Zeitschrift: L'Enseignement Mathématique
Herausgeber: Commission Internationale de l'Enseignement Mathématique
Band: 54 (2008)
Heft: 3-4

Artikel: Normal surfaces in topologically finite 3-manifolds
Autor: Tillmann, Stephan
DOI: https://doi.org/10.5169/seals-109942

L'Enseignement Mathématique (2) 54 (2008), 329–380

NORMAL SURFACES IN TOPOLOGICALLY FINITE 3-MANIFOLDS

by Stephan TILLMANN *)

INTRODUCTION

Spun-normal surfaces made their first appearance in unpublished work by Thurston. He described essential surfaces in the figure eight knot complement as spun-normal surfaces with respect to the ideal triangulation by two regular ideal hyperbolic tetrahedra. There have since been interesting applications of spun-normal surfaces; for instance in work by Weeks [14], Kang [8], Kang and Rubinstein [9]. They are also part of the repertoire of many people working in the field. There does not however appear to be a basic reference for fundamental properties of these surfaces. This paper attempts to fill this gap.

Let M be the interior of a compact 3-manifold with non-empty boundary. Then M admits an ideal triangulation and this gives the end-compactification of M the structure of a triangulated, closed pseudo-manifold P (see Proposition 1.2). The complement of the 0-skeleton in P is identified with M, and spun-normal surfaces in M can be studied using the triangulation of P. The remainder of this introduction restricts to this setting; Figure 1 shows a spun-normal surface in a pseudo-manifold with boundary and indicates the more general situation.

A spun-normal surface meets each 3-simplex in a union of pairwise disjoint normal discs such that there are at most finitely many normal quadrilaterals but possibly infinitely many normal triangles which accumulate at the 0-skeleton of P. A spun-normal surface is therefore properly embedded in M. Two spun-normal surfaces are regarded as equivalent if there is an isotopy taking one to the other whilst leaving all simplices invariant (normal

*) This work was supported by a Postdoctoral Fellowship by the CRM/ISM in Montréal and under the Australian Research Council's Discovery funding scheme (project number DP0664276).

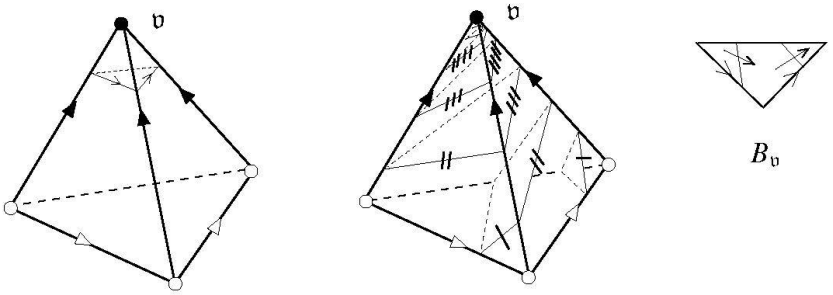

FIGURE 1

A spun-normal surface in a non-orientable pseudo-manifold with boundary. The triangulation has a single 3-simplex with two faces identified as indicated by the arrows. There are two vertices; the link of v is a Möbius band, and the link of the other vertex is a disc. The shown spun-normal surface spins into v; it is a disc minus a point on its boundary and it meets a suitable vertex linking surface, B_v, in a 2-sided, non-separating arc.

isotopy). One may choose a small open neighbourhood of each 0-simplex v in P such that its boundary, B_v, is a closed (possibly non-orientable) normal surface. The topology of spun-normal surfaces is analysed in Section 1, and the main observations are the following:

1. A spun-normal surface without vertex linking components is up to normal isotopy uniquely determined by its quadrilateral discs (see Lemma 1.19).
2. If two spun-normal surfaces meet every 3-simplex in quadrilateral discs of the same type, then their *geometric sum* is well-defined (see Lemma 1.28).
3. For each spun-normal surface S and each 0-simplex v there is a well-defined element $\partial_v(S) \in H_1(B_v; \mathbf{Z})$ which determines S in a neighbourhood of v uniquely up to normal isotopy (see Lemma 1.33 and Lemma 1.35).
4. A spun-normal surface S is topologically finite unless there is some 0-simplex v such that $\partial_v(S) \neq 0$ and $\chi(B_v) < 0$ (see Corollary 1.36).
5. (Kneser–Haken finiteness) If S is a spun-normal surface no two components of which are normally isotopic, then the number of components of S which are not vertex linking is at most $12t$, where t is the number of 3-simplices in the triangulation (see Lemma 1.37 which includes a stronger bound).

Section 2 analyses algebraic properties of a spun-normal surface S derived from the quadrilaterals in its cell structure. Recording the number of quadrilaterals of each type gives the *normal Q-coordinate of S*. The *Q-matching equations* described by Tollefson [12] for compact 3-manifolds result from the fact that a properly embedded surface meets a small regular neighbourhood of an edge in a collection of discs, and that each of these discs is uniquely determined by its intersection with the boundary of the

neighbourhood. The equations are shown to give a necessary and sufficient condition on a collection of quadrilateral discs to be (up to normal isotopy) contained in the cell structure of some spun-normal surface. A spun-normal surface without vertex linking components is thus uniquely determined by its normal Q-coordinate (see Theorem 2.4).

The triangulation of P is denoted by $\mathcal{T}$ and the solution space of the Q-matching equations by $Q(\mathcal{T})$. This is a real vector subspace of $\mathbf{R}^{3t}$. The projective solution space $PQ(\mathcal{T})$ is the intersection of $Q(\mathcal{T})$ with the unit simplex; it is a convex rational polytope. Rational points are therefore dense in $PQ(\mathcal{T})$ and correspond to the projective classes not merely of spun-normal surfaces (which are embedded), but also of *immersed* and *branched immersed spun-normal surfaces* (see Proposition 2.7).

Section 3 introduces the *boundary curve map*

$$\partial\colon Q(\mathcal{T}) \to \oplus_{\mathfrak{v}} H_1(B_{\mathfrak{v}}; \mathbf{R}),$$

which is linear and generalises the above map $\partial_{\mathfrak{v}}$. It is thus possible to compute $\partial_{\mathfrak{v}}(S)$ directly from the normal Q-coordinate of S. The definition of ∂ is motivated by the boundary map used by Kang and Rubinstein [9] in the case where each vertex linking surface is a torus or Klein bottle, and their methods are generalised to prove the following:

THEOREM 0.1. *Let P be a closed 3-dimensional pseudo-manifold with triangulation $\mathcal{T}$. Then*

1. *$Q(\mathcal{T})$ has dimension $v_o - e + 3t = \chi(P) + 2t - v_n$;*
2. *$\partial\colon Q(\mathcal{T}) \to \oplus_{\mathfrak{v}} H_1(B_{\mathfrak{v}}; \mathbf{R})$ is onto, and its restriction to integer lattice points in $Q(\mathcal{T})$ has image of finite index in $\oplus_{\mathfrak{v}} H_1(B_{\mathfrak{v}}; \mathbf{Z})$; and*
3. *$PQ(\mathcal{T})$ is non-empty and hence of dimension $\dim Q(\mathcal{T}) - 1$.*

Here t, e, v is the number of 3-simplices, 1-simplices, 0-simplices respectively; so $\chi(P) = v - e + t$. Moreover, $v = v_o + v_n$, where v_n is the number of 0-simplices with non-orientable linking surface.

Note that $\chi(P) = 0 = v_n$ if P is a manifold, and $\chi(P) = v$ if the link of each vertex is a torus or Klein bottle. The above dimension of $Q(\mathcal{T})$ and target of ∂ correct the main result of [9] in the case of pseudo-manifolds with non-orientable vertex links.

The set of projective classes of (embedded) spun-normal surfaces has a natural compactification, $PF(\mathcal{T}) \subset PQ(\mathcal{T})$, points of which correspond to projective classes of certain transversely measured singular codimension-one foliations of the complement of the 0-skeleton, M. The restriction

$\partial\colon PF(\mathcal{T}) \to \oplus_{\mathfrak{v}} H_1(B_{\mathfrak{v}}; \mathbf{R})$ determines the (possibly singular) foliation of each $B_{\mathfrak{v}}$. If P is orientable, then the intersection pairing on $\oplus_{\mathfrak{v}} H_1(B_{\mathfrak{v}}; \mathbf{R})$ defines a bi-linear skew-symmetric form on $Q(\mathcal{T})$; for any $N, L \in Q(\mathcal{T})$:

$$\langle N, L \rangle = \sum \iota\Big(\partial_{\mathfrak{v}}(N), \partial_{\mathfrak{v}}(L)\Big).$$

This is used to analyse the structure of the components of $PF(\mathcal{T})$.

THEOREM 0.2. *Let P be a closed 3-dimensional pseudo-manifold with triangulation $\mathcal{T}$. The set $PF(\mathcal{T})$ is a finite union of convex rational polytopes. Each maximal convex polytope R in this union satisfies* $\dim R \leq t - 1$. *Moreover, if P is orientable then:*

$$\chi(P) - 1 \leq \dim R \leq \chi(P) + \dim(R \cap \ker \partial), \tag{0.1}$$

and if P is non-orientable then:

$$\chi(P) - v_n - 1 \leq \dim R \leq 2\chi(P) - v_n + \dim(R \cap \ker \partial), \tag{0.2}$$

where $\dim \varnothing = -1$ *throughout.*

The lower bounds in Equations (0.1) and (0.2) are shown to be sharp by the examples given in Section 4, the complement of the figure eight knot and the Gieseking manifold.

This article does not address many topics which are standard in algorithmic topology. For instance, work by Jaco and Oertel [3] implies that there is an algorithm to determine whether a topologically finite manifold is *large*, i.e. whether it contains a closed, incompressible, 2-sided surface distinct from a 2-sphere or a vertex linking surface. The extension of such algorithms to spun-normal surfaces (for instance, to determine whether there are essential discs or annuli) is left for future research; it presupposes the ability to put incompressible surfaces into spun-normal form and to decide whether a spun-normal surface is incompressible — neither of which is addressed in this paper.

ACKNOWLEDGEMENTS. The author thanks Daryl Cooper for many enlightening conversations about spun-normal surfaces. He also thanks Steven Boyer, Craig Hodgson and Hyam Rubinstein for helpful comments. Thanks to Saul Schleimer for pointing out that there does not seem to be a complete treatment of the geometric sum operation in the literature.

CONTENTS

1. SPUN-NORMAL SURFACES

This section defines spun-normal surfaces in a 3-dimensional pseudo-manifold (possibly with boundary) and analyses their basic properties. It is also shown that a topologically finite 3-manifold is homeomorphic to the complement of the 0-skeleton in a closed, triangulated pseudo-manifold.

1.1 PSEUDO-MANIFOLD

Let $\widetilde{\Delta} = \{\widetilde{\Delta}^3_1, \ldots, \widetilde{\Delta}^3_t\}$ be a disjoint union of 3-simplices, each of which is a regular Euclidean 3-simplex of edge length one. The subsimplices of the elements in $\widetilde{\Delta}$ are referred to as simplices in $\widetilde{\Delta}$. Let Φ be a collection of Euclidean isometries between 2-simplices in $\widetilde{\Delta}$ with the property that for each 2-simplex $\widetilde{\Delta}^2$ there is at most one isometry having $\widetilde{\Delta}^2$ as its range or domain (but not both). There is a natural quotient map $p\colon \widetilde{\Delta} \to \widetilde{\Delta}/\Phi$. Let $P = \widetilde{\Delta}/\Phi$ and note that P inherits a singular PL cone-metric. Since p restricts to an embedding of the interior of every 2-simplex and every 3-simplex, the only possible non-manifold points in P are at images of vertices and barycentres of 1-simplices.

DEFINITION 1.1 (Pseudo-manifold). The quotient space $P = \widetilde{\Delta}/\Phi$ is a *pseudo-manifold* if p restricts to an embedding of the interior of every 1-simplex in $\widetilde{\Delta}$.

A pseudo-manifold is *closed* if each 2-simplex in $\widetilde{\Delta}$ is the range or domain of an element in Φ, otherwise it is a *pseudo-manifold with boundary*. A pseudo-manifold is *orientable* if all 3-simplices in $\widetilde{\Delta}$ can be oriented such that all elements in Φ are orientation reversing. For the remainder of this paper, it is assumed that P is a pseudo-manifold and a quotient map $p\colon \widetilde{\Delta} \to P$ is fixed; additional hypotheses will be added.

1.2 TRIANGULATION

The image under p of an n-simplex in $\widetilde{\Delta}$ is called an *n-singlex* in P. The resulting combinatorial cell-decomposition of P is termed a *triangulation* of P and denoted by $\mathcal{T} = (\widetilde{\Delta}, \Phi, p)$. If p restricts to an embedding of each simplex, then $\mathcal{T}$ is a triangulation in the traditional sense; it will then be referred to as a *simplicial triangulation*. The image in P of the set of i-simplices in $\widetilde{\Delta}$ is denoted by $P^{(i)}$ and termed the *i-skeleton of P*.

The number of 3-singlices in P is t, and we denote by f the number of 2-singlices, e the number of 1-singlices and v the number of 0-singlices.

1.3 IDEAL TRIANGULATION

A manifold is *topologically finite* if it is homeomorphic to the interior of a compact 3-manifold. If P is closed, then removing the 0-skeleton yields a topologically finite 3-manifold $M = P \setminus P^{(0)}$ and P is called the *end-compactification* of M. This motivates the following terminology.

An *ideal i-simplex*, $i \in \{1,2,3\}$, is an i-simplex with its vertices removed; the vertices of the i-simplex are referred to as the *ideal vertices* of the ideal i-simplex. Similarly for singlices. The restriction $(\widetilde{\Delta} \setminus \widetilde{\Delta}^{(0)}, \Phi|_{\widetilde{\Delta}^{(2)} \setminus \widetilde{\Delta}^{(0)}}, p|_{\widetilde{\Delta} \setminus \widetilde{\Delta}^{(0)}})$ of the triangulation of P to $P \setminus P^{(0)}$ is an *ideal triangulation* of $P \setminus P^{(0)}$. A 0-simplex of P is referred to as an *ideal vertex* of the ideal triangulation or of M.

If M is the interior of a compact 3-manifold $\overline{M}$ with non-empty boundary and M is homeomorphic to $P \setminus P^{(0)}$ for some pseudo-manifold P, then M is said to admit an ideal triangulation. The following result is implicit in Matveev [10] and shows that every topologically finite 3-manifold arises in this way from a closed pseudo-manifold:

PROPOSITION 1.2 (Topologically finite manifold has ideal triangulation). *If M is the interior of a compact 3-manifold $\overline{M}$ with non-empty boundary, then M admits an ideal triangulation. The ideal vertices of the ideal triangulation are in one-to-one correspondence with the boundary components of $\overline{M}$.*

Proof. This derives from the following results in [10]. Theorem 1.1.13 due to Casler asserts that $\overline{M}$ possesses a *special spine* Σ; it has the property that $\overline{M}$ is homeomorphic to a regular neighbourhood of Σ in M. Theorem 1.1.26 implies that Σ is dual to an ideal triangulation of M with the property that the ideal vertices are in one-to-one correspondence with the boundary components of $\overline{M}$. □

1.4 NORMAL DISCS, ARCS AND CORNERS

A *normal corner* is an interior point of a 1-simplex. A *normal arc* is a properly embedded straight line segment on a 2-simplex with boundary consisting of normal corners.

A *normal disc* is a properly embedded disc in a 3-simplex whose boundary consists of normal arcs no two of which are contained on the same face of the 3-simplex; moreover, the normal disc is the cone over its boundary with cone point the barycentre of its normal corners. It follows that the boundary of a normal disc consists of either three or four normal arcs, and it is accordingly called a *normal triangle* or a *normal quadrilateral*. A normal disc is uniquely determined by its intersection with the 1-skeleton.

A *normal isotopy* is an isotopy of a simplex which preserves all its subsimplices. There are exactly three normal isotopy classes of normal arcs in a 2-simplex. A normal arc is called *$\mathfrak{v}$-type* if it separates the vertex $\mathfrak{v}$ from the other two vertices of the 2-simplex. There are exactly seven normal

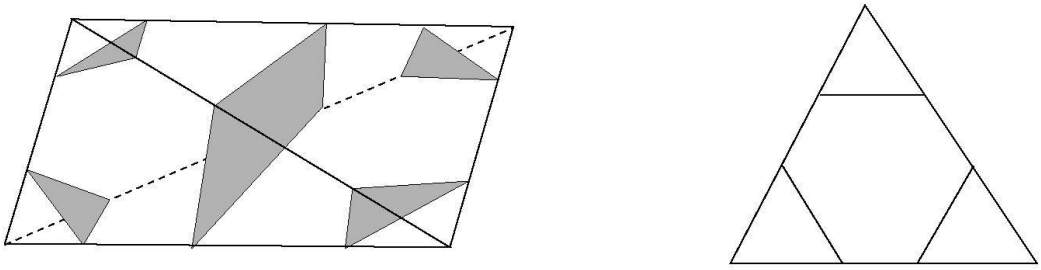

FIGURE 2
Normal discs in a 3-simplex and normal arcs in a 2-simplex

isotopy classes of normal discs in a 3-simplex (four of normal triangles and three of normal quadrilaterals). Given a normal triangle $\mathfrak{t}$ in a 3-simplex, there is a unique complementary region containing precisely one 0-simplex $\mathfrak{v}$; $\mathfrak{t}$ is said to be *dual* to $\mathfrak{v}$.

A piecewise linear arc in a 2-simplex is termed an *elementary arc* if it is normally isotopic to a normal arc. A piecewise linear disc in a 3-simplex is termed an *elementary disc* if it is normally isotopic to a normal disc. The following lemma provides a useful normal isotopy for any collection of elementary cells.

LEMMA 1.3 (Straightening). *Given an arbitrary collection of pairwise disjoint elementary arcs in a 2-simplex, there is a normal isotopy keeping the 1-skeleton fixed and taking each elementary arc to a normal arc. Given an arbitrary collection of pairwise disjoint elementary discs in a 3-simplex, there is a normal isotopy keeping the 1-skeleton fixed and taking each elementary disc to a normal disc.*

Proof. Let Δ^2 be a 2-simplex containing an arbitrary collection of pairwise disjoint elementary arcs. There is a homeomorphism $\Delta^2 \to \Delta^2$ which fixes the boundary and takes each elementary arc to a normal arc. The first part of the lemma now follows from the fact that a homeomorphism of a closed ball which fixes the boundary is isotopic to the identity.

Given an arbitrary collection of pairwise disjoint elementary discs in a 3-simplex Δ^3, there is a homeomorphism $h\colon \Delta^3 \to \Delta^3$ which fixes the 1-skeleton and takes each elementary disc to a normal disc. The first part of the lemma shows that the restriction of h to $\partial\Delta^3$ is isotopic to the identity. One can therefore isotope h in a neighbourhood of $\partial\Delta^3$ such that the boundary of each elementary disc is straight and the result follows as above. □

The image under p of a normal disc (respectively arc, corner) in $\widetilde{\Delta}$ is called a *normal disc (respectively arc, corner) in* P. Note that a normal disc in P may have fewer arcs or corners in its boundary than its pre-image in $\widetilde{\Delta}$. Two normal discs in P are said to be *of the same type* if they are the images of normally isotopic normal discs in $\widetilde{\Delta}$. A *normal isotopy of* P is an isotopy of P which leaves all singlices invariant.

