of this construction produces strongly essential ideal triangulations of the interiors of compact irreducible atoroidal 3-manifolds with boundary consisting of incompressible tori.

The third construction shows how strongly essential triangulations of non-compact complete hyperbolic 3-manifolds can be obtained by subdividing the polyhedral decompositions of Epstein and Penner [7].

In the fourth construction, a one-vertex strongly essential triangulation of a closed 3-manifold with a Riemannian metric of constant negative or zero curvature is constructed from the dual of the cut locus of an arbitrary point. The same method applies to some classes of elliptic closed 3-manifolds, i.e. those with constant positive curvature. However in the latter case, there are restrictions on the fundamental groups.

We expect that these constructions will extend to higher dimensional manifolds, but in this paper we focus on dimension three.

We also will show that an ideal triangulation of a 3-manifold admitting a strict angle structure must be strongly essential and if a triangulation has a semi-angle structure, then it must be essential. The class of semi-angle structures contains the class of taut structures. Moreover, examples of taut triangulations which are not strongly essential are given, as well as an example of a taut triangulation which is strongly essential, but does not support a strict angle structure.

In the last section, we show that there are algorithms to decide whether a given triangulation of a closed 3-manifold is essential or strongly essential, based on the fact that the word problem is uniformly solvable in the class of fundamental groups of compact, connected 3-manifolds. Similarly, building on work of Friedl and Wilton [9], we show that there is an algorithm to test whether an ideal triangulation is essential based on the fact that the subgroup membership problem is solvable in the same generality. However, to test whether an ideal triangulation is strongly essential, we require the double coset membership problem to be solvable for pairs of peripheral subgroups. Building on work of Aschenbrenner, Friedl and Wilton [2], we restrict this algorithm to the following three classes: topologically finite

hyperbolic 3-manifolds, compact Seifert fibred spaces, and compact, irreducible, orientable, connected 3-manifolds with incompressible boundary consisting of a disjoint union of tori.

We would like to thank the referee for many detailed suggestions to improve the exposition.

2. Preliminaries

The purpose of this section is to establish the basic definitions and to recall basic properties of 3-manifolds. Our set-up is somewhat rigid in order to describe objects, such as the compact core, in an algorithmically constructible fashion.

2.1. Triangulations

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

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

Consider the quotient space $\widehat{M} = \tilde{\Delta}/\Phi$ with the quotient topology, and denote the quotient map $p: \tilde{\Delta} \rightarrow \widehat{M}$. We will make the additional assumption on Φ that for every k -simplex τ in $\tilde{\Delta}$ 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-Threlfall [33].) We will also always assume that \widehat{M} is connected. The triple $\mathcal{T} = (\tilde{\Delta}, \Phi, p)$ is a (*singular*) *triangulation* of \widehat{M} , but for brevity we will often simply say that \widehat{M} is given with the structure of a triangulation.

Let $\widehat{M}^{(k)} = \{p(\tau^k) \mid \tau^k \subseteq \tilde{\Delta}\}$ denote the set of images of the k -simplices of $\tilde{\Delta}$ under the projection map. Then the elements of $\widehat{M}^{(k)}$ are precisely the equivalence classes of k -simplices in $\tilde{\Delta}$. 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 $\gamma_e: [0, 1] \rightarrow \widehat{M}$ with endpoints in $\widehat{M}^{(0)}$ via the composition of maps $[0, 1] \rightarrow \widetilde{\Delta} \rightarrow \widehat{M}$, 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 except that we often need to take care of the fact that we have chosen an arbitrary orientation of each edge.

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 *compact core* M^c of M is the result of taking the second barycentric subdivision of \widehat{M} and deleting the open star of each vertex. The components of $\widehat{M} \setminus M^c$ are referred to as the *ends* of M . The dual 2–skeleton X in \widehat{M} is a *spine* for M^c (see §2.7). Hence $\pi_1(X) = \pi_1(M^c) = \pi_1(M)$, whilst the fundamental group of \widehat{M} is a quotient thereof.

2.4. Hats off

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

2.5. Surfaces in 3–manifolds

Throughout, 3–manifolds will be compact, connected and we will work in the PL category. A 3–manifold N is called *irreducible* if every embedded 2–sphere bounds a 3–ball; it is P^2 –irreducible if N is irreducible and

does not contain any embedded 2-sided projective planes. An embedded 2-sided surface S in N is called *incompressible* if S is not simply connected and the inclusion map induces an injection $\pi_1(S) \rightarrow \pi_1(N)$. If S is properly embedded in N then S is called *∂ -incompressible* if the induced map $\pi_1(S, \partial S) \rightarrow \pi_1(N, \partial N)$ is one-to-one. A closed 3-manifold is called *Haken* if it is P^2 -irreducible and contains an incompressible surface. If N has non-empty boundary, then it is Haken if it is P^2 -irreducible and each component of ∂N is incompressible. Moreover, N is *atoroidal* if any incompressible torus or Klein bottle is isotopic to a component of ∂N ; N is *anannular* if it has non-empty boundary and contains no incompressible and ∂ -incompressible annuli or Möbius bands.

2.6. Heegaard splittings

A *handlebody* is a regular neighbourhood of an embedded graph in a 3-manifold N and a *compression body* is a regular neighbourhood of the connected union of a finite collection of graphs and some boundary components of N .

A *Heegaard surface* in a compact 3-manifold N is a 2-sided, properly embedded, connected, closed surface S in N , which splits N into two handlebodies or compression bodies H_1, H_2 . The decomposition $N = H_1 \cup_S H_2$ is called a *Heegaard splitting* of N . For emphasis, we will sometimes say that this is a *2-sided* Heegaard splitting. Note that if N is non-orientable, so are H_1, H_2 .

A *1-sided Heegaard surface* is a 1-sided, properly embedded, connected, closed surface S in N such that the complement of an open regular neighbourhood $U(S) \subset N$ is a handlebody or compression body H . The decomposition $N = U(S) \cup_{\partial U(S)} H$ is called a *1-sided Heegaard splitting* of N . In this case, N can be viewed as the result of identifying points of one of the boundary surfaces of a compression body or handlebody under a free involution. Note that if N is non-orientable, then so are S and H .

A (1- or 2-sided) Heegaard surface S in N is *irreducible* if there is no separating embedded 2-sphere in N , which intersects S in a single loop that is essential on S . We will also say that the associated Heegaard splitting is irreducible. If N is irreducible and there is such a 2-sphere Σ meeting S in a single essential loop, then Σ bounds a 3-ball B in N , which meets S in a non-simply connected region. By [37], the Heegaard splitting of B given by $B \cap S$ is standard, consisting of trivial handles. Equivalently, $B \cap S$ is isotopic to the boundary of a regular neighbourhood of the union of ∂B and an unknotted collection of properly embedded arcs in B .

2.7. Spines

To construct triangulations, we will make use of *spines* of 3-manifolds. We follow the terminology of the monograph by Matveev [24]. A spine of the compact manifold N with non-empty boundary is a compact subpolyhedron P of N with the property that N collapses to P . Typically $\dim P = \dim N - 1$. We will usually arrange for $P \subset \text{Int}(N)$, so that P is a spine if and only if $N \setminus P$ is homeomorphic with $\partial N \times [0, 1)$ (see [24], Theorem 1.1.7). A spine of a closed 3-manifold N is a spine of the complement of an open ball in N . A spine of a topologically finite non-compact 3-manifold N is a spine of a compact core of N .

If $M = \widehat{M} \setminus \widehat{M}^{(0)}$ is given with the structure of an ideal triangulation, there is a natural spine dual to the triangulation obtained by inserting a so-called *butterfly* in each tetrahedron. The butterfly is the intersection of the dual 2-skeleton with the tetrahedron (see Figure 1).




Figure 1. — A butterfly in a tetrahedron.

A spine is *simple* if every point on the spine has a neighbourhood homeomorphic to a neighbourhood of a point in the interior of the butterfly. The points on a simple spine fall into three categories: (*true*) *vertices* have arbitrarily small neighbourhoods homeomorphic to the butterfly, *non-singular points* have a neighbourhood homeomorphic to a disc, and all other points are *triple points* (see Figure 2). A connected component of the set of all non-singular points is called a *2-stratum*; a connected component of the set of all triple points is called a *1-stratum*. A simple spine is *special* if each *2-stratum* is an open 2-cell and each *1-stratum* is an open 1-cell.




Figure 2. — From left to right: a non-singular point, a triple point, and a true vertex.

In our existence proofs for triangulations with certain properties, we will often first construct a special spine and then a triangulation dual to the spine. A general theory for constructing triangulations dual to *thickenable special polyhedra* is given in [24], §1.1.5. In our setting, we start with a simple spine of a manifold, so know from the outset that we have a thickenable polyhedron, and only need to show that it is special in order to justify the existence of a dual triangulation. As pointed out in [24] (Remark 1.1.11), given a simple polyhedron with at least one true vertex, the condition on the 1-strata for a special spine follows from the condition on the 2-strata.

3. Essential and parallel edges

The purpose of this section is to establish the definitions of essential and strongly essential triangulations of 3-manifolds. We are interested in two distinct cases in this paper: 1-vertex triangulations of a closed pseudo-manifold \widehat{M} (which will generally be a closed manifold), and ideal triangulations of a topologically finite manifold M with any number of ideal vertices. This section continues in the notation of §2, and we will always assume that \widehat{M} is triangulated and $M = \widehat{M} \setminus \widehat{M}^{(0)}$ is ideally triangulated.

3.1. Closed pseudo-manifolds

The definitions for a closed pseudo-manifold \widehat{M} are straightforward:

DEFINITION 3.1 (Essential and parallel edges). — *An edge loop in \widehat{M} is essential if it is not null-homotopic. The edge paths γ, δ in \widehat{M} are parallel if there is a path homotopy between γ and one of $\pm\delta$.*

DEFINITION 3.2 (Essential triangulation). — *The triangulation of \widehat{M} is essential if it has exactly one vertex and every edge loop in \widehat{M} is essential. It is strongly essential if, in addition, no two edge paths are parallel.*

In the case of an ideal triangulation, we will restrict the allowable path homotopies. For instance, Thurston’s ideal triangulation of the figure eight knot complement M yields a simply connected pseudo-manifold \widehat{M} since the fundamental group is normally generated by meridians, and hence the triangulation of \widehat{M} is not essential (and hence not strongly essential). However, the *geometry* of the figure eight knot complement begs for a definition that renders this triangulation strongly essential.

We will first give a definition that displays the similarity between the ideal case and the closed case, and then give equivalent conditions in terms of the compact core, which are easier to verify in practice.

3.2. Ideal triangulations via pseudo-manifolds

The key to our definition in the ideal case is that intermediate paths in the homotopy are not allowed to pass through the ideal vertices, except for their endpoints. The reader is reminded that an ideal edge of M is an edge of \widehat{M} .

DEFINITION 3.3 (Admissible path homotopy). — *A path homotopy¹ $H: [0, 1] \times [0, 1] \rightarrow \widehat{M}$ 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 if there is a path homotopy $H: [0, 1] \times [0, 1] \rightarrow \widehat{M}$ with $H(x, 0) = \gamma(x)$, $H(x, 1) = \gamma(0)$ for all $x \in [0, 1]$ and $H((0, 1) \times [0, 1]) \subset M$.*

DEFINITION 3.4 (Essential and parallel ideal edges). — *An ideal edge of M is essential if it is not admissibly null-homotopic in \widehat{M} . Two ideal edges e, f of M are parallel if there is an admissible path homotopy between γ_e and one of $\pm\gamma_f$.*

DEFINITION 3.5 (ESSENTIAL IDEAL TRIANGULATION). — *The ideal triangulation of M is essential if every ideal edge of M is essential. It is strongly essential if, in addition, no two ideal edges of M are parallel.*

We remark that whilst the definition of an essential triangulation asks for exactly one vertex, the corresponding “minimality” condition is built into the definition of an ideal triangulation—having one ideal vertex for each end of M .

