

TRAVERSING THREE-MANIFOLD TRIANGULATIONS AND SPINES

J. HYAM RUBINSTEIN, HENRY SEGERMAN AND STEPHAN TILLMANN

ABSTRACT. A celebrated result concerning triangulations of a given closed 3–manifold is that any two triangulations with the same number of vertices are connected by a sequence of so-called 2-3 and 3-2 moves. A similar result is known for ideal triangulations of topologically finite non-compact 3–manifolds. These results build on classical work that goes back to Alexander, Newman, Moise, and Pachner. The key special case of 1–vertex triangulations of closed 3–manifolds was independently proven by Matveev and Piergallini. The general result for closed 3–manifolds can be found in work of Benedetti and Petronio, and Amendola gives a proof for topologically finite non-compact 3–manifolds. These results (and their proofs) are phrased in the dual language of spines.

The purpose of this note is threefold. We wish to popularise Amendola’s result; we give a combined proof for both closed and non-compact manifolds that emphasises the dual viewpoints of triangulations and spines; and we give a proof replacing a key general position argument due to Matveev with a more combinatorial argument inspired by the theory of subdivisions.

1. INTRODUCTION

Suppose you have a finite number of triangles. If you identify edges in pairs such that no edge remains unglued, then the resulting identification space looks locally like a plane and one obtains a closed surface, a two-dimensional manifold without boundary. The classification theorem for surfaces, which has its roots in work of Camille Jordan and August Möbius in the 1860s, states that each closed surface is topologically equivalent to a sphere with some number of handles or crosscaps. For example, the torus is a sphere with a single handle, and the projective plane is a sphere with a single crosscap. A modern proof of the classification theorem is due to John Conway and presented by Francis and Weeks [21].

Now suppose you have a finite number of tetrahedra and identify triangular faces in pairs such that no face remains unglued. In this case, the resulting identification space is again closed (it has no boundary). The resulting space looks everywhere like three-dimensional euclidean space *except* possibly at the vertices and at mid-points of edges. This can be seen as follows.

The local picture near a vertex in the identification space can be understood by tracing the small triangles in the tetrahedra cutting off the vertices that are identified under the face gluings (see Figure 1). If these small triangles globally glue up to a sphere, then in the identification space the sphere formed by the triangles bounds a ball and the vertex in the identification space looks like the centre of a ball. In this case, the vertex is called *material*. However, if the triangles glue up to a sphere with at least one handle or crosscap, then the space near this vertex looks like a cone on this surface. Such a vertex is called *ideal*.

Only the cone on a sphere is a ball; the cone on a different surface is more difficult to visualise. In general one can imagine this cone as obtained as follows. First take a product of the surface S with an interval I . The resulting three-dimensional space has two boundary components, each a

2000 *Mathematics Subject Classification.* 57Q15, 57Q25, 52B70, 57N10, 57M50, 57M27.

Key words and phrases. triangulation, 3-manifold, surface, combinatorial topology, triangulated manifolds, bistellar moves, standard spines.

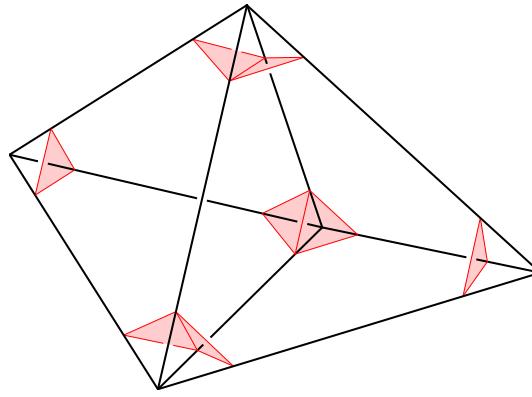


FIGURE 1. Small triangles near the vertices of the tetrahedra glue up to form a surface.

copy of S . One then collapses one of these boundary components to a point, giving the cone on the surface. If S is orientable, one can picture the product $S \times I$ in three-dimensional Euclidean space. If S is a sphere, it is now also easy to see that after collapsing one of the boundary components, one obtains a ball.

The situation at the midpoints of edges arises from the issue that an edge may be identified with itself but in the opposite direction. In this case, a small neighbourhood of the midpoint of such an edge is bounded by a projective plane (see Figure 2). If this happens, we can turn the mid-point into an ideal vertex by subdividing the tetrahedra.

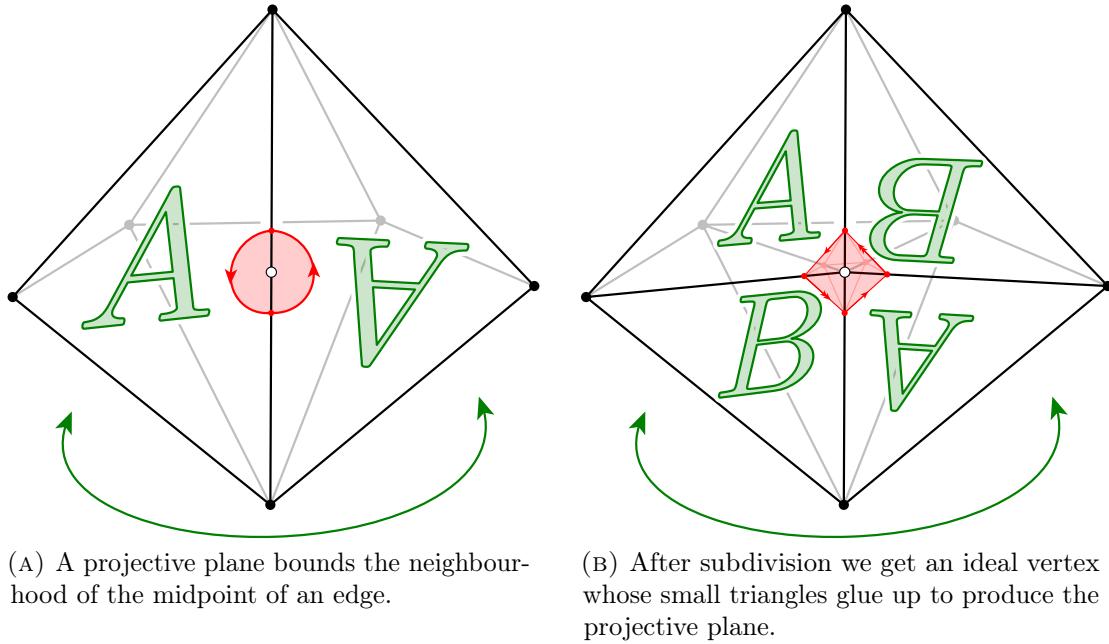


FIGURE 2

Hence we will assume that when we identify faces of a collection of tetrahedra, then the resulting space looks *everywhere* like three-dimensional Euclidean space *except* possibly at the vertices. If all vertices are material, we have a *three-dimensional manifold* and if there is at least one ideal

vertex, we have a *three-dimensional pseudo-manifold*. In each case, the space also comes with a *triangulation*, namely the collection of tetrahedra and face pairings with the property that the only non-manifold points are at the vertices.

Why think about triangulations? As Bill Thurston [66] puts it, “Manifolds are around us in many guises.” We will not give the technical definition of a manifold here. A key result of Moise [49] states that every three-dimensional manifold can be triangulated in the way described above. If one is given a compact three-dimensional manifold with non-empty boundary and one cones each boundary component to a point (one cone point for each boundary component), then it is a closed three-dimensional pseudo-manifold. A relative version of Moise’s theorem allows the compact three-dimensional manifold with boundary to be triangulated, and coning each triangulation of a boundary component to a point gives a triangulation of the pseudo-manifold. In fact, there is a more general definition of three-dimensional pseudo-manifolds than that considered in this paper, where not only the vertices but also points at edges have neighbourhoods that don’t look like balls. Such spaces arise naturally in algebraic geometry. Banagl and Friedman [6] showed that these more general three-dimensional pseudo-manifolds can also be triangulated.

This is in contrast with higher dimensions. There is a four-dimensional manifold, called Freedman’s E_8 manifold, which was shown in the mid 1980s through work of Casson [3] to have no triangulation. The existence of such manifolds in all dimensions greater or equal to five was recently shown by Manolescu [46]. An overview of the so-called triangulation conjecture can be found in [44, 45].

The interest in triangulations lies in their combinatorial framework that allows the study of geometric and topological properties of a manifold both in theory and in practice. In three-dimensional geometry and topology this is completely general due to Moise’s result, and hence gives a framework for proving important theoretical results [40, 24, 26, 34, 35, 36], design algorithms for decisions problems [23, 63, 25, 32, 65, 61, 41, 18], and allows the analysis of the computational complexity of such decision problems [28, 2, 15, 12, 62]. Triangulations also carry the geometry of a three-manifold [52, 55, 53, 22] and allow rigorous computation of geometric invariants [54, 14, 31], such as the volume or the length of a shortest curve that cannot be contracted to a point. In addition, there are also numerous topological or geometric invariants that can be computed from triangulations, see, for example, [7, 39, 13, 38]. Many of these algorithms are implemented in publicly available software packages [16, 11].

Turaev and Viro [69, 68] showed that this can also be turned around. They *define* powerful invariants of closed 3–manifolds *using* triangulations, and there is currently much activity in trying to link these invariants to known topological or geometric invariants, see [51, 17] and the references therein.

To show that a quantity computed using a triangulation is indeed independent of the triangulation requires means to compare the quantities associated to different triangulations. A similar situation arises in *knot theory*. In this case, there are quantities, properties or mathematical objects that can be computed from a diagram, which is a particular way of drawing the knot in the plane. These are computed using a prescribed set of rules. Reidemeister showed that any two diagrams of the same knot are related by a particular set of moves (see Figure 3). So if the result computed before applying any one of Reidemeister’s moves is the same as the result computed after performing that move, then one has shown that the result is independent of the particular diagram and hence an invariant of the knot. Examples of this are the property of tricolourability, the Jones polynomial or the fundamental group. An excellent, elementary introduction to these ideas is given by Adams [1].

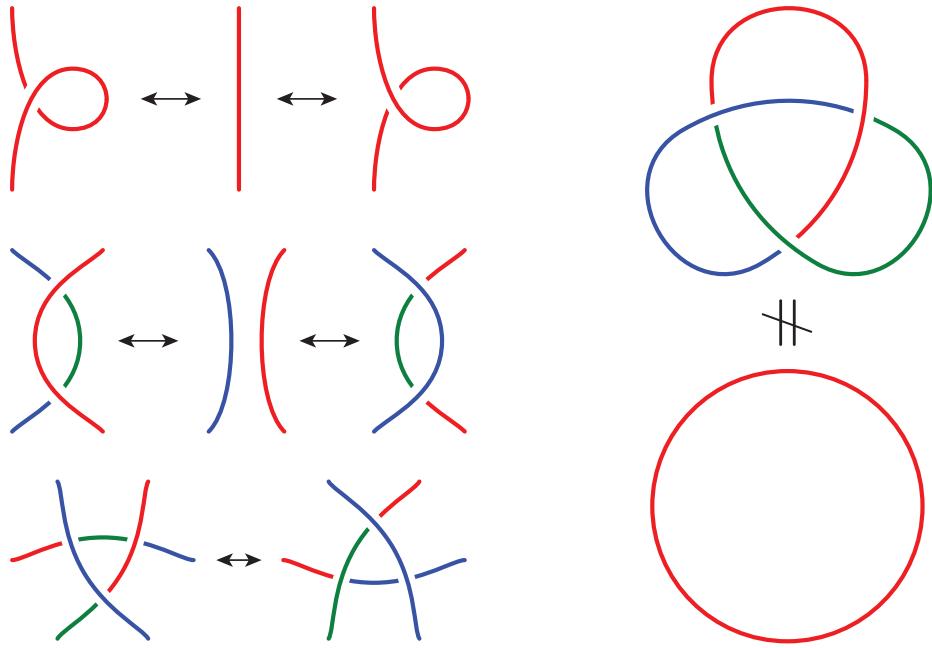


FIGURE 3. Reidemeister moves. The colouring indicates that they preserve tricolourability. In particular, the trefoil knot cannot be deformed into the unknot.

The invariants of Turaev and Viro work similarly. Suitable moves that change one triangulation to another are the *bistellar moves* that have been popularised by Pachner [57]. The bistellar moves play the role of Reidemeister's moves in this setting. We drop down by one dimension to introduce the bistellar moves and highlight some subtle points.