The above definitions apply verbatim to define normal cells in ideal simplices and ideal singlices, and normal isotopies in $P \setminus P^{(0)}$.

1.5 Normal surfaces and normal curves

This subsection collects some well-known facts about normal surfaces; the reader should consult Jaco and Rubinstein [4] and Thompson [11] for details and proofs.

Definition 1.4 (Normal surface). A subset S of P is a *normal surface* in P if it meets every 3-singlex in P in a (possibly empty) finite union of pairwise disjoint normal discs.

Note that a normal surface is a properly embedded (not necessarily connected) compact surface in $(P \setminus P^{(0)}) \subset P$.

Lemma 1.5. *Let A be a finite union of pairwise disjoint normal arcs in Δ^2. Then A is uniquely determined by its intersection with the 1-skeleton up to normal isotopy.*

A curve on a compact, triangulated surface is *normal* if it meets each 2-singlex in a finite collection of pairwise disjoint normal arcs. (A curve will always be connected, whilst a surface may not be connected.)

Lemma 1.6. *Every non-empty normal curve in $\partial\Delta^3$ is simple and closed. Moreover, it bounds some properly embedded disc in Δ^3.*

Consequences of Lemma 1.5 are the following:

Corollary 1.7. *A finite union of pairwise disjoint normal curves on $\partial\Delta^3$ is uniquely determined by its intersection with the 1-skeleton up to normal isotopy.*

COROLLARY 1.8. *Let S be a normal surface in P. Then S is uniquely determined by its intersection with the 1-skeleton of P up to normal isotopy.*

The *length* of a normal curve is the number of intersections with the 1-skeleton (equivalently: the number of normal arcs it is composed of).

LEMMA 1.9. *Every normal curve γ bounds a disc in $\partial\Delta^3$ containing one or two vertices. In the former case, γ has length three, in the latter, its length is a multiple of four.*

LEMMA 1.10. *If a non-empty normal curve γ does not meet some edge of Δ^3, then γ has length three or four.*

1.6 SPUN-NORMAL SURFACES

The definition of a normal surface is now extended; the following notions are defined with respect to the PL cone-metric on P.

DEFINITION 1.11 (Normal subset). Let P be a pseudo-manifold (possibly with boundary). A subset S of P is *normal* if

1. S intersects each 3-singlex in P in a (possibly empty) countable union of pairwise disjoint normal discs;
2. the set of accumulation points of $S \cap P^{(1)}$ is contained in $P^{(0)}$;
3. if $\{x_i\} \subset S$ has accumulation point $x \in P$, then $x \in S$ or $x \in P^{(0)}$.

Note that any normal subset of P is contained in $P \setminus P^{(0)}$. The following two facts follow from the observation that whenever a normal subset S of P meets a 3-singlex Δ^3 in a point $x \in \partial\Delta^3$, then it meets Δ^3 in a normal disc containing x.

LEMMA 1.12. *Let S be a normal subset of P. Then $p^{-1}(S)$ is a normal subset of $\widetilde{\Delta}$.*

LEMMA 1.13 (Normal subset is a surface). *Let S be a normal subset of P. Then every point on S has a small closed neighbourhood in S which is homeomorphic to a disc.*

For each 0-singlex $\mathfrak{v} \in P^0$ choose a small, open neighbourhood $N_\mathfrak{v}$ with the property that the neighbourhoods are pairwise disjoint and $\overline{\partial N_\mathfrak{v} \setminus \partial P} = B_\mathfrak{v}$

is a normal surface in P. Note that all normal discs in $B_\mathfrak{v}$ are normal triangles no two of which are normally isotopic. A normal subset in P is called a *vertex linking surface* if it is normally isotopic to $B_\mathfrak{v}$ for some 0-singlex $\mathfrak{v}$. Let $P^c = P \setminus \bigcup N_\mathfrak{v}$.

REMARK 1.14. Condition (2) in Definition 1.11 is not redundant. For instance, it rules out that countably many pairwise disjoint copies of $B_\mathfrak{v}$ accumulate on $B_\mathfrak{v}$.

Each surface $B_\mathfrak{v}$ consists of normal triangles and hence inherits an induced triangulation $\mathcal{T}_\mathfrak{v}$. A path in $B_\mathfrak{v}$ is called *normal* if it intersects each triangle in $\mathcal{T}_\mathfrak{v}$ in a (possibly empty) disjoint union of normal arcs.

LEMMA 1.15 (Vertex linking surface). *If S is a normal subset of P consisting only of normal triangles, then each connected component of S is a vertex linking surface.*

Proof. Without loss of generality, it may be assumed that S is connected. Given a normal triangle in S dual to $\mathfrak{v} \in P^{(0)}$, the normal triangles it meets along its boundary arcs are all dual to $\mathfrak{v}$. Whence all normal triangles in S are dual to a fixed 0-singlex $\mathfrak{v}$. Since $S \cap P^c \cap P^{(1)}$ consists of finitely many points, one may (up to normal isotopy) assume that S is contained in $N_\mathfrak{v}$. Since S meets every 3-singlex in a union of pairwise disjoint normal discs, it follows that it meets every 3-singlex incident with $\mathfrak{v}$ in at least one normal triangle contained in $N_\mathfrak{v}$. This forces S to be normally isotopic to $B_\mathfrak{v}$. □

LEMMA 1.16 (Finiteness). *A normal subset S of P contains at most finitely many normal quadrilaterals.*

Proof. If S contains infinitely many normal quadrilaterals, then it contains infinitely many quadrilaterals of the same type. Using condition (2), there are finitely many possibilities for the limit points of the normal corners of these quadrilaterals. In each case, one can construct a sequence of points on S which accumulates on a point not on S but in the interior of a 1-simplex in P, contradicting the third condition. □

LEMMA 1.17. *Every connected component of a normal subset of P is a properly embedded surface in $P \setminus P^{(0)}$.*

Proof. It suffices to assume that S is connected and contains a normal quadrilateral. Then S contains finitely many normal quadrilaterals. If S contains only finitely many normal triangles, then S is a normal surface and therefore a closed, properly embedded surface in $P \setminus P^{(0)}$. Hence assume that S contains infinitely many normal triangles. We may assume that P^c contains all normal quadrilaterals contained in S and that $\partial P^c \cap S$ is transverse. Conditions (2) and (3) in Definition 1.11 imply that $P^c \cap P^{(1)} \cap S$ contains finitely many points. Since normal triangles are flat, it follows that $S \cap \partial P^c$ consists of finitely many pairwise disjoint simple closed curves and arcs. Hence $S \cap P^c$ is a properly embedded surface in P^c. Since $N_{\mathfrak{v}}$ can be made arbitrarily small, the conclusion follows. □

It follows from Lemmas 1.15 and 1.16 that a normal subset may have infinitely many connected components at most finitely many of which are not vertex linking.

DEFINITION 1.18 (Spun-normal surface). A normal subset of P is termed a *spun-normal surface* in P if it has finitely many connected components. If a spun-normal surface contains infinitely many normal triangles dual to $\mathfrak{v} \in P^{(0)}$, then it is said to *spin into* $\mathfrak{v}$. In particular, a normal surface is a spun-normal surface which spins into no vertex.

LEMMA 1.19 (Uniqueness). *Let S_0 and S_1 be normal subsets in P with the property that for any type of normal disc, S_0 and S_1 contain the same (not necessarily finite) number of normal discs of this type. Then S_0 and S_1 are normally isotopic.*

Proof. There is a normal isotopy in $\widetilde{\Delta}$ which takes $p^{-1}(S_0)$ to $p^{-1}(S_1)$ and, in particular, takes the intersection of $p^{-1}(S_0)$ with the 1-skeleton to the intersection of $p^{-1}(S_1)$ with the 1-skeleton. From the latter, a normal isotopy taking $p^{-1}(S_0)$ to $p^{-1}(S_1)$ can be constructed which descends to P. □

1.7 REGULAR EXCHANGE AND GEOMETRIC SUM: NORMAL ARCS

Given two normal subsets S and F of P, one may perturb S by an arbitrarily small normal isotopy so that S and F have no common point of intersection with the 1-skeleton, and one can attempt to produce a new normal subset from $S \cup F$ by performing a canonical cut and paste operation along $S \cap F$. Such a *geometric sum of S and F* was first used by Haken. Since a

normal disc is uniquely determined by its boundary, it is convenient to start with the geometric sum of normal arcs.

DEFINITION 1.20 (Normal subset of Δ^2). A subset A of a 2-simplex Δ is *normal* if

1. A is a (possibly empty) countable union of pairwise disjoint normal arcs;
2. the set of accumulation points of $A \cap \Delta^{(1)}$ is contained in $\Delta^{(0)}$;
3. if $\{x_i\} \subset A$ has accumulation point $x \in \Delta$, then $x \in A$ or $x \in \Delta^{(0)}$.

DEFINITION 1.21 (General position). Two normal subsets of a 2-simplex Δ^2 are said to be in *general position* if their intersection is contained in the interior of Δ^2.

Let $\mathfrak{a}$ be a normal arc on a 2-simplex Δ^2 dual to vertex $\mathfrak{v}$. Denote by $C_\mathfrak{a}$ the complementary region of $\Delta^2 \setminus \mathfrak{a}$ containing $\mathfrak{v}$. Given normal arcs $\mathfrak{a}, \mathfrak{b}$ on Δ^2 in general position, denote by D a small (piecewise linear) disc in $\operatorname{int}\Delta^2$ with centre $\mathfrak{a} \cap \mathfrak{b}$. Then define $\mathfrak{a} \uplus \mathfrak{b}$ to be the union of two normal arcs obtained as follows. Remove from $\mathfrak{a} \cup \mathfrak{b}$ the intersection with D and adjoin to it two arcs on ∂D (termed *circular arcs*) as follows. If $\mathfrak{a}$ and $\mathfrak{b}$ are of different types, then add the intersection of ∂D with $(C_\mathfrak{a} \cup C_\mathfrak{b}) \setminus (C_\mathfrak{a} \cap C_\mathfrak{b})$. Otherwise take the intersection with the complement thereof. This procedure is termed a *regular exchange at the intersection point,* and the position of the circular arcs is referred to as the *switch condition at the intersection point.* The result is a disjoint union of two elementary arcs which can therefore be straightened to a disjoint union of two normal arcs which is denoted by $\mathfrak{a} \uplus \mathfrak{b}$ and termed the *geometric sum of* $\mathfrak{a}$ *and* $\mathfrak{b}$. This is illustrated in Figure 3.

Let A and B be two normal subsets of a 2-simplex Δ^2 which are in general position. For each intersection point $p \in A \cap B$, one can choose a sufficiently small disc D with centre p so that if $p = \mathfrak{a} \cap \mathfrak{b}$ for $\mathfrak{a} \subset A$ and $\mathfrak{b} \subset B$, then $D \cap (A \cup B) = D \cap (\mathfrak{a} \cup \mathfrak{b})$, and any two such discs are disjoint. One can then perform a regular exchange at each intersection point and the result, $\mathcal{C}$, is (a priori) a collection of embedded arcs (possibly open or half-open) and circles in Δ^2. This is termed *the regular exchange at* $A \cap B$ since $\mathcal{C}$ is uniquely determined up to normal isotopy fixing the 1-skeleton.

The definition of regular exchange extends to any two *elementary* arcs meeting transversely in a single point in the interior of Δ^2. Since the resulting normal arcs are uniquely determined by the relative position of the intersection points of the elementary arcs with $\partial\Delta^2$, it follows that performing a regular

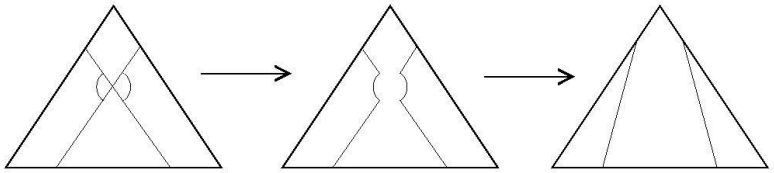

(a) Regular exchange of normal arcs of different types

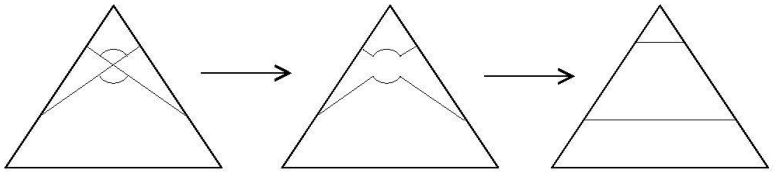

(b) Regular exchange of normally isotopic arcs

FIGURE 3

exchange between two elementary arcs and then straightening gives the same result as first straightening the elementary arcs, then taking the geometric sum. The following argument is similar to the one given in lecture notes by Cameron Gordon [2].

LEMMA 1.22 (Geometric sum of normal arcs). *Let A and B be two normal subsets of a 2-simplex Δ^2 which are in general position. Then the regular exchange at $A \cap B$ yields a family of pairwise disjoint elementary arcs. It can be straightened to a normal subset of Δ^2; this is denoted by $A \uplus B$ and called the* geometric sum of A and B. *Moreover, if $|\cdot|_{\mathfrak{v}} \in \mathbf{N} \cup \{\infty\}$ denotes the number of normal arcs dual to vertex $\mathfrak{v}$ of Δ^2, then $|A \uplus B|_{\mathfrak{v}} = |A|_{\mathfrak{v}} + |B|_{\mathfrak{v}}$.*

Proof. Choose discs for the regular exchange at $A \cap B$, and hence two circular arcs at each intersection point determined by the switch condition. These discs and arcs will be fixed throughout the proof; the set of circular arcs is denoted by C. Since A and B are normal subsets, the normal subset A_0 consisting of all normal arcs $\mathfrak{a} \subset A$ with the property that $\mathfrak{a}$ meets an arc in B not normally isotopic to $\mathfrak{a}$ is finite. Define $B_0 \subset B$ analogously. Assume that some arc in A_0 meets an arc in B_0 which is not normally isotopic to it. The properties that will be used are the following: (1) each connected component of A_0 (resp. B_0) is an elementary arc which is made up of *straight subarcs* of arcs in $A \cup B$ and circular arcs in C; (2) the arcs

in A_0 (resp. B_0) are pairwise disjoint; (3) whenever an elementary arc in A_0 meets an elementary arc in B_0, then they meet in a single point, and the corresponding regular exchange can be done by deleting subarcs from $A \cup B$ and inserting arcs from C.

Denote the vertices of Δ^2 by x, y, z and the opposite edges by s_x, s_y, s_z respectively. Label the normal corners of elementary arc $\mathfrak{a}$ dual to $\mathfrak{v}$ by $\mathfrak{v}$. For each side of Δ^2, this gives a finite word in the labels. Consider the side s_z. If some arc dual to x meets an arc dual to y, then on s_z, read from x to y, the word has the form $w_0(x,y) \cdot y \cdot w_1(x,y) \cdot x \cdot w_2(x,y)$. Consequently, if the word is of the form $x^n y^m$, then there are no intersections between arcs dual to x and arcs dual to y. Assume that there is some point of intersection. Then the word contains the subword yx. Since arcs in A_0 are pairwise disjoint, the two labels do not correspond to endpoints of two normal arcs in A_0. Similarly for B_0. Without loss of generality, assume that y is the label of an endpoint of $\mathfrak{a} \subset A_0$ and x is the label of an endpoint of $\mathfrak{b} \subset B_0$. Perform the regular exchange at $\mathfrak{a} \cap \mathfrak{b}$ using circular arcs from C. This yields two elementary arcs, one containing a portion of $\mathfrak{b}$ between s_y and the intersection point as well as the corner $\mathfrak{a} \cap s_y$. This elementary arc, denoted by $\mathfrak{b}_0$, has the property that $\mathfrak{b}_0 \cap (B_0 \setminus \mathfrak{b}) = \varnothing$ as well as $\mathfrak{b}_0 \cap A_0 = (\mathfrak{b} \cap A_0) \setminus (\mathfrak{a} \cap \mathfrak{b})$. Moreover, the switch condition at any point of $\mathfrak{b}_0 \cap A_0$ coincides with the switch condition at that point with respect to $\mathfrak{b}$, since $\mathfrak{b}_0$ and $\mathfrak{b}$ are normally isotopic. In particular, it can be realised using circular arcs from C. Moreover, any elementary arc met by $\mathfrak{b}_0$ has at most one point of intersection with it since none of the subarcs in which $\mathfrak{b}_0$ and $\mathfrak{b}$ differ meet any other arc. The same discussion applies to the other elementary arc, $\mathfrak{a}_0$, arising from the regular exchange. Let A_1 be the collection of elementary arcs obtained from A_0 by deleting $\mathfrak{a}$ and adding $\mathfrak{a}_0$, and define B_1 likewise. Then $|A_0|_{\mathfrak{v}} + |B_0|_{\mathfrak{v}} = |A_1|_{\mathfrak{v}} + |B_1|_{\mathfrak{v}}$ for each $\mathfrak{v} \in \{x, y, z\}$, $|A_1| = |A_0|$, $|B_1| = |B_0|$, $|A_1 \cap B_1| = |A_0 \cap B_0| - 1$ and A_1 and B_1 satisfy the above properties (1)–(3). It follows that this process can be iterated to yield two normal subsets A_k and B_k with the property that $|A_k| = |A_0|$, $|B_k| = |B_0|$, and whenever arcs in A_k and B_k meet, then they are normally isotopic. Moreover, the number of arcs dual to x (resp. y, z) in $A_k \cup B_k$ equals the number of normal arcs in A_0 dual to x (resp. y, z) plus the number of normal arcs in B_0 dual to x (resp. y, z).

Letting $A' = (A \setminus A_0) \cup A_k$ and $B' = (B \setminus B_0) \cup B_k$ gives sets with the property that each point in $A' \cap B'$ is the intersection point of normally isotopic arcs; A' and B' can be viewed as obtained by performing all regular exchanges at all points in $A \cap B$ corresponding to intersection points of arcs of different normal types. Let A_x and B_x be the sets of all elementary

arcs dual to x in A' and B' respectively. The elementary arcs in A_x can be labelled $\mathfrak{a}_0, \mathfrak{a}_1, \ldots$ with the property that $C_{\mathfrak{a}_0} \supset C_{\mathfrak{a}_1} \supset \ldots$. Similarly for B_x.

Since A and B are in general position, the intersection points of $A_x \cup B_x$ with the side s_y (resp. s_z) can be labelled $y_1, y_2, \ldots$ (resp. $z_1, z_2, \ldots$) such that $(y_i, x) \subset (y_{i-1}, x)$ (resp. $(z_i, x) \subset (z_{i-1}, x)$) for each $i > 1$. It now follows inductively that performing all regular exchanges at $A_x \cap B_x$ yields a union of pairwise disjoint elementary arcs $\{\mathfrak{n}_i\}$ with $\mathfrak{n}_i$ having endpoints y_i and z_i. The resulting number of normal arcs dual to x is thus uniquely determined by the intersection points with the 1-skeleton. Whence the claimed relationship.