⁽¹⁾ i.e. a homotopy keeping endpoints fixed

3.3. Ideal triangulations via compact core

We now give equivalent conditions using the intersection of edges with the compact core M^c of M . This alternative viewpoint for essential edges can be found in [32].

LEMMA 3.6. — *An ideal edge e of M is essential if and only if there is no path homotopy $H: [0, 1] \times [0, 1] \rightarrow M^c$ between $\gamma_e \cap M^c$ (suitably parametrised) and a path in ∂M^c .*

Proof. — We write $\gamma = \gamma_e$. Suppose there is a path homotopy $H: [0, 1] \times [0, 1] \rightarrow M^c$ between $\gamma \cap M^c$ and a path β in ∂M^c . Coning the endpoints of β to the vertex corresponding to the component of ∂M^c containing β gives a loop β' in \widehat{M} , which is admissibly null-homotopic. The path homotopy between $\gamma \cap M^c$ and β extends to an admissible path homotopy between γ and β' , hence showing that γ is admissibly null-homotopic.

Conversely, suppose that $\widehat{H}: [0, 1] \times [0, 1] \rightarrow \widehat{M}$ is an admissible path homotopy between γ and a constant loop at a vertex of \widehat{M} . This restricts to a homotopy $H: E = (0, 1) \times [0, 1] \rightarrow M$ and we can assume that H is transverse to the surface $\partial M^c \subset M$. Then $H^{-1}(\partial M^c)$ is a compact, properly embedded 1-submanifold of E , and $H^{-1}(\partial M^c) \cap \partial E$ consists of the two points $(a, 0), (b, 0)$ with $\gamma(a), \gamma(b) \in \partial M^c$. Let α be the arc in $H^{-1}(\partial M^c)$ joining these two points, and let D be the disc in E bounded by α and the line segment in ∂E between these points. Then the restriction $H|_D: D \rightarrow M$ gives a path homotopy in M between $\gamma \cap M^c$ and a path in ∂M^c . Composing this with a deformation retraction $r: M \rightarrow M^c$, defined using the product structure on the ends of M , gives a path homotopy in M^c between $\gamma \cap M^c$ and a path in ∂M^c . \square

LEMMA 3.7. — *Two ideal edges e and f of M are parallel if and only if there is a homotopy $H: [0, 1] \times [0, 1] \rightarrow M^c$ between $\gamma_e \cap M^c$ and $\gamma_f \cap M^c$ (suitably parametrised) with the property that $H(0, t), H(1, t) \in \partial M^c$ for all $t \in [0, 1]$, i.e. the endpoints of all intermediate paths remain on the boundary.*

Proof. — Suppose $\gamma = \gamma_e$ and $\delta = \gamma_f$ are edge paths in \widehat{M} with the property that they have the same initial vertex v and the same terminal vertex w . If γ and δ are related by an admissible path homotopy of \widehat{M} , then using the conical structure of the ends, one can construct a homotopy with image in M^c taking $\gamma \cap M^c$ to $\delta \cap M^c$ whilst keeping the endpoints of all intermediate paths on the boundary. Conversely, if there is such a homotopy of M^c taking $\gamma \cap M^c$ to $\delta \cap M^c$, then coning the resulting arcs on ∂M^c to the vertices gives an admissible homotopy of \widehat{M} taking γ to δ . \square

As a final remark, we emphasise that “essential” and “strongly essential” apply to the edges of \widehat{M} in two different ways—depending on whether they are considered as edges of \widehat{M} or ideal edges of M . We hope that this will not cause any confusion, and we regret that in [32] “essential in \widehat{M} ” stands for “essential in M .”

4. Heegaard splittings

Our first construction takes as its launching point an irreducible one-sided Heegaard splitting:

THEOREM 4.1. — *Assume that M is a closed P^2 -irreducible 3-manifold M with the property that $H_1(M, \mathbb{Z}_2) \neq 0$ and $M \neq \mathbb{R}P^3$. Then M admits an essential one-vertex triangulation.*

In particular, any closed non-orientable P^2 -irreducible 3-manifold has such a triangulation.

Proof. — We first claim that M admits an irreducible one-sided Heegaard splitting. We know from [29] that M admits a one-sided Heegaard splitting $M = N(K) \cup Y$, where Y is a handlebody and $N(K)$ is a regular neighbourhood of a one-sided surface. (In [29], the focus is on orientable 3-manifolds, but the same argument works in the non-orientable case.) To prove the existence of an irreducible one-sided Heegaard splitting, consider a one-sided Heegaard splitting of smallest genus and suppose that it is not irreducible. Then there is a 2-sphere S meeting the Heegaard surface K in a single essential loop on K . Moreover S bounds a 3-ball B in M and the splitting $K \cap B$ of B consists of trivial handles. We can then replace the splitting surface K by $K' = (K \setminus K \cap B) \cup D$, where D is a disk bounded by $K \cap S$ on S . Then K' gives a smaller genus one-sided splitting, a contradiction.

Hence assume that $M = N(K) \cup Y$ is an irreducible one-sided Heegaard splitting of M . We next note that there is no 2-sphere or 2-sided projective plane meeting K in a single essential curve. Indeed, since M is P^2 -irreducible there is no such 2-sided projective plane in M , and since the Heegaard splitting is irreducible, there is no such sphere.

We choose a bouquet of circles Γ which is a spine of Y and has the property that every circle is essential in M . To achieve this, we first choose an appropriate free generating set for $\pi_1(Y)$. This is non-empty so long as $\pi_1(M)$ is strictly larger than \mathbb{Z}_2 . However, if $\pi_1(M) \cong \mathbb{Z}_2$, then M is homotopy equivalent to $\mathbb{R}P^3$, and Perelman’s proof of the Geometrisation Theorem implies that M is homeomorphic to $\mathbb{R}P^3$, contradicting our hypothesis. Consequently it follows that $\pi_1(M)$ is strictly larger than \mathbb{Z}_2 , and

so Γ is non-empty but some circles may be inessential. To achieve that the core of every one-handle is essential, we perform handle slides in the handlebody Y , using the fact that the image of $\pi_1(Y)$ is a non trivial subgroup of $\pi_1(M)$. This constructs the graph Γ .

Next, we pick a non-separating family of meridian discs D_1, \dots, D_k for Y dual to each of the k circles of Γ . Each boundary curve $C_i = \partial D_i$ is a component of the boundary of an annulus A_i or a Möbius band B_i in $N(K)$. The annulus A_i satisfies $\partial A_i = C_i \cup C'_i$ where $C'_i = A_i \cap K$. The Möbius band B_i meets K in an essential (possibly singular) orientation-reversing curve and $B_i \cup D_i$ is an immersed 2-sided projective plane. The projective plane theorem in [6] allows us to replace this immersed 2-sided projective plane by an embedded one. But as we are assuming that M is P^2 -irreducible, there is no such projective plane in M . Note that A_i is embedded except that C'_i can have isolated double points. We now form the 2-complex $X = K \cup D_1 \cup A_1 \cup \dots \cup D_k \cup A_k$ and claim this is a special spine of M . See Figure 3.




Figure 3. — Two discs D_i and D_j with their associated annuli A_i and A_j near the one-sided Heegaard surface K . Here, i may be equal to j .

It is clear that $M \setminus X$ is a 3-cell, and hence X is a spine of M . Moreover, we may assume that the curves C'_i on K meet transversely, and hence the spine X is simple. If X has no true vertex, i.e. no two curves C'_i and C'_j on K meet, we may perform a *lune move* (see Figure 4 and [24], §1.2) to arrange for the existence of true vertices. Denote the resulting simple spine with at least one true vertex again by X . We wish to show that X is special. As discussed in §2.7, by [24] (Remark 1.1.11) it now remains to show that the 2-strata of X are open 2-cells.




Figure 4. — The lune move introduces two true vertices.

Since the interior of each $A_i \cup D_i$ is a disc, we need to show that each 2-stratum on K is also a 2-cell. For the sake of contradiction, suppose that we can find a 2-stratum F on K , which is not a disc. Choose an essential simple closed curve C in F . Assume first that C is 2-sided in $N(K)$, i.e. that the lift to ∂Y under the natural map $\partial Y \rightarrow K$ has two components. Denote these components by C_+ and C_- . An annulus A properly embedded in $N(K)$ can be found with $\partial A = C_+ \cup C_-$ and $A \cap K = C$. Since $C \subset F$, it follows that C_+, C_- are both disjoint from D_1, \dots, D_k . Hence C_+, C_- bound disjoint discs D_+, D_- properly embedded in Y . We obtain an embedded 2-sphere $S = D_+ \cup A \cup D_-$ intersecting K in a single essential curve, contrary to assumption. If C is one-sided in $N(K)$ then there is a Möbius band B properly embedded in $N(K)$ with $B \cap K = C$. Then ∂B bounds a properly embedded disc D in Y and $B \cup D$ is a 2-sided projective plane, contrary to assumption.

As in [24] §1.1.5, there is a triangulation \mathcal{T} of M dual to X with precisely one vertex. We claim that this triangulation is essential. The edge loops are either dual to the discs D_1, \dots, D_k (and hence homotopic to the circles of Γ) or they are dual to faces F of K in X (and hence homologically non-trivial). In either case they are homotopically non-trivial and therefore essential. This completes the proof of the main statement.

The last sentence follows since, as is well-known, $H_1(M, \mathbb{Z}_2) \neq 0$ for any closed non-orientable 3-manifold. \square

Example 4.2 (Quaternionic space). — The minimal two tetrahedron triangulation of quaternionic space S^3/Q_8 (see Figure 2(d) in [19]) illustrates the construction in the above proof. Here, the incompressible one-sided Heegaard surface is a Klein bottle K , and sits as the union of two normal quadrilaterals in the triangulation. There is precisely one edge in the complement

of K , corresponding to the single compression disc D , and there are two more edges in the triangulation arising from the components of $K \setminus D$.

We next give a variation on the above construction, which applies to the interior of a compact orientable irreducible atoroidal 3-manifold M , which has boundary consisting of two incompressible tori.

THEOREM 4.3. — *Suppose M is a compact orientable irreducible annular 3-manifold with boundary having two tori components which are incompressible. Also assume that M is not homeomorphic to $T^2 \times [0, 1]$. Then there is an essential ideal triangulation \mathcal{T} , which is dual to a spine X of M coming from a compression body decomposition.*

Proof. — Choose an irreducible Heegaard decomposition of M so that $M = Y \cup Y'$, where both Y and Y' are compression bodies with a common boundary component S and one other boundary component being an incompressible torus in ∂M . Since the decomposition is irreducible, there cannot be a 2-sphere intersecting S in a single essential curve. A spine X for M can be built by $X = S \cup D_1 \cup \dots \cup D_k \cup D'_1 \cup \dots \cup D'_m$ where D_1, \dots, D_k is a collection of compressing discs for Y and D'_1, \dots, D'_m is a similar collection for Y' . Moreover, each disc collection is chosen such that it cuts its compression body up into a collar of the associated boundary torus, and, as in the previous proof, the boundaries of the discs can be assumed transverse on S and having at least one intersection point. Hence X is a simple spine of M with at least one true vertex.