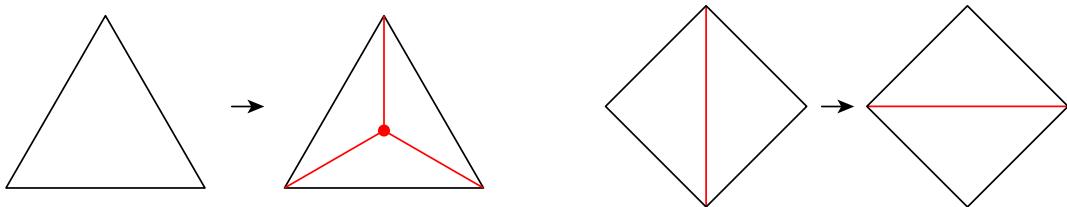


FIGURE 4. The 1-3 and 2-2 moves.

An interlude about surfaces. For surfaces, the corresponding moves are shown in Figure 4. The 1-3 move introduces a vertex at the centre of a triangle and connects this to the three vertices of the triangle, thus dividing it into three triangles. Hence the name 1-3 move. Its inverse is called a 3-1 move, and can be performed on any vertex at which three different triangles meet. The second move is the 2-2 move. Two triangles sharing an edge can be thought of as a quadrilateral with a diagonal drawn in. The 2-2 move changes this diagonal to the other diagonal of the quadrilateral. Note that the 1-3 and 3-1 moves change the number of vertices, edges and triangles, whilst the 2-2 move leaves all of these quantities unchanged. There is a conceptually nice way of thinking about these moves. The boundary of a tetrahedron has a natural triangulation with four triangles. The above moves can be thought of as swapping any subset of triangles on the boundary of the tetrahedron with their complementary triangles. The three bistellar moves we have just described can be used to change a given triangulation into a different triangulation.

However, in the setting we are interested in, one is given two different triangulations of the same space. What does this mean? One way to interpret this is that our topological space, in this case a surface S , comes with two sets of triangles drawn on it, and we'd like to apply the moves to change one set to another. The immediate idea is to superimpose both triangulations. This introduces new vertices at the places where edges of the two triangulations meet. The resulting cells may not all be triangles, but can be easily be subdivided into triangles. So the result is a third triangulation that is a *common refinement* of both triangulations. It is therefore enough to show that one can transform a triangulation to an arbitrary refinement using the given set of moves. There is a hidden assumption in this argument, though, namely that the edges of both triangulations only meet in finitely many points. Think of a small disc on the surface as parameterised as the euclidean plane. It might be that an edge in one triangulation looks like $y = 0$ whilst an edge in the other triangulation looks like $y = x \sin(\frac{1}{x})$ for $x > 0$ and includes the origin as a vertex. These edges meet in infinitely many points! We appear to have gone down the rabbit hole further: given a triangulation of a surface, one can define a smooth structure on the surface that makes the triangulation smooth. A nice account of this is given by Thurston [66, §3.10]. So now each triangulation gives us a smooth structure on the surface. A theorem of Whitehead says that one can indeed perturb the triangulations in order to arrange the finite intersection property. In applying this theorem, there is yet again a hidden assumption: that one of the triangulations may be curvy but at least sufficiently nice with respect to the smooth structure induced on the surface by the other triangulation. To concoct something unpleasant, imagine that an edge with respect to one triangulation looks like a fractal curve with respect to the smooth structure induced by the other. That both triangulations are piecewise smooth with respect to the same structure can also be arranged by a perturbation. This was shown by Epstein [19] in 1966. Hatcher [29] recently gave a very nice proof of this fact using the so-called Kirby torus trick. So we may (and will) assume that the triangulations have the finite intersection property after performing a small perturbation.

How do we transform a given triangulation \mathcal{T}_0 of the closed surface S to an arbitrary subdivision \mathcal{T}_1 of \mathcal{T}_0 only using 1-3, 2-2 and 3-1 moves? First apply a small isotopy of the triangulation \mathcal{T}_0 that keeps all vertices fixed and moves each edge to a position where it does not meet any of the additional vertices of \mathcal{T}_1 . Then repeatedly apply 1-3 moves to the isotopic copy of \mathcal{T}_0 to introduce vertices precisely where the vertices of \mathcal{T}_1 are. Denote the resulting triangulation \mathcal{T}'_0 . We now have two triangulations of the surface that share the same set of vertices. It is now a pleasant exercise to show that there is a sequence of 2-2 moves that transforms \mathcal{T}'_0 into a triangulation that is isotopic to \mathcal{T}_1 , again keeping the vertices fixed. So we have argued that any triangulation of a surface can be transformed to a subdivision by first applying an isotopy fixing the vertices of the initial triangulation, then applying a sequence of 1-3 moves followed by a sequence of 2-2 moves, and completing with an isotopy fixing the vertices of the target triangulation. The isotopies are merely for cosmetic reasons: they put edges neatly on top of each other, rather than roughly in the right place. In conclusion, given two arbitrary triangulations, we can transform one to the other by applying an isotopy, a sequence of 1-3 moves, a sequence of 2-2 moves, an isotopy, a sequence of 2-2 moves, a sequence of 3-1 moves, and finishing with an isotopy. Topological proofs of the existence of the sequence of 2-2 moves can be found in work of Mosher [50] and Hatcher [30]. A geometric proof using the Epstein-Penner decomposition [20] transforms each triangulation with the same number of vertices via 2-2 moves to a *canonical triangulation* by first choosing a geometric structure on the surface. This can be found in [67].

Many authors also allow combinatorial isomorphisms when moving between different triangulations. These are maps from S to S that take vertices to vertices, edges to edges, and triangles to

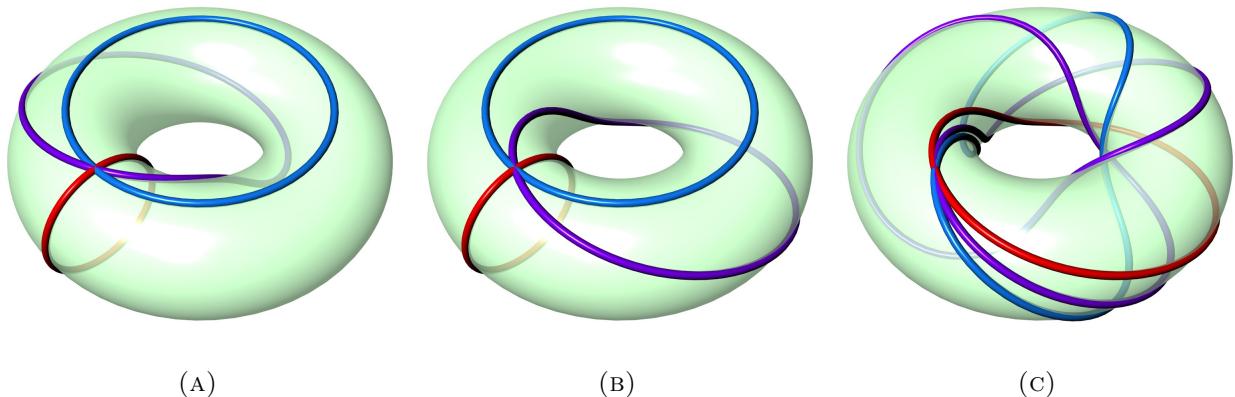


FIGURE 5. Two one-vertex triangulations of the torus related by a single 2-2 move, and another one-vertex triangulation obtained by performing two more flips.

triangles. This has the following effect. Combinatorially, there is a unique one-vertex triangulation of the torus: it has two triangles, three edges and one vertex. If one allows combinatorial isomorphism as a move to go from one triangulation to another, then for any two one-vertex triangulations of the torus, no bistellar moves would be required at all, but only a single combinatorial isomorphism to go from one to another. If one is not allowed to perform any combinatorial isomorphisms but only bistellar moves and isotopies, then there are infinitely many inequivalent one-vertex triangulations of a torus. However, one can apply a sequence of 2-2 moves and isotopies to connect any one-vertex triangulation to any other. See Figure 5 for an example of two such triangulations related by a single 2-2 move. The set of all isotopy classes of one-vertex triangulations of the two-torus connected by 2-2 moves thus has the structure of an infinite trivalent tree. In particular, for any two such triangulations, there is a unique minimal sequence of moves connecting them. All of this has beautiful stories in itself, of which we mention three. One story links triangulations to geometry and the Farey tessellation of the hyperbolic plane [50, 30]. In particular, we can not only tell any two of them apart, but can easily work out the minimal number of moves needed to transform one to another. Another story gives a geometric interpretation of Markov triples and diophantine equations [58, 59], see [64] for a nice account of this. A third story uses these triangulations of the torus to construct special triangulations of the so-called lens spaces, a special class of 3-manifolds [37, 33].

Here is another reason why one might not want to allow combinatorial isomorphisms. To return to the analogy with knot theory, there are knots that cannot be transformed to their mirror image by a sequence of Reidemeister moves. An example is the trefoil knot. If we allowed the operation that takes a knot to its mirror image as an additional transformation, then the distinction between a right-handed and a left-handed trefoil knot is lost.

Back to 3-dimensional spaces. The situation for 3-dimensional manifolds regarding common refinements of triangulations is similar to surfaces, albeit slightly more difficult to visualise. Details and references are given in §2.6. The moves on 3-dimensional triangulations are also completely analogous. The 1-4 move introduces a vertex inside a tetrahedron and connects it to the four vertices of the tetrahedron with four edges. This creates six new triangular faces, spanned by the edges of the tetrahedron and the new vertex, and four new tetrahedra. The inverse move is called a

4-1 move. Then there is the 2-3 move, which is performed on two different tetrahedra meeting in a triangular face. The 2-3 move deletes this face by introducing a new edge connecting the opposite corners of the tetrahedra. Its inverse is the 3-2 move. Each of these moves changes the number of tetrahedra, faces and edges. However, the 2-3 and 3-2 moves do not affect the number of vertices in the triangulation. The 1-4 and 4-1 moves only change the number of material vertices, but not the number of ideal vertices. One may again imagine all of these moves as swapping subsets of tetrahedra on the boundary of a 4-dimensional simplex!

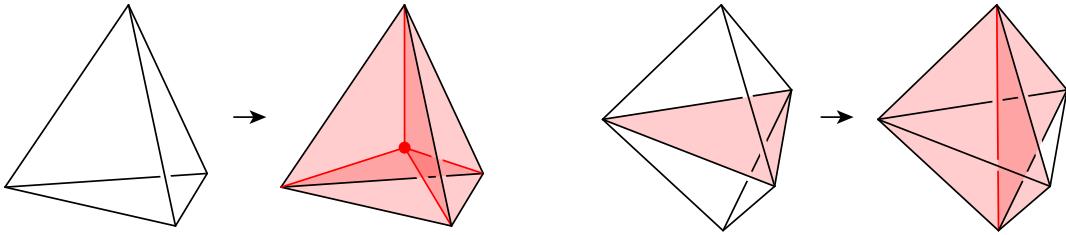


FIGURE 6. The 1-4 and 2-3 moves.

The use of bistellar moves in the work of Pachner has its roots in the stellar moves used by Alexander [4] and Newman [56]. The bistellar moves are equivalent to moves on a dual structure, called a spine, due to Matveev [47]. Classically, the stellar and bistellar moves were only defined, and results involving them only proved, for special types of triangulations that are less general than the ones we defined above. Vertices, edges, triangular faces and tetrahedra are called *simplices*, and one indicates the dimension d of a simplex by saying that it is a d -simplex. A simplex contained in another simplex σ is said to be a *face* of σ . A *simplicial* triangulation requires any two simplices to meet in either a face or not at all. A *combinatorial* triangulation has the additional requirement that the manifold structure is completely evident from the combinatorics — this is technically made precise by requiring that the so-called link of every simplex is a sphere. For surfaces and 3-dimensional manifolds, every simplicial triangulation is combinatorial, and the distinction is only of relevance in higher dimensions (which are of no concern for this paper). Our triangulations are sometimes called *singular* or *semi-simplicial* in the literature. They can be turned into simplicial triangulations by performing at most two *barycentric subdivisions*. We define these in §2.5 and show that barycentric subdivision can be achieved using the bistellar moves 1-4, 2-3 and 3-2. We also show that stellar moves can be achieved using these moves—the caveat here is that one stellar move may turn into an arbitrarily long sequence of bistellar moves. In this sense, there are infinitely many stellar moves but only four bistellar moves. This discussion connects our triangulations to the general theory of simplicial or combinatorial triangulations via bistellar moves, and we therefore use the word *triangulation* throughout this paper without further qualification. We also always allow the implicit use of isotopy — so our statements are really statements about triangulations up to isotopy, rather than fixed triangulations. We summarise this discussion in the following theorem.