If follows that performing all regular exchanges at $A \cap B$ yields a union of pairwise disjoint elementary; it can be pulled straight to a union of pairwise disjoint normal arcs, $A \uplus B$. The other parts of Definition 1.20 are verified as follows. We have $(A \uplus B) \cap \Delta^{2(1)} = (A \cap \Delta^{2(1)}) \cup (B \cap \Delta^{2(1)})$, whence the set of limit points of $(A \uplus B) \cap \Delta^{2(1)}$ is contained in $\Delta^{2(0)}$.

It follows from Definition 1.20 that for each vertex $\mathfrak{v}$ there is a neighbourhood $N_{\mathfrak{v}}$ with the property that $N_{\mathfrak{v}}$ only meets normal arcs dual to $\mathfrak{v}$ in $A \cup B$. It follows from the construction that $N_{\mathfrak{v}}$ can be chosen such that it meets only normal arcs dual to $\mathfrak{v}$ in $A \uplus B$. Thus, $A \uplus B$ is a normal subset of Δ^2. □

The method of pre-assigning switch conditions at intersection points fails for certain immersed sets of normal arcs. Assume $A = \{\mathfrak{a}_i\}$ is a set consisting of three normal arcs on a 2-simplex Δ^2 which are pairwise in general position and no two of which are normally isotopic. Then for any distinct $\mathfrak{a}_i, \mathfrak{a}_j \in A$, the regular exchange at $\mathfrak{a}_i \cap \mathfrak{a}_j$ is defined.

However, if all these regular exchanges are performed, then the result may not be elementary; Figure 4(a) shows a constellation where this fails. However, we have $(\mathfrak{a}_0 \uplus \mathfrak{a}_1) \uplus \mathfrak{a}_2 = \mathfrak{a}_0 \uplus (\mathfrak{a}_1 \uplus \mathfrak{a}_2) = (\mathfrak{a}_0 \uplus \mathfrak{a}_2) \uplus \mathfrak{a}_1$ even though intersection points are possibly resolved differently. See Figure 4(b) and notice that the last picture is invariant under rotation by $\frac{2\pi}{3}$ whilst the first and second are not.

LEMMA 1.23. *Let A, B, C be three normal subsets of a 2-simplex which are pairwise in general position. Then $A \uplus (B \uplus C) = (A \uplus B) \uplus C$.*

Proof. Let A' be the set of all normal arcs in A which meet normal arcs of different types in B or C; define B' and C' similarly. Then the previous lemma applies to show that $A' \uplus (B' \uplus C')$ and $(A' \uplus B') \uplus C'$ are normal

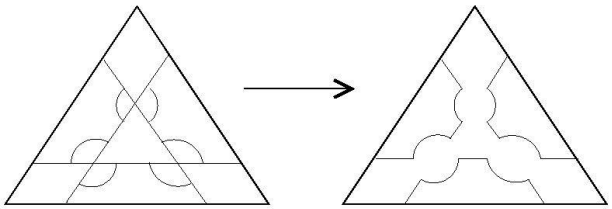

(a) Regular exchange of immersed may not be normal

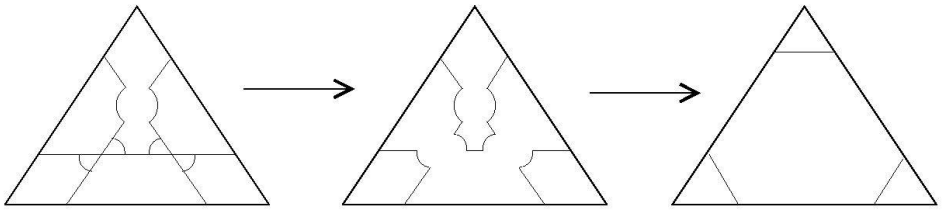

(b) Geometric sum is associative

FIGURE 4

subsets with the property that for every vertex $\mathfrak{v}$,

$$\begin{aligned} |A' \uplus (B' \uplus C')|_{\mathfrak{v}} &= |A'|_{\mathfrak{v}} + |B' \uplus C'|_{\mathfrak{v}} \\ &= |A'|_{\mathfrak{v}} + |B'|_{\mathfrak{v}} + |C'|_{\mathfrak{v}} \\ &= |A' \uplus B'|_{\mathfrak{v}} + |C'|_{\mathfrak{v}} = |(A' \uplus B') \uplus C'|_{\mathfrak{v}} . \end{aligned}$$

Whence $A' \uplus (B' \uplus C')$ and $(A' \uplus B') \uplus C'$ are identical, since they have identical intersection with the 1-skeleton.

It remains to analyse arcs which are pairwise in general position and dual to a common vertex. Here the switch conditions at intersection points are uniquely determined regardless of the order in which regular exchanges are performed. The lemma follows. □

Given two normal subsets A and B on a triangulated, compact surface, let

$$A \uplus B = \bigcup_{\Delta^2} \left((A \cap \Delta^2) \uplus (B \cap \Delta^2) \right),$$

where the union is taken over all 2-singlices in the triangulation. Then $A \uplus B$ is a well-defined normal subset.

LEMMA 1.24. *Let A and B be two finite unions of pairwise disjoint normal curves on $\partial\Delta^3$ with the property that for each face Δ^2, the sets $A \cap \Delta^2$ and $B \cap \Delta^2$ are in general position. Then $C = A \uplus B$ is a finite collection of normal curves. Moreover, if A and B only contain curves of length three or four, and each curve of length four does not meet the same edge Δ^1 of Δ^3, then C contains only curves of length three or four and each curve of length four does not meet Δ^1.*

Proof. The statement follows from Corollary 1.7 and Lemma 1.22. By assumption, one may normally isotope A to a normal subset A' which is disjoint from B. Then the intersection of $A' \cup B$ with the 1-skeleton is normally isotopic to the intersection of $A \cup B$ with the 1-skeleton. It follows that $A' \cup B$ is normally isotopic to $A \uplus B$. □

1.8 REGULAR EXCHANGE AND GEOMETRIC SUM: NORMAL DISCS

DEFINITION 1.25 (General position). Two normal subsets S and F of a 3-simplex Δ^3 are said to be *in general position* if $S \cap F \cap \Delta^{3(1)} = \varnothing$. In particular, for every face Δ^2 of Δ^3, the normal subsets $S \cap \Delta^2$ and $F \cap \Delta^2$ are in general position.

Let $\mathfrak{d}_0$ and $\mathfrak{d}_1$ be two normal discs in a 3-simplex Δ^3 which are in general position. A geometric sum of $\mathfrak{d}_0$ and $\mathfrak{d}_1$ should restrict to the geometric sum of $\partial\mathfrak{d}_0$ and $\partial\mathfrak{d}_1$ and yield two normal discs. Whence each normal curve in $\partial\mathfrak{d}_0 \uplus \partial\mathfrak{d}_1$ should have length three or four. It follows that $\mathfrak{d}_0$ and $\mathfrak{d}_1$ cannot be quadrilateral discs of different types since otherwise $\partial\mathfrak{d}_0 \uplus \partial\mathfrak{d}_1$ is a single curve of length eight.

Hence assume that if both $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are normal quadrilaterals, then they are normally isotopic. Also assume that $\alpha = \mathfrak{d}_0 \cap \mathfrak{d}_1 \neq \varnothing$. Recall the definition of a normal disc as the cone to the barycentre of its vertices. Any normal triangle is a flat Euclidean triangle, and a normal quadrilateral is made up of (at most) four Euclidean triangles. The incidence between two normal discs does not change under a normal isotopy of their union, so to determine the possibilities for α, it may be assumed that if $\mathfrak{d}_0$ is a normal quadrilateral, then it is flat.

It follows that if not both $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are normal quadrilaterals, then α is a properly embedded arc in each disc and necessarily has endpoints on distinct faces of Δ^3. If both discs are normal quadrilaterals, then α cannot contain a connected component not meeting $\partial\Delta^3$ for otherwise $\mathfrak{d}_0$ would separate the vertices of $\mathfrak{d}_1$ from the barycentre of $\mathfrak{d}_1$. Moreover, α meets

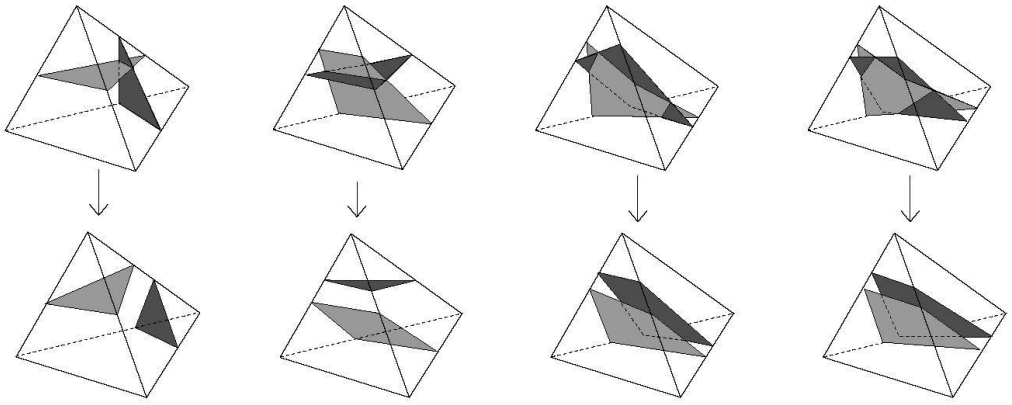

FIGURE 5

Regular exchange of normal discs (not shown are the two possible intersections of two quadrilateral discs in one arc)

$\partial\Delta^3$ in either two or four points. It follows that α is either a single arc, the disjoint union of two arcs, or a cross; each properly embedded in Δ^3. In particular, the above definition allows a piecewise linear analogue of saddle tangencies.

Let N_α be a regular neighbourhood of α in Δ^3. A *regular exchange of* $\mathfrak{d}_0$ *and* $\mathfrak{d}_1$ *at* α consists of deleting the portion of $\mathfrak{d}_0 \cup \mathfrak{d}_1$ in N_α, and adding discs contained in the complement of $\mathfrak{d}_0 \cup \mathfrak{d}_1$ on $\partial(N_\alpha)$ as follows.

If $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are normal quadrilaterals, then let X_0 and X_1 be the connected components of $\Delta^2 \setminus (\mathfrak{d}_0 \cup \mathfrak{d}_1)$ containing 1-simplices. If $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are a normal quadrilateral and a normal triangle, let X_0 (resp. X_1) be the component containing a 1-simplex (resp. the vertex dual to the triangle). If $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are normal triangles of the same type, let X_0 (resp. X_1) be the component containing a 2-simplex (resp. the vertex dual to the triangles). If $\mathfrak{d}_0$ and $\mathfrak{d}_1$ are normal triangles of the different types, let X_0 and X_1 be the components containing the vertices dual to the triangles. Then let $C = (X_0 \cup X_1) \cap \partial N_\alpha$.

It follows that $((\mathfrak{d}_0 \cup \mathfrak{d}_1) \setminus N_\alpha) \cup C$ is a union of two disjoint elementary discs which can be straightened to give two disjoint normal discs, denoted by $\mathfrak{d}_0 \uplus \mathfrak{d}_1$, and the restriction of C to each face coincides with the switch condition at each endpoint of α on that face. For each connected component α' of α, the placement of the discs in $\partial N_{\alpha'}$ is termed the *switch condition at* α'.

DEFINITION 1.26 (Compatible normal subsets). Two normal subsets S and F of P are *compatible* if they are in general position and if the normal quadrilaterals in $S \cap \Delta^3$ and $F \cap \Delta^3$ are all of the same type for each 3-singlex Δ^3 in P.

The above definition of regular exchange extends to a regular exchange at $S \cap F$ where S and F are compatible normal subsets. It remains to show that the result (after straightening) is a normal subset.

For the proof of the lemma below, it is necessary to extend the above definition of regular exchange to elementary discs $\mathfrak{d}_0$ and $\mathfrak{d}_1$ in the 3-simplex Δ^3 subject to (1) $\mathfrak{d}_0 \cap \mathfrak{d}_1 \cap \Delta^{3(1)} = \varnothing$, (2) if one of the discs has three corners, then they have precisely one arc of intersection with endpoints on different faces of Δ^3, and (3) if both elementary discs meet the 1-skeleton in four normal corners, then the discs are normally isotopic and have intersection consisting of either one or two arcs or a cross and no two endpoints of $\mathfrak{d}_0 \cap \mathfrak{d}_1$ are on the same face of Δ^3. This generalisation is straightforward; it should be noted that in the realm of elementary discs, saddle tangencies can be eliminated by an arbitrarily small isotopy fixing $\partial\Delta^3$.

LEMMA 1.27 (Geometric sum of compatible normal subsets). *Let S and F be compatible normal subsets of a 3-simplex Δ^3. Then the regular exchange at $S \cap F$ yields a family of pairwise disjoint elementary discs. It can be straightened to a normal subset of Δ^3; this is denoted by $S \uplus F$. Moreover, the number of discs in $S \uplus F$ of any type equals the sum of the number (possibly infinite) of the discs of that type in S and F.*

Proof. The structure of the proof is as in Lemma 1.22, and only an outline is given. Let S_0 be the set of all normal discs in S meeting normal discs in F of a different normal type; define F_0 analogously. Then S_0 and F_0 are both finite (possibly empty) sets of normal discs.

Let Δ^1 be an edge of Δ^3 and denote its vertices by $\mathfrak{v}$ and $\mathfrak{w}$. Choose an ordering of the faces Δ^2_1 and Δ^2_2 incident with Δ^1. A normal triangle dual to $\mathfrak{v}$ (resp. $\mathfrak{w}$) meets each face in a $\mathfrak{v}$-type arc (resp. $\mathfrak{w}$-type arc), and a normal quadrilateral with a corner on Δ^1 meets one of the faces in a $\mathfrak{v}$-type and the other in a $\mathfrak{w}$-type arc. With respect to the ordering, the corner of a normal disc on Δ^1 is labelled by $(\mathfrak{v},\mathfrak{v})$, $(\mathfrak{v},\mathfrak{w})$, $(\mathfrak{w},\mathfrak{v})$ or $(\mathfrak{w},\mathfrak{w})$ accordingly. Since S and F are compatible, the corners of all quadrilaterals have the same label, say $(\mathfrak{v},\mathfrak{w})$. Read from $\mathfrak{v}$ to $\mathfrak{w}$, this

gives a word of finite length in these tuples, and if no disc meets a disc of different type, then the word has the form $(\mathfrak{v},\mathfrak{v})^n(\mathfrak{v},\mathfrak{w})^m(\mathfrak{w},\mathfrak{w})^l$, where $n,m,l \geq 0$.

Given a word not of this form, as before, an inductive argument allows to successively perform regular exchanges which move a pair $(\mathfrak{v},\mathfrak{v})$ to the left or a pair $(\mathfrak{w},\mathfrak{w})$ to the right. Having done this for all edges, it follows that whenever two discs meet, then they are of the same type. The proof is now completed similarly to the proof of Lemma 1.22. □

LEMMA 1.28. *The geometric sum of compatible normal subsets of P is well-defined and associative.*

Proof. Let S and F be compatible normal subsets of P. Then for each 3-simplex $\widetilde{\Delta}^3$ in $\widetilde{\Delta}$, $p^{-1}(S)\cap\widetilde{\Delta}^3$ and $p^{-1}(F)\cap\widetilde{\Delta}^3$ are compatible. It follows that the geometric sum $p^{-1}(S)\uplus p^{-1}(F)$ is well-defined. Since the definition of regular exchange between normal subsets of a 2-simplex is defined without reference to a 3-simplex, it follows that $p\left(p^{-1}(S)\uplus p^{-1}(F)\right)$ is a well-defined normal subset of P. □

This has the following, classical, consequence:

COROLLARY 1.29. *Let S and F be two normal subsets of P which are in general position, and let γ be a simple closed curve in $S\cap F$. Then γ is either 2-sided in both S and F or it is 1-sided in both S and F.*

Proof. If γ is non-separating in S, but not in F, then the switch condition at an intersection point of γ with the 2-skeleton changes upon a full traverse of γ. But this is not possible. □

LEMMA 1.30 (Additive identity). *If a normal subset S in P contains infinitely many normal triangles dual to $\mathfrak{v}\in P^{(0)}$, then $S\uplus B_{\mathfrak{v}}$ is normally isotopic to S.*

Proof. The intersection $S\cap B_{\mathfrak{v}}$ consists of at most finitely many pairwise disjoint simple closed curves and arcs. Since $B_{\mathfrak{v}}$ contains only normal triangles, it follows that $S\uplus B_{\mathfrak{v}}$ and S satisfy the hypothesis of Lemma 1.19. □

1.9 Boundary curves of spun-normal surfaces

Let S be a normal subset in P which is transverse to $B_\mathfrak{v}$ for some $\mathfrak{v} \in P^{(0)}$. It may be assumed that $N_\mathfrak{v}$ meets only normal triangles in S and that S and $B_\mathfrak{v}$ are in general position. Note that S and $B_\mathfrak{v}$ are compatible, and that $S \cap B_\mathfrak{v}$ is a finite union of pairwise disjoint simple closed curves or properly embedded arcs. Since $B_\mathfrak{v}$ is 2-sided in P, it follows that $S \cap B_\mathfrak{v}$ is 2-sided in S and hence, by Corollary 1.29, $S \cap B_\mathfrak{v}$ is 2-sided in $B_\mathfrak{v}$. This implies that $S \cap B_\mathfrak{v}$ can be given a transverse orientation in $B_\mathfrak{v}$, and we will make the following canonical choice.

Since $\overline{N}_\mathfrak{v}$ is the cone over $B_\mathfrak{v}$ on $\mathfrak{v}$, it inherits a triangulation which is the cone on $\mathfrak{v}$ over the triangulation of $B_\mathfrak{v}$. Moreover, $S \cap \overline{N}_\mathfrak{v}$ is a normal subset in $\overline{N}_\mathfrak{v}$ with respect to this triangulation. Any normal disc $\mathfrak{n}$ in $\overline{N}_\mathfrak{v}$ contained in the 3-singlex $\Delta^3_\mathfrak{v}$ in $\overline{N}_\mathfrak{v}$ inherits a well-defined transverse orientation by assigning $+1$ to the component of $\operatorname{int}(\Delta^3_\mathfrak{v}) \setminus \mathfrak{n}$ containing $\mathfrak{v}$. These orientations match up to give $S \cap \overline{N}_\mathfrak{v}$ a transverse orientation, and $S \cap B_\mathfrak{v}$ is given the induced transverse orientation. In particular, this shows that $S \cap \overline{N}_\mathfrak{v}$ is 2-sided in $\overline{N}_\mathfrak{v}$.