Suppose there is a 2-stratum F of X , which is not a disc. Then $F \subset S$ and there must be either a 2-sphere meeting S in a single essential curve or an essential annulus between the two boundary components of M intersecting S in a single loop. Both of these are ruled out since the splitting is irreducible and M is anannular. So X is a special spine of M .

Let N be the open manifold obtained from M by adding a collar $\partial M \times [0, 1)$ to its boundary. Then N is homeomorphic to the interior of M , X is a special spine of N with at least one true vertex, and we denote by \mathcal{T} the ideal triangulation of N dual to X . An ideal edge e meets M in a properly embedded arc λ_e , and e is essential if and only if λ_e is not path homotopic into ∂M .

Now e is either dual to a face of S or to a disc in one of the families $\{D_1, \dots, D_k\}$, $\{D'_1, \dots, D'_m\}$. If e is dual to a 2-stratum contained in S , then λ_e has endpoints on different boundary tori and hence e is essential. If e is dual to one of the discs, then it may not be essential and we will show that (analogous to the proof of the previous theorem) one may perform a handle slide to arrive at an essential ideal triangulation. Taking advantage

of the symmetry of the situation, it suffices to show this for the edges dual to D_1, \dots, D_k .

We first show that there is at least one disc D_i for Y with the property that the arc λ_i dual to D_i is not path homotopic into a boundary torus. By way of contradiction, suppose this is not the case. Then the image of $\pi_1(Y)$ in $\pi_1(M)$ would have to be the same as the image of $\pi_1(T)$ in $\pi_1(M)$. But then the inclusion $T \subset M$ is a homotopy equivalence and so $M = T \times [0, 1]$, contrary to the assumption (see [34]). Hence Y has at least one non-trivial dual arc to a disc.

Next suppose some disc D_j has a dual arc λ_j path homotopic into the boundary torus. We can perform a handle slide of the handle dual to D_j over the handle dual to D_i . This replaces D_j, λ_j by a new disc/arc pair which is non-trivial, in the sense that the arc is not path homotopic into the boundary torus.

So by induction on the number of pairs D_j, λ_j for which λ_j is path homotopic into the boundary torus, we conclude that \mathcal{T} is essential after doing handle slides. \square




Figure 5

Example 4.4 (The Whitehead link). — Refer to Figure 5. Diagram (a) shows the Whitehead link. Diagrams (b) and (c) show a special spine for the complement of the Whitehead link. In (c) we have a horizontal plane, with two tubes, one following one of the link components above the plane, and the other following the other link component below the plane. In (b) we have two discs, one above and one below the plane. Triple points for the spine are shown with a solid line, while the apparent boundary of our view of the surface is shown dotted, and the cut out borders of the surface around infinity are shown dashed. The true vertices of the spine are shown with dots. Note that there are four true vertices, so the corresponding triangulation has four ideal tetrahedra. It can be checked that this is the same triangulation as listed in SnapPy [4] for the link S^2_1 .

5. Hierarchies

THEOREM 5.1. — *Assume that M is a closed Haken 3-manifold M . Then M admits a strongly essential one-vertex triangulation.*

In particular, every closed P^2 -irreducible non-orientable 3-manifold has a strongly essential triangulation.

Proof. — We build a hierarchy for M whose complement consists of either one or two open 3-cells. This is closely related to the construction in [27] of 0-efficient triangulations of Haken 3-manifolds. There are two cases, depending on whether $H_1(M, \mathbb{Q}) \neq 0$ or $H_1(M, \mathbb{Q}) = 0$. The dual of the hierarchy will be a one-vertex or two-vertex triangulation respectively. In the first case, we show that the triangulation is strongly essential and in the second case, we crush an edge which joins the two vertices to a point and then establish that the resulting triangulation is strongly essential.

5.1. Case 1: $H_1(M, \mathbb{Q}) \neq 0$

Begin with a closed incompressible non-separating 2-sided surface S_1 , which exists since $H_1(M, \mathbb{Q}) \neq 0$, using the well-known method of [34]. Cut open along S_1 to give a compact connected manifold M_1 with incompressible boundary consisting of two copies of S_1 . Construct a non-separating 2-sided incompressible and ∂ -incompressible proper surface S_2 in M_1 (again using [34]) and cut open along S_2 . Note that M_2 may not have incompressible boundary. Iterating this construction, we can find a hierarchy $\mathcal{H} = \{S_1, \dots, S_k\}$ with the property that M cut open along \mathcal{H} is a single ball. Note that surfaces S_i in the hierarchy may be compressing disks.

We require that:

- (1) S_i is a non-separating 2-sided incompressible and ∂ -incompressible proper surface (including the possibility of compressing disks) in M_{i-1} .
- (2) S_i is in general position with respect to $\{S_1, \dots, S_{i-1}\}$.
- (3) For each $i = 2, \dots, k$, S_i is chosen amongst all surfaces in M_{i-1} that satisfy (1) and (2) so that ∂S_i meets the 1-skeleton of the previous surfaces and itself minimally relative to a lexicographical complexity function. That is, we look at the vector

$$(\partial S_i \star \partial S_2, \partial S_i \star \partial S_3, \dots, \partial S_i \star \partial S_i)$$

where $\partial S_i \star \partial S_j$ counts the number of transverse intersection points between ∂S_i and ∂S_j . We choose S_i so that this is smallest in the lexicographical ordering.

Note that Waldhausen [38] uses a different but related criterion, choosing a surface S_i of maximal Euler characteristic. Conditions (2) and (3) imply that the hierarchy is a special spine and so the dual cellulation to the hierarchy is a triangulation. The dual triangulation \mathcal{T} to \mathcal{H} is then a one-vertex triangulation.

5.1.1. We claim that a triangulation dual to a hierarchy satisfying the above three conditions has the property that every edge loop is essential and no two loops are homotopic, keeping the vertex base point fixed.

5.1.2. Suppose that α is an edge loop. Then by definition of the dual triangulation, α is dual to a surface S_i of the hierarchy, i.e meets this surface in a single point and is disjoint from the other surfaces of the hierarchy. If we cut M open along surfaces S_1, \dots, S_{i-1} , the result is a compact 3-manifold M_{i-1} where S_i is a non-separating 2-sided incompressible and ∂ -incompressible proper surface. If α is contractible, then there is a map $f : D \rightarrow M$, where D is a disc, satisfying $f(\partial D) = \alpha$. Since all the surfaces S_1, \dots, S_{i-1} are incompressible, we can homotope f so that $f(D)$ misses all these surfaces, starting inductively with a surface of smallest index. But then $f : D \rightarrow M_{i-1}$ implies that α is contractible in M_{i-1} , which contradicts invariance of algebraic intersection number between S_i and α . Hence we have established that the triangulation \mathcal{T} is essential.

5.1.3. We will now establish the second property, i.e that \mathcal{T} is strongly essential. Suppose that α_1, α_2 are different edge loops. These are dual to surfaces S_i, S_h in the hierarchy. There are two cases, $i \neq h$ and $i = h$. The

first case is relatively easy. We can assume without loss of generality that $i < h$. If we cut M open along surfaces S_1, \dots, S_{i-1} , the result is a compact 3-manifold M_{i-1} where S_i is a non-separating 2-sided incompressible and ∂ -incompressible proper surface. By assumption, α_1 is a dual loop to S_i , whereas α_2 is disjoint from S_i . Hence the loops are not homotopic in M_{i-1} . Suppose there was a homotopy $H : D \rightarrow M$ between these two edges loops, where D is a disc. Hence $H(\partial D) = \alpha_1 \cup \alpha_2$. As in 5.1.2, we can homotope H off the surfaces S_1, \dots, S_{i-1} since these surfaces are incompressible and ∂ -incompressible. But this would then give a homotopy between α_1, α_2 in M_{i-1} , which as we noted is impossible. This completes the proof in the first case.

5.1.4. For the second case $i = h$ and we again cut open along S_1, \dots, S_{i-1} and examine a homotopy $H : D \rightarrow M$ between α_1, α_2 . As in the previous case, such a homotopy can be pulled off S_1, \dots, S_{i-1} so is contained in M_{i-1} . By definition, the edge loops α_1 and α_2 intersect S_i exactly once and are disjoint from S_1, \dots, S_{i-1} , so the pullback of S_i under the homotopy is an arc with two endpoints, one on each of the boundary arcs, which are pre-images of α_1, α_2 , together with some simple closed curves. By altering the homotopy, we can arrange that the pullback of S_i under the homotopy is a single arc, since any loops in the pullback of S_i can be eliminated using the incompressibility of S_i .

5.1.5. Consider now the pullback of the remainder of the hierarchy, i.e the surfaces S_{i+1}, \dots, S_k . Our plan is to change the homotopy so that it misses all of these surfaces, and is still disjoint from S_1, \dots, S_{i-1} . Let $j \in [i+1, k]$ be the least integer such that S_j has non-empty intersection with the image of the homotopy H . The pullback of S_j under H will have arcs with ends on the pullback of S_i and loops in the interior of D . No arcs of the pullback of S_j can have ends on ∂D , since the image of ∂D is the two edges α_1 and α_2 , whose only intersection with the hierarchy is with S_i . We are not yet concerned with the intersections of D with surfaces deeper in the hierarchy.

5.1.6. Loops can be eliminated in the usual way, using the incompressibility of S_j in M_{j-1} . Removing the arcs is more complicated. See Figure 6. Note that the figure shows $H(D)$ embedded in M , although in general it will be immersed and the endpoints of α_1 and α_2 are identified. Also in the argument following, we refer to discs, which in general are images of discs, rather than embeddings. To simplify matters, we will first deal with the special case of $H(D)$ embedded and then indicate how to proceed in the general case. We choose an innermost arc of $H(D) \cap S_j$ on one of the two sides of S_i ; in the figure this arc is γ . This bounds a disc D_1 on D , with $\partial D_1 = \gamma \cup \delta$, and $\delta \subset S_i$. Since S_j is ∂ -incompressible in M_{j-1} , the disc

D_1 in M_{j-1} with boundary the union of $\gamma \subset \partial S_j$ and $\delta \subset \partial M_{j-1}$ (where $\gamma \cap \delta = \partial\gamma = \partial\delta$) implies the existence of a disc D_2 in S_j bounded by γ and an arc ϵ in ∂M_{j-1} .




Figure 6. — An example showing how the image of the homotopy disc can intersect the hierarchy. The diagrams to the right are a close-up view of the intersection arcs on the left, and the way to alter the homotopy to avoid the arc γ .

5.1.7. In fact, the arc ϵ is contained only in S_i . To see this, first note that the only intersections we see with δ are with arcs of intersection with S_j (i.e. there are no arcs of intersection with $S_{j'}$ for $i < j' < j$). This is because we have already chosen $H(D)$ to miss all of these surfaces. If ϵ were not contained entirely in S_i , then it would cross either the boundary of S_i , or the boundary of some other surface $S_{j'}$ with $j' < j$. In either case, we could modify S_j by replacing D_2 with D_1 . The resulting surface would still be incompressible and ∂ -incompressible, and would have smaller lexicographical complexity since replacing ϵ with δ would remove intersections of ∂S_j with the boundary of some surface of the hierarchy with smaller index than j .