Theorem 1.1 (Alexander, Newman, Moise, Pachner). *The set of all triangulations of a closed three-dimensional manifold M is connected under 1-4, 2-3, 3-2 and 4-1 moves.*

An excellent account of the history and proof of a similar result that holds in all dimensions was recently given by Lickorish [42]. A more general definition of triangulations, and a different set of moves, is used by Ludwig and Reitzner [43].

Banagl and Friedman [6, Proposition 2.16] give a version of this theorem for closed three-dimensional pseudo-manifolds starting with a more general definition of a pseudo-manifold, and use this to extend the Turaev-Viro invariants to three-dimensional pseudo-manifolds:

Theorem 1.2 (Banagl-Friedman). *The set of all triangulations of a closed three-dimensional pseudo-manifold is connected under 1-4, 2-3, 3-2 and 4-1 moves.*

We remark that in the above, the number of ideal vertices remains unchanged, but the number of material vertices may vary. Proofs of Theorems 1.1 and 1.2 are sketched in §2.6.

The interest of results such as the above is indeed the construction of invariants. This is made even easier by limiting the number of moves required, which is achieved by the following refinement of Theorem 1.1 of Matveev [48] and Piergallini [60].

Theorem 1.3 (Matveev, Piergallini). *The set of all triangulations of a closed three-dimensional manifold M with exactly one material vertex is connected under 2-3 and 3-2 moves, excepting triangulations with a single tetrahedron.*

This theorem was originally stated in the dual language of spines, which we discuss in §2.7. There are precisely three closed three-manifolds that admit a triangulation with a single tetrahedron: the three-sphere and two other lens spaces, namely $L(4, 1)$ and $L(5, 2)$. See Figure 7. All of these manifolds are orientable. Moreover, up to combinatorial equivalence, there are precisely four triangulations with a single tetrahedron. The three-sphere has two such triangulations, but one of them has two vertices. So the exclusion only concerns three triangulations of three three-dimensional manifolds. No 2-3 or 3-2 moves can be applied to these triangulations.

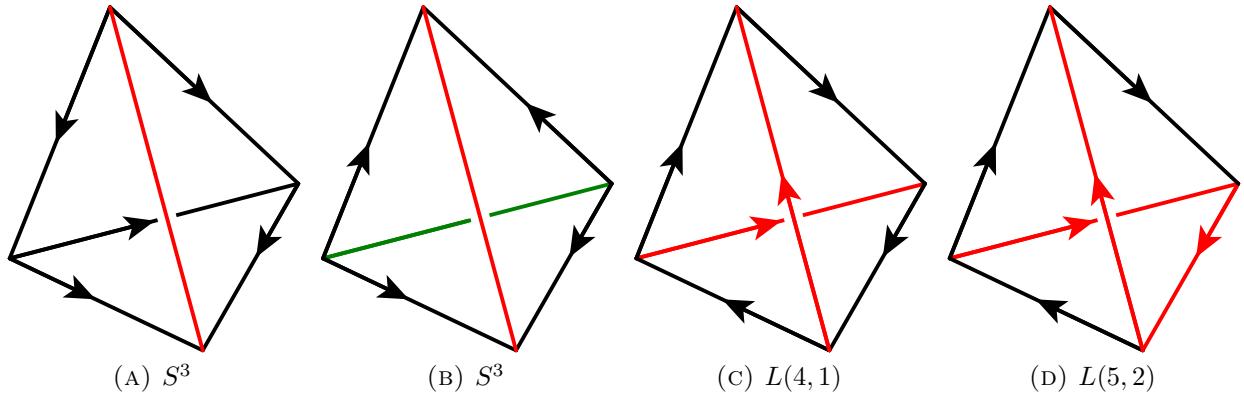


FIGURE 7. Single tetrahedron triangulations of closed manifolds.

Benedetti and Petronio [8, 9] outlined a proof of the fact that the above result also holds for any prescribed number of material vertices in a triangulation of a 3-dimensional manifold, and Amendola [5] shows that this extends to 3-dimensional pseudo-manifolds. In each case, key results of Matveev [48] are used and the techniques involve the dual viewpoint of spines. The main purpose of this note is threefold. We wish to popularise Amendola's result; we give a proof that emphasises the dual viewpoints of triangulations and spines; and we give a proof replacing a key general position argument of [48] with a more combinatorial argument inspired by the theory of subdivisions.

Theorem 1.4 (Amendola). *Let M be a three-dimensional manifold or pseudo-manifold. The set of triangulations of M with a fixed number (possibly zero) of material vertices is connected under 2-3 and 3-2 moves, excepting triangulations with a single tetrahedron.*

As for the additional exceptions to the theorem: The only closed orientable three-dimensional pseudo-manifolds having a triangulation with a single tetrahedron are the manifolds given above. The only non-orientable pseudo-manifold having a triangulation with a single tetrahedron is the so-called Gieseking manifold shown in Figure 8. Here, the vertex link is a Klein bottle.

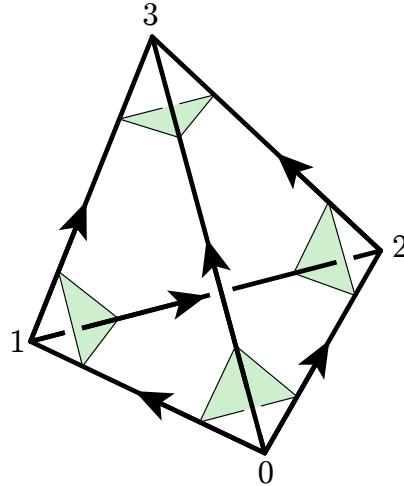


FIGURE 8. The Gieseking manifold. Vertices are labelled by the number of arrows pointing towards the vertex. The faces are identified by $012 \leftrightarrow 023$ and $013 \leftrightarrow 123$.

Theorems 1.1 and 1.2 are important connectivity results. They are useful for defining invariants in terms of triangulations, and for building censuses of triangulations. Theorem 1.3 is a great improvement over Theorem 1.1 in the view of the theory of *0-efficient triangulations* of Jaco and Rubinstein [36], which require only a single vertex for the most important classes of closed three-dimensional manifolds. In the same vein, Theorem 1.4 allows us to pass between any two triangulations with no material vertices without introducing material vertices along the way. We also remark that Amendola [5] proves a more general theorem that keeps track of finer information on the manifold than Theorem 1.4.

In the literature, there are a number of papers that attribute Theorem 1.4 to Matveev or Piergallini. Personal communication with both of these authors revealed that both think they only proved Theorem 1.3, but that their techniques may extend to that effect. This paper affirms this belief.

We give a complete proof of Theorem 1.4. We will assume Theorems 1.1 and 1.2 (proofs are sketched in §2.6), but we do not assume Theorem 1.3. Our main task is thus to convert any appearance of 1-4 or 4-1 moves into sequences of 2-3 and 3-2 moves. Our contribution in reproving Theorem 1.3 in the proof of Theorem 1.4 is to view the passage between triangulations from two perspectives: the triangulations on the one hand and the dual notion of *special spines* on the other.

The structure of this paper is as follows. More examples, as well as the definitions of triangulations and spines are given in Section 2. Here, we also describe classical constructions, such as stellar

moves and barycentric subdivision, and outline proofs of classical results going back to Alexander and Newman.

We set the scene in Section 3, proving a lemma that simplifies the set-up for the proof of the main result. Whilst the main ideas for the main proof will be present, the interplay between triangulations and spines is not yet required. The strategy for the proof of the main result is presented in Section 4. The remainder of the paper aims at filling in the missing details. This starts in Section 5 by discussing Matveev’s arch, and describing arch constructions as moves on triangulations. The first serious steps in providing details are taken in Section 6, which can be viewed as a pre-processing step for the two triangulations, reducing the problem of connecting them to the problem of *sweeping a membrane across a ball*. This last task is performed in detail in Section 7, thus completing this paper.

In writing this paper, we have considered other strategies to prove this particular connectivity result for triangulations. Inevitably, the problem was always reduced to a situation that could not easily be swept under the rug. In the end, we settled for sweeping across a ball, as here the generic situation is easy to deal with and hence the overall strategy easy to follow.

2. TRIANGULATIONS AND SPINES

In the hope of providing a reference that is useful for novices in the field, we provide this section to make this paper as self-contained as possible. We give further examples (§2.1); define manifolds and pseudo-manifolds (§2.2) and the triangulations we consider (§2.3). We then discuss classical moves on triangulations, namely stellar moves (§2.4) and barycentric subdivision (§2.5) and explain how these can be performed using bistellar flips. We use these in §2.6 with sketches of proofs of Theorems 1.1 and 1.2. The dual notion of spines (§2.7) completes the summary of background material required for this paper.

2.1. Examples. Before we formalise the definition of a triangulation given in the introduction, we give some examples. In the introduction, we have already seen triangulations just involving one tetrahedron. Figure 7 shows that we can build the 3–dimensional sphere with just one tetrahedron. A simplicial triangulation of the 3–dimensional sphere requires at least five tetrahedra. The advantage of the more general triangulations thus lies in their efficiency.

Ideal triangulations similarly turn out to be an efficient way of encoding manifolds. An example is given by Thurston’s ideal triangulation of the complement of the figure eight knot shown in Figure 9. This is a space obtained by deleting the figure eight knot from the three–dimensional sphere. One can build a pseudo-manifold with just two tetrahedra with the property that the complement of the vertex in this pseudo-manifold is homeomorphic to the complement of the figure eight knot. If one considers the space obtained by deleting a small regular neighbourhood of the knot from the 3–sphere, one obtains a compact 3–manifold with boundary a torus. To divide this into (material) tetrahedra requires ten — where now we allow some faces of tetrahedra not to be identified so as to form the boundary.

2.2. Manifolds and pseudo-manifolds. A euclidean 3–simplex σ is the convex hull of four points p_0, p_1, p_2, p_3 in euclidean 3–space that are in general position. What matters in our discussion is not the euclidean, but just the affine structure. A unified discussion of simplices and subsimplices is aided by the following definition. An n –simplex σ is the convex hull of $n + 1$ affinely independent points p_0, \dots, p_n in some affine space of dimension at least n . The convex hull of a subset of

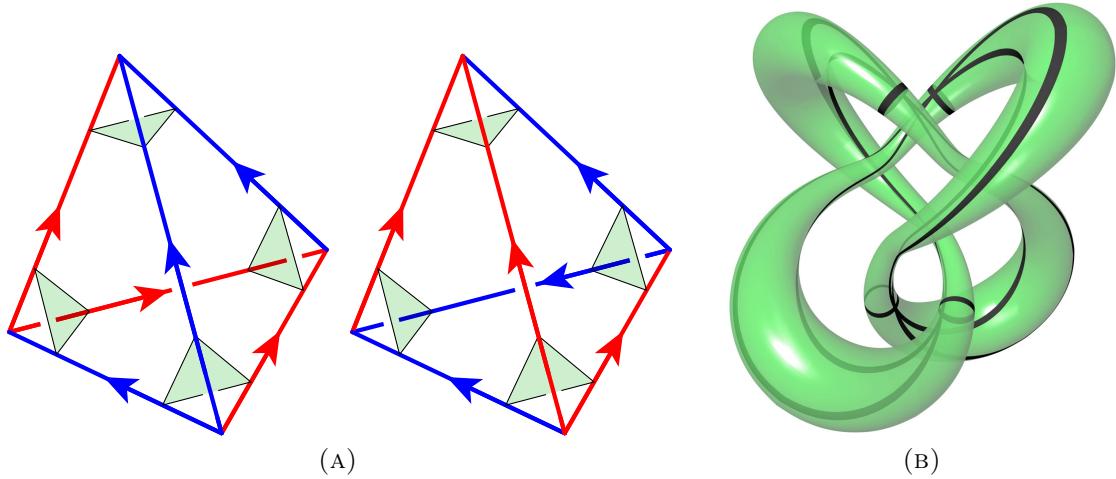


FIGURE 9. The ideal triangulation of the complement of the figure eight knot in the three-sphere. The shown triangles form a tube enclosing the knot.

$\{p_0, \dots, p_n\}$ is a *face* or *subsimplex* of σ . We refer to 3–simplices as tetrahedra, 2–simplices as triangles, 1–simplices as edges and 0–simplices as vertices.