Lemma 1.31 (Boundary curves are non-separating and 2-sided). *Let S be a normal subset in P. Each curve or arc in $S \cap B_\mathfrak{v}$ is 2-sided in $B_\mathfrak{v}$ and in S. Up to normal isotopy of S, one may assume that $S \cap B_\mathfrak{v}$ is a (possibly empty) union of pairwise disjoint, non-separating, 2-sided simple closed curves or properly embedded arcs on $B_\mathfrak{v}$.*

Proof. It remains to show that each component of $S \cap B_\mathfrak{v}$ is non-separating in $B_\mathfrak{v}$. If S contains at most finitely many normal discs dual to $\mathfrak{v}$, then there is a normal isotopy of S making S disjoint from $B_\mathfrak{v}$, and there is nothing to prove. Hence assume that S contains infinitely many normal triangles dual to $\mathfrak{v}$. To simplify notation, let $B = B_\mathfrak{v}$.

Let c be a connected component of $S \cap B$. Perform the regular exchange at c and denote the result by $S(c)$. If c is separating in B, then there are subsets of $S(c)$ corresponding to the components B_+ and B_- of $B \setminus c$ and there is a normal homotopy of $S(c)$ which pushes one of these subsets towards v and the other out of $\overline{N}_\mathfrak{v}$ so that the result, also denoted by $S(c)$, is transverse to B and meets B in one separating curve fewer than S.

Performing regular exchanges along all singular curves in $S(c)$ yields a normal subset $S'(c)$ with the properties that $S'(c) \cap B = S(c) \cap B$ and $S'(c)$ is normally isotopic to $S \uplus B$. Hence $S'(c)$ is normally isotopic to S by Lemma 1.30. Since $S \cap B$ has finitely many connected components, one may, by induction, assume that no component of $S \cap B$ is separating. □

COROLLARY 1.32 (Spun-normal surface = normal surface in a manifold). *If P is a 3-manifold, then every connected component of a normal subset is a normal surface.*

Proof. This follows from the above lemma since if P is a 3-manifold, then every vertex linking surface is a sphere or disc. □

1.10 THE ORIENTABLE DOUBLE COVER

If P is non-orientable, denote by $\widetilde{P}$ the orientable double cover of P and by $\widetilde{B}_\mathfrak{v}$ a connected component of the pre-image of $B_\mathfrak{v}$ for each $\mathfrak{v} \in P^{(0)}$. A spun-normal surface $S \subset P$ lifts to a spun-normal surface $\widetilde{S} \subset \widetilde{P}$ which is invariant under the non-trivial deck transformation. Either P or $\widetilde{P}$ is given a fixed orientation.

Each boundary component can be given a transverse orientation by choosing a normal vector pointing towards the dual 0-singlex. It is then oriented such that the tuple *(orientation, transverse orientation)* agrees with the orientation of the ambient pseudo-manifold.

1.11 THE BOUNDARY CURVE MAP

A transversely oriented normal curve on $B_\mathfrak{v}$ defines an element in $H_1(B_\mathfrak{v}, \partial B_\mathfrak{v}; \mathbf{Z})$ as follows. The triangulation of $B_\mathfrak{v}$ can be given the structure of a Δ-*complex* (see [7], Section 2.1). Thus, each edge in the triangulation of $B_\mathfrak{v}$ is given a well-defined direction, and can be used for a well-defined simplicial homology theory. If $B_\mathfrak{v}$ is non-orientable, denote by $\widetilde{B}_\mathfrak{v}$ the chosen lift in the double cover; otherwise let $\widetilde{B}_\mathfrak{v} = B_\mathfrak{v}$. Recall that $\widetilde{B}_\mathfrak{v}$ is oriented.

Let $\mathfrak{a}$ be a transversely oriented normal curve on $B_\mathfrak{v}$, and consider its pre-image $\widetilde{\mathfrak{a}}$ in $\widetilde{B}_\mathfrak{v}$. Using the transverse orientation and the orientation of $\widetilde{B}_\mathfrak{v}$, give $\widetilde{\mathfrak{a}}$ an orientation such that the tuple *(orientation, transverse orientation)* associated to $\widetilde{\mathfrak{a}}$ agrees with the orientation of $\widetilde{B}_\mathfrak{v}$. Then homotope $\widetilde{\mathfrak{a}}$ in the direction of the transverse orientation into the 1-skeleton of $\widetilde{B}_\mathfrak{v}$. Using the orientation gives a well-defined element of $Z_1(\widetilde{B}_\mathfrak{v}, \partial\widetilde{B}_\mathfrak{v})$, and hence an element of $H_1(\widetilde{B}_\mathfrak{v}, \partial\widetilde{B}_\mathfrak{v})$. This maps to a non-torsion element $\overline{\mathfrak{a}} \in H_1(B_\mathfrak{v}, \partial B_\mathfrak{v})$ since $\mathfrak{a}$ is 2-sided on $B_\mathfrak{v}$. One may therefore view $\overline{\mathfrak{a}} \in H_1(B_\mathfrak{v}, \partial B_\mathfrak{v}; \mathbf{Z})$ and let

$$\partial_\mathfrak{v}(S) = \sum \overline{\mathfrak{a}} \in H_1(B_\mathfrak{v}, \partial B_\mathfrak{v}; \mathbf{Z}),$$

where the sum is taken over all connected components of $B_\mathfrak{v} \cap S$.

LEMMA 1.33 (Boundary curves are well-defined). *Let S_0 and S_1 be normally isotopic normal subsets of P. Then $\partial_{\mathfrak{v}}(S_0) = \partial_{\mathfrak{v}}(S_1)$. Moreover, $\partial_{\mathfrak{v}}(S)$ is trivial if and only if $S \cap B_{\mathfrak{v}}$ may be assumed empty.*

Proof. A dual statement will be proved. It suffices to assume that $B_{\mathfrak{v}}$ is orientable. Let B_0 and B_1 be two vertex linking surfaces normally isotopic to $B_{\mathfrak{v}}$ and meeting S transversely in normal triangles. The normal isotopy gives a canonical identification $H_1(B_{\mathfrak{v}}, \partial B_{\mathfrak{v}}; \mathbf{Z}) \cong H_1(B_k, \partial B_k; \mathbf{Z})$. Denote the corresponding families of transversely oriented intersection curves by $\mathcal{C}_k$. Then B_0 and B_1 bound a submanifold X of P homeomorphic to $B_{\mathfrak{v}} \times [0,1]$. The surface $S \cap X$ is given a well-defined transverse orientation using the construction in Subsection 1.9 which proves that $S \cap N_{\mathfrak{v}}$ is 2-sided. It follows that its set of transversely oriented boundary curves is precisely the union $\mathcal{C}_0 \cup \mathcal{C}_1$, which therefore determines an element of $B_1(B_{\mathfrak{v}}, \partial B_{\mathfrak{v}}; \mathbf{Z})$.

The argument given in the last paragraph in the proof of Lemma 1.31 is now adapted to prove the second part. Let $B = B_{\mathfrak{v}}$. Assume there is a subset $\mathcal{C}$ of $S \cap B$ with the property that there is a subsurface B' of B with $\partial B' = \mathcal{C}$ and B' is on the positive side of each component of $\mathcal{C}$. Then the cited argument can be applied with $B_+ = B'$ and $B_- = B \setminus B'$ to show that the components in $\mathcal{C}$ can be deleted from $S \cap B$ by a normal isotopy. This completes the proof. □

COROLLARY 1.34. *A spun-normal surface S spins into $\mathfrak{v} \in P^{(0)}$ if and only if $\partial_{\mathfrak{v}}(S) \neq 0$.*

1.12 THE SPINNING CONSTRUCTION

An interesting normal subset in $\overline{N}_{\mathfrak{v}}$ is constructed as follows (this generalises the construction of an infinite normal annulus in [8]).

Let $\mathcal{C}$ be a finite collection of pairwise disjoint, transversely oriented, non-separating normal curves on $B_{\mathfrak{v}}$ with the property that the intersection with every normal triangle in $B_{\mathfrak{v}}$ consists of k copies of a normal arc dual to a corner $\mathfrak{c}_0$ with transverse orientation pointing towards $\mathfrak{c}_0$, and l copies of a normal arc dual to $\mathfrak{c}_1 \neq \mathfrak{c}_0$ with transverse orientations pointing away from $\mathfrak{c}_1$; both k, l are non-negative integers. For instance, the system of transversely oriented normal curves in $S \cap B_{\mathfrak{v}}$ is of this form.

Let $\{B_i\}$, $i \in \mathbf{N}$, be a countably infinite family of vertex linking surfaces dual to v in $N_{\mathfrak{v}}$ such that $\bigcup B_i$ is a normal subset. On each B_i, there is a copy, $\mathcal{C}_i$, of $\mathcal{C}$. For each i and for each $c_i \in \mathcal{C}$, delete a small open

neighbourhood of c_i in B_i. Since c_i separates its neighbourhood, label the boundary component in the positive direction by c_i^+, and the other by c_i^-. Since there is a normal isotopy of $N_\mathfrak{v}$ taking c_i^- to c_{i+1}^+ for each i and each c, they can be joined by an annulus in $N_\mathfrak{v}$. Similarly, annuli can be attached between c and c_1^+ for each $c \in \mathcal{C}$. There is a normal isotopy fixing $\overline{N}_\mathfrak{v}^{(1)}$ such that the result is a normal subset of $\overline{N}_\mathfrak{v}$; this is said to be obtained by *spinning on* $\mathcal{C}$.

LEMMA 1.35 (Controlled spinning). *Deleting $S \cap N_\mathfrak{v}$ from S and attaching the result of spinning on $S \cap B_\mathfrak{v}$ (with the induced transverse orientation) yields a normal subset which is normally isotopic to S.*

Proof. Delete the portion of S in $N_\mathfrak{v}$ and replace it by a normal subset in $N_\mathfrak{v}$ obtained from the spinning construction on $S \cap B_\mathfrak{v}$, where each component is spun in the $+1$ direction on $B_\mathfrak{v}$ according to the convention in the proof of Lemma 1.31. The result meets each 3-singlex in elementary discs and straightening gives a normal subset S' in P with the property that S and S' satisfy the hypothesis of Lemma 1.19. It follows that S and S' are normally isotopic. □

COROLLARY 1.36. *Let S be a spun-normal surface which spins into $\mathfrak{v} \in P^{(0)}$. Then $\chi(B_\mathfrak{v}) \leq 0$. If $\chi(B_\mathfrak{v}) < 0$, then S is topologically infinite. If $\chi(B_\mathfrak{v}) = 0$, then S is homeomorphic to the interior of $S \cap (P \setminus N_\mathfrak{v})$ for suitably chosen $N_\mathfrak{v}$.*

Proof. Assume that S spins into $\mathfrak{v}$. If $\chi(B_\mathfrak{v}) > 0$, there are no 2-sided, non-separating curves or arcs on $B_\mathfrak{v}$ and S cannot be spun into $\mathfrak{v}$. Hence $\chi(B_\mathfrak{v}) \leq 0$. Since S can be viewed as obtained from the spinning construction, one has

$$\chi(S) \leq \chi(S \cap P^c) + \sum_{i=1}^{\infty} \chi(B_\mathfrak{v})$$

for suitably chosen P^c. So if $\chi(B_\mathfrak{v}) < 0$, then S is topologically infinite.

If $\chi(B_\mathfrak{v}) = 0$, then $B_\mathfrak{v}$ may be chosen such that $S \cap B_\mathfrak{v}$ contains no collection of curves which is homotopically trivial. It follows that $S \cap \overline{N}_\mathfrak{v}$ consists of properly embedded *half-open annuli* (i.e. annuli having one boundary curve removed) or discs with one point on the boundary removed meeting $B_\mathfrak{v}$ along a boundary arc. □

1.13 Kneser's lemma for spun-normal surfaces

LEMMA 1.37 (Kneser–Haken finiteness). *Let P be a triangulated pseudo-manifold (possibly with boundary), and S be a spun-normal surface such that no two components of S are normally isotopic and no component is a vertex linking surface. Then the number $|S|$ of components of S satisfies $|S| \leq 12t$ and*

$$|S| \leq 3t + \frac{3}{2}\dim H_2(P^c, \partial P^c; \mathbf{Z}_2) = 3t + \frac{3}{2}\dim H_1(P^c; \mathbf{Z}_2). \tag{1.1}$$

Moreover, if S is 2-sided, then $|S| \leq 6t$ and $|S| \leq 3t + \frac{1}{2}\dim H_2(P^c, \partial P^c; \mathbf{Z}_2)$.

Proof. Whenever S does not spin into $\mathfrak{v}$, one may add $B_{\mathfrak{v}}$ to S disjointly since no component of S is a vertex linking surface; the result is again denoted by S. The normal discs of $p^{-1}(S)$ divide a 3-simplex in $\widetilde{\Delta}$ into the following types of regions:

1. *slabs*: trivial I-bundles over normal discs,
2. *thick regions*: truncated tetrahedra and truncated triangular prisms,
3. *vertex regions*: "small" tetrahedra which contain a vertex of $\mathcal{T}$.

In total, there are at most $v + 2t$ components of $P \setminus S$ which contain the image under p of at least one of the latter types of regions. The remaining components of $P \setminus S$ are entirely made up of slabs. Each such slab component is either a trivial or a twisted I-bundle over a spun-normal surface.

If it is a trivial I-bundle, then there is a 2-sided surface in S which is the boundary of a twisted I-bundle over a 1-sided surface in S (since otherwise there would be two normally isotopic components of S). All other slab components are twisted I-bundles over a spun-normal surface not in S. Since this core surface is not normally isotopic to any of the components of S, we may add it disjointly to S. Let F be a 1-sided component of S. A small regular neighbourhood of F is a twisted I-bundle with boundary a 2-sided spun-normal surface, N, in P. If N is not normally isotopic to a component of S, we join it to S.

Let $\{V_i\}$ be the set of all vertex linking surfaces in S, $\{F_i\}$ be the set of 1-sided surfaces, the set of corresponding 2-sided surfaces be $\{N_i\}$, and the set of all remaining 2-sided surfaces be $\{S_i\}$. Then $|\{F_i\}| = |\{N_i\}|$. Note that P^c may be chosen such that the intersection of each component of S with it is connected. Let $S^c = S \cap P^c$, and similarly for its components. Each F_i^c is a 1-sided and therefore non-separating surface in P^c. Since F_i^c is 1-sided, there is a closed loop in a small regular neighbourhood of F_i^c meeting F_i^c transversely in a single point. Using the intersection pairing with $\mathbf{Z}_2$ coefficients, it follows

that F_i^c determines a non-zero element of $H_2(P^c, \partial P^c; \mathbf{Z}_2)$ and that $\{F_i^c\}$ is in bijective correspondence with a linearly independent set of elements of $H_2(P^c, \partial P^c; \mathbf{Z}_2)$.

Assume S' is a 2-sided component of S which is not a vertex linking surface and meets no thick region. Then S' bounds an I-bundle to either side. If one of them is trivial, then there is a component in S which is normal isotopic to S'. Hence both of these I-bundles are twisted. But then P is decomposed into two twisted I-bundles glued along their boundaries, and $S = S'$ does not meet any thick region. This is not possible. So each element of $\{N_i\} \cup \{S_i\}$ meets at least one thick region. There are at most $6t$ normal discs in the boundaries of thick regions. Each surface in $\{N_i\} \cup \{S_i\}$ meets at least one thick region in at least one disc. This implies that $|\{N_i\}| + |\{S_i\}| \leq 6t$. Thus $|\{F_i\}| \leq 6t$, and the first bound is obtained.

We now use a doubling trick from [1]. Each N_i meets thick regions only on one side, and each S_i meets at least one thick region on each side. Push each N_i off itself away from the twisted I-bundle, and call the resulting copy N_i', and push each S_i off itself and call the resulting copy S_i'. This can be done such that all surfaces are still pairwise disjoint. Each surface in $\{N_i'\} \cup \{S_i\} \cup \{S_i'\}$ meets at least one thick region in at least one disc. Thus, $2|S| = 2|\{V_i\}| + 2|\{F_i\}| + 2|\{N_i\}| + 2|\{S_i\}| = 2|\{V_i\}| + 3|\{F_i\}| + |\{N_i\}| + 2|\{S_i\}| \leq 2|\{V_i\}| + 3 \dim H_2(P^c, \partial P^c; \mathbf{Z}_2) + 6t$. Dividing by two and subtracting the vertex linking surfaces gives the second inequality. The stated equality follows since Lefschetz duality and the universal coefficient theorem yield $H_2(P^c, \partial P^c; \mathbf{Z}_2) \cong H_1(P^c; \mathbf{Z}_2)$. □

1.14 Normalising properly embedded surfaces

Haken's approach to normalising properly embedded *closed* surfaces as described in Chapter 3 of [10] applies to ideal triangulations without change. However, it does not apply to properly embedded *non-compact* surfaces; an obstruction for putting such a surface into spun-normal form is, for instance, given by Corollary 1.36. Moreover, it is shown by Kang [8] that there is an ideal triangulation of the complement of the figure eight knot with two ideal 3-simplices, such that no Seifert surface for the knot can be put into spun-normal form (see also Section 4.2). A result due to Thurston (see Walsh [13]) states that essential surfaces other than virtual fibres can be normalised in any hyperbolic 3-manifold with torus cusps and ideal triangulation with essential edges. A general theory of normalisation with respect to ideal triangulations is implicit in the work of Brittenham and Gabai [6].

2. Matching equations

A normal subset S is (up to vertex linking components and normal isotopy) uniquely determined by its quadrilateral discs. Recording the number of quadrilaterals of each type gives the normal Q-coordinate of S. This satisfies two necessary conditions: (1) it is admissible, and (2) it satisfies a *Q-matching equation* for each edge not contained in the boundary of P. This equation is given by Tollefson [12] for compact 3-manifolds. The main result of this section (Theorem 2.4) states that, conversely, any admissible solution to the Q-matching equation is realised by a spun-normal surface which is unique up to normal isotopy.

2.1 Normal Q-coordinate

There are t 3-singlices in P. Let $\{q_1, \ldots, q_{3t}\}$ be the set of all quadrilateral types in P. A normal subset S in P meets each 3-singlex in at most finitely many quadrilateral discs of each type. If x_i is the number of normal quadrilaterals of type q_i, then let $N(S) = (x_1, \ldots, x_{3t})$ be the *normal Q-coordinate* of S. Thus, $N(S)$ is a well-defined point in $\mathbf{R}^{3t}$ with the coordinate axes labelled by the quadrilateral types. A point $(x_1, \ldots, x_{3t}) \in \mathbf{R}^{3t}$ is called *admissible* if each $x_i \geq 0$ and if q_i, q_j, q_k are the three distinct quadrilateral types contained in any 3-singlex, then at most one of x_i, x_j, x_k is non-zero. The cell structure of S allows normal quadrilaterals of at most one quadrilateral type in each ideal 3-singlex, and hence $N(S)$ is admissible.