5.1.8. Thus $\delta \cup \epsilon$ is a loop in S_i . Since S_i is incompressible, $\delta \cup \epsilon$ bounds a disc in S_i . We rule out the possibility that D_3 is on the “outside”; that is that α_1 and α_2 intersect D_3 . The disc D_3 together with D_1 and D_2 forms a sphere, which bounds a ball B in M since M is irreducible. Either this ball is above or below the sphere as seen in Figure 6. If it is above then S_j is contained in B , if below then S_i is contained in B . In either case, a surface contained in a ball is not incompressible. Therefore the disc D_3 is on the “inside”, as labelled in Figure 6.

5.1.9. Now, we can modify the homotopy, eliminating the arc γ , by replacing H on a small neighbourhood of the pullback of D_1 with a map to the union of D_2 and D_3 , pushed slightly off as in the lower right diagram of Figure 6. It is possible that parts of S_j connect onto the other side of S_i from the arc γ (such a part of S_j is shown in the figure), and we should be concerned that in removing the (pullback of the) arc γ , we do not introduce more arcs. There are no arcs on the part of the disc that is pushed off of D_2 , since such an arc would correspond to a piece of surface with boundary on S_j . An arc on the part of the disc that is pushed off of D_3 has a corresponding arc of the boundary of S_j in D_3 itself.

5.1.10. We claim that the only arcs that are possible in D_3 are arcs of intersection with S_j that intersect each of δ and ϵ once, or δ twice. (In Figure 6 only one arc that hits both δ and ϵ is shown.) First, there are no arcs of intersection of D_3 with surfaces $S_{j'}$ for $i < j' < j$. To see this, note that there are no such intersections with δ by our choice of $H(D)$, and any such intersections with ϵ would contradict the minimality of the lexicographical complexity function for S_j . If there were any such arcs of intersection entirely in the interior of D_3 , then the first such surface in the hierarchy, $S_{j'}$, would intersect S_i in a loop bounding a disc, which would contradict incompressibility of $S_{j'}$ (since $S_{j'}$ is not a disc). Second, any arcs of intersection with S_j that intersect ϵ twice could be isotoped over to reduce the lexicographical complexity function for S_j . Therefore each arc has at least one of its endpoints on δ , which means that it connects to a pre-existing arc on $H(D)$. So this modification of the homotopy removes the pullback of γ from D , and does not introduce any new arcs. Continuing in this way, we can remove all arcs of intersection between S_j and $H(D)$.

5.1.11. We can then repeat the process for $S_{j''}$ where $j < j'' \leq k$. Eventually, we end up with the graph of intersection between $H(D)$ and the hierarchy consisting of just a single arc, which means that the edge loops α_1, α_2 are identical in the triangulation. The proof of the first case is now complete assuming that $H(D)$ is embedded.

5.1.12. We now indicate the modifications required to deal with the general case when $H(D)$ is singular. Firstly in 4.1.7 we need to explain why the arc ϵ must be contained in S_i . Note that the subdisk D_1 of D , mapped by H , must be homotopic into S_j , since S_j is ∂ -incompressible in M_{j-1} . Hence δ is homotopic to ϵ in ∂M_{j-1} . We claim that this contradicts our choice of S_j satisfying property (3) that ∂S_j has a minimal number of transverse intersection points with boundary curves $\partial S_{j'}$ for $j' < j$ with these numbers lexicographically ordered.

To prove the claim, notice that the curves $\partial S_{j'}$ for $j' < j$ form a trivalent graph Γ (often called a boundary pattern) in ∂M_{j-1} . The homotopy between δ and ϵ can be easily shown to imply there is an innermost bigon between a subarc of ϵ and a subarc of Γ in case ϵ is not contained in S_i . But then ϵ can be isotoped across this bigon producing a proper isotopy of S_j in M_{j-1} reducing the number of intersection points of ∂S_j with the boundary pattern, with the lexicographic ordering of complexity. So this establishes the claim and establishes the step in 4.1.7.

The remainder of the argument in 4.1.8 to 4.1.11 follows as in the embedded case, requiring only the observation that the ball B in 4.1.8 is the image of a ball, rather than an embedded ball. But this does not affect the arguments and so the first case is established.

5.2. Case 2: $H_1(M, \mathbb{Q}) = 0$

In the second case, the first surface S_1 in the hierarchy will be a closed separating incompressible 2-sided surface. The remaining surfaces for the two components of M split open along S_1 will be properly embedded non-separating incompressible and ∂ -incompressible 2-sided surfaces. After cutting open along these we get two 3-cells, so the dual triangulation \mathcal{T} to this hierarchy \mathcal{H} has two vertices. We claim that the result of crushing any edge E joining the two vertices is a well-defined one-vertex triangulation \mathcal{T}^* of M . Here the method of [17] for crushing suitable 3-cells with normal 2-sphere boundary will be followed. To prove this claim, we first establish that the boundary B of a regular neighbourhood of E is a normal 2-sphere.

5.2.1. Suppose F is a face of the triangulation \mathcal{T} with at least two edges in ∂F identified with E . Then F is either a cone, a Möbius band, a 3-fold or a dunce hat (see [17] for the terminology). The last three cases cannot occur since then the ends of E would have to be the same vertex. Hence assume F is a cone. Then the third edge E' of ∂F is a loop which is contractible via F . We can now use the same method as in Case 1 to produce a contradiction: E' is dual to some surface S_i of the hierarchy. The contracting disc F can be pulled off of all the previous surfaces S_1, \dots, S_{i-1} of the hierarchy. We then see that E' is contractible in M_{i-1} , and this gives a contradiction to the fact that E' is dual to S_i . So this shows that F cannot be a cone either. It follows that each face of the triangulation contains E at most once, and hence the surface B is normal.




Figure 7. — The possible ways in which the edge E can be incident to a tetrahedron.




Figure 8. — The two possible arrangements of surfaces giving a tetrahedron incident with the two distinct vertices in \mathcal{T} .

5.2.2. Figure 7 shows the possible ways in which the edge E can be incident to a tetrahedron and the resulting normal discs in S : the edge E either meets a tetrahedron only in the vertices, or in one edge, or in a pair of opposite edges. We will now show that the last case does not happen. Suppose there is a tetrahedron of \mathcal{T} with two opposite edges identified to E . This tetrahedron is dual to a vertex x of the hierarchy which is at the intersection of S_1, S_i, S_j , three surfaces of the hierarchy, with the first being the closed separating 2-sided surface. If S_i, S_j were on the same side of S_1 , we would have E dual to faces on different surfaces of the hierarchy, which cannot happen. Since S_i, S_j must be on opposite sides of S_1 , x is a crossing point of curves $C_i \subset \partial S_i$ and $C_j \subset \partial S_j$ in S_1 . The edge E will then be dual to faces of the hierarchy on S_1 which have x as a vertex and are diagonally opposite at x (since the edges of the tetrahedron corresponding to E are opposite each other). Since each edge is dual to one face, we conclude that there is a face F of the hierarchy with two vertices at x . Then there is a loop C^* which is contained in F and meets both S_i and S_j transversely only at x . We claim that C^* is an element of infinite order in $H_1(M)$, contrary to assumption. To

see this, first cut M open along S_1 , the surface containing the loop. We now use the Mayer-Vietoris sequence to compute $H_1(M)$ as the result of gluing back along S_1 . We can push C^* into both sides of the surface. By duality with S_i on one side and S_j on the other, we see that both copies of the loop are elements of infinite order in the homology of the cut open manifold M_1 . Hence the Mayer-Vietoris sequence shows that gluing produces an element of infinite order coming from gluing the two copies of C^* together. So we conclude this case does not occur.

5.2.3. It follows that each tetrahedron contains E at most once, and since the endpoints of E are distinct vertices in the triangulation, we can crush E and the surrounding tetrahedra to obtain a triangulation \mathcal{T}^* as claimed.

5.2.4. To complete the proof of the theorem (showing that \mathcal{T}^* is essential and strongly essential), we follow a similar argument to the case of infinite homology. First suppose that α is an edge loop in \mathcal{T}^* . Then there is an edge β in \mathcal{T} which becomes α after crushing. If β is also an edge loop, then we can follow exactly the same argument as in the first case to get a contradiction. So assume that β joins the two vertices and so $\beta \cup E$ is a loop. Clearly α being contractible is equivalent to $\beta \cup E$ being contractible. Since β joins the two vertices, it must be dual to S_1 , and so $\beta \cup E$ is contractible if and only if there is a homotopy between β and E . But now we can use the same argument as in case 1 to show that E and β are dual to the same face of S_1 and so are in fact the same edge. But then α cannot have come from $\beta = E$, and we have our contradiction. This finishes the first part showing that \mathcal{T}^* is essential. Finally suppose that α_1, α_2 are homotopic edge loops in \mathcal{T}^* . Again choose edges β_1, β_2 which map to α_1, α_2 under the crushing map from \mathcal{T} to \mathcal{T}^* . There are three cases:

5.2.5. In the first case, β_1, β_2 have the same endpoints and are homotopic keeping their endpoints fixed. As before, the same argument as in the case of infinite homology starting at 5.1.3 shows this cannot happen. (Note that this case covers both the endpoints of each edge being the same or different, and the argument is the same in each case.)

5.2.6. In the second case β_1 has both endpoints on one of the two vertices, and β_2 has both endpoints on the other vertex, and the loop formed by following β_1 then E , then β_2 in reverse then E in reverse, is contractible. In this case, both β_1, β_2 are loops. We can now follow the previous argument involving Mayer Vietoris to show that the homology class of the homotopic loops β_1, β_2 is of infinite order, since these loops are dual to spanning surfaces on either side of S_1 . So this contradicts our assumption that the first homology is finite.

5.2.7. In the third case, without loss of generality β_2 has both endpoints at the same vertex, while β_1 has them at different vertices, and the loop $\beta_1 \cup \beta_2 \cup E$ is contractible, giving us a homotopy $H: D \rightarrow M$. Since it has different vertices at its endpoints, β_1 must be dual to S_1 , and since it has the same vertex at its ends, β_2 is dual to some $S_i, i > 1$. As before, our plan is to pull the homotopy off most of the hierarchy, while fixing the boundary. This time, our goal is that the intersection of $H(D)$ with the hierarchy consists of a ‘T’ pattern: an arc of intersection with S_1 joining a point of E with a point of β_1 , together with an arc of intersection with S_i joining a point of β_2 with a point of the first arc. Then, collapsing E corresponds to deleting the subarc of the S_1 arc joining the triple point to the E section of the boundary, then collapsing the E part of the boundary. The result is a homotopy showing that α_1 and α_2 are dual to the same face of the hierarchy and so are in fact the same edge.

5.2.8. To achieve our goal ‘T’ intersection pattern, first we pull $H(D)$ off all surfaces of the hierarchy from S_1 down to S_{i-1} other than the single arc of intersection with S_1 . Loops are removed using incompressibility as usual, and arcs of intersection that meet the intersection arc with S_1 are removed as in 5.1.6. When we come to S_i , we find that one of our arcs of intersection hits β_2 at one end, and the S_1 arc at the other, since there is nowhere else that the arc can go to. We remove any other intersections of S_i with $H(D)$ as usual. At this stage then we have our ‘T’ intersection pattern, together with some intersections with surfaces from S_{i+1} down to S_k . Let $j \in [i+1, k]$ be the least integer such that S_j has non-empty intersection with the image of the homotopy H . There are two cases, depending on whether an arc of S_j has both endpoints on the same one of the three “legs” of the ‘T’, or on different legs. The former case is dealt with innermost first, again as in 5.1.6, and we perform all of these homotopies first. Now for the second case, there are two subcases depending on which one of the two components formed by the initial separating surface S_1 contains S_j . If S_j is in the component not containing S_i then the S_j arc runs over the “top” of the ‘T’, and no parts of S_j can appear on the other side of S_1 . This case is again dealt with as in 5.1.6.