Let $\tilde{\Delta}$ be a finite union of pairwise disjoint euclidean 3–simplices. Every k –simplex τ in $\tilde{\Delta}$ is contained in a unique 3–simplex σ_τ . A 2–simplex in $\tilde{\Delta}$ is called a *facet*. Let Φ be a family of affine isomorphisms pairing the facets in $\tilde{\Delta}$, with the properties that $\varphi \in \Phi$ if and only if $\varphi^{-1} \in \Phi$, and every facet is the domain of a unique element of Φ . The elements of Φ are termed *face pairings*.

The quotient space $P = \tilde{\Delta}/\Phi$ with the quotient topology is then a *closed 3–dimensional pseudo-manifold*, and the quotient map is denoted $p: \tilde{\Delta} \rightarrow P$. We will always assume that P is connected. In the case where P is not connected, the results of this paper apply to its connected components.

If one considers any point $p \in P$ that is not a vertex or a mid-point of an edge, then a small neighbourhood of p looks like the neighbourhood of a point in 3–dimensional euclidean space. This is easy to see for points in the interior of a tetrahedron, not too difficult for points in faces, and a little more involved for points that lie in edges. So what makes P a *pseudo*-manifold is the fact that we have no control over small neighbourhoods of its vertices, and over small neighbourhoods of mid-points of edges. Place a small triangle in the corner of each tetrahedron so that these triangles glue up to form a closed surface. This surface is called the *link* of the vertex, and the neighbourhood of the vertex is a cone over this surface, with cone point the vertex. This cone is a ball precisely when the surface is a sphere. We thus call a vertex *material* when its link is a sphere and *ideal* otherwise. To see what happens at midpoints of edges, take the intersection of the tetrahedron with a small sphere based at the midpoint. This gives a lune in the tetrahedron. Now trace the chain of tetrahedra around this edge. With each tetrahedron, this adds another lune until we get to the last tetrahedron. Here again, a lune is added but this is also identified to the initial lune. There are two different ways of making the identification: one gives a sphere, and the other gives a projective plane. In the second case, the edge is identified with itself but in opposite orientation. To rule out this case and to align the structure of the triangulation with the structure of our pseudo-manifolds, we require that

(1) the restriction of the map $p: \tilde{\Delta} \rightarrow P$ to the interior of each k -simplex in $\tilde{\Delta}$ is a homeomorphism for all $k \in \{0, 1, 2, 3\}$.

This is only a restriction on the triangulations we allow, but not on the pseudo-manifolds. The restriction has consequences only for $k = 1$, indeed it only avoids the situation where an edge may fold back onto itself, but it allows extra identifications along the boundary of a simplex. With this requirement, the pseudo-manifold P is a *manifold* if and only if every vertex is material.

We still refer to the image under p of a triangle in $\tilde{\Delta}$ as a triangle in P , but the reader should bear in mind that this may not be an embedded triangle and have identifications along its boundary. The same applies to tetrahedra or edges in P .

The quotient space P is studied via the map $p: \tilde{\Delta} \rightarrow P$. The *degree* of the triangle, edge or vertex τ in P is the number of preimages it has in $\tilde{\Delta}$; equivalently, the degree of the restriction $p: p^{-1}(\tau) \rightarrow \tau$.

There is a more restrictive class, where such confusion cannot arise. We say that $(\tilde{\Delta}, \Phi)$ is *simplicial* if

(1) the restriction of the map $p: \tilde{\Delta} \rightarrow P$ to each k -simplex $\tilde{\Delta}$ is a homeomorphism for all $k \in \{0, 1, 2, 3\}$, and

(2) for any two simplices τ_0 and τ_1 in $\tilde{\Delta}$, $p(\tau_0) \cap p(\tau_1)$ is either empty or a single simplex.

2.3. Triangulations. In the previous section, we have defined a 3–dimensional manifold or pseudo-manifold using tetrahedra. Now suppose you are given such a manifold or pseudo-manifold M . We say that the triple $\mathcal{T} = (\tilde{\Delta}, \Phi, h)$ is a *triangulation* of M if

$h: P \rightarrow M$ is a piecewise linear homeomorphism.

The images of vertices, edges, triangles and tetrahedra under the composition of the map $p: \tilde{\Delta} \rightarrow P$ with h give us the vertices, edges, triangles and tetrahedra of the triangulation in M . Again, these generally have self-identifications along their boundaries. The triangulation $\mathcal{T} = (\tilde{\Delta}, \Phi, h)$ is a *simplicial triangulation* of M if $(\tilde{\Delta}, \Phi)$ is simplicial.

We say that two triangulations $\mathcal{T}_0 = (\tilde{\Delta}_0, \Phi_0, h_0)$ and $\mathcal{T}_1 = (\tilde{\Delta}_1, \Phi_1, h_1)$ are *equivalent* or *the same* if there is a simplicial bijection $s: \tilde{\Delta}_0 \rightarrow \tilde{\Delta}_1$ such that $h_0 \circ p_0 \circ s$ is isotopic to $h_1 \circ p_1$. This in particular allows the same union of tetrahedra and set of face pairings to give inequivalent triangulations that are combinatorially equivalent.

We will be concerned with changing a given triangulation. On a formal level, this amounts to changing a triangulation $\mathcal{T}_0 = (\tilde{\Delta}_0, \Phi_0, h_0)$ of M to another triangulation $\mathcal{T}_1 = (\tilde{\Delta}_1, \Phi_1, h_1)$. Such a change can be very local, such as replacing one simplex and the associated face pairings in $(\tilde{\Delta}_0, \Phi_0)$ with a small number of simplices and corresponding face-pairings. Or it could be a change that replaces all simplices and face-pairings. As an example, we mention that it is always possible to transform a triangulation into a simplicial triangulation by applying at most two *barycentric subdivisions*. This replaces each 3–simplex by $24 \times 24 = 576$ 3–simplices. We give the definition in §2.5 and show how to achieve this using 1-4, 2-3 and 3-2 moves.

In the remainder of this paper, we try to keep notation to a minimum in order to not distract from the main ideas. A precise notation is, of course, relevant for an actual implementation in software. There is a natural way to track vertices, edges, and so forth through a sequence of moves on a triangulation. Our moves are typically performed topologically inside the pseudo-manifold, and the simplices can then be disassembled to form a triangulation together with a preferred quotient map $\tilde{\Delta} \rightarrow M$.

2.4. Stellar moves. There are three types of stellar moves. The first type is simply the 1-4 move. See Figure 6. It places a vertex in the interior of a tetrahedron and divides the tetrahedron into four tetrahedra. Each of these can be viewed as a cone from a boundary face of the original tetrahedron to the central vertex.

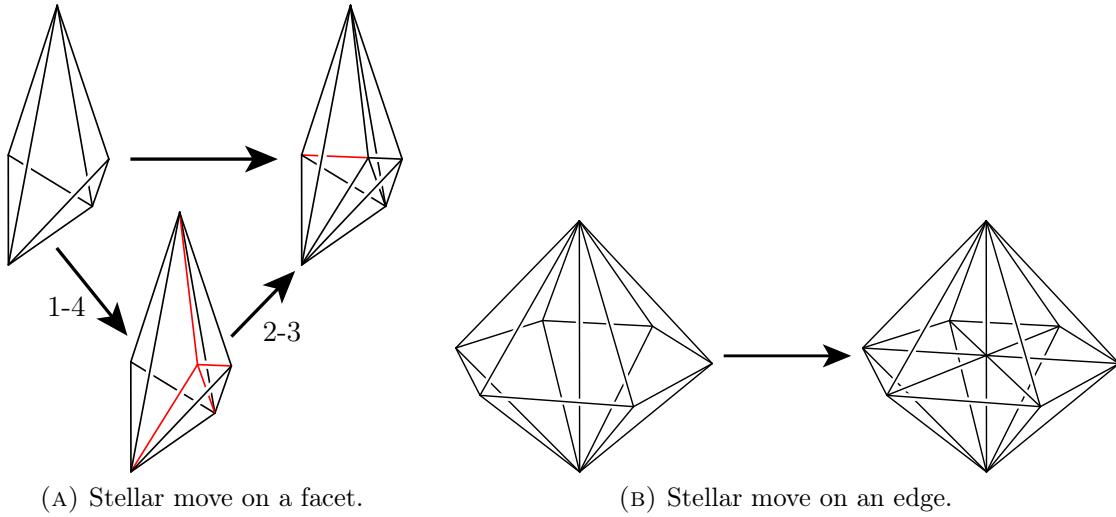


FIGURE 10. Stellar moves.

The second type of stellar move is conceptually similar, this time performed on a facet τ that is contained in two distinct tetrahedra σ_0 and σ_1 . Here, one first performs a 1-3 move on τ . One then subdivides each of the two tetrahedra into three by coning the triangulation of the subdivided face to the remaining vertex. See Figure 10a. This operation can be described as an operation on $\tilde{\Delta}$, so it is irrelevant that some of the edges or vertices of the tetrahedra may be identified. To achieve this move by our four bistellar moves, one can first apply a 1-4 move on σ_0 . The face τ is not affected by this, and one can now perform a 2-3 move that replaces the two tetrahedra meeting in τ with three tetrahedra meeting in an edge. The result is the same as the stellar move on the face τ .

The third move is performed on an edge e that is only contained at most once in any tetrahedron and of degree at least two. We place a new vertex on e . In each tetrahedron, one now creates a new triangle by coning the edge opposite e to the new vertex, thus cutting the tetrahedron into two. See Figure 10b. Again, this operation can be done consistently in $\tilde{\Delta}$, and so is well-defined. We need to show how to realise this move by bistellar moves. Choose a cyclic order of the tetrahedra abutting e , and starting with one tetrahedron label them $\sigma_0, \dots, \sigma_k$. Suppose σ_i and σ_{i+1} meet in e along τ_i , and σ_k and σ_0 meet in e along τ_k . First apply a 1-4 move on σ_0 . This increases the degree of e by one. Then apply a 2-3 move along the face τ_0 (interpreted as a face in the new triangulation). This decreases the degree of e by one. Now we iteratively apply 2-3 moves along τ_1, τ_2, \dots , until only τ_{k-1} and τ_k are left. At this stage, e (interpreted as an edge in the resulting triangulation) has degree three. Hence apply a 3-2 move (noting that all three tetrahedra meeting in e are distinct). See Figure 15.

2.5. Barycentric subdivision. A *barycentric subdivision* can be defined inductively. A barycentric subdivision of a 1-simplex introduces a vertex at its midpoint and divides it into two 1-simplices. A barycentric subdivision of a 2-simplex introduces a vertex at its centre and cones this to the

barycentric subdivision of its boundary. This divides it into six 2-simplices. For a 3-simplex, we again introduce a vertex at its centre and then cone to the barycentric subdivision of its boundary, thus dividing it into 24 3-simplices. See Figure 11.

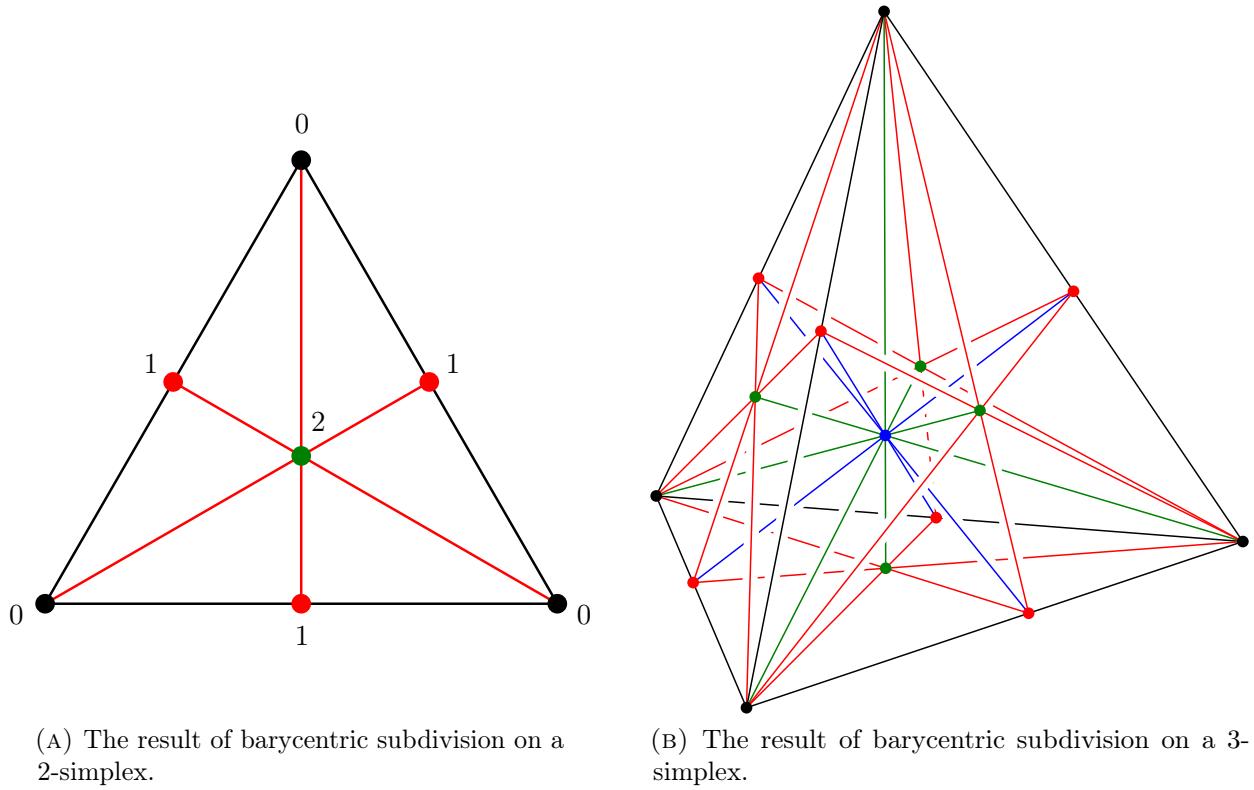


FIGURE 11. Barycentric subdivision.