2.2 Convention for oriented 3-simplices

Let Δ^3 be an oriented, regular Euclidean 3-simplex. The edges of Δ^3 are labelled with parameters z, z', z'', such that opposite edges have the same parameter, and the ordering z, z', z'' agrees with a right-handed screw orientation of Δ^3; this is pictured in Figure 6. It follows that the labelling is uniquely determined once the parameter z is assigned to any edge of Δ^3. The vertices of a normal triangle dual to a vertex of Δ^3 inherit moduli from the edge parameters; this labelling is always viewed from the vertex.

There are three quadrilateral types in Δ^3; let $q^{(k)}$ denote the quadrilateral type which does not meet the edges labelled $z^{(k)}$ for $k \in \{0, 1, 2\}$. The symmetry group of Δ^3 is the alternating group on its four vertices. It contains a normal Kleinian four group which leaves pairs of opposite edges invariant and hence fixes quadrilateral types. The quotient group is $\mathbf{Z}_3$; it may be identified with the group of even permutations of the three quadrilateral types,

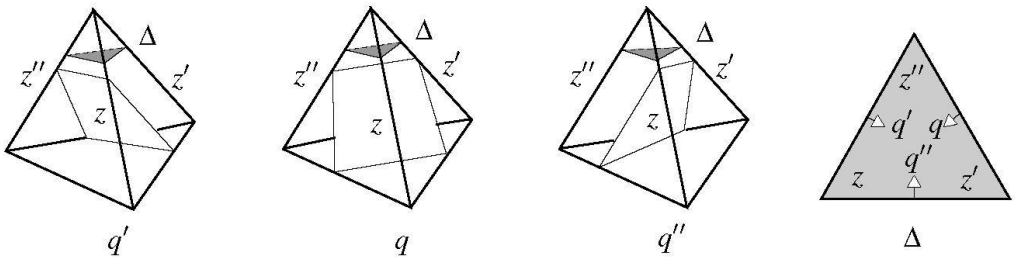

FIGURE 6
Edge labels, quadrilateral types and Q-moduli

and is hence generated by the cycle $\pi = (q, q', q'')$. If p and q are quadrilateral types in Δ^3, let:

$$s_p(q) = k \text{ if } \pi^k(p) = q \text{ for } k \in \{-1, 0, 1\}.$$

An isometry $\varphi\colon \Delta_0^3 \to \Delta_1^3$ between oriented 3-simplices induces a bijection between quadrilateral types, and one has $s_p(q) = s_{\varphi(p)}(\varphi(q))$ if φ is orientation preserving, and $s_p(q) = -s_{\varphi(p)}(\varphi(q))$ otherwise. The above extends to normal quadrilateral types in an oriented 3-singlex since the quotient map from a 3-simplex to a 3-singlex induces a bijection between normal quadrilateral types.

2.3 DEGREE AND ABSTRACT NEIGHBOURHOOD

The *degree of a 1-singlex* Δ^1 in P, $\deg(\Delta^1)$, is the number of 1-simplices in $\widetilde{\Delta}$ which map to Δ^1. Given a 1-singlex Δ^1 in P, there is an associated *abstract neighbourhood* $B(\Delta^1)$ of Δ^1 which is a ball triangulated by $\deg(\Delta^1)$ 3-simplices with the property that there is a well-defined simplicial quotient map $p_{\Delta^1}\colon B(\Delta^1) \to P$ taking $\widetilde{\Delta}^1$ to Δ^1.

If Δ^1 has at most one pre-image in each 3-simplex in $\widetilde{\Delta}$, then $B(\Delta^1)$ is obtained as the quotient of the collection $\widetilde{\Delta}_{\Delta^1}$ of all 3-simplices in $\widetilde{\Delta}$ containing a pre-image of Δ^1 by the set Φ_{Δ^1} of all face pairings in Φ between faces containing a pre-image of Δ^1. There is an obvious quotient map $b_{\Delta^1}\colon B(\Delta^1) \to P$ which takes into account the remaining identifications on the boundary of $B(\Delta^1)$.

We now describe $B(\Delta^1)$ if Δ^1 has more than one pre-image in some 3-simplex. In this case $\deg(\Delta^1) > 1$. For each 3-simplex $\widetilde{\Delta}_j^3$, if Δ^1 has k pre-images in $\widetilde{\Delta}_j^3$, take k copies of this 3-simplex, $\widetilde{\Delta}_{j,1}^3, \ldots, \widetilde{\Delta}_{j,k}^3$, each with an isometry $i_{j,h}\colon \widetilde{\Delta}_j^3 \to \widetilde{\Delta}_{j,h}^3$. To each pre-image $\widetilde{\Delta}^1$ of Δ^1 in $\widetilde{\Delta}_j^3$ assign one of

the maps $i_{j,h}$ and call $i_{j,h}(\widetilde{\Delta}^1)$ the *marked edge* of $\widetilde{\Delta}^3_{j,h}$. This gives a collection $\widetilde{\Delta}_{\Delta^1}$ of $\deg(\Delta^1)$ 3-simplices.

A face pairing φ of faces of $\widetilde{\Delta}^3_j$ and $\widetilde{\Delta}^3_h$ gives rise to a face pairing between faces of $\widetilde{\Delta}^3_{j,i}$ and $\widetilde{\Delta}^3_{h,k}$ if and only if it identifies the edges of $\widetilde{\Delta}^3_j$ and $\widetilde{\Delta}^3_h$ corresponding to the marked edges of $\widetilde{\Delta}^3_{j,i}$ and $\widetilde{\Delta}^3_{h,k}$. Denote the resulting set of face pairings by Φ_{Δ^1}.

The quotient map is denoted by $p_{\Delta^1}\colon \widetilde{\Delta}_{\Delta^1} \to \widetilde{\Delta}_{\Delta^1}/\Phi_{\Delta^1} = B(\Delta^1)$. If $\deg(\Delta^1) > 1$ or if $\Delta^1 \subset \partial P$, then the triangulation of $B(\Delta^1)$ is simplicial.

We now describe the map b_{Δ^1} in the case where Δ^1 has more than one pre-image in some 3-simplex. Since $\deg(\Delta^1) > 1$, the given triangulation of $B(\Delta^1)$ is simplicial and hence for each simplex $\widetilde{\Delta}^3_{j,h}$ in $B(\Delta^1)$ the inverse $p^{-1}_{\Delta^1}$ is well-defined. Since the face pairings defining $B(\Delta^1)$ arise from the face pairings in Φ, the map $p_{\Delta^1}\colon B(\Delta^1) \to P$ defined by $p \circ i^{-1}_{j,h} \circ p^{-1}_{\Delta^1}\colon \widetilde{\Delta}^3_{j,h} \to P$ for each 3-simplex is well-defined.

The unique edge in $B(\Delta^1)$ which is the image of all marked edges is termed the *axis* of $B(\Delta^1)$. One of the endpoints of the axis is termed the *north pole* and the other the *south pole*. The set of all 1-singlices not containing a pole is termed *the equator* of $B(\Delta^1)$. See Figure 7(a).

LEMMA 2.1. *Let S be a normal subset of P. Then $p^{-1}_{\Delta^1}(S)$ is a normal subset of $B(\Delta^1)$.*

Proof. The set $p^{-1}(S)$ is a normal subset of $\widetilde{\Delta}$ with the property that whenever $\mathfrak{a} \subset p^{-1}(S)$ is a normal arc on a 2-simplex in the range of some face pairing $\varphi \in \Phi$, then $\varphi(\mathfrak{a}) \subset p^{-1}(S)$. The result now follows from the definitions of $\widetilde{\Delta}_{\Delta^1}$ and Φ_{Δ^1}. □

LEMMA 2.2. *Any normal subset in $B(\Delta^1)$ is a (possibly infinite) union of pairwise disjoint properly embedded discs in $B(\Delta^1)$.*

Proof. This follows from the fact that $B(\Delta^1)$ is a manifold and from Corollary 1.32. □

2.4 SLOPES OF QUADRILATERAL TYPES

Let Δ^1 be a 1-singlex not contained in ∂P. The abstract neighbourhood $B(\Delta^1)$ of a 1-singlex Δ^1 in P is a ball and hence the 3-singlices in $B(\Delta^1)$ may be oriented coherently. In particular, the convention of Subsection 2.2 applies to each oriented 3-singlex in $B(\Delta^1)$. To simplify notation, write $e = \Delta^1$. If p is the type of a normal quadrilateral in the oriented abstract neighbourhood

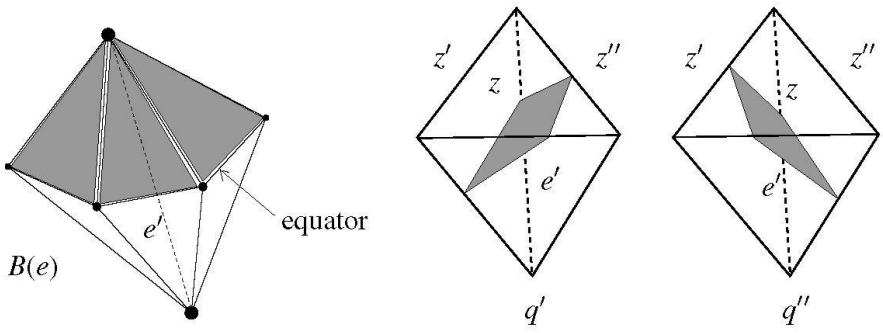

(a) Abstract neighbourhood (b) Positive (left) and negative (right) slope

FIGURE 7

$B(e)$ which does not meet the axis e', let $s_e(q) = s_p(q)$ for any quadrilateral type q in the same 3-singlex as p. This gives a function s_e defined on the set of all quadrilateral types in $B(e)$, and the value $s_e(q)$ is termed *the slope of q with respect to e*. The two quadrilateral types in a 3-singlex meeting e' have slopes of opposite signs, and a quadrilateral type which does not meet e' has slope zero. See Figure 7(b), where the shown 3-singlex is oriented using a right-handed screw orientation.

The orientation of $B(e)$ induces an orientation of its hemispheres, and the equator will be given the orientation induced from the southern hemisphere.

If $D \subset B(e)$ is a properly embedded disc such that the poles are contained in different components of $\partial B(e) \setminus \partial D$, then ∂D is given the orientation induced from the component containing the south pole. If P is a point in the intersection of ∂D and the equator, let $s(P) = 1$ if ∂D crosses the equator from the southern to the northern hemisphere at P, and let $s(P) = -1$ otherwise. Thus $0 = \sum s(P)$, where the sum is taken over all points in the intersection of ∂D and the equator; the empty sum is defined to be zero throughout.

Now assume that D is a connected normal surface in $B(e)$ which meets e'. Then D is a disc; it is the join of ∂D to $D \cap e'$. In particular, ∂D is a normal curve separating the poles on $\partial B(e)$ and D meets each 3-singlex in exactly one normal disc. The boundary of a normal triangle in D does not meet the equator, and a quadrilateral disc of type q in D contributes two normal arcs to its boundary which meet in a point P at the equator. The above conventions imply that $s(P) = s_e(q)$. Thus, $\sum s_e(q) = 0$, where the sum is taken over all quadrilateral types in D. If D is any normal subset in $B(e)$ which does not

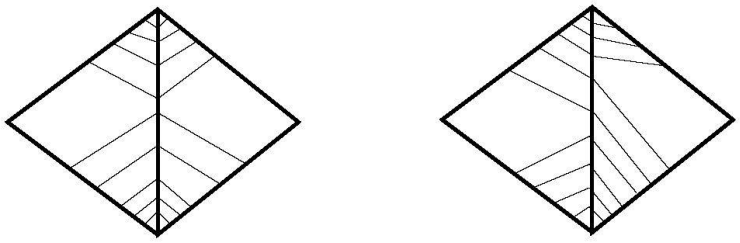

$\sum s_{\Delta^1}(q)x_q = 0$ $\sum s_{\Delta^1}(q)x_q \neq 0$

FIGURE 8
$B(\Delta^1)$ cut open along $p_{\Delta^1}(\widetilde{\Delta}^2)$

meet e', this equation is trivially satisfied since a quadrilateral disc meets the equator if and only if it has a vertex on e'.

Note that the equation is well-defined up to sign since reversing the orientation of $B(\Delta^1)$ changes the sign of the slope of each quadrilateral meeting Δ^1. If Δ^1 is contained in ∂P, then define $s_{\Delta^1}(q) = 0$ for each quadrilateral type in $B(\Delta^1)$.

LEMMA 2.3. *A normal subset $\widetilde{S} \subset \widetilde{\Delta}_{\Delta^1}$ which contains infinitely many normal triangles of each type can be normally isotoped such that $p_{\Delta^1}(\widetilde{S})$ is a normal subset of $B(\Delta^1)$ if and only if $\sum s_{\Delta^1}(q)x_q = 0$, where x_q is the number of normal quadrilaterals of type q in $p_{\Delta^1}(\widetilde{S})$.*

Proof. It follows from the above discussion that the equation is a necessary condition, and it remains to show that it is sufficient. Given two 3-simplices in $\widetilde{\Delta}_{\Delta^1}$, the pattern on each face is the same: there are three countably infinite stacks of parallel copies of normal arcs. Fix the discs of $\widetilde{S}$ in one 3-simplex $\widetilde{\Delta}^3$. If there are no face pairings for faces of $\widetilde{\Delta}^3$, then $B(\Delta^1) = \widetilde{\Delta}^3$ and there is nothing to prove.

Next assume that there is precisely one face pairing, φ_1, involving faces of $\widetilde{\Delta}^3$. Then the equation is trivial. Let $\widetilde{\Delta}^3_2$ be a simplex with one face $\widetilde{\Delta}^2$ in the range or domain of φ_1. Then there is an ideal normal isotopy of $\widetilde{\Delta}^3_2$ which fixes all subsimplices not in $\widetilde{\Delta}^2$ such that φ_1 matches the normal arcs on $\widetilde{\Delta}^2$ bijectively with the normal arcs on its counterpart. If there are no other face pairings involving $\widetilde{\Delta}^3_2$, we are done. Otherwise there is exactly one other face pairing, and the above procedure can be iterated. It terminates after $\deg(\Delta^1) - 1$ steps.

Hence assume that there are two face pairings for the faces of every 3-singlex. Then the equation is non-trivial. Ignoring one of the face pairings,

φ_0, for faces of $\widetilde{\Delta}^3$, one can proceed as in the previous paragraph. Denote by $\widetilde{\Delta}^3_k$ the last simplex in the iteration, and assume that $\varphi_0(\widetilde{\Delta}^2) = \widetilde{\Delta}^2_k$, where $\widetilde{\Delta}^2$ is a face of φ_0. It follows from the construction that $\widetilde{S} \cap \widetilde{\Delta}^3_k$ can only be altered by a normal isotopy which fixes all edges of $\widetilde{\Delta}^3_k$ other than the non-marked edges of $\widetilde{\Delta}^2_k$; see Figure 8. To obtain a normal subset of $B(\Delta^1)$, the normal arcs dual to the vertices not contained on the marked edges on $\widetilde{\Delta}^2$ and $\widetilde{\Delta}^2_k$ have to be identified in pairs, and this determines a unique bijection between the normal corners on the edges which are not marked. This bijection determines a unique normal isotopy of $\widetilde{S} \cap \widetilde{\Delta}^3_k$ such that equivalent corners have the same image under p_{Δ^1}. Note that $p_{\Delta^1}(\widetilde{S}) \cap \partial B(\Delta^1)$ is a normal subset.

If $p_{\Delta^1}(\widetilde{S})$ is a normal subset, then $p_{\Delta^1}(\widetilde{S}) \cap \partial B(\Delta^1)$ is a countably infinite union of circles and the equation $\sum s_{\Delta^1}(q)x_q = 0$ is satisfied. Assume that $p_{\Delta^1}(\widetilde{S})$ is not a normal subset. Then some component of $p_{\Delta^1}(\widetilde{S}) \cap \partial B(\Delta^1)$ crosses the equator an odd number of times, and, since $p_{\Delta^1}(\widetilde{S}) \cap \partial B(\Delta^1)$ is properly embedded in $\partial B(\Delta^1)$ minus the poles, it follows that $\sum s_{\Delta^1}(q)x_q \neq 0$. This proves the lemma. □

2.5 THE Q-MATCHING EQUATION

For each 1-singlex Δ^1 in P, the *total slope with respect to* Δ^1, $s_{\Delta^1}(q)$, of a quadrilateral type q in P is defined to be the sum over all pre-images of q in $B(\Delta^1)$ of the signs of their slopes. If S is a normal subset of P, then it follows from Lemma 2.3 that the normal Q-coordinate $N(S) = (x_1, \ldots, x_{3t})$ of S satisfies a linear equation for Δ^1, called the *Q-matching equation of* Δ^1:

$$0 = \sum_{i=1}^{3t} s_{\Delta^1}(q_i)x_i .$$

This equation is well-defined up to sign since reversing the orientation of $B(\Delta^1)$ changes the sign of the slope of each quadrilateral meeting Δ^1.

THEOREM 2.4. *Let P be a triangulated pseudo-manifold. For each admissible integer solution N of the Q-matching equations there exists a spun-normal surface S in P with no vertex linking components such that $N(S) = N$. Moreover, S is unique up to normal isotopy.*

Proof. It suffices to show existence; uniqueness follows from Lemmata 1.15 and 1.19.

Recall that there is a bijective correspondence between normal disc types in $\widetilde{\Delta}$ and P. Given an admissible integer solution $(x_1, \ldots, x_{3t})$, place x_i pairwise disjoint normal discs in $\widetilde{\Delta}$ each of which maps to a normal disc of type

q_i in P. Then place infinitely many normal triangles of each type in each 3-simplex in $\widetilde{\Delta}$ such that the resulting collection of normal discs is a normal subset, denoted by $\widetilde{S}$, of $\widetilde{\Delta}$. It suffices to show that there is a normal isotopy i of $\widetilde{\Delta}$ such that $p(i(\widetilde{S}))$ is a normal subset of P.

The normal isotopy is uniquely determined by defining it on the 1-skeleton. The normal subset $\widetilde{S}$ determines a normal subset $\widetilde{S}_{\Delta^1}$ in $\widetilde{\Delta}_{\Delta^1}$. It follows from Lemma 2.3 that this may be normally isotoped so that $p_{\Delta^1}(\widetilde{S}_{\Delta^1})$ is normal in $B(\Delta^1)$. For each 1-simplex $\widetilde{\Delta}^1$ in $\widetilde{\Delta}$ there exists a unique marked edge $\widetilde{\Delta}^1_m$ in some $\widetilde{\Delta}_{h,k}$. Define i on $\widetilde{\Delta}^1$ to be an normal isotopy taking the normal corners in $\widetilde{S} \cap \widetilde{\Delta}^1$ to the corresponding normal corners in $i^{-1}_{h,k}(\widetilde{S}_{\Delta^1} \cap \widetilde{\Delta}^1_m)$. Since each 1-simplex in $\widetilde{\Delta}$ corresponds to a unique marked edge, this is well defined and extends to the desired normal isotopy. □

REMARK 2.5. There are proofs for two special cases of Theorem 2.4 in the literature. The proofs by Tollefson [12] (for compact manifolds) and Kang [8] (for topologically finite manifolds with torus ends) construct an explicit surface. A different proof for 3-manifolds with torus cusps is sketched by Weeks in the documentation of `SnapPea` [14]. The above proof is inspired by the latter.