5.2.9. Finally, we deal with S_j arcs that connect between an S_1 leg and the S_i leg. We order these by the distance of the S_i end of the arc from the triple point of the ‘T’, and deal with the S_j arc closest to the triple point first. See Figure 9. Here the arc of intersection with S_j is labelled γ . Since γ is the innermost arc, there are no other arcs of intersection inside the disc D_1 . As in the situation in 5.1.6, boundary incompressibility of S_j in M_{j-1} together with minimal lexicographical complexity means that there is a disk D_2 in S_j bounded by γ and another arc, here the union of $\epsilon_1 \subset S_1$

and $\epsilon_i \subset S_i$. This arc must be contained entirely in S_1 and S_i (with one connected component in each) by minimal lexicographical complexity for S_j . The arcs ϵ_1 and ϵ_i have no intersections with $S_{j'}$ for $i < j' < j$, for the same reasons.




Figure 9. — The image of the homotopy disc intersecting the hierarchy. The diagrams to the right are a close-up view of the intersection arcs on the left, and the way to alter the homotopy to avoid the arc of intersection with S_j .

5.2.10. Now consider η , the part of the boundary curve of S_i between $\delta_i \cap \delta_1$ and $\epsilon_i \cap \epsilon_1$. This must be a single arc on S_i which does not intersect any other surface boundary curve, by minimal complexity for S_i , comparing it with the curve $\delta_1 \cup \epsilon_1$. The curves η, δ_i and ϵ_i together form a loop in S_i , and bound a disc D_4 in S_i by the incompressibility of S_i . Here there is no question of which side of the loop contains the disc since η is on the boundary of S_i . Also, the curves η, δ_1 and ϵ_1 together form a loop in S_1 , and bound a disc D_3 in S_1 by the incompressibility of S_1 . If the disc were “outside” the loop, so that E and β_1 intersect it, then similarly to as in 5.1.8, the union of D_3 with the discs D_1, D_2 and D_4 gives a sphere, so by irreducibility of M , one or the other side is a ball. But then either S_i or S_1 would be contained in the ball (depending on which side of the sphere it is), which is impossible because of the incompressibility of the surfaces. Therefore D_3 is on the “inside” of the loop $\eta \cup \delta_1 \cup \epsilon_1$ as shown in Figure 9.

5.2.11. We claim that there are no curves of intersection between the interiors of any of D_2, D_3, D_4 , and any of the boundary arcs of the hierarchy. First, $D_2 \subset S_j$, so none of the surfaces with index less than or equal to j can have boundary on it. Second, D_3 has no loops of intersection in its interior since it is a disc. There are no intersections with arcs of S_j , since S_j is contained in the other component of $M \setminus S_1$. For intersections with $S_{j'}$ for $j' < j$, D_3 has boundary δ_1, η and ϵ_1 . The arc δ_1 has no intersections by our choice of $H(D)$. The arcs ϵ_1 and η have no intersections with $S_{j'}$ for $j' < j$ by our previous arguments in 5.2.9 and 5.2.10 respectively. So there are no intersections with D_3 . Lastly, D_4 again has no loops of intersection in its interior since it is a disc. For intersections with $S_{j'}$ for $i < j' < j$, there are no intersections with δ_i by our choice of $H(D)$, and no intersections with ϵ_i or η again by our arguments in 5.2.9 and 5.2.10 respectively. For intersections of S_j with D_4 , the above arguments also rule out any arcs that do not intersect δ_i . Any arc that does intersect δ_i implies that there is an arc of intersection in $H(D) \setminus D_1$ starting at δ_i . However, we have already removed all such arcs that have both endpoints on the S_i arc, and the only other option would be an arc connecting an S_1 arc and the S_i arc that is closer to the triple point of the ‘T’ than the supposedly innermost such arc we are considering. So there are no arcs of intersection with any of D_2, D_3 or D_4 . Therefore we can modify the homotopy, eliminating the arc γ , by replacing H on a small neighbourhood of the pullback of D_1 with a map to the union of D_2, D_3 and D_4 , pushed slightly off as in the lower right diagram of Figure 9. Continuing in this way, we can remove all arcs of intersection between S_j and $H(D)$.

5.2.12. Again, we repeat the process for $S_{j''}$ where $j < j'' \leq k$. Eventually, we end up with the graph of intersection between $H(D)$ and the hierarchy consisting of the desired ‘T’ graph, and we are done with this case, and the proof of the main statement is complete. The last sentence of the theorem follows since a closed P^2 -irreducible non-orientable 3-manifold is Haken. \square

Remark 5.2. — It is also possible to extend this technique to cusped hyperbolic 3-manifolds as illustrated by the following examples. We will not pursue this but use a different technique to produce essential and strongly essential triangulations for cusped and closed hyperbolic n -manifolds in Section 7. In the cusped case, the technique is to subdivide the canonical decomposition into ideal hyperbolic polyhedra, using [7], into ideal simplices. In the closed case, we follow a similar approach starting with a Delaunay cell decomposition.

Example 5.3 (The Borromean rings). — See Figure 10. The three components of the Borromean rings are labelled A, B and C . The first two surfaces, S_1 and S_2 in the hierarchy are tori, parallel to A and B . The third, S_3 is

a sphere with three holes, one boundary component is a longitude of S_1 , and the other two are meridians of S_2 . The fourth, S_4 is a disc, which can be described by starting with a sphere with three holes, this time with one boundary component a longitude of S_2 and the other two meridians of S_1 , and then cutting this surface along the intersection with S_3 . The result of this as shown in Figure 10 is not a hierarchy, as the singular edges at the intersection between S_3 and S_4 are degree 4. We can fix this by altering S_4 slightly, pushing the components of $\partial S_4 \cap S_3$ off itself by moving it along S_3 . There are four different ways to do this, as we can push each of the two pairs of components in two different directions. The cellulation dual to the complex we have before pushing off has square faces – pushing off breaks each square into two triangles. The triangulations obtained by pushing off in one direction versus another are therefore related by flipping the diagonal of that square. This is sometimes referred to as a 4-4 move on the triangulation.




Figure 10. — A hierarchy for the Borromean rings.

Example 5.4 (The figure 8 knot complement). — See Figure 11. We view the figure 8 knot complement as the punctured torus bundle with monodromy

$$\begin{pmatrix} 2 & 1 \\ 1 & 1 \end{pmatrix}.$$

A fundamental domain for the manifold is a cube with its vertical edges deleted, the vertical faces of the cube identified by horizontal Euclidean translation, and the bottom face of the cube identified to the top via the monodromy. In the figure we see this cube, together with parts of surfaces in

images of the cube under the monodromy and its inverse. The first surface of the hierarchy, S_1 is a boundary parallel torus. The second, S_2 is the punctured torus fibre, cut off at S_1 . The third and fourth, S_3 and S_4 are vertical discs parallel to generators of the fundamental group of the punctured torus fibre.




Figure 11. — A hierarchy for the figure 8 knot complement.

The triangulation dual to this hierarchy has one finite vertex and one ideal vertex. We crush an edge of this triangulation to produce an essential and strongly essential triangulation. See Figure 12. On the left we see the hierarchy again. We need to crush an edge joining the finite and ideal vertices, or equivalently, remove a face of S_1 . One choice is to remove one of the faces of S_1 incident to the edge marked e . However, the same face is incident to e twice, so we must remove a face of S_2 as well, otherwise the resulting hierarchy would have a face with boundary. On the right, after removing these faces, many of the remaining faces coalesce into larger faces as there are no longer curves of triple points between them. In the dual picture, this corresponds to edges of the triangulation being identified after the crushing. There are only two quadruple points remaining, labelled A and B . Thus the

corresponding triangulation has only two tetrahedra, and in this case we get the canonical ideal triangulation of the figure 8 knot complement.




Figure 12. — Crushing an edge of the triangulation dual to the hierarchy corresponds to removing the faces of the hierarchy incident to an edge (here labelled e , note that we see two translates of it). After removing these faces, edges with only two faces incident are drawn with dashed lines.

6. Angle structures

We give a quick summary of strict angle structures and semi-angle structures. A detailed discussion may be found in e.g. [23]. Given an ideal triangulation, each ideal tetrahedron has a real number, called an *angle*, associated with each of its edges, so that opposite edges have the same angles. The angle equations are then that the sum of the three angles at an ideal vertex for each such tetrahedron is π and the sum of the angles around each edge in the triangulation is 2π . A strict angle structure is a solution of the angle equations for which each angle is strictly positive, whereas a semi-angle structure is a solution with each angle non-negative.

THEOREM 6.1. — *Suppose that M is a 3-manifold with an ideal triangulation \mathcal{T} . If \mathcal{T} admits a strict angle structure, then \mathcal{T} is strongly essential. If \mathcal{T} admits a semi-angle structure, then it is essential.*

Proof. — Consider the case of a strict angle structure on \mathcal{T} . We establish first that \mathcal{T} is essential. So suppose not, i.e. there is an edge α of \mathcal{T} which is homotopic, keeping its ends on a boundary torus T , into T . Construct the covering \overline{M} corresponding to the image of the subgroup $\pi_1(T)$ in $\pi_1(M)$. Then there are homeomorphic lifts \overline{T} of T and $\overline{\alpha}$ of α in \overline{M} , so that ends

of $\bar{\alpha}$ lie in \bar{T} . There is a collar of the boundary component \bar{T} of \bar{M} which contains $\bar{\alpha}$. We can push \bar{T} across $\bar{\alpha}$ to give a new surface \bar{T}' that is possibly not normal. Using the barrier arguments of [17], the edge $\bar{\alpha}$ is a barrier for normalising the incompressible surface \bar{T}' , and hence we shrink \bar{T}' to obtain a normal surface T^* in \bar{M} , which is isotopic but not normally isotopic to \bar{T} since T^* is disjoint from $\bar{\alpha}$.

According to the discussion of the result of Casson and Rivin in [22] (see also [23], Theorem 3), there cannot be any normal torus T^* which is not peripheral in a triangulation with a strict angle structure. A direct argument is as follows. Note that \bar{M} has a lifted strict angle structure. We apply a formal Gauss-Bonnet argument (cf. [23], Proposition 13). Namely if we put the dihedral angles of edges at the vertices of the normal cells of T^* , then a triangle has angle sum π and a quadrilateral has angle sum $2\alpha + 2\beta < 2\pi$, where α, β, γ are the 3 angles at edges of the ideal tetrahedron containing the quadrilateral. Adding zero for each triangle and $2\alpha + 2\beta - 2\pi = -2\gamma$ for each quadrilateral gives $2\pi\chi(T^*)$. We conclude that the torus T^* has negative Euler characteristic if it has any quadrilateral disks. But since T^* is not normally isotopic to a peripheral torus in the ideal triangulation, it must have quadrilaterals. So this gives a contradiction to the existence of the homotopy of an edge into a boundary torus. This completes the proof that \mathcal{T} is essential.