It follows from the definition of a triangulated pseudo-manifold that one can apply this procedure to each 3-simplex in the triangulation, and hence break the triangulation up into a finer triangulation having 24-times as many tetrahedra. It turns out that this can be achieved using stellar moves. Hence, by the previous section, barycentric subdivision can be realised using 1-4, 2-3, 3-2 moves on the original triangulation.

To achieve barycentric subdivision, first apply a 1-4 move to each 3-simplex in the original triangulation Δ to obtain a new triangulation Δ_0 with four times as many tetrahedra. There are facets and edges that persist from Δ to Δ_0 . On each of these facets, perform a stellar move, giving a new triangulation Δ_1 with three times as many tetrahedra as Δ_0 . The edges that persisted from Δ to Δ_0 are still present in Δ_1 . Now perform stellar moves on all of these edges, giving a triangulation Δ_2 with twice as many tetrahedra as Δ_1 . So in total we have $4 \times 3 \times 2 = 24$ times as many tetrahedra as we started with, and we hope the reader is convinced that the overall effect is the same as applying a barycentric subdivision to each tetrahedron! The moves are illustrated in Figures 12–15.

2.6. Common refinements. Even though we stated in the introduction that we will prove our main result assuming Theorems 1.1 and 1.2, we cannot resist sketching proofs of these results.

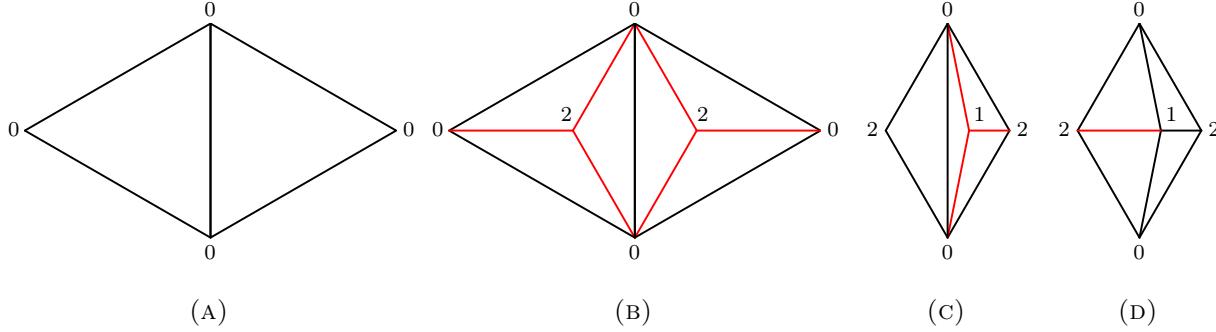


FIGURE 12. Barycentric subdivision for two-dimensional triangulations via bistellar moves.

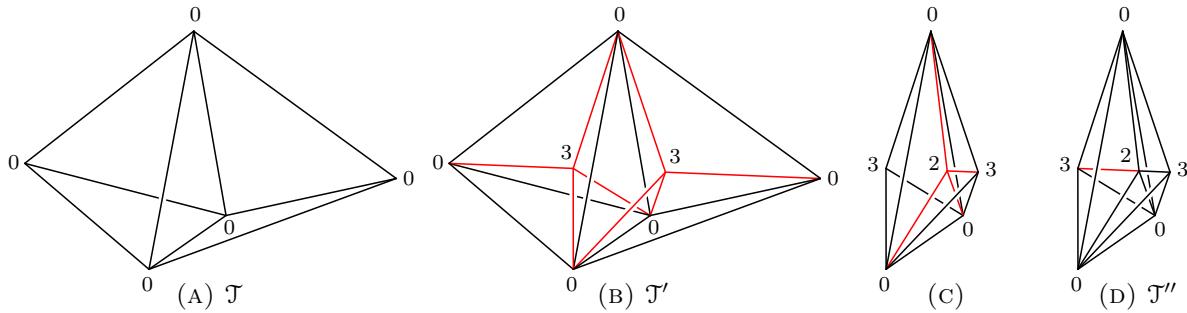


FIGURE 13. Introducing the vertices corresponding to the three- and two-dimensional simplices for the barycentric subdivision of a three-dimensional triangulation.

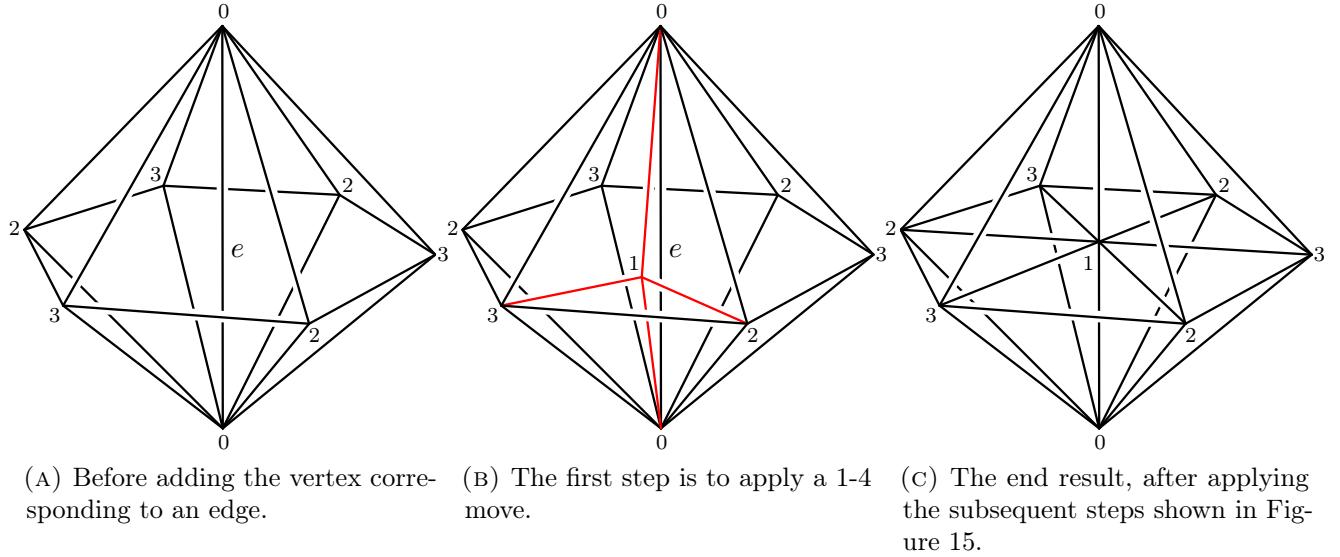


FIGURE 14. Introducing the vertex corresponding to a one-dimensional simplex for the barycentric subdivision of a three-dimensional triangulation.

We are given triangulations \mathcal{T}_0 and \mathcal{T}_1 of a 3-manifold M . To apply the classical results, we first apply two barycentric subdivisions to each in order to obtain the simplicial triangulations $\mathcal{T}'_0 = (\tilde{\Delta}_0, \Phi_0, p_0)$ and $\mathcal{T}'_1 = (\tilde{\Delta}_1, \Phi_1, p_1)$ respectively. We have quotient maps $p_0: \tilde{\Delta}_0 \rightarrow M$ and

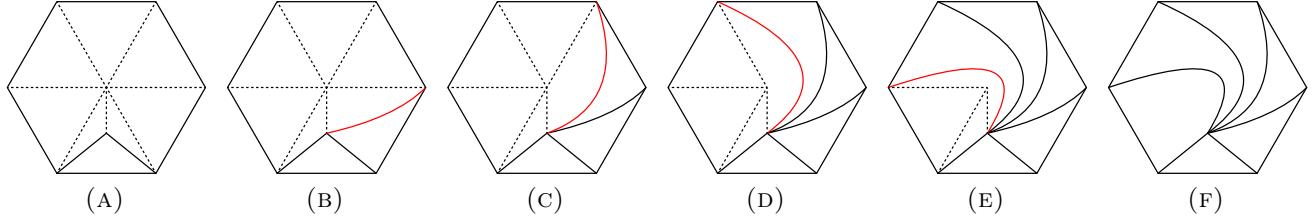


FIGURE 15. The subsequent steps in adding the vertices corresponding to the one-dimensional simplices for the barycentric subdivision of a three-dimensional triangulation. The horizontal cross-section through the tetrahedra is shown. Edges are drawn with a solid line, while intersections of faces with the cross-section are drawn with a dotted line.

$p_1: \tilde{\Delta}_1 \rightarrow M$. We view M as the quotient space Δ_0/Φ_0 , so that p_0 is a linear map on simplices. The work of Moise (see also Brown [10] and Hamilton [27]) now allows us to compose p_1 with an isotopy $\iota: M \rightarrow M$ so that the restriction of $\iota \circ p_1$ is also linear on each simplex. In particular, Alexander's theorem [4] (see also Lickorish [42] and Turaev-Viro [69]) now allows us to transform \mathcal{T}'_0 and \mathcal{T}'_1 to a common refinement by applying stellar moves.

We have shown above that barycentric subdivisions and stellar moves can be performed through sequences of bistellar moves. This completes our sketch of the proof of Theorem 1.1.

The class of pseudo-manifolds considered in this paper is less general than that considered by Banagl and Friedman [6]. Our pseudo-manifolds are merely compact manifolds with boundary, where we have added a cone to each boundary component. It is therefore straightforward to sketch a proof of Theorem 1.2 analogous to the above. We have the same set-up as before, and apply two barycentric subdivisions, giving quotient maps $p_0: \tilde{\Delta}_0 \rightarrow M$ and $p_1: \tilde{\Delta}_1 \rightarrow M$. The work of Brown [10] (see [6, Theorem 2.13]) now allows us to compose p_1 with an isotopy $\iota: M \rightarrow M$ so that the restriction of $\iota \circ p_1$ is also linear on each simplex. Due to the topology of the pseudo-manifold, the isolated ideal vertices are necessarily fixed under the isotopy. Alexander's theorem again provides stellar moves resulting in a common refinement, and we conclude as above that there is the desired sequence of bistellar moves.

2.7. Special spines. We now describe a 2-dimensional object, a *spine*, that encodes M , a 3-dimensional manifold or pseudo-manifold. To begin with, suppose you have a triangulation Δ of M , and pass to the second barycentric subdivision Δ'' of Δ . This has the following nice property. Let v be a vertex of Δ'' that corresponds to a vertex of Δ . Then the collection $\text{Lk}(v)$ of all tetrahedra in Δ'' containing v is a cone on a triangulated surface with cone point v , and for any two distinct vertices of Δ , the two corresponding surfaces are disjoint. Hence if we remove $\text{Lk}(v)$ for each vertex v of Δ from Δ'' , we obtain the triangulation of a compact 3-dimensional manifold M^c with boundary, and M is obtained from M^c by adding the cones on its boundary components. Hence we can reconstruct M from M^c .