2.6 PROJECTIVE SOLUTION SPACE

Let B be the coefficient matrix of the Q-matching equations. Considering B as a linear transformation $\mathbf{R}^{3t} \to \mathbf{R}^e$, the set of all solutions to the Q-matching equations is $\ker B$. This will be denoted by $Q(\mathcal{T})$. One often considers the *projective solution space* $PQ(\mathcal{T})$, which consists of all elements of $Q(\mathcal{T})$ with the property that the sum of the coordinates equals one. This is a convex polytope, and its vertices are called *vertex solutions*. Since the entries in B are integers, the vertices of $PQ(\mathcal{T})$ are rational solutions, and hence the subset of all rational points is dense in $PQ(\mathcal{T})$. Such a polytope is termed *rational*. The vector in $\mathbf{R}^{3t}$ with each coordinate equal to one is contained in $Q(\mathcal{T})$. Thus, $\dim_{\mathbf{R}} PQ(\mathcal{T}) = \dim_{\mathbf{R}} Q(\mathcal{T}) - 1$.

The subset of all admissible solutions in $PQ(\mathcal{T})$ is denoted by $PF(\mathcal{T})$. If $N \in PF(\mathcal{T})$ is rational, then there is a countable family of spun-normal surfaces S_i without vertex linking components such that $N(S_i) = \alpha_i N$ for some $\alpha_i \in \mathbf{R}$. A spun-normal surface is a *minimal representative* for N, if the corresponding scaling factor is minimal. Since the admissible solutions can be found by setting in turn two coordinates from each tetrahedron equal to zero, it follows that $PF(\mathcal{T})$ is a finite union of convex rational polytopes,

each of which is contained in a $(t-1)$-dimensional rational polytope. Addition of solutions within a convex cell of $PF(\mathcal{T})$ corresponds to the geometric sum operation.

2.7 Spun-normal branched immersions

Definition 2.6 (Spun-normal branched immersion). Let P be a pseudo-manifold (possibly with boundary). Let S be obtained from a compact (not necessarily connected or closed) surface by removing some boundary components. Suppose that S has a cell structure with 2-cells which are either triangles and quadrilaterals, and that $f\colon S \to P$ is a piecewise linear map such that the interior of every 2-cell of S is mapped homeomorphically onto the interior of a normal disc in P satisfying the following extra conditions.

1. f is normally isotopic to an immersion in the complement of the 0-skeleton of S;
2. f is transverse to the 2-singlices in P;
3. the set of accumulation points of $f(S) \cap P^{(1)}$ is contained in $P^{(0)}$;
4. if $\{x_i\} \subset f(S)$ has accumulation point $x \in P$, then $x \in f(S)$ or $x \in P^{(0)}$.

Then S contains finitely many quadrilaterals and f is proper. The map f is called a *spun-normal branched immersion*, and if f is in general position, then $f(S)$ is termed a *branched immersed spun-normal surface* in P. Let $x(f) \in \mathbf{Z}^{3t}$ be the point defined as follows: the coefficient for each normal quadrilateral type is the number of 2-cells in S which map to discs of this type.

Proposition 2.7. *Let P be a 3-dimensional pseudo-manifold with triangulation $\mathcal{T}$. If $f\colon S \to P$ is a spun-normal branched immersion, then $x(f) \in Q(\mathcal{T})$. Every non-zero point in $Q(\mathcal{T})$ with non-negative integral coordinates is represented by a (not necessarily unique) spun-normal branched immersion with finitely many branch points.*

It follows from the definition that $x(f) \in Q(\mathcal{T})$; the remainder of the proposition is proved in Subsection 3.2.

2.8 Traditional normal surface theory

To the $3t$ quadrilateral coordinates $(q_1, \ldots, q_t'')$ adjoin $4t$ further coordinates, one for each normal triangle type. If S is a closed normal surface in P, let $N_\Delta(S) \in \mathbf{R}^{7t}$ be its *normal coordinate*. The normal coordinate of a closed normal surface satisfies three *compatibility equations* for each 2-singlex not

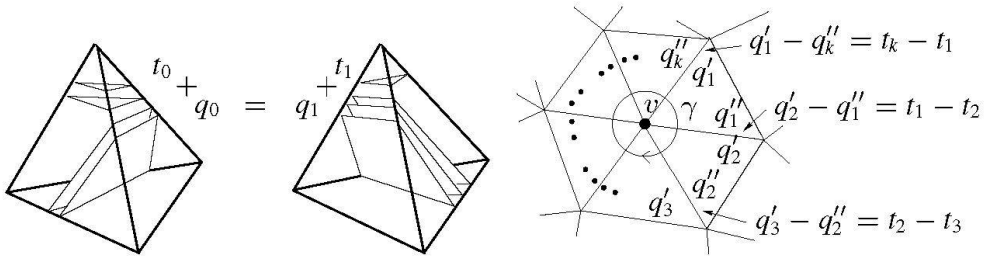

(a) Compatibility equations (b) Compatibility and Q-matching

FIGURE 9
Normal surface coordinates

contained in ∂P, arising from the fact that the total numbers of normal arcs "on either side" have to match up. Denote by $C(\mathcal{T})$ the space of all solutions to the compatibility equations.

Since the compatibility equations can be viewed as conditions associated to normal arc types, there is precisely one such equation for each edge in the induced triangulation of ∂P^c not contained in its boundary. Let $\mathfrak{c} \in B_{\mathfrak{v}}^{(0)}$ be contained in the interior of $B_{\mathfrak{v}}$. Then each element of $C(\mathcal{T})$ satisfies one compatibility equation for each normal arc in $B_{\mathfrak{v}}$ which ends in $\mathfrak{c}$ (see Figure 9(b)). Pick one of the arcs, and write the compatibility equation in the form $q'_{i+1} - q''_i = t_i - t_{i+1}$. Then, proceeding to the next arc, one obtains $q'_{i+2} - q''_{i+1} = t_{i+1} - t_{i+2}$, etc. Summing all these equations gives on the left hand side (up to sign) the Q-matching equation of the 1-singlex in P containing $\mathfrak{c}$, and the right hand side equals zero. Hence there is a well-defined linear map

$$pr\colon C(\mathcal{T}) \to Q(\mathcal{T})$$

defined by projection onto the quadrilateral coordinates.

A canonical basis for $C(\mathcal{T})$ consisting of t *tetrahedral solutions* and e *edge solutions* is given in [9]. Let Δ^3 be a 3-singlex in P, and let $\{\mathfrak{q}_i, \mathfrak{t}_j \mid i = 1, \ldots, 3,\ j = 1, \ldots, 4\}$ be the set of normal discs types in Δ^3. The *tetrahedral solution* associated to Δ^3 is:

$$W_{\Delta^3} = \mathfrak{t}_1 + \mathfrak{t}_2 + \mathfrak{t}_3 + \mathfrak{t}_4 - \mathfrak{q}_1 - \mathfrak{q}_2 - \mathfrak{q}_3\,.$$

For each edge $\widetilde{\Delta}^1$ contained in a 3-simplex $\widetilde{\Delta}^3$ in $\widetilde{\Delta}$ there is a unique normal quadrilateral type $\mathfrak{q}(\widetilde{\Delta}^1)$ in $\widetilde{\Delta}^3$ disjoint from it, and there are two normal triangle types $\mathfrak{t}_1(\widetilde{\Delta}^1)$ and $\mathfrak{t}_2(\widetilde{\Delta}^1)$ meeting it. The *edge solution* associated to a 1-singlex Δ^1 in P is:

$$W_{\Delta^1} = p\Big(\sum_{\widetilde{\Delta}^1 \subset p^{-1}(\Delta^1)} \mathfrak{t}_1(\widetilde{\Delta}^1) + \mathfrak{t}_2(\widetilde{\Delta}^1) - \mathfrak{q}(\widetilde{\Delta}^1) \Big) .$$

Using signed intersection numbers with edges, it follows from work in [9] that these solutions form a basis for $C(\mathcal{T})$ (even though this is not stated in [9] in this generality):

PROPOSITION 2.8 (Kang–Rubinstein). *Let P be a pseudo-manifold with triangulation $\mathcal{T}$. The set of all tetrahedral and edge solutions is a basis for $C(\mathcal{T})$ as a vector space over* $\mathbf{R}$. *In particular, $C(\mathcal{T})$ has dimension $t + e$.*

2.9 CONVENTION FOR ORIENTABLE PSEUDO-MANIFOLDS

If P is oriented, it is possible to fix a convention for the quadrilateral labels which determines the Q-matching equations without reference to the abstract neighbourhoods as follows. The orientation of P is pulled back to $\widetilde{\Delta}$. Assign the parameter z_i to any edge of $\widetilde{\Delta}^3_i$, and label the remaining edges according to the convention in Section 2.2 in the unique resulting way with z_i, z_i', z_i'', and denote the quadrilateral types in $\widetilde{\Delta}^3_i$ by q_i, q_i', q_i''. There is a 1–1 correspondence between quadrilateral types in $\widetilde{\Delta}^3_i$ and in $\Delta^3_i = p(\widetilde{\Delta}^3_i)$, and the quadrilateral types in Δ^3_i are denoted by the same symbols.

The orientation of $B(\Delta^1)$ can be chosen such that the quotient map $p_{\Delta^1} : B(\Delta^1) \to P$ is orientation preserving. It can be deduced from Figure 7 that if a 1-simplex in $p^{-1}(\Delta^1)$ has parameter z_i, then the normal quadrilaterals of type q_i' have positive slope on $\partial B(\Delta^1)$, and the normal quadrilaterals of type q_i'' have negative slope. Label the 1-singlices in P by $\Delta^1_1, \ldots, \Delta^1_e$. It follows that if a_{ij} is the number of pre-images of Δ^1_j with label z_i, then the contribution to the Q-matching equation is $a_{ij}(q_i' - q_i'')$. Defining a_{ij}' and a_{ij}'' accordingly, one obtains the *Q-matching equation of Δ^1_j* :

$$0 = \sum_{i=1}^{t} (a_{ij}'' - a_{ij}')q_i + (a_{ij} - a_{ij}'')q_i' + (a_{ij}' - a_{ij})q_i'' , \tag{2.1}$$

and the coefficient matrix of the system of Q-matching equations is given by

$$B = \begin{pmatrix} a_{11}'' - a_{11}' & a_{11} - a_{11}'' & a_{11}' - a_{11} & \ldots & a_{t1}' - a_{t1} \\ \vdots & & & & \vdots \\ a_{1e}'' - a_{1e}' & \ldots & & & a_{te}' - a_{te} \end{pmatrix} . \tag{2.2}$$

3. Boundary matters

It was shown in the first section that a spun-normal surface is uniquely determined by the quadrilateral discs in its cell decomposition, and that it has well-defined boundary curves. This section introduces an algebraic version of the boundary curve map which is defined on the whole solution space to the Q-matching equations. The discussion is restricted to closed pseudo-manifolds merely for pragmatic reasons: more notation and words are required if P is not closed.

3.1 The boundary curve map revisited

Let $\mathfrak{t}$ be a normal triangle in $B_\mathfrak{v}$ and Δ^3 be the 3-singlex which contains $\mathfrak{t}$. Let $\mathfrak{a}$ be a normal arc in $\partial\mathfrak{t}$, and Δ^2 be the face of Δ^3 containing $\mathfrak{a}$. There is a unique quadrilateral type q in Δ^3 such that q and $\mathfrak{t}$ have the same arc type on Δ^2; see Figure 6. Let q be the *Q-modulus of $\mathfrak{a}$ (with respect to $\mathfrak{t}$)*, and give $\mathfrak{a}$ a transverse orientation (with respect to $\mathfrak{t}$) by attaching a little arrow pointing into the interior of $\mathfrak{t}$. This construction is dual to the labelling of the vertices. The orientation conventions of Subsections 1.10 and 1.11 will be used throughout this section. Lift the labelling to $\widetilde{B}_\mathfrak{v}$ if $B_\mathfrak{v}$ is non-orientable; otherwise write $\widetilde{B}_\mathfrak{v} = B_\mathfrak{v}$.

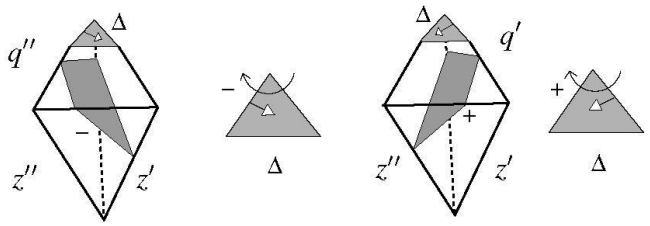

FIGURE 10
Quadrilateral slopes and transverse orientation

Let γ be an oriented path in $\widetilde{B}_\mathfrak{v}$ which is disjoint from the 0-skeleton, has endpoints in the 1-skeleton, and whose interior meets the 1-skeleton transversely. To γ one can associate a linear functional $\nu(\gamma)$ in the quadrilateral types by taking the positive Q-modulus of an edge if it crosses with the transverse orientation, and by taking the negative Q-modulus if it crosses against it (where each edge in $\widetilde{B}_\mathfrak{v}$ is counted twice — using the two adjacent triangles). Evaluating $\nu(\gamma)$ at a solution N to the Q-matching equations, gives a real number $\nu_N(\gamma)$.

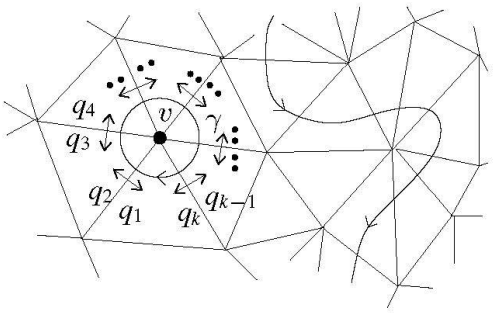

FIGURE 11
$0 = \nu(\gamma) = \sum_{i=1}^{k}(-1)^i q_i$ is the Q-matching equation

Let $\mathfrak{c}$ be a vertex in the triangulation of $B_\mathfrak{v}$ contained in the 1-singlex Δ^1 in P, and consider a pre-image $\tilde{\mathfrak{c}}$ of $\mathfrak{c}$ in $\widetilde{B}_\mathfrak{v}$. Then the linear functional $\nu(\gamma)$ associated to a small circle γ with clockwise orientation around $\tilde{\mathfrak{c}}$ gives the Q-matching equation of Δ^1 by setting $\nu(\gamma) = 0$. This can be deduced from Figure 11.

LEMMA 3.1. *Let γ be an oriented, closed loop in $\widetilde{B}_\mathfrak{v}$ and $N \in Q(\mathcal{T})$. The number $\nu_N(\gamma) \in \mathbf{R}$ depends only on the homotopy class of γ and defines a homomorphism $\nu_{N,\mathfrak{v}}\colon \pi_1(\widetilde{B}_\mathfrak{v}) \to (\mathbf{R}, +)$.*

Proof. Since $\nu_N(\gamma) = 0$ if γ is a small circle about a vertex of $\widetilde{B}_\mathfrak{v}$, it follows that ν_N is well-defined for homotopy classes of loops which intersect the 1-skeleton transversely away from the 0-skeleton. □

Assume that $\widetilde{B}_\mathfrak{v}$ is a closed, orientable surface of genus $g_\mathfrak{v}$. Corresponding to a basis of $H_1(\widetilde{B}_\mathfrak{v})$ choose the system of $2g_\mathfrak{v}$ closed oriented curves shown in Figure 12. Then one obtains a linear map $\nu_\mathfrak{v}\colon Q(\mathcal{T}) \to \mathbf{R}^{2g_\mathfrak{v}}$ defined by

$$\nu_\mathfrak{v}(N) = \left(-\nu_N(\lambda_1), \nu_N(\mu_1), \ldots, -\nu_N(\lambda_{g_\mathfrak{v}}), \nu_N(\mu_{g_\mathfrak{v}})\right). \tag{3.1}$$

Let $g = \sum g(\widetilde{B}_\mathfrak{v})$, and let $\nu\colon Q(\mathcal{T}) \to \mathbf{R}^{2g}$ be the map defined by $\nu = \oplus\nu_\mathfrak{v}$.

There is the following canonical isomorphism $\mathbf{Z}^{2g_\mathfrak{v}} \cong H_1(\widetilde{B}_\mathfrak{v}; \mathbf{Z})$. Let γ be an oriented closed curve on $\widetilde{B}_\mathfrak{v}$. Then $(x_1, y_1, \ldots, x_{g_\mathfrak{v}}, y_{g_\mathfrak{v}}) \in \mathbf{Z}^{2g_\mathfrak{v}}$ maps to $[\gamma] \in H_1(\widetilde{B}_\mathfrak{v}; \mathbf{Z})$ if and only if $\iota(\gamma, \lambda_i) = -x_i$ and $\iota(\gamma, \mu_i) = y_i$ for each i, where ι denotes the algebraic intersection number. This convention identifies the standard system of curves with the standard basis of $\mathbf{Z}^{2g_\mathfrak{v}}$. This extends to a unique homomorphism

$$h_\mathfrak{v}\colon \mathbf{R}^{2g_\mathfrak{v}} \cong H_1(\widetilde{B}_\mathfrak{v}; \mathbf{R}) \to H_1(B_\mathfrak{v}; \mathbf{R}).$$

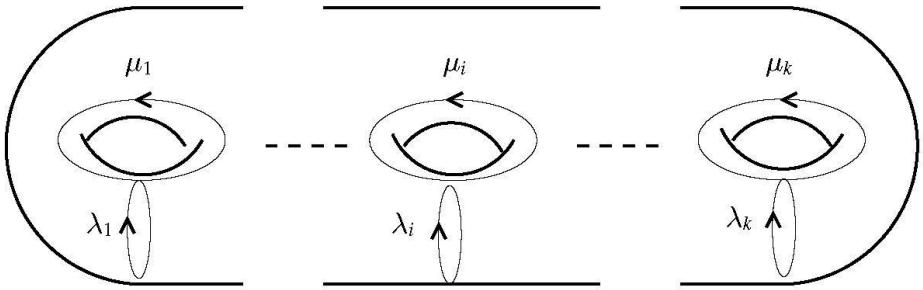

FIGURE 12

The shown surface $\widetilde{S}$ of genus k, $k \geq 0$ even or odd, is placed in $\mathbf{R}^3$ such that it is invariant under reflection in any of the three coordinate planes. It follows that the antipodal map $(x, y, z) \to (-x, -y, -z)$ restricts to an orientation reversing homeomorphism with quotient S so that the quotient map $\widetilde{S} \to S$ is a covering of degree two.

It follows from the definitions that $h_\mathfrak{v}(\mathbf{Z}^{2g_\mathfrak{v}}) \subseteq H_1(B_\mathfrak{v}; \mathbf{Z})$.