To show that \mathcal{T} is strongly essential, assume to the contrary that there are edges α_1, α_2 which are homotopic keeping their ends on peripheral tori. We can truncate M and \mathcal{T} , and finish the proof with the following argument of Lackenby [22]. Consider the homotopy as a mapped-in quadrilateral disc with two opposite boundary arcs running along peripheral tori and two arcs running along truncated copies of α_1, α_2 . Making this map transverse to the truncated triangulation and least weight, we can classify all the possible disc types in the complement of the pullback of the 2-skeleton of \mathcal{T} . Lackenby shows that all the disc types have non-positive formal contribution to the Euler characteristic in a combinatorial Gauss-Bonnet formula. Moreover there are some negative contributions so this contradicts the fact that the Euler characteristic is one. So this completes the proof that the triangulation is strongly essential if it supports a strict angle structure.

Suppose next that \mathcal{T} has a semi-angle structure. As in the first paragraph, if \mathcal{T} was not essential, there would be a normal torus T^* embedded in the covering space \bar{M} which has a lifted taut triangulation. A standard argument using thin position (which goes back to Thompson [36] and Stocking [35], with a detailed discussion given in Schleimer [30], Theorem 6.2.2) shows that in the product region between the vertex linking torus and the

normal torus T^* is an almost normal torus T' which has exactly one octagon. However, this octagon makes a negative contribution to the formal Euler characteristic of T' (see Figure 13), thus contradicting the combinatorial Gauss-Bonnet formula. This completes the proof of this case. \square




Figure 13. — An octagonal disc from an almost normal surface meeting the edge with angle α twice contributes negative combinatorial area $2(\alpha - 2\pi)$ to the combinatorial Gauss-Bonnet formula.

Example 6.2. — We first give simple examples of taut triangulations which are not strongly essential. Suppose we take any example of an ideal triangulation with a taut structure and choose a pair of adjacent triangular faces, sharing an edge e . Assume also that the π angles at e are on opposite sides of the two faces. Unglue the triangulation along e and these two faces, and glue onto the resulting four faces a pillow consisting of two tetrahedra with two faces and a degree two edge in common (sometimes this is called a 0–2 move). It is then easy to see that the taut structure extends over this new triangulation. However the degree two edge is an obstruction to the existence of a strict angle structure (see [21]). Moreover this triangulation clearly has two edges which are homotopic keeping their ends on the boundary tori, namely the two edges that the edge e is split into. So we see explicitly that this triangulation is not strongly essential.

Example 6.3 (m136). — We now give an example of a taut triangulation which is strongly essential, but does not support a strict angle structure. This is Example 7.7 of [16], which is a triangulation of the manifold m136 from the SnapPea [4] census. For convenience we reproduce the combinatorial data of this example in Tables 6.3 and 6.3.

To reconstruct the triangulation, take 7 tetrahedra with vertices labelled 0 through 3, all with consistent orientations, and make the appropriate

gluings given in Table 6.3. For example, the top left entry in the table tells us to glue the face of tetrahedron #0 with vertices labelled 0,1,2 to the face of tetrahedron #1 with vertices labelled 3,1,2 in the orientation that matches the vertices in the order given. The identifications of edges can be deduced from this information, and the resulting edge cycles are given in Table 2.

Table 1. — Gluing data for a triangulation of m136.

Tetrahedron	Face 012	Face 013	Face 023	Face 123	Shape parameter at 01
0	1 (312)	4 (302)	6 (130)	4 (132)	$2i$
1	3 (102)	2 (012)	2 (203)	0 (120)	$-1 + 2i$
2	1 (013)	6 (321)	1 (203)	4 (031)	$\frac{3}{5} + \frac{1}{5}i$
3	1 (102)	6 (230)	5 (021)	5 (023)	-1
4	5 (312)	2 (132)	0 (130)	0 (132)	$\frac{1}{5} + \frac{2}{5}i$
5	3 (032)	6 (012)	3 (123)	4 (120)	2
6	5 (013)	0 (302)	3 (301)	2 (310)	$\frac{1}{2} + \frac{1}{2}i$

Table 2. — Edge cycles in the triangulation.

Edge	Degree	Tetrahedron (Vertex numbers)
0	4	0 (01), 4 (30), 2 (21), 1 (31)
1	4	0 (02), 1 (32), 2 (30), 6 (13)
2	10	0 (03), 6 (10), 5 (10), 3 (30), 6 (02), 5 (03), 3 (13), 6 (30), 0 (23), 4 (32)
3	10	0 (12), 4 (13), 2 (32), 1 (30), 2 (20), 1 (02), 3 (12), 5 (02), 3 (02), 1 (12)
4	6	0 (13), 4 (02), 5 (32), 3 (32), 5 (12), 4 (12)
5	4	1 (01), 2 (01), 6 (32), 3 (10)
6	4	2 (13), 6 (21), 5 (31), 4 (01)

It is shown in [16] that the triangulation does not admit a strict angle structure. However, it can be checked (for instance, using Regina [3]) that the triangulation does admit a taut angle structure. We now show that the triangulation is strongly essential. This shows that unfortunately, having a strongly essential triangulation is necessary but not sufficient to produce strict angle structures, even starting with a taut structure.

Although the property of a triangulation being strongly essential is a homotopic property, here we use the geometric structure. The last column of Table 1 gives shape parameters for the seven tetrahedra that form a solution to Thurston’s gluing equations corresponding to the complete hyperbolic structure. The corresponding developing map $D: \widetilde{M} \rightarrow \mathbb{H}^3$ has the

property that distinct cusps of \widetilde{M} map to distinct points of $\partial\mathbb{H}^3$. If there are two homotopic edges in \widetilde{M} , then the corresponding geodesics in \mathbb{H}^3 will coincide. If we had a triangulation with all positive volume tetrahedron shapes corresponding to the complete structure (so a geometric triangulation), then we would be done. This follows because the geodesic image g of an edge in \mathbb{H}^3 is surrounded by positive volume tetrahedra, and no other tetrahedra in the developing map can intersect these other than via the face identifications with their neighbours. In particular, no other tetrahedron in \mathbb{H}^3 can have g as an edge. (Alternatively, a geometric triangulation gives a strict angle structure, so we can apply Theorem 6.1.)

For the shapes that we do have, each positive volume tetrahedron in \mathbb{H}^3 again intersects other tetrahedra only via face identifications. So if there were a pair of homotopic edges, we would be able to walk from one to the other through a sequence of flat tetrahedra. In our case the flat tetrahedra are tetrahedra 3 and 5. Figure 14 shows the way in which the two tetrahedra are glued to each other. Notice that they are glued to each other on the bottom two faces of tetrahedron 3 and the top two faces of tetrahedron 5. Moreover, as we develop through copies of these tetrahedra in \widetilde{M} , each new pair of tetrahedra are parabolic translates of the previous pair (around the vertex labelled 2 in both tetrahedra). Thus, no pair of edges are mapped to the same geodesic in \mathbb{H}^3 .

We use a similar argument in the proof of Theorem 7.1.




Figure 14. — The two flat tetrahedra in the triangulation with shapes corresponding to the complete structure. Edges are numbered in circles with arrows indicating the orientation of each edge relative to the tetrahedra. The numbering and orientations match the ordering given in Table 2. Vertex numbers are also shown. The flat shapes for these tetrahedra are such that the diagonal edges have dihedral angles of π .

7. Metrics with constant curvature

THEOREM 7.1. — *Assume that M is a compact orientable hyperbolic 3-manifold M with $h > 0$ incompressible torus boundary components. Then M has a strongly essential ideal triangulation.*

Proof. — This follows relatively easily from techniques in [16]. Given our manifold M , we construct an ideal triangulation using Algorithm 5.2 of [16]. Briefly, this uses the Epstein-Penner decomposition [7] of M into convex ideal hyperbolic polyhedra. These polyhedra are further subdivided into ideal pyramids using a coning procedure. For our purposes, the cone vertex on each polyhedron can be chosen arbitrarily. The pyramids are further subdivided into ideal tetrahedra. Where the induced triangulations on the glued faces of adjacent polyhedra disagree, we insert layered triangulations to bridge between them. This construction produces an ideal triangulation of M that comes with ideal hyperbolic shapes for all of the tetrahedra, those coming from the polyhedral cells having positive volume while the tetrahedra in the bridging layered triangulations have volume zero. Thus there is a pseudo-developing map $D: \widetilde{M} \rightarrow \mathbb{H}^3$ corresponding to the complete hyperbolic structure on M . The edges of the triangulation all map to geodesics in \mathbb{H}^3 . Since a geodesic must have distinct endpoints in \mathbb{H}^3 , this means that each edge in \widetilde{M} must also have distinct endpoints, and so none of them are null-homotopic, and so the edges are all essential.

A similar argument shows that two edges are homotopic fixing their endpoints if and only if they map to the same geodesic in \mathbb{H}^3 . Our pseudo-developing map shows that this could only possibly happen within one of the layered triangulation regions. However, an analysis of the construction of the triangulations in these regions (Algorithm 5.1 of [16]) shows that no pair of edges in such a region join the same vertices of the associated ideal polygon. This shows that the triangulation is also strongly essential. \square

Remark 7.2. — Theorem 6.1 gives a slightly different construction proving Theorem 7.1 in the case of a hyperbolic link complement in S^3 . Corollary 1.2 of [16] states that such a manifold admits an ideal triangulation with a strict angle structure. Then by Theorem 6.1, this triangulation is strongly essential.

Given a closed Riemannian 3-manifold with constant curvature, our next aim is to mimic the proof of Theorem 7.1, starting with the Delaunay cell decomposition dual to the cut locus of a point, to obtain a strongly essential triangulation with a single vertex. Without loss of generality, we may normalise the metric to have constant curvature -1 , 0 or $+1$. An extra condition on the diameter is required in the case of constant positive curvature.

For constant zero or negative curvature, the cut locus from any point is always a spine, obtained from the boundary of a Dirichlet domain, but is not always generic so its dual need not be a triangulation. As an example, the flat metric on the torus obtained by gluing up a rectangle gives a non-generic spine but gluing up a centrally symmetric hexagon gives a generic spine. For the cut locus in the hexagon case is a theta curve and the dual is the usual essential one-vertex triangulation of a torus.

THEOREM 7.3. — *Suppose that M is a closed Riemannian manifold of dimension 3 with a metric with constant curvature -1 , 0 or $+1$. In the case of curvature $+1$, also assume that the diameter of the manifold is smaller than π . Then the cell decomposition dual to the cut locus of any point x_0 can be subdivided to give a strongly essential one-vertex triangulation.*

Proof. — We first follow the well-known construction of Dirichlet domains and Delaunay cell decompositions in the case of constant zero or negative curvature. The modification required to deal with the case of constant positive curvature is given at the end. For basic information on hyperbolic manifolds, see [28].

Consider the picture in the universal covering \widetilde{M} of M . Since \widetilde{M} is isometric to Euclidean or hyperbolic space, there is a unique geodesic segment between any two points of \widetilde{M} . The Dirichlet construction defines a region \mathcal{D} in \widetilde{M} given by choosing a lift x_1 of the base point x_0 in M and letting $\mathcal{D} = \{x \in \widetilde{M} : d(x, x_1) \leq d(x, x_j) \text{ for all } j > 1\}$, where the lifts of x_0 are denoted x_1, x_2, \dots and d is the Riemannian distance between points in \widetilde{M} . Then the Dirichlet domain \mathcal{D} is a convex geometric polyhedron with totally geodesic faces, and the boundary of \mathcal{D} projects to the cut locus of x_0 in M , which gives a spine for M .