One can now collapse M^c onto a 2-dimensional complex S that is the union of all triangles in the first barycentric subdivision of Δ with the property that they do not contain any vertex of Δ . For each tetrahedron of Δ this is the butterfly shown in Figure 16a. Note that M^c is homeomorphic to a regular neighbourhood of S in M^c . We claim that S determines M^c (and hence M) *independent* of its embedding in M^c . The important properties of S that make this possible are the following.

Each point in S has a neighbourhood that is homeomorphic to one of the three models shown in Figure 16b. A point is *singular* if it does not have a neighbourhood homeomorphic to a disc. The

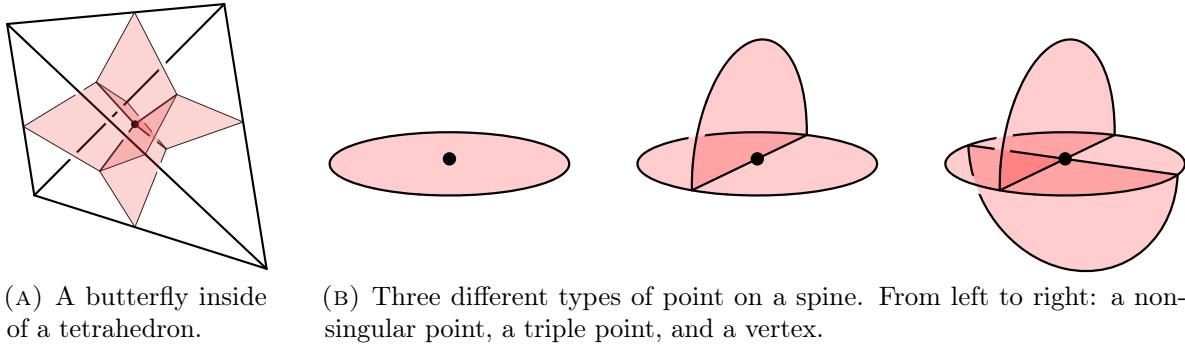


FIGURE 16. Spines.

set of all singular points consists of the triple points and vertices, and is naturally a graph. The nodes of this graph are called vertices of the spine. The complement of the set of singular points in S is a surface, and we claim that it is a disjoint union of open discs. This comes from the local structure: each connected component of the complement is the union of all triangles in Δ' having a vertex on the barycentre of an edge in Δ . If one removes the vertices from the singular set, then each connected component is an interval—this is clear because we have a graph. The relationship with the triangulation is again evident: the interval connects the barycentres of adjacent tetrahedra through the barycentre of a common face. These properties are summarised by saying that S is a *special spine* of M .

Having observed this, one can show that if S_0 and S_1 are special spines of M_0 and M_1 , with the property that S_0 and S_1 are homeomorphic, then M_0 and M_1 are homeomorphic. The reader may find pleasure in proving this result, or can find it in Matveev's book [48, Theorem 1.1.17].

Special spines and triangulations are dual objects: each vertex of a spine is dual to a tetrahedron; each edge of a spine is dual to a triangle; each face of a spine is dual to an edge of the triangulation. The complementary regions of a spine in the pseudo-manifold are dual to the vertices of the triangulation. Indeed, each such region is the cone over a surface, where the surface is homeomorphic to the vertex link. In particular, material vertices correspond to three-balls.

2.8. Moves on triangulations and spines. Using this duality, we should be able to view moves performed on triangulations in the dual spine picture. For example, the bistellar moves we introduced earlier can be seen as moves performed on spines, as shown in Figure 17. Pictures of spines are often more flexible, with many ways to draw the same spine. Figure 18 shows an alternative picture for the 2-3 move.

Similarly, we can see our barycentric subdivision moves (as implemented using bistellar moves) in the dual spine picture. These are shown in Figures 19 and 20.

The spine picture is sometimes clearer than the triangulation picture when illustrating certain sequences of moves. For example, Matveev introduces the *V-move*, which replaces one vertex of a spine with three vertices. Matveev shows that the V-move can be implemented using 2-3 and 3-2 moves, assuming that a spine has more than one vertex (or dually that the triangulation has more than one tetrahedron), see Proposition 1.2.8 of [48]. For the convenience of the reader, we reproduce the argument in Figure 21. This sequence of moves would be confusing and difficult to draw in the triangulation picture. (Try it!)

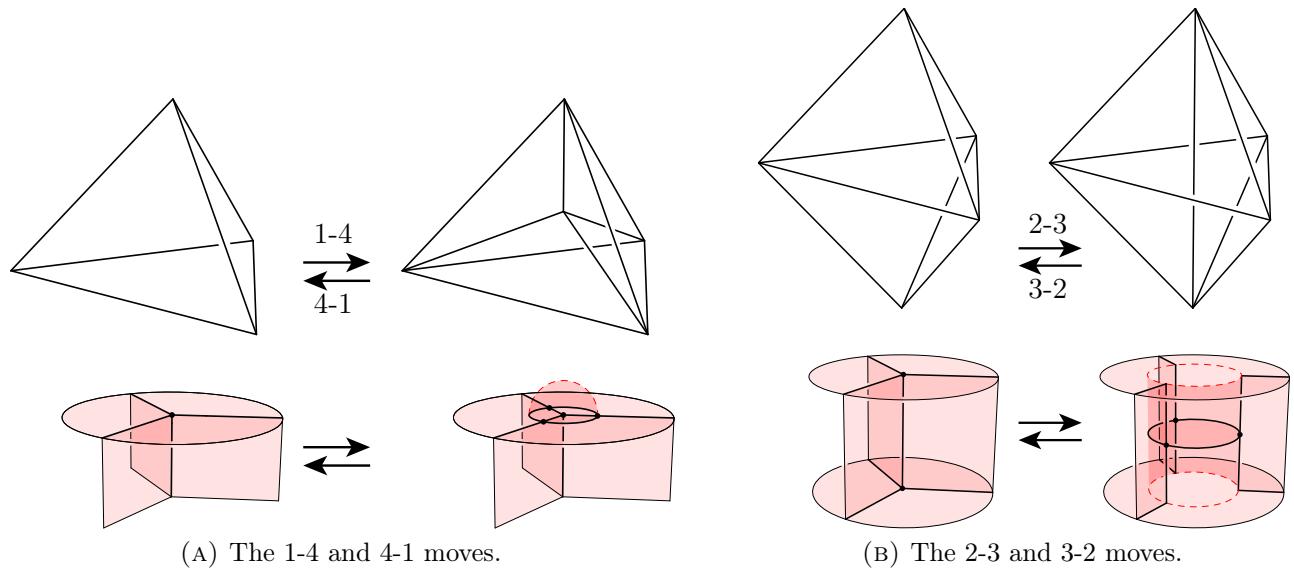


FIGURE 17. Bistellar moves on triangulations and their dual special spines.

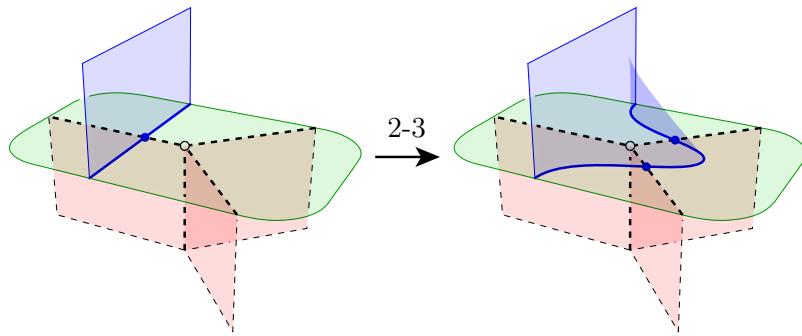


FIGURE 18. An alternative picture of the 2-3 move on a spine.

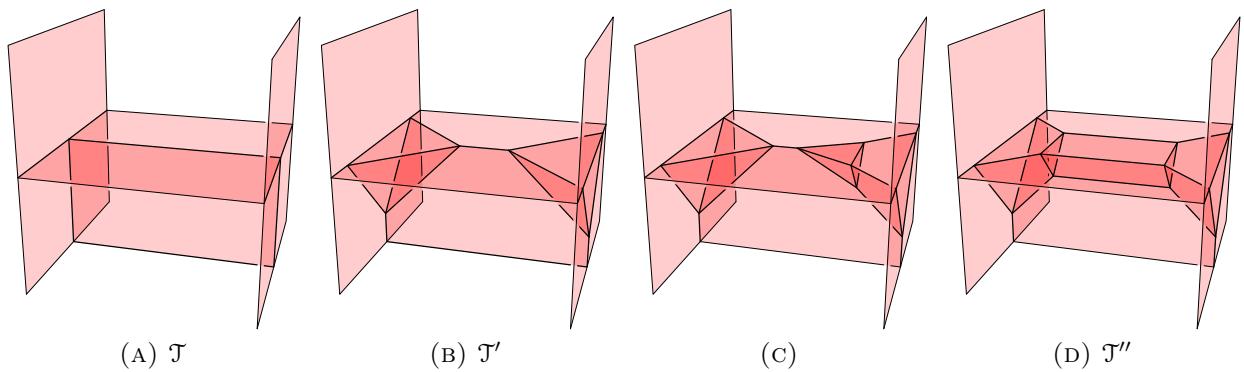


FIGURE 19. Figure 13, drawn in the dual spine picture.

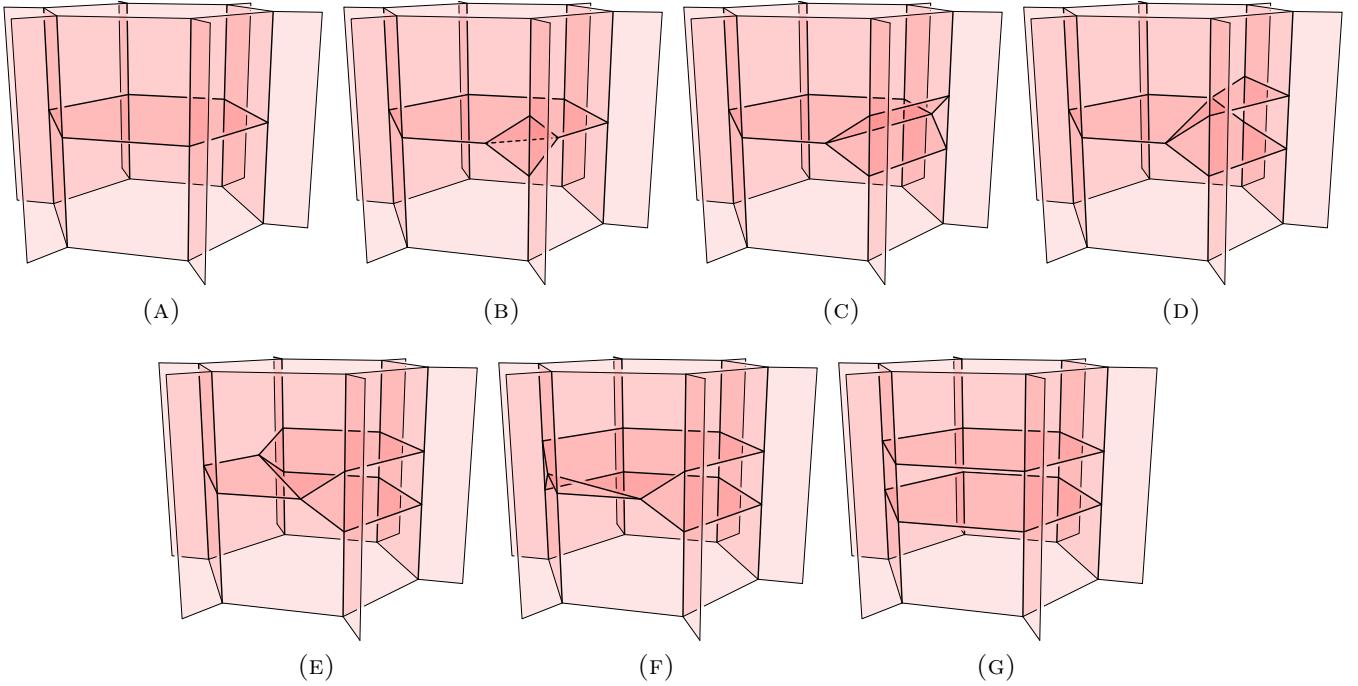


FIGURE 20. Figures 14 and 15, drawn in the dual spine picture.

Note that by applying symmetries of the spine in the initial drawing, we can apply any of the three possible V-moves to any vertex of a spine that contains more than one vertex. The procedure requires an auxiliary neighbouring vertex – any vertex in a spine has such a neighbour unless it is incident to itself along all four arcs of triple points, but then it would be the only vertex in the spine.