DEFINITION 3.2 (Boundary curve map). Let $\partial_\mathfrak{v} = h_\mathfrak{v} \circ \nu_\mathfrak{v}$, and let

$$\partial\colon Q(\mathcal{T}) \to \oplus_\mathfrak{v} H_1(B_\mathfrak{v}; \mathbf{R})$$

be the map defined by $\partial = \oplus_\mathfrak{v} \partial_\mathfrak{v}$.

The map $\nu_\mathfrak{v}$, and hence the boundary curve map, has the following geometric property:

PROPOSITION 3.3. *If N is the normal Q-coordinate of the spun-normal surface S, then $h_\mathfrak{v} \circ \nu_\mathfrak{v}(N) = \partial_\mathfrak{v}(S)$.*

Proof. It suffices to assume that P is oriented. Let S be a spun-normal surface in P, and denote by S' the subcomplex consisting of all normal quadrilaterals in S. Orient the edges of normal quadrilaterals such that the orientation of the normal arc, the corresponding normal vector $n_\mathfrak{v}$ pointing towards the dual vertex and the transverse orientation form a positively oriented basis (see Figure 13). It follows that whenever the quadrilateral disc $\tilde{\mathfrak{q}}$ has positive slope with respect to $\widetilde{\Delta}^1$, then its oriented edges point away from $\tilde{\mathfrak{q}} \cap \widetilde{\Delta}^1$, and if it has negative slope then they point towards the intersection point. It follows that at each point of intersection of S with a 1-singlex, the number of oriented edges pointing towards it equals the number pointing away from it. In particular, one obtains an oriented 1-chain c of quadrilateral edges in $\partial S'$ along which subcomplexes made up of triangles meet S' in S.

The normal vector $n_\mathfrak{v}$ and the orientation of P determine an orientation of each vertex linking surface $B_\mathfrak{v}$. Give each transversely oriented curve in $B_\mathfrak{v} \cap S$

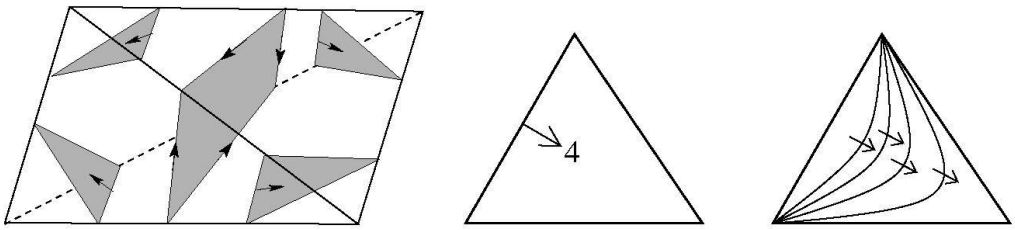

(a) Orientation of quadrilateral sides (b) Construction and labelling of oriented arcs

FIGURE 13
Oriented 1-chains of quadrilateral edges

the orientation induced from the orientation of B_v. Then $[B_v \cap S] = \partial_v(S) \in H_1(B_v; \mathbf{Z})$. It suffices to show that c can be viewed as a union of oriented closed curves $\bigcup c_i$ with the property that each c_i is homotopic in $S \cup B_v$ to an oriented simple closed curve h_i in B_v such that $[B_v \cap S] = \sum [h_i]$.

The following construction determines $\bigcup c_i$ and $\bigcup h_i$. Given $N = N(S)$, associate a modulus and a transverse orientation to 1-singlices in B_v as follows. Identify Q-moduli with the values given by N. Let s be an edge in B_v and let t and t' be the triangles in B_v meeting in s, and n and n' be the associated Q-moduli. If $n = n'$, then give s the modulus zero and no transverse orientation. If $n > n'$, then give s the modulus $n - n'$ and the transverse orientation inherited from t, and if $n < n'$, then give s the modulus $n' - n$ and the transverse orientation inherited from t'. The modulus of s determines the number of quadrilateral edges of the same arc type as s which are identified with edges of normal triangles in S, and the transverse orientation determines the 3-singlex containing the corresponding quadrilaterals (see Figure 14).

Let c' be the union of edges in B_v with non-zero modulus. A unique collection of transversely oriented simple closed curves on B_v is derived from c' as follows. If s is an edge in c' with modulus n and transverse orientation pointing into t, place n arcs in t, each of which has endpoints identical to s, any two of which only meet in their endpoints, and give all of them the induced transverse orientation (pointing towards the vertex of t opposite s; see Figure 13). These arcs are in a natural bijective correspondence with edges of quadrilaterals in S meeting edges of triangles. A bijection between endpoints of arcs at a vertex v of B_v is now defined which corresponds to the identifications of quadrilaterals and triangles in S. Please refer to Figure 14(b) and 14(d). Let H_v be the union of all triangles

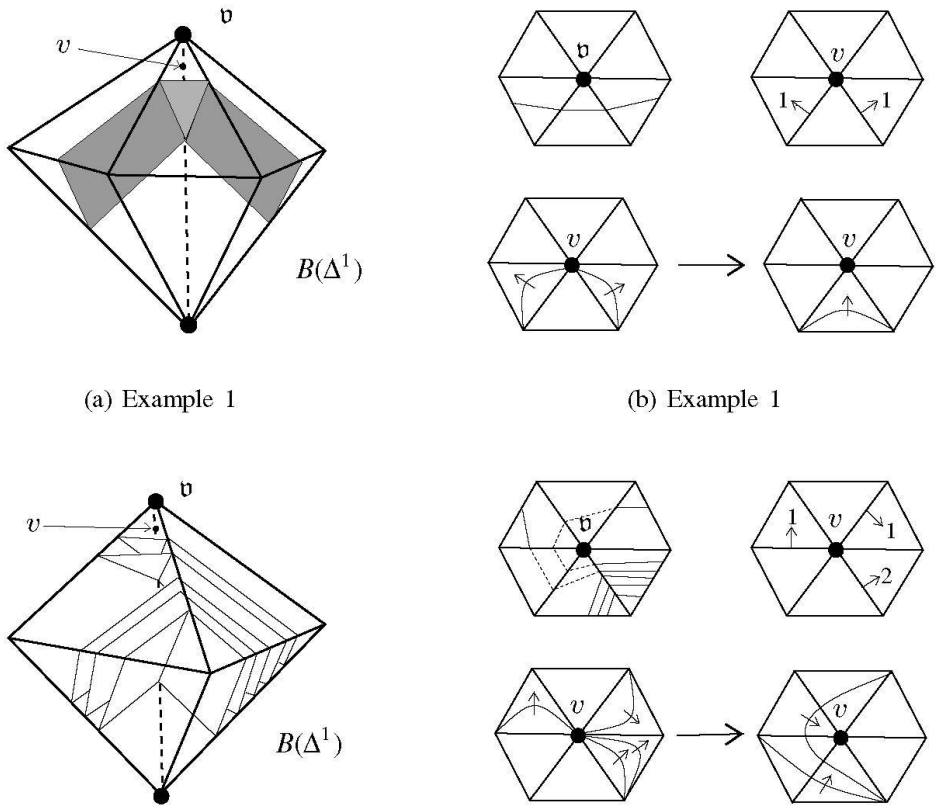

(a) Example 1 (b) Example 1

(c) Example 2 (d) Example 2

FIGURE 14

Two examples; Figures (b) and (d) show: (Top left) The pattern on the upper hemisphere of $B(\Delta^1)$; (Top right) The edge labelling on $B_\mathfrak{v}$; (Bottom) Introducing arcs and resolving intersection points at v

in $B_\mathfrak{v}$ containing v, and let A be the set of arcs which have not been paired yet. An *outermost pair* γ and δ satisfies the following two criteria (this is a variant of constructions in [12, 8]):

1. If the endpoints of γ and δ are identified at v, then the transverse orientations of γ and δ match.
2. The disc cut out from $H_\mathfrak{v}$ by $\gamma \cup \delta$ which the transverse orientation of δ points away from does not contain any arcs in A.

If an outermost pair is found, identify their endpoints and isotope the union away from v (opposite to their transverse orientation). Then repeat the above by replacing A by $A - \{\gamma, \delta\}$. Thus, after glueing and isotopy, one obtains a unique disjoint collection $\mathcal{C}$ of simple closed curves, each with a transverse

orientation. Orient these curves using the above convention for orientation of Q-edges. It follows from the construction that:

1. Each $c_i \in \mathcal{C}$ is homotopic in $S \cup B_\mathfrak{v}$ to a unique closed (possibly not simple) curve c_i'' of quadrilateral edges in $\partial S'$ (e.g. introduce labels for the arcs specifying a 1- and a 2-simplex of $B_\mathfrak{v}$ and a normal quadrilateral). In particular, c_i'' inherits a well-defined orientation from c_i which agrees with the orientation of the quadrilateral edges.

2. The element in $H_1(B_\mathfrak{v}; \mathbf{Z})$ represented by $\mathcal{C}$ is $\partial_\mathfrak{v}(S)$.

This completes the proof. □

3.2 Spun-normal branched immersions

Proof of Proposition 2.7. Let $N = (x_i)$ be a point in $Q(\mathcal{T})$ with non-negative integral coordinates. If P is not closed, one may double P along its boundary and construct a spun-normal branched immersion corresponding to the doubled solution (noting that doubling reverses signs of corners). The restriction of this to either half gives the desired map. Hence assume that P is closed. At first, a normal branched immersion is constructed which is not in general position. For each quadrilateral type q_i place x_i copies of this type in P such that all vertices are barycentres of 1-singlices; this is viewed as a map $f_q\colon S_q \to P$, where S_q is an abstract union of $\sum x_i$ pairwise disjoint quadrilateral discs. Choose each $B_\mathfrak{v}$ such that its intersection with the 1-skeleton consists of barycentres.

Using the construction in the proof of Proposition 3.3, a finite family, $\mathcal{C}_\mathfrak{v}$, of pairwise disjoint transversely oriented simple closed curves in $B_\mathfrak{v}$ can be constructed with the property that each curve in $\mathcal{C}_\mathfrak{v}$ is homotopic to a unique curve in the 1-skeleton of $B_\mathfrak{v}$; denote the resulting homotopy taking $\mathcal{C}_\mathfrak{v}$ into the 1-skeleton of $B_\mathfrak{v}$ by $h_\mathfrak{v}\colon B_\mathfrak{v} \times I \to B_\mathfrak{v}$.

Note that the spinning construction can be applied to any finite collection of pairwise disjoint, transversely oriented, non-separating (not necessarily normal) curves on $B_\mathfrak{v}$. After performing the straightening, the result is a union, $S_\mathfrak{v}''$, of normal discs in $\overline{N}_\mathfrak{v}$ and normal discs in P which meet $B_\mathfrak{v}$ in normal cells; $S_\mathfrak{v}'' \cap N_\mathfrak{v}$ is properly embedded in $N_\mathfrak{v}$. Applying this spinning construction to $\mathcal{C}_\mathfrak{v}$ yields a subset $F_\mathfrak{v}$ of $\overline{N}_\mathfrak{v}$. Then $h_\mathfrak{v}$ can be extended in a neighbourhood of $B_\mathfrak{v}$ in $N_\mathfrak{v}$ such that $h_\mathfrak{v}(F_\mathfrak{v}, 1)$ meets P in normal triangles and $F_\mathfrak{v} \cap N_\mathfrak{v} = h_\mathfrak{v}(F_\mathfrak{v}, 0) \cap N_\mathfrak{v}$ is isotopic to $h_\mathfrak{v}(F_\mathfrak{v}, 1) \cap N_\mathfrak{v}$. Let $S_\mathfrak{v}'$ be the union of all connected components of $h_\mathfrak{v}(F_\mathfrak{v}, 1)$ which are not vertex linking. There is a triangulated (possibly non-compact) surface $S_\mathfrak{v}$ and a simplicial isomorphism $f_\mathfrak{v}\colon S_\mathfrak{v} \to S_\mathfrak{v}'$.

Lift $S' = f_q(S_q) \cup_{\mathfrak{v}} f_{\mathfrak{v}}(S_{\mathfrak{v}})$ to $\widetilde{\Delta}$. Let $\mathfrak{a}$ be a normal arc in P with endpoints barycentres of 1-singlices, and denote its pre-images in $\widetilde{\Delta}$ by $\mathfrak{a}_0$ and $\mathfrak{a}_1$. It follows from the definition of moduli that there are as many normal discs in $p^{-1}(S')$ containing $\mathfrak{a}_0$ as there are normal discs containing $\mathfrak{a}_1$. Moreover, these numbers are finite. One can therefore define an arbitrary bijection between these two sets which determines a unique bijection between the corresponding boundary arcs in $S_q \cup_{\mathfrak{v}} S_{\mathfrak{v}}$. Do this for all such normal arcs and denote the resulting quotient by S. Then S is a triangulated (possibly non-compact) surface S and there is a well-defined normal branched immersion $f\colon S \to P$ obtained from glueing the maps f_q and $f_{\mathfrak{v}}$. The set of branched points is contained in the pre-image of all barycentres; hence there are finitely many branched points. It follows from the construction that $x(f) = N$, and the map f can be homotoped into general position if desired. □

3.3 Analysis of the linear maps and Proof of Theorem 0.1

If $B_{\mathfrak{v}}$ is non-orientable, then it may be assumed that the non-trivial deck transformation $\widetilde{P} \to \widetilde{P}$ induces an involution σ on $\widetilde{B}_{\mathfrak{v}}$ which coincides with the antipodal map in Figure 12. Whence $\sigma\mu_i = \mu_{g_{\mathfrak{v}}+1-i}^{-1}$ and $\sigma\lambda_i = \lambda_{g_{\mathfrak{v}}+1-i}$. Denote the resulting involution on $H_1(\widetilde{B}_{\mathfrak{v}})$ by σ also. It follows that $\nu_N(\mu_i) = \nu_N(\sigma\mu_i) = -\nu_N(\mu_{g_{\mathfrak{v}}+1-i})$, and $\nu_N(\lambda_i) = \nu_N(\sigma\lambda_i) = \nu_N(\lambda_{g_{\mathfrak{v}}+1-i})$. This observation yields the following upper bound:

Lemma 3.4. $\dim(\operatorname{im}\nu) \leq 2\chi(P) - v_n$.

Proof. If $B_{\mathfrak{v}}$ is orientable, then $2g_{\mathfrak{v}} = 2 - \chi(\widetilde{B}_{\mathfrak{v}}) = 2 - \chi(B_{\mathfrak{v}})$.

If $B_{\mathfrak{v}}$ is non-orientable, then $g_{\mathfrak{v}} = 1 - \chi(B_{\mathfrak{v}})$. It follows from the description of τ that the kernel of $\nu_{\cdot,\mathfrak{v}}$ has dimension at least $g_{\mathfrak{v}}$ if $g_{\mathfrak{v}}$ is even. If $g_{\mathfrak{v}} = 2k+1$ is odd, then $\nu_N(\mu_{k+1}) = \nu_N(\sigma\mu_{k+1}) = -\nu_N(\mu_{k+1})$ implies $\nu_N(\mu_{k+1}) = 0$, and the kernel of $\nu_{\cdot,\mathfrak{v}}$ has again dimension at least $g_{\mathfrak{v}}$.

Thus:

$$\begin{aligned}\dim(\operatorname{im}\nu) &\leq \sum_{\text{or'ble}} 2 - \chi(B_{\mathfrak{v}}) + \sum_{\text{non-or'ble}} 1 - \chi(B_{\mathfrak{v}}) \\ &= 2v_o + v_n - \chi(\partial P^c) = 2\chi(P) - v_n,\end{aligned}$$

since P is closed and hence $\chi(\partial P^c) = 2v - 2\chi(P)$. □

Lemma 3.5. *One has* $\ker\nu = \operatorname{im} pr$. *Moreover, if* $N \in \ker\nu$, *then* $pr^{-1}(N) = L + \mathit{span}\{N_\Delta(B_1), \ldots, N_\Delta(B_v)\}$ *for some* $L \in C(\mathcal{T})$. *In particular, the image of* $pr\colon C(\mathcal{T}) \to Q(\mathcal{T})$ *has dimension* $t + e - v$.

Proof. First assume that $B_\mathfrak{v}$ is orientable. Let $N \in Q(\mathcal{T})$. Assume that the induced triangulation $\mathcal{T}_\mathfrak{v}$ of $B_\mathfrak{v}$ supports the triangle coordinates $t_0, \ldots, t_k$. If a value is assigned to t_0, then the Q-moduli of the common edges uniquely determine values for the adjacent triangle coordinates such that the matching equations are satisfied. It needs to be shown that these assignments are well defined globally if and only if $N \in \ker \nu$.

Indeed, choose a closed, oriented, simple path γ in $B_\mathfrak{v}$, disjoint from the 0-skeleton, transverse to the 1-skeleton and meeting 2-cells in at most one normal arc. Let t be a 2-cell of $B_\mathfrak{v}$ which γ passes through, and assume that t_0 is the associated triangle coordinate. If t_0 is given any value, then the adjacent coordinate t_i of the next triangle that γ passes through must take the value $t_i = (q_0 - q_i) + t_0$, where $-(q_0 - q_i)$ is the oriented sum of the Q-moduli associated to the edge crossed by γ. In this way, uniquely determined values are given to the coordinates of triangles traversed by γ, and upon returning to the initial triangle, one has the equation $t_0 = -\nu_N(\gamma) + t_0$, which is satisfied if and only if $\nu_N(\gamma) = 0$. Similarly, values can be assigned to the remaining triangle coordinates supported by $\mathcal{T}_\mathfrak{v}$. This procedure is well defined, if and only if $\nu_{N,\mathfrak{v}} \colon \pi_1(B_\mathfrak{v}) \to (\mathbf{R}, +)$ is the trivial homomorphism. Changing the initial value assigned to t_0 changes the triangle coordinates by a multiple of $N_\Delta(B_\mathfrak{v})$.

Now assume that $B_\mathfrak{v}$ is non-orientable. Then the above can be done for the lift of the triangulation of P to the orientable double cover $\widetilde{P}$. Let $\tilde{t}_0$ and $\tilde{t}_1$ be normal triangles on $\widetilde{B}_\mathfrak{v}$ which are exchanged by the non-trivial deck transformation σ. Let γ be a closed normal path in $\widetilde{B}_\mathfrak{v}$ passing through $\tilde{t}_0$ and $\tilde{t}_1$ which is invariant under σ. Then γ can be written as the union of arcs a and $\sigma(a)$ with $\nu_N(a) = \nu_N(\sigma(a))$. Then $0 = \nu_N(\gamma) = \nu_N(a) + \nu_N(\sigma(a)) = 2\nu_N(a)$. Whence the above procedure assigns the same values to $\tilde{t}_0$ and $\tilde{t}_1$. It follows that the normal triangles constructed in $\widetilde{P}$ can be isotoped to be invariant under σ. □

LEMMA 3.6. *The matrix B has rank $e - v_o$, and we have* $\dim Q(\mathcal{T}) = 3t - e + v_o = 2t + \chi(P) - v_n$ *and* $\dim \operatorname{im} \nu = 2\chi(P) - v_n$.