Now \widetilde{M} is tiled by copies of the fundamental domain \mathcal{D} under the action of the covering transformation group. The geometric dual to this tiling is the Delaunay cell decomposition \mathcal{C} of \widetilde{M} into convex polyhedra with vertices at the points x_1, x_2, \dots . In particular, each edge of \mathcal{C} is a geodesic in \widetilde{M} joining two distinct lifts of x_0 . Then \mathcal{C} projects to a decomposition of M into convex geometric polyhedra with a single vertex x_0 . We can now repeat the argument from the proof of Theorem 7.1 to obtain a subdivision into topological tetrahedra giving a strongly essential triangulation.

In the case of constant positive curvature $+1$, if the diameter is strictly less than π , then the Dirichlet domain lies in a ball of radius at most $\frac{\pi}{2}$. In this case, the Dirichlet domain is strictly convex and so the same argument works as in the constant zero or negative curvature cases. \square

Remark 7.4. — It seems likely that both Theorem 7.1 and Theorem 7.3 should also apply in higher dimensions. If there are no coincidences of geometry then the Epstein-Penner decomposition consists only of simplices of the appropriate dimension, and then, by the arguments above, this triangulation is strongly essential. In the non-generic case we would need to subdivide convex ideal hyperbolic polytopes into ideal simplices and bridge between incompatible induced triangulations of the polytopes’ facets. The former is easy enough, either using coning constructions or regular triangulations (as used in [11], see also [10]), but it isn’t immediately obvious how to do the latter by layering simplices.

Similar observations apply to the cut locus and its dual cell decomposition.

Example 7.5 (Lens spaces). — It is shown by Anisov ([1], Theorems 2.4 and 3.1) and Guéritaud ([13], proof of Theorem 3) that we get the minimal layered triangulation of a lens space from the above construction, i.e. the layered triangulations with the minimal number of tetrahedra (see [18]). In this case, the cut locus of a generic point is the spine dual to the minimal layered triangulation of $L(p, q)$, unless $q = \pm 1 \pmod p$, in which case one needs to perturb the cut locus slightly.

Remark 7.6. — Ehrlich-Im Hof [8, p650] mention examples of Riemannian metrics with no conjugate points on manifolds of dimension ≥ 3 for which the dual triangulation to the cut locus is not strongly essential.

8. Algorithms

The previous sections have given various constructions to obtain essential or strongly essential triangulations. We now show that there are algorithms to decide whether a given triangulation of a closed 3-manifold is essential or strongly essential, based on the fact that the word problem is uniformly solvable in the class of fundamental groups of compact, connected 3-manifolds. Similarly, we show that there is an algorithm to test whether an ideal triangulation is essential based on the fact that the subgroup membership problem is solvable in the same generality. However, to test whether an ideal triangulation is strongly essential, we require the double coset membership problem to be solvable for pairs of peripheral subgroups—here, we restrict the discussion to the most interesting classes of manifolds.

We address the closed case in §8.1 and the ideal case in §8.2. In each, we first determine group theoretic decision problems that are equivalent to deciding whether a triangulation is essential or strongly essential, and then determine conditions on the manifold that ensure that the problems can be

solved. The reader is referred to the excellent survey of decision problems for 3–manifold groups by Aschenbrenner, Friedl and Wilton [2] for definitions and further references for the results we use.

8.1. The closed case

First suppose that \widehat{M} is a closed, connected 3–manifold with 1–vertex triangulation, and write $M = \widehat{M}$. Then a finite presentation for $\pi_1(M, v)$, where v is the vertex, is obtained from the triangulation by arbitrarily orienting the edges. Every generator of the presentation corresponds to an (oriented) edge loop γ , and every relator corresponds to a face. The triangulation is essential if and only if for each edge loop γ , the corresponding element in the fundamental group is non-trivial, written $[\gamma] \neq 1 \in \pi_1(M, v)$. It therefore suffices to solve the word problem for each generator.

The triangulation is strongly essential if, in addition, for any two (oriented) edge loops γ and δ , we have $[\gamma] \neq [\delta]$ and $[\gamma] \neq [\delta]^{-1}$. This, again, reduces to solving the word problem.

We conclude that there is an algorithm to test whether a triangulation of M is essential or strongly essential if there is a solution to the word problem in $\pi_1(M)$.

The Geometrisation Theorem of Thurston and Perelman together with work of Hempel [15] implies that 3–manifold groups are residually finite. The word problem in the larger class of finitely presented, residually finite groups has been known to be uniformly solvable since the 1960s due to work of Dyson [5] and Mostowski [25]. The existence of the desired algorithms therefore follows in full generality from the following:

THEOREM 8.1. — *The Word Problem is uniformly solvable in the class of fundamental groups of compact, connected 3–manifolds.*

COROLLARY 8.2. — *Given a 1–vertex triangulation of a closed, connected 3–manifold, there is an algorithm to test whether it is essential, and if so, whether it is strongly essential.*

8.2. The ideal case

The main ingredient to test whether an ideal triangulation is essential is the following:

THEOREM 8.3 (FRIEDL AND WILTON [9]). — *The Membership Problem is uniformly solvable in the class of fundamental groups of compact, connected 3–manifolds.*

COROLLARY 8.4. — *Given an ideal triangulation of the interior of a compact, connected 3-manifold, there is an algorithm to test whether it is essential.*

Proof. — We are given $M = \widehat{M} \setminus \widehat{M}^{(0)}$, where \widehat{M} is a triangulated 3-dimensional closed pseudo-manifold (with an arbitrary number of vertices). Taking the second barycentric subdivision gives the compact core M^c as described in Section 2, and M^c has a triangulation arising from the second barycentric subdivision of $\widehat{M}^{(0)}$.

Given any edge path γ of \widehat{M} with the same initial and terminal vertex, take the intersection of γ with M^c and connect the endpoints a and b of $\gamma \cap \partial M^c$ by an arc contained in the 1-skeleton of ∂M^c to give a loop γ^c in M^c based at a (this can be done in an algorithmic fashion). The triangulation of M^c yields a finite presentation of $\pi_1(M^c, a)$ as well as a finite generating set for $H = \text{im}(\pi_1(B, a) \rightarrow \pi_1(M^c, a))$, where B is the component of ∂M^c containing a . One now needs to test whether $[\gamma^c] \in H$. The ideal triangulation of M is essential if and only if the answer to this membership test is negative for each edge γ of \widehat{M} with the same initial and terminal vertex. \square

Our test whether a triangulation is strongly essential requires a solution to the so-called *Double Coset Membership Problem*. A class of groups for which this can be solved is provided in the discussion after Corollary 4.16 in [2]. There it is argued that given the fundamental group of a compact Seifert fibred space or a topologically finite hyperbolic 3-manifold, the product of any two finitely generated subgroups is separable. In the case of Seifert fibred manifolds, this is due to Niblo [26] building on work of Scott [31]. In particular, the following result holds by applying the argument of the proof of Lemma 4.3 in [2].

PROPOSITION 8.5. — *The Double Coset Membership Problem is uniformly solvable for pairs of finitely generated subgroups in the class of fundamental groups of compact Seifert fibred spaces or topologically finite hyperbolic 3-manifolds.*

COROLLARY 8.6. — *Given an ideal triangulation of a topologically finite hyperbolic 3-manifold or the interior of a compact Seifert fibred space with boundary, there is an algorithm to test whether it is strongly essential.*

Proof. — We follow verbatim the proof of Corollary 8.4, where M is now known to be a topologically finite hyperbolic 3-manifold or the interior of a compact Seifert fibred space with boundary. If the triangulation is not essential, it is also not strongly essential.

Hence suppose the algorithm shows the triangulation to be essential. To test whether the triangulation is strongly essential, we need to perform the following additional test. Suppose γ and δ are edge paths in \widehat{M} , with γ running from vertex $i(\gamma) = v_1$ to vertex $t(\gamma) = v_2$, and δ running from $i(\delta) = v_2$ to $t(\delta) = v_1$. Again consider the intersections of γ and δ with M^c and denote the first point of intersection of γ with ∂M^c by $\gamma(a_\gamma)$ and the second by $\gamma(b_\gamma)$. Similarly for δ .

Then choose a path ρ_2 on ∂M^c from $\gamma(b_\gamma)$ to $\delta(a_\delta)$ and a path ρ_1 from $\delta(b_\delta)$ to $\gamma(a_\gamma)$. Then $\gamma' = (\gamma|_{[a_\gamma, b_\gamma]}) \star \rho_2 \star (\delta|_{[a_\delta, b_\delta]}) \star \rho_1$ is a loop in the 1-skeleton of M^c based at $\gamma(a_\gamma)$. Denote the boundary component of M^c containing $\gamma(a_\gamma)$ by B_1 and the component containing $\gamma(b_\gamma)$ by B_2 (they are possibly the same). Then let $H_1 = \text{im}(\pi_1(B_1, \gamma(a_\gamma)) \rightarrow \pi_1(M^c, \gamma(a_\gamma)))$ and $H_2 = \text{im}(\pi_1(B_2, \gamma(b_\gamma)) \rightarrow \pi_1(M^c, \gamma(a_\gamma)))$, where for the second inclusion map the path $\gamma|_{[a_\gamma, b_\gamma]}$ joining $\gamma(a_\gamma)$ and $\gamma(b_\gamma)$ is chosen.

Lemma 3.7 implies that γ and $-\delta$ are related by an admissible path homotopy if and only if $[\gamma'] \in H_2H_1$. So we have reduced the problem of deciding whether there is an admissible path homotopy between γ and $-\delta$ to testing whether a certain, algorithmically constructed loop represents an element in the double coset H_2H_1 . Since the groups H_2 and H_1 are finitely generated subgroups of $\pi_1(M)$, our hypotheses on M allow us to solve this problem algorithmically (Proposition 8.5).

For any two edge paths γ, δ in the triangulation, we now need to apply the above algorithm to γ and δ if $i(\gamma) = t(\delta)$ and $i(\delta) = t(\gamma)$; and to γ and $-\delta$ if $i(\gamma) = t(-\delta)$ and $i(-\delta) = t(\gamma)$. Then the triangulation is strongly essential if and only if none of the constructed loops represents an element in the associated double coset. \square

We will now bootstrap the above results to obtain a more general result.

COROLLARY 8.7. — *Given an ideal triangulation of the interior of a compact, irreducible, orientable, connected 3-manifold with incompressible boundary consisting of a disjoint union of tori, there is an algorithm to test whether it is strongly essential.*

Proof. — We first determine whether the triangulation is essential using Corollary 8.4. Suppose it is. The next step determines a system of tori giving the JSJ decomposition of the 3-manifold M . This is done using work of Jaco and Tollefson [20] as follows. The compact core M^c is triangulated by construction. Algorithm 8.2 of [20] determines the collection of JSJ tori (noting that annuli do not arise because the ends of M are all tori) as normal surfaces with respect to the finer triangulation. In particular, any edge in

the ideal triangulation of M meets these JSJ tori transversely. The JSJ tori divide the tetrahedra in M^c into polytopes, and this finer cell decomposition of M^c gives natural inclusion maps of the fundamental groups of the JSJ tori and the JSJ pieces into $\pi_1(M)$ (using appropriately chosen basepoints in the 0-skeleton and paths in the 1-skeleton).