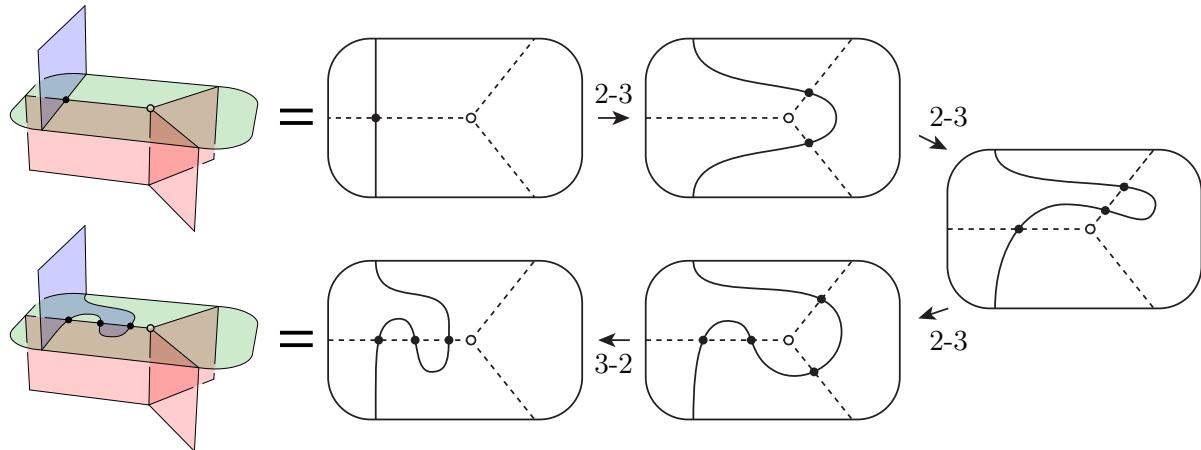


FIGURE 21. The V-move and its construction from 2-3 and 3-2 moves.

3. A WARM-UP EXERCISE

As a warm-up exercise for our proof of Theorem 1.4, we prove the following lemma. This introduces some of the ideas of the main proof in an easier context, and will be useful later. In the proof, we introduce the *triangular 0-2 move* and *pillow marks*. They give a first taste for what is to come.

Lemma 3.1. *Let \mathcal{T}_A and \mathcal{T}_B be two triangulations of a given 3-manifold M with the same number, k , of material vertices. Then there is a sequence of triangulations $\mathcal{T}_A = \mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n = \mathcal{T}_B$, with adjacent triangulations related by 2-3, 3-2, 1-4 and 4-1 moves, such that each triangulation \mathcal{T}_i has at least k material vertices.*

Proof. The statement of this lemma is essentially the statement of Theorem 1.2, except that we also require that the number of material vertices does not go below the number, k , of material vertices that \mathcal{T}_A and \mathcal{T}_B have. We will work from a sequence \mathcal{S} of triangulations $\mathcal{T}_A = \mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n = \mathcal{T}_B$ given to us by Theorem 1.2, and will modify it if and when the number of material vertices goes below k .

In fact, we may assume that $\mathcal{T}_2, \dots, \mathcal{T}_{n-1}$ have precisely $k - 1$ vertices, and thus \mathcal{T}_2 is obtained from \mathcal{T}_1 by a 1-4 move, \mathcal{T}_n is obtained from \mathcal{T}_{n-1} by a 4-1 move, and all other moves are 2-3 and 3-2 moves. If we are able to replace this by a sequence of triangulations for which the number of vertices is always at least k , then by induction we obtain the result.

Our tool to increase the number of vertices in a triangulation is the *triangular 0-2 move*. Its reverse is the *triangular 2-0 move*. See Figure 22. The triangular 0-2 move is performed by ungluing one of the triangular gluings of the triangulation, and inserting into the resulting gap a *triangular pillow* – a three-ball formed from two tetrahedra, glued to each other along three faces, so that the three-ball has an internal vertex. The triangular 0-2 move can be performed on any triangular face of the triangulation, and implemented using a 1-4 move followed by a 3-2 move, as shown in Figure 22. The inverse move is called the *triangular 2-0 move*, and can be performed on any triangular pillow as long as the two outer faces of the pillow are not glued to each other.¹

We modify our sequence of triangulations $\mathcal{T}_A = \mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n = \mathcal{T}_B$ as follows. Immediately before the initial 4-1 move takes us to $k - 1$ material vertices, we perform a triangular 0-2 move on some triangle Δ of the triangulation, increasing the number of material vertices to $k + 1$. If this triangle Δ is not deleted by any of the subsequent moves taking us to \mathcal{T}_n , then we perform those moves as before, with the resulting triangulations modified from the original sequence by our triangular 0-2 move. After performing the final 1-4 move (taking us to $k + 1$ material vertices), we remove the triangular pillow with the triangular 2-0 move, taking us back to k material vertices, and we have arrived at \mathcal{T}_B .

If there is no triangle of \mathcal{T}_A that persists throughout the sequence of moves then the above process will not work. Whichever choice we make for Δ , at some point a 2-3 or 3-2 move in the sequence (or the final 4-1 move) is supposed to delete Δ . However, we will be unable to perform this move because the triangular pillow blocks it. In these cases, we must first move the triangular pillow elsewhere.

In order to organise the argument, for each triangulation of \mathcal{S} , we will mark a triangle with a *pillow mark*, showing where we want the triangular pillow to be. We denote by \mathcal{T}'_i the *marked triangulation* resulting from adding a pillow mark to the triangulation \mathcal{T}_i . Denote the result of

¹If they were glued then the entire triangulation would consist of only the triangular pillow, so performing the 2-0 move would result in a triangulation with no tetrahedra!

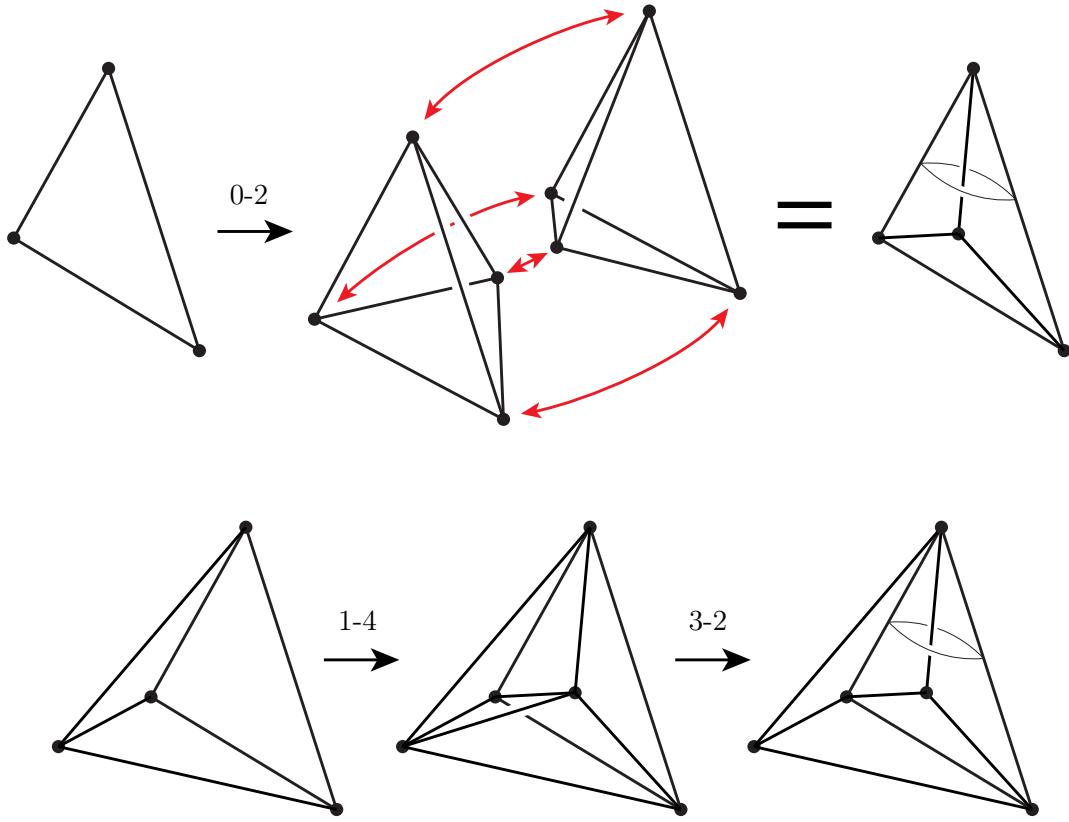


FIGURE 22. Above: the triangular 0-2 move adds a vertex and two tetrahedra to a triangulation. Below: the triangular 0-2 move can be implemented using a 1-4 move applied to a tetrahedron adjacent to the triangle, followed by a 3-2 move. We draw a bigon to show the two external faces of the triangular pillow.

inserting a triangular pillow into a marked triangulation \mathcal{T}'_i according to the pillow mark by $\widehat{\mathcal{T}}'_i$. We refer to these triangulations $\widehat{\mathcal{T}}'_i$ as *waypoint triangulations*. Our plan then is to build a new sequence of triangulations $\widehat{\mathcal{S}}$, which again connects \mathcal{T}_A to \mathcal{T}_B , but for which the number of material vertices never goes below k . We do this by going from $\mathcal{T}_A = \mathcal{T}_1$ to $\widehat{\mathcal{T}}'_1$ via a triangular 0-2 move (i.e. a 1-4 followed by a 3-2), then connecting each adjacent pair $\widehat{\mathcal{T}}'_i$ and $\widehat{\mathcal{T}}'_{i+1}$ of waypoint triangulations using 1-4, 2-3, 3-2 and 4-1 moves, then finally going from $\widehat{\mathcal{T}}'_n$ to $\mathcal{T}_n = \mathcal{T}_B$ via a triangular 2-0 move (i.e. a 2-3 followed by a 4-1).

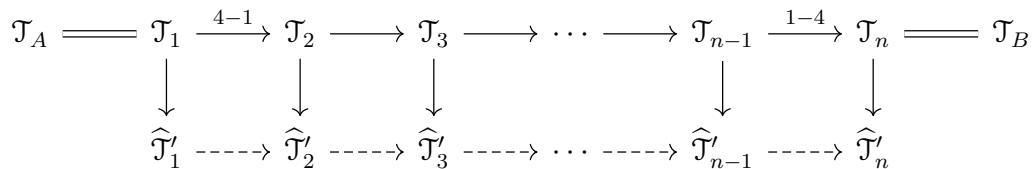


FIGURE 23. The original sequence of triangulations \mathcal{S} is on the upper row. The modified waypoint triangulations are below. Dashed arrows indicate sequences of 1-4, 2-3, 3-2 and 4-1 moves. All vertical arrows indicate triangular 0-2 moves.

First, we describe where the pillow mark is on each \mathcal{T}'_i . The idea is to have the pillow mark of \mathcal{T}'_i move when it is involved in the bistellar move relating \mathcal{T}'_i and \mathcal{T}'_{i+1} . Having done this, we will describe sequences of 1-4, 2-3, 3-2 and 4-1 moves to take each waypoint triangulation $\widehat{\mathcal{T}}'_i$ to the subsequent $\widehat{\mathcal{T}}'_{i+1}$. See Figure 23.

In \mathcal{T}'_1 , we choose the pillow mark arbitrarily. There are two cases under which we determine the new pillow mark on \mathcal{T}'_{i+1} based on the location of the pillow mark on \mathcal{T}'_i :

If a pillow mark is on a triangle of \mathcal{T}'_i that is not deleted by the bistellar move relating \mathcal{T}_i and \mathcal{T}_{i+1} , then we leave it where it is, duplicating it on \mathcal{T}'_{i+1} . If a pillow mark is on a triangle of \mathcal{T}'_i that is deleted by the bistellar move, then we must first move it to any triangle of \mathcal{T}_i that is not deleted by the move. We denote the resulting marked triangulation by \mathcal{T}''_i . This done, we perform the bistellar move and get \mathcal{T}'_{i+1} , where the pillow mark doesn't move between \mathcal{T}''_i and \mathcal{T}'_{i+1} . Note that there is always a triangle of \mathcal{T}_i that is not deleted by the bistellar move, since every triangle on the boundary of the region altered by a bistellar move remains after the move.