Proof. It will first be shown that the rank of B is at most $e - v_o$. The rank of B equals $3t - \dim \ker(B) = e - \dim \ker B^T$. It therefore suffices to find v_o linearly independent elements in $\ker B^T$. The dual system of equations described by the transpose has one equation for each quadrilateral type and one variable, x_{Δ^1}, for each 1-singlex Δ^1 in P. The equation associated to q is $0 = \sum s_{\Delta^1}(q) x_{\Delta^1}$, where the sum is taken over all 1-singlices in P and $s_{\Delta^1}(q)$ is the total slope of q with respect to Δ^1 introduced in Section 2.1.

Assume that P is orientable. In this case, q has the same slope at opposite normal corners, and opposite signs at normal corners sharing a normal arc. Let S be a vertex linking surface. Then the assignment $x_{\Delta^1}(S) = \#(S \cap \Delta^1)$ defines an element in $\ker B^T$. Since each triangle disc is uniquely determined by its intersection with the 1-simplices, it follows that the set of v solutions thus obtained from the v vertex linking surfaces is linearly independent, giving rank at most $e - v$.

Hence assume that P is non-orientable. Considering its lift to the orientable double cover, one finds that if the edge of a quadrilateral disc in P is normally isotopic into an orientable vertex linking surface, then the signs at the respective corners are opposite. This implies that the above argument can be applied to an orientable component of ∂P^c. It does not, however, work for non-orientable components. This gives rank at most $e - v_o$ in general.

Thus, $\dim Q(\mathcal{T}) \geq 3t - e + v_o = 2t + \chi(P) - v_n$. It needs to be shown that this is also an upper bound. One has $2t + \chi(P) - v_n \leq \dim Q(\mathcal{T}) = \dim \ker \nu + \dim \operatorname{im} \nu = \dim \operatorname{im} pr + \dim \operatorname{im} \nu = t + e - v + \dim \operatorname{im} \nu$. Thus $2\chi(P) - v_n \geq \dim \operatorname{im} \nu \geq 2\chi(P) - v_n$. This forces equality and hence the conclusion. □

LEMMA 3.7. $\ker \partial = \ker \nu$.

Proof. It follows from the definition that $\ker \nu \subset \ker \partial$. Assume that $\partial(N) = 0$. If $B_\mathfrak{v}$ is orientable, then $h_\mathfrak{v}$ is an isomorphism which implies $\nu_\mathfrak{v}(N) = 0$. Hence assume that $B_\mathfrak{v}$ is non-orientable.

Since $\dim \operatorname{im} \nu = 2\chi(P) - v_n$, it suffices to show that whenever δ represents the homotopy class of an oriented 1-sided simple closed curve in $B_\mathfrak{v}$ with the property that $\nu_\mathfrak{v}(N)$ determines $\tilde{\delta}$, then $\iota(\gamma, \tilde{\delta}) = 0$ for each standard generator.

Indeed, since δ is 1-sided, we have $\sigma(\tilde{\delta}) = \tilde{\delta}$. Thus, $\iota(\tilde{\delta}, \gamma) = \nu_N(\gamma) = \nu_N(\sigma(\gamma)) = \iota(\tilde{\delta}, \sigma(\gamma)) = \iota(\sigma(\tilde{\delta}), \sigma(\gamma)) = -\iota(\tilde{\delta}, \gamma)$. Whence $\iota(\tilde{\delta}, \gamma) = 0$ for each standard generator, which gives $\nu_\mathfrak{v}(N) = 0$. □

Proof of Theorem 0.1. The first part of the statement follows from the above lemmata in conjunction with the discussion in Subsection 2.6.

Since $\ker \partial = \ker \nu$, it follows that $\dim \operatorname{im} \nu = \dim \operatorname{im} \partial$, and hence that ∂ is surjective. Since ν is defined over the integers, its restriction to integer lattice points in $Q(\mathcal{T})$ has image in $\mathbf{Z}^{2g}$, whence ∂ has image of finite index in $\oplus_\mathfrak{v} H_1(B_\mathfrak{v}; \mathbf{Z})$. □

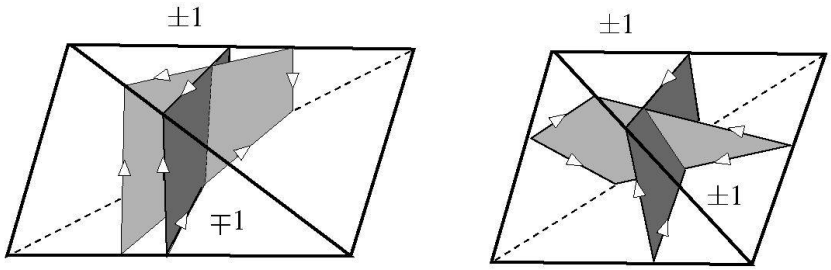

FIGURE 15
Orientation of quadrilateral sides and intersection numbers

3.4 INTERSECTION NUMBERS AND PROOF OF THEOREM 0.2

Assume that P is oriented. The material in this subsection is based on the observation that one can determine the algebraic intersection number of the oriented boundary curves of two spun-normal surfaces from their normal Q-coordinates. The intersection pairing on $H_1(B_\mathfrak{v};\mathbf{R})$ pulls back to a bi-linear, skew-symmetric pairing $\star$ on $\mathbf{R}^{2g_\mathfrak{v}}$. Taking sums gives a pairing $\star$ on $\mathbf{R}^{2g}$ with the property that for any $N, L \in Q(\mathcal{T})$, we have:

$$\nu(N) \star \nu(L) = \sum_{\mathfrak{v} \in P^{(0)}} \nu_\mathfrak{v}(N) \star \nu_\mathfrak{v}(L) = \sum_{\mathfrak{v} \in P^{(0)}} \iota\Big(\partial_\mathfrak{v}(N), \partial_\mathfrak{v}(L)\Big) .$$

To obtain a corresponding bi-linear, skew-symmetric form on $Q(\mathcal{T})$, let

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

and let C_t be the $(3t \times 3t)$ block diagonal matrix with t copies of C on its diagonal. Then for any $N, L \in Q(\mathcal{T})$, define:

$$\langle N, L \rangle = N^T C_n L . \tag{3.2}$$

LEMMA 3.8. *Let P be an oriented, closed pseudo-manifold. For all $N, L \in Q(\mathcal{T})$:*

$$\langle N, L \rangle = \nu(N) \star \nu(L) .$$

Proof. One can compute $\langle N, L \rangle$ by summing intersection numbers obtained for normal quadrilaterals over all 3-singlices as follows (see Figure 15). If N and L have different non-zero Q-coordinates in some 3-singlex, the associated normal quadrilaterals are naturally transverse, and the contribution to $\langle N, L \rangle$ is (up to sign) twice the product of the Q-coordinates, and the sign

is determined by the order of the surfaces. If they have the same non-zero Q-coordinate in some tetrahedron, the contribution is zero, independent of the chosen transverse intersection. Given the map ν and the orientation conventions, this gives precisely the intersection pairing in homology of vertex links. □

Proof of Theorem 0.2. If R is a maximal convex polytope in $PF(\mathcal{T})$, then it is the intersection of a $(\dim R + 1)$-dimensional vector space R' with the unit simplex. The set $PF(\mathcal{T})$ is obtained by setting in turn two coordinates from each 3-singlex equal to zero. Thus R' is the intersection of $Q(\mathcal{T})$ with a t-dimensional subspace of $\mathbf{R}^{3t}$, and its dimension is at least $(2t + \chi(P) - v_n) + t - 3t = \chi(P) - v_n$.

Assume that P is oriented. For any $N, L \in R'$, one has $\nu(N) \star \nu(L) = \langle N, L \rangle = 0$. The image $\nu(R')$ therefore lies in a self-annihilating subspace of $\mathbf{R}^{2g}$. Its dimension is thus at most $g = \chi(P)$, giving $\dim R' \leq \chi(P) + \dim(R' \cap \ker \nu)$. Thus $\dim R = \dim R' - 1 \leq \chi(P) + \dim(R' \cap \ker \nu) - 1 = \chi(P) + \dim(R \cap \ker \nu)$.

If P is non-orientable, the previous paragraph applies to $\widetilde{P}$, and we have $\chi(\widetilde{P}) = 2\chi(P) - v_n$. This completes the proof. □

4. Examples

A general procedure to determine the boundary slope of a spun-normal surface on a vertex linking torus is given. Then the figure eight knot complement and the Gieseking manifold are discussed. These examples can be found in [9].

4.1 Boundary slopes on tori

If $B_\mathfrak{v}$ is a torus, then the boundary curve map gives a convenient way of determining the boundary curves of a spun-normal surface S from its normal Q-coordinate $N = N(S)$. In this case $B_\mathfrak{v}$ may be chosen such that it meets S in a finite family of parallel simple closed curves, all with the same transverse orientation, and it meets $N_\mathfrak{v}$ in the same number of parallel open annuli. The homology class of such a curve is termed the *slope* of S on $B_\mathfrak{v}$. Since $N_\mathfrak{v} \setminus \{\mathfrak{v}\}$ is orientable, one can choose generators $\{\lambda, \mu\}$ for $\pi_1(B_\mathfrak{v})$ such that together with a normal vector n pointing into $N_\mathfrak{v}$, $\{\lambda, \mu, n\}$ is a positively oriented basis for $TN_\mathfrak{v}$. This corresponds to the orientation

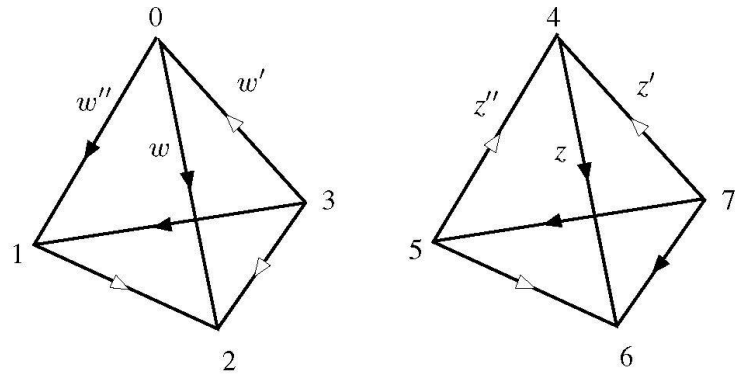

FIGURE 16
An ideal triangulation of the figure eight knot complement

convention of a standard meridian-longitude pair of a knot complement. Since ν_N is a homomorphism, the following holds under the above assumptions.

1. If $\nu_N(\mu) = \nu_N(\lambda) = 0$, then S is disjoint[1]) from $B_\mathfrak{v}$.
2. If $\nu_N(\mu) \neq 0$ or $\nu_N(\lambda) \neq 0$, then let $d > 0$ denote the greatest common divisor of the numbers $|\nu_N(\mu)|$ and $|\nu_N(\lambda)|$. Put $p = -\nu_N(\lambda)/d$ and $q = \nu_N(\mu)/d$. Then $s = \mu^p\lambda^q$ has $\nu_N(s) = 0$ and hence it is a boundary slope of S. Furthermore, S has d boundary curves on $B_\mathfrak{v}$.

Now assume that each vertex linking surface is a torus; label the surfaces $B_\mathfrak{v}$ by $T_1, \ldots, T_v$. Choose generators $\{\lambda_i, \mu_i\}$ for each T_i as above. The oriented boundary curves of a spun-normal surface S with normal Q-coordinate $N = N(S)$ are then determined by the vector:

$$(-\nu_N(\lambda_1), \nu_N(\mu_1), \ldots, -\nu_N(\lambda_v), \nu_N(\mu_v)) \in \mathbf{Z}^{2v}, \tag{4.1}$$

and boundary curves of linear combinations of compatible surfaces can be determined using linear combinations of vectors of the form (4.1) since ν is additive. Moreover, the intersection pairing has the following form:

$$\nu(N) \star \nu(L) = \frac{1}{2}\sum_{i=1}^{v} \nu_N(\mu_i)\nu_L(\lambda_i) - \nu_N(\lambda_i)\nu_L(\mu_i). \tag{4.2}$$

4.2 THE FIGURE EIGHT KNOT COMPLEMENT

Let M denote the complement of the figure eight knot. An oriented, ideal triangulation of M is encoded in Figure 16. Since M is oriented, the

[1]) This is also shown in the file normal_surface_construction.c in [14].

convention given in Subsection 2.9 will be used, with the quadrilateral types dual to $w^{(k)}$ and $z^{(k)}$ denoted by $p^{(k)}$ and $q^{(k)}$ respectively. One computes:

$$B = \begin{pmatrix} -1 & -1 & 2 & -1 & -1 & 2 \\ 1 & 1 & -2 & 1 & 1 & -2 \end{pmatrix},$$

which determines a single Q-matching equation:

$$0 = p + p' - 2p'' + q + q' - 2q'' .$$

This implies that the space $PQ(\mathcal{T})$ is four-dimensional.

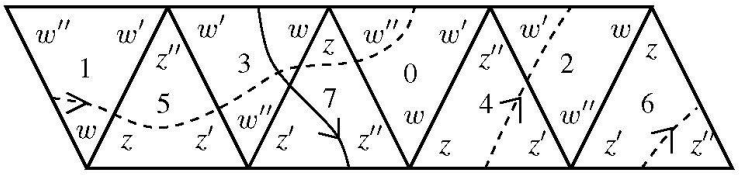

FIGURE 17

The induced triangulation of the vertex linking torus, where the sides of the rectangle are identified by translations parallel to its sides and triangle i is dual to vertex i in Figure 16. The shown elementary curves are the standard meridian (solid) and longitude (dashed).

The induced triangulation of the vertex linking surface is shown in Figure 17 and used to determine the linear functionals associated to the standard peripheral curves:

$$\nu(\lambda) = 2p + 2p' - 4p'' ,$$
$$\nu(\mu) = -p' + p'' - q + q'' .$$

The normal Q-coordinates of closed normal surfaces satisfy $\nu(\lambda) = \nu(\mu) = 0$, whence any closed (embedded) normal surface is vertex linking. It follows that $PF(\mathcal{T})$ is zero-dimensional. A direct calculation reveals that there are four projective classes of admissible solutions; all have minimal representative a once-punctured Klein bottle. Their normal Q-coordinates and boundary slopes are listed in Table 1. This calculation in particular shows that no spun-normal surface is a Seifert surface for the knot.

4.3 THE GIESEKING MANIFOLD

The Gieseking manifold, M', is double covered by the figure eight knot complement, M, and the covering transformation is encoded by the involution (05)(14)(26)(37) on the vertices in Figure 16. The resulting ideal triangulation $\mathcal{T}'$ of M' has one ideal 3-singlex and one ideal 1-singlex. Denote the three

TABLE 1
Normal surface in the figure eight knot complement

solution	$\nu(\mu)$	$\nu(\lambda)$	slope
$(2,0,0,0,0,1)$	1	4	-4
$(0,2,0,0,0,1)$	-1	4	4
$(0,0,1,2,0,0)$	-1	-4	-4
$(0,0,1,0,2,0)$	1	-4	4

quadrilateral types in M' by r, r', r'', where $r^{(i)}$ lifts to $p^{(i)}$. The Q-matching equation is $r + r' - 2r'' = 0$. It can be worked out from the triangulation or by observing that the induced involution on quadrilateral types in M is $(p\ q')(p'\ q)(p''\ q'')$. Thus, $\dim PQ(\mathcal{T}') = 1$ and $PF(\mathcal{T}') = \varnothing$.

The boundary curve map is defined via the induced triangulation of the double cover of the Klein bottle; using the generators from the above section, one has: $\nu(\lambda) = -2r - 2r' + 4r'' = 0$ and $\nu(\mu) = -2r' + 2r''$. Generators λ', μ' can be chosen for $H_1(B_\mathfrak{v}) \cong \mathbf{Z} \oplus \mathbf{Z}_2$ such that the map $H_1(\widetilde{B}_\mathfrak{v}) \to H_1(B_\mathfrak{v})$ is given by $\lambda \to \lambda'$ and $\mu \to (\mu')^2 = 0$. The composition

$$Q(\mathcal{T}') \to \mathbf{Z}^2 \to H_1(\widetilde{B}_\mathfrak{v}; \mathbf{R}) \to H_1(B_\mathfrak{v}; \mathbf{R})$$

is then

$$N \to (-\nu_N(\lambda), \nu_N(\mu)) = (0, \nu_N(\mu)) \to \nu_N(\mu)\lambda \to \nu_N(\mu)\lambda' .$$

Since $\nu(\mu) = -2r' + 2r''$, it follows that the map $\partial\colon Q(\mathcal{T}') \to H_1(B_\mathfrak{v}; \mathbf{R})$ is surjective. Its restriction to integral points in $Q(\mathcal{T}')$ has image of index two in $H_1(B_\mathfrak{v}; \mathbf{Z})$, which gives a subgroup of index four in $H_1(B_\mathfrak{v})$.

REFERENCES

[1] BACHMAN, D. A note on Kneser-Haken finiteness. *Proc. Amer. Math. Soc. 132* (2004), 899–902

[2] GORDON, C. McA. *The Theory of Normal Surfaces.* Lecture notes typeset by Richard Kent, available at http: //www.math.brown.edu/∼rkent/ (2001).

[3] JACO, W. and U. OERTEL. An algorithm to decide if a 3-manifold is a Haken manifold. *Topology 23* (1984), 195–209.

[4] JACO, W. and J. H. RUBINSTEIN. PL equivariant surgery and invariant decompositions of 3-manifolds. *Adv. in Math. 73* (1989), 149–191.

[5] JACO, W. and J. H. RUBINSTEIN. 0-efficient triangulations of 3-manifolds. *J. Differential Geom. 65* (2003), 61–168.

[6] GABAI, D. Essential laminations and Kneser normal form. *J. Differential Geom. 53* (1999), 517–574 .

[7] HATCHER, A. *Algebraic Topology.* Cambridge University Press, Cambridge, 2002.

[8] KANG, E. Normal surfaces in non-compact 3-manifolds. *J. Aust. Math. Soc. 78* (2005), 305–321.

[9] KANG, E. and J.H. RUBINSTEIN Ideal triangulations of 3-manifolds I: spun normal surface theory. *Geom. Topol. Monogr. 7*, Proceedings of the Casson Fest (2004), 235–265.

[10] MATVEEV, S. *Algorithmic Topology and Classification of 3-Manifolds.* Algorithms and Computation in Mathematics *9*. Springer-Verlag, Berlin, 2003.

[11] THOMPSON, A. Thin position and the recognition problem for S^3. *Math. Res. Lett. 1* (1994), 613–630.

[12] TOLLEFSON, J.L. Normal surface Q-theory. *Pacific J. Math. 183* (1998), 359–374.

[13] WALSH, G. Incompressible surfaces and spunnormal form. Preprint arXiv: math.GT/0503027 (2005).

[14] WEEKS, J. *SnapPea*, available at www.geometrygames.org.

(Reçu le 14 juin 2007)

Stephan Tillmann

Department of Mathematics and Statistics
The University of Melbourne
VIC 3010,
Australia
e-mail: tillmann@ms.unimelb.edu.au