Suppose, as above, that γ and δ are edge paths with respect to the initial triangulation of \widehat{M} , with γ running from vertex v_1 to vertex v_2 , and δ running from v_2 to v_1 . Denote the intersections of γ and δ with M^c with the same symbols, and as above denote the boundary component linking v_k by B_k . Choose an orientation for each torus T in the JSJ decomposition, and check whether the algebraic intersection number of γ with T is the negative of the intersection number of δ with T . If the answer is no for any JSJ torus, then there is no admissible homotopy between γ and $-\delta$. Otherwise, we will first analyse the intersections of the edges with the JSJ pieces M_j .

We start with some preliminary observations. The conditions determining the properties essential and strongly essential with respect to the compact core apply to the intersection of γ (resp. δ) with each piece in the JSJ decomposition of M^c . We will count these components of γ starting at the boundary component B_1 of M_c , and write $\gamma = \gamma_1 \star \gamma_2 \star \dots$ with $\gamma_j \subset M_j$, where M_j is a JSJ piece (these pieces are not necessarily distinct). Note that γ_j may have endpoints on the same JSJ torus, but in this case, they are distinct as otherwise γ_j would be a simple closed loop in M^c .

If the connected component γ_j is not essential in M_j , then it can be homotoped into ∂M_j . Using the methods of [2], this can be done algorithmically using the solution of the word problem in $\pi_1(M_j)$, where we choose appropriate basepoints in the 0-skeleta and generating sets as well as paths for the inclusion maps in the 1-skeleta. Moreover, using the chosen generating sets for the JSJ tori, one can algorithmically replace γ_j by the corresponding path τ in the 1-skeleton of the JSJ tori. Let γ^* denote the path obtained from γ by replacing γ_j with τ . Then γ and γ^* are homotopic keeping their endpoints fixed, and we define the components of γ^* to be the components of γ , except that $\gamma_{j-1} \star \tau \star \gamma_{j+1}$ is counted as one component since in this case $M_{j-1} = M_{j+1}$ and $\tau \subset \partial M_{j-1}$. Hence we decrease the total number of components by two. Moreover, the intersection of γ^* with the JSJ tori is still transverse at the endpoints of each component of γ^* (even though γ^* is no longer transverse to them). We then repeat the above step with γ^* , and denote the resulting properly embedded arc in M^c again by γ^* , until each component of γ^* is essential. Similarly for δ , resulting in δ^* .

If there is an admissible homotopy taking γ to $-\delta$ in \widehat{M} , then there is an admissible homotopy between γ^* and $-\delta^*$ in M^c and they have the same number of components. Moreover, the k th component of γ^* has endpoints on the same tori and with the same respective signs as the k th component of $-\delta^*$. Hence we check this property of γ^* and δ^* first, and if it is not satisfied, then there is no admissible homotopy taking γ to $-\delta$.

Hence assume that the intersections match up. For each k , we check whether the k th component of γ^* is parallel to the k th component of δ^* in the JSJ piece that contains them. If the answer is no for any k , then there is no admissible homotopy taking γ to $-\delta$. Hence assume that the components are pairwise parallel.

Then there is a homotopy taking γ^* to the union of δ^* with possibly some loops on the JSJ tori inserted. This is because in checking parallelism in the JSJ pieces, we are free to move the ends of arcs on the JSJ tori, which is no longer possible when considering the edges in M^c . However, choosing basepoints appropriately, the Double Coset Membership Test in each JSJ piece will return the corresponding elements in the JSJ tori. Since for each k the k th components are parallel, we can choose the same paths (but with opposite orientation) on the JSJ tori to connect the terminal points of a pair of parallel components as well as the initial points of the next pair of parallel components, and hence write the loop $\gamma^* \star \rho_1 \star \delta^* \star \rho_2$ based at a_γ as a product of elements in the fundamental groups of the JSJ tori and the boundary components B_1 and B_2 (by choosing appropriate base points and inclusion homomorphisms). It then suffices to check that the resulting element is contained in the double coset H_2H_1 as above. There is an admissible homotopy between γ and $-\delta$ if and only if the answer to this test is positive.

Finally, if $v_1 = v_2$ and there is no admissible path homotopy found between γ and $-\delta$, we also need to run the whole algorithm with δ and $-\delta$ interchanged. \square

Bibliography

- [1] ANISOV (S.). — Cut loci in lens manifolds. *C. R. Math. Acad. Sci. Paris* 342, no. 8, p. 595-600 (2006).
- [2] ASCHENBRENNER (M.), FRIEDL (S.) and WILTON (H.). — Decision problems for 3-manifolds and their fundamental groups, arXiv:1405.6274 (2014).
- [3] BURTON (B. A.), BUDNEY (R.), PETTERSSON (W.) et al. — Regina: Software for 3-manifold topology and normal surface theory, <http://regina.sourceforge.net/>, (1999-2012).

- [4] CULLER (M.), DUNFIELD (N. M.), and WEEKS (J. R.). — SnapPy, a computer program for studying the geometry and topology of 3-manifolds, <http://snappy.computop.org/>.
- [5] DYSON (V. H.). — The word problem and residually finite groups, *Notices Amer. Math. Soc.* 11, p. 734 (1964).
- [6] EPSTEIN (D. B. A.). — Projective planes in 3-manifolds, *Proc. London Math. Soc.* (2) 76, p. 180-184 (1962).
- [7] EPSTEIN (D. B. A.) and PENNER (R. C.). — Euclidean decompositions of noncompact hyperbolic manifolds, *Journal of Differential Geometry* 27, no. 1, p. 67-80 (1988).
- [8] EHRLICH (P.) and IM HOF (H.-C.). — Dirichlet regions in manifolds without conjugate points, *Commentarii Math. Helv.* 54, p. 642-658 (1979).
- [9] FRIEDL (S.) and WILTON (H.). — The membership problem for 3-manifold groups is solvable, arXiv:1401.2648, 2014, to appear in *Algebraic & Geometric Topology*.
- [10] GELFAND (I. M.), KAPRANOV (M. M.), and ZELEVINSKY (A. V.). — Discriminants, resultants, and multidimensional determinants, *Mathematics: Theory & Applications*, Birkhäuser Boston Inc., Boston, MA (1994).
- [11] GAROUFALIDIS (S.), HODGSON (C. D.), RUBINSTEIN (J. H.) and SEGERMAN (H.). — 1-efficient triangulations and the index of a cusped hyperbolic 3-manifold, *Geometry & Topology* 19, p. 2619-2689 (2015).
- [12] GUÉRITAUD (F.). — On canonical triangulations of once-punctured torus bundles and two-bridge link complements, with an appendix by David Futer, *Geometry & Topology* 10, p. 1239-1284 (2006).
- [13] GUÉRITAUD (F.). — Delaunay triangulations of lens spaces, arXiv:0901.2738v2, (2009).
- [14] HATCHER (A.). — Notes on Basic 3-manifold Topology, available from <http://www.math.cornell.edu/hatcher/3M/3Mdownloads.html>.
- [15] HEMPEL (J.). — Residual finiteness for 3-manifolds, in: S. M. Gersten and John R. Stallings (eds.), *Combinatorial Group Theory and Topology*, pp. 379-396, *Ann. of Math. Stud.*, vol. 111, Princeton Univ. Press, Princeton, NJ (1987).
- [16] HODGSON (C. D.), RUBINSTEIN (J. H.) and SEGERMAN (H.). — Triangulations of hyperbolic 3-manifolds admitting strict angle structures, *Journal of Topology*. 5, p. 887-908 (2012).
- [17] JACO (W.) and RUBINSTEIN (J. H.). — 0-efficient triangulations of 3-manifolds, *J. Differential Geometry* 65, p. 61-168 (2003).
- [18] JACO (W.) and RUBINSTEIN (J. H.). — Layered-triangulations of 3-manifolds, arXiv:math/0603601.
- [19] JACO (W.), RUBINSTEIN (J. H.) and TILLMANN (S.). — Minimal triangulations for an infinite family of lens spaces, *Journal of Topology* 2 p. 157-180 (2009).
- [20] JACO (W.) and TOLLEFSON (J. L.). — Algorithms for the complete decomposition of a closed 3-manifold, *Illinois J. Math.* 39, no. 3, p. 358-406 (1995).
- [21] KANG (E.) and RUBINSTEIN (J. H.). — Ideal triangulations of 3-manifolds II; taut and angle structures, *Algebraic & Geometric Topology* 5, p. 1505-1533 (2005).
- [22] LACKENBY (M.). — Taut ideal triangulations of 3-manifolds, *Geom. Topol.* 4, p. 369-395 (2000).
- [23] LUO (F.) and TILLMANN (S.). — Angle structures and normal surfaces, *Trans. Amer. Math. Soc.* 360, no. 6, p. 2849-2866 (2008).
- [24] MATVEEV (S.). — Algorithmic topology and classification of 3-manifolds. *Algorithms and Computation in Mathematics*, 9. Springer, Berlin (2003).

- [25] MOSTOWSKI (A. W.). — On the decidability of some problems in special classes of groups, *Fund. Math.* 59, p. 123-135 (1966).
- [26] NIBLO (G. A.). — Separability properties of free groups and surface groups, *J. Pure Appl. Algebra* 78, p. 77-84 (1992).
- [27] RAMANAYAKE (S.). — 0-efficient triangulations of Haken three-manifolds, Ph.D. thesis, submitted, University of Melbourne (2015).
- [28] RATCLIFFE (J.). — *Foundations of Hyperbolic Manifolds*, Graduate Texts in Mathematics, Springer, Second Edition (2006).
- [29] RUBINSTEIN (J. H.). — One-sided Heegaard splittings of 3-manifolds, *Pacific J. of Math.* 76, p. 185-200 (1978).
- [30] SCHLEIMER (S.). — Almost normal Heegaard splittings, Ph.D. Thesis, University of California, Berkeley. 2001. 81 pp.
Available from <http://homepages.warwick.ac.uk/~masgar/index.html>.
- [31] SCOTT (P.). — Subgroups of surface groups are almost geometric, *J. London Math. Soc.* 17, p. 555-565 (1978).
- [32] SEGERMAN (H.) and TILLMANN (S.). — Pseudo-developing maps for ideal triangulations I: Essential edges and generalised hyperbolic gluing equations, *Topology and Geometry in Dimension Three: Triangulations, Invariants, and Geometric Structures* (Proceedings of the Jacofest conference), *AMS Contemporary Mathematics* 560, p. 85-102 (2011).
- [33] SEIFERT (H.) and THRELFALL (W.). — *Lehrbuch der Topologie*, AMS Chelsea Publishing, 31, (1934) (reprinted 2003).
- [34] STALLINGS (J. R.). — On fibering certain 3-manifolds, *Topology of 3-manifolds and related topics* (Proc. The Univ. of Georgia Institute, 1961), Prentice Hall, p. 95-100
- [35] STOCKING (M.). — Almost normal surfaces in 3-manifolds, *Trans. Amer. Math. Soc.* 352, no. 1, p. 171-207 (2000).
- [36] THOMPSON (A.). — Thin position and the recognition problem for S^3 , *Math. Res. Lett.* 1, no. 5, p. 613-630 (1994).
- [37] WALDHAUSEN (F.). — Heegaard-Zerlegungen der 3-Sphäre, *Topology* 7 no. 2, p. 195-203 (1968).
- [38] WALDHAUSEN (F.). — On irreducible 3-manifolds which are sufficiently large, *Annals of Math.* 87, p. 56-88 (1968).