Having constructed these markings on the triangulations of \mathcal{S} , the resulting waypoint triangulations $\widehat{\mathcal{T}}'_i$ never have fewer than k material vertices. In fact, $\widehat{\mathcal{T}}'_1$ and $\widehat{\mathcal{T}}'_n$ have $k+1$ vertices, while all others have k vertices. All that remains is to show how to connect \mathcal{T}_1 to $\widehat{\mathcal{T}}'_1$, how to connect $\widehat{\mathcal{T}}'_i$ to $\widehat{\mathcal{T}}'_{i+1}$ for each subsequent pair of waypoint triangulations, and how to connect $\widehat{\mathcal{T}}'_n$ to \mathcal{T}_n by sequences of 1-4, 2-3, 3-2 and 4-1 moves, without reducing the number of material vertices below k .

To connect \mathcal{T}_1 to $\widehat{\mathcal{T}}'_1$, we apply a triangular 0-2 move. This is implemented by a 1-4 move followed by a 3-2 move, as shown in Figure 22, and takes us from a triangulation with k vertices to one with $k+1$ vertices. The connection between $\widehat{\mathcal{T}}'_n$ to \mathcal{T}_n is the same but in reverse.

In the case that we did not move the pillow mark between \mathcal{T}'_i and \mathcal{T}'_{i+1} , then the triangular pillow is not in the way of the corresponding 2-3 or 3-2 move on $\widehat{\mathcal{T}}'_i$, and so we simply apply the move to obtain $\widehat{\mathcal{T}}'_{i+1}$, which doesn't alter the number of vertices. In the case that we did move the pillow mark, we get from $\widehat{\mathcal{T}}'_i$ to $\widehat{\mathcal{T}}''_i$ by first performing a triangular 0-2 move in the location of the new pillow mark, and then perform a triangular 2-0 move to remove the old triangular pillow and obtain $\widehat{\mathcal{T}}'_{i+1}$. Thus the number of material vertices goes up by one to $k+1$ and then down by one to k in this process, but never goes below k . Having moved the pillow mark out of the way of the 2-3 or 3-2 move that converts \mathcal{T}_i into \mathcal{T}_{i+1} , we apply the move to $\widehat{\mathcal{T}}''_i$ to obtain $\widehat{\mathcal{T}}'_{i+1}$, which again doesn't alter the number of vertices.

Making these moves connects the waypoint triangulations together, and so connects \mathcal{T}_A to \mathcal{T}_B by a sequence of bistellar moves, in such a way that all triangulations in the sequence have at least k material vertices. \square

Remark 3.2. One might ask why we use the triangular 0-2 move to increase the number of vertices instead of the simpler 1-4 move. The answer is that the triangular pillow is easier to “hide” from the other Pachner moves than the result of the 1-4 move. For example, the sequence \mathcal{S} might have a triangulation with only two tetrahedra on which we are supposed to perform a 2-3 move. Neither of the two tetrahedra would be valid locations to hide the extra vertex if we used the 1-4 move.

4. PROOF STRATEGY

In order to prove Theorem 1.4, suppose that we are given two triangulations, $\mathcal{T}_A, \mathcal{T}_B$ of a given manifold M with the same number, k say, of material vertices. By Theorems 1.1 and 1.2, we have a sequence \mathcal{S} of triangulations $\mathcal{T}_A = \mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_n = \mathcal{T}_B$, with adjacent triangulations related by

2-3, 3-2, 1-4 and 4-1 moves. We convert this sequence \mathcal{S} of triangulations into another sequence $\widehat{\mathcal{S}}$ that connects \mathcal{T}_A to \mathcal{T}_B , but that does not use 1-4 or 4-1 moves.

Note that every triangulation of the desired sequence $\widehat{\mathcal{S}}$ must have k material vertices, since the number of vertices only changes under 1-4 and 4-1 moves. By Lemma 3.1, we may assume that every triangulation in our sequence \mathcal{S} has at least k material vertices, so we must avoid intermediate triangulations with more than k material vertices. Again, as in the proof of the lemma, it suffices to consider an innermost pair of triangulations with the property that the two triangulations have k material vertices, and the triangulations between them $k+1$. In particular, we may assume that the first bistellar move is a 1-4 move applied to \mathcal{T}_A , followed by a sequence of 2-3 and 3-2 moves, and the sequence is then completed with a 4-1 move resulting in \mathcal{T}_B .

We need a number of tools to achieve this, primarily the *arch*. The arch is a structure we can add in to a spine (or dually a triangulation). The move of *inserting an arch* has the effect of connecting together two three-dimensional regions in the complement of the spine, while only altering the spine in a very localised fashion. Dually, this move unglues a triangular face Δ at which two (or possibly one) tetrahedra are glued, and inserts the arch (made from a single tetrahedron). The effect is to identify two of the three vertices incident to Δ . This move cannot be achieved by 2-3 and 3-2 moves, because these moves do not change the number of vertices of the triangulation. We discuss the arch in detail in Section 5.1.

Since the triangulations in our sequence \mathcal{S} have at most $k+1$ material vertices, we will need to insert at most one arch in each triangulation. In order to organise the argument, we will mark the triangle of each triangulation where we wish to insert an arch. Each such *arch mark* shows both the triangle for an arch to be introduced into, and the pair of incident vertices to be connected. See Figure 24. We denote by \mathcal{T}'_i the *marked triangulation* resulting from adding an arch mark to the triangulation \mathcal{T}_i . Dually, an arch mark is associated to an edge of the spine, and indicates which pair of incident three-dimensional regions are to be connected by inserting the arch.

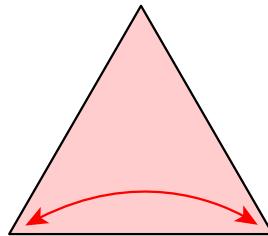


FIGURE 24. An arch mark on a triangle shades the triangle and adds arrows showing the vertices to be identified by the arch.

We will choose where to put each arch mark in such a way that after performing the arch insertions, the resulting triangulations have the same number of material vertices as \mathcal{T}_A and \mathcal{T}_B . Denote the result of inserting an arch into a marked triangulation \mathcal{T}'_i by $\widehat{\mathcal{T}}'_i$. We refer to these triangulations $\widehat{\mathcal{T}}'_i$ as *waypoint triangulations*. Our plan then is to connect each adjacent pair $\widehat{\mathcal{T}}'_i$ and $\widehat{\mathcal{T}}'_{i+1}$ of waypoint triangulations using only 2-3 and 3-2 moves, producing a new sequence of triangulations $\widehat{\mathcal{S}}$, which again connects \mathcal{T}_A to \mathcal{T}_B , but that only uses 2-3 and 3-2 moves. See Figure 25. We give the details of the proof in Section 6.

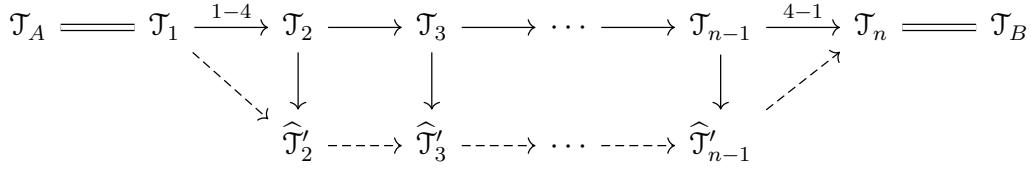


FIGURE 25. The original sequence of triangulations \mathcal{S} is on the upper row. The waypoint triangulations are below. Dashed arrows indicate sequences of 2-3 and 3-2 moves. All vertical arrows indicate arch insertion moves.

5. THE ARCH AND ASSOCIATED CONSTRUCTIONS

Two of the tools that are key to the visual simplicity and conceptual elegance of working with spines are Matveev's *arch* and *arch-with-membrane*. This section describes these constructions and their equivalent moves on triangulations.

5.1. The arch [48, page 7]. Figure 26 shows the arch, and the related structure of an *arch-with-membrane* in spine form. Figure 27 shows the arch in the dual triangulation picture. Also see Fig. 1.29 of [48].

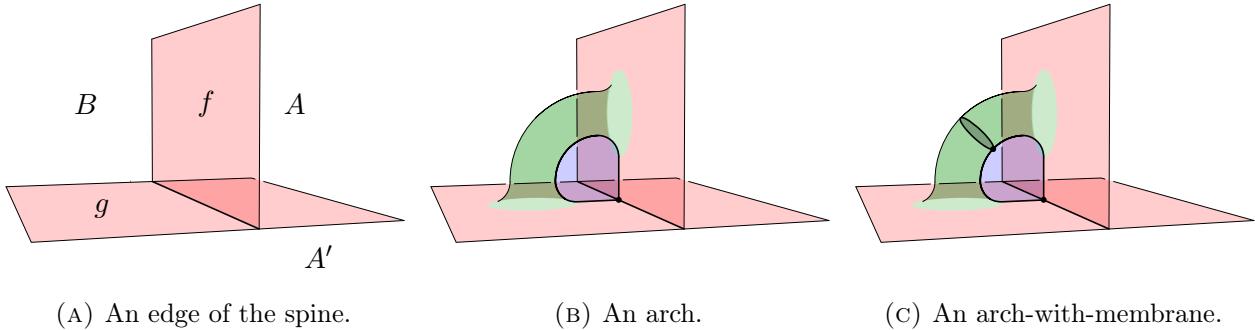


FIGURE 26. Introducing an arch (i.e. going from Figure 26a to 26b) connects together the regions A and A' . Figure 26c shows the related structure of an arch-with-membrane.

Proposition 5.1. *The move of inserting an arch into a spine, connecting two distinct three-dimensional regions, at least one of which is a three-ball, results in a spine of the same pseudo-manifold.*

Proof. We first consider the set of all triple points. It is easy to see that this is a union of open intervals: no loops have been introduced. It is slightly more subtle to see that all of the two-dimensional surfaces (the non-singular points) are again disks. The monogon region in Figure 26b is a disk. The surface containing the tube is formed by starting with the two disks f and g as shown in Figure 26a, taking their connect sum, and then cutting the result along the boundary of the monogon. This cuts open the surface to form a disk, as long as f and g were distinct. If $f = g$, then we have two cases. Consider the transverse coorientation on f pointing from A into B . Following this coorientation around, if we find that it also points from A' into B , then $A = A'$, contradicting the fact that the two regions to be connected by the arch are distinct. If instead we

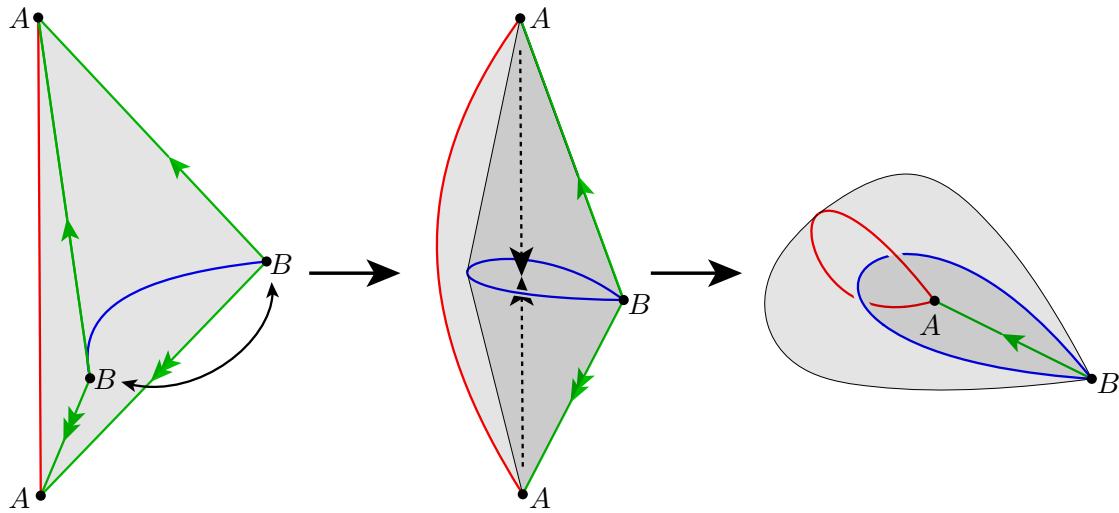


FIGURE 27. The arch is harder to see in the dual picture of the triangulation. Here we show the shape of the single tetrahedron. First we identify the two edges marked with a single arrow, and separately the two edges marked with a double arrow. This produces the middle diagram: a shape with two triangular faces on the exterior, and two cone-shaped faces on the interior. Second, we glue the two interior faces together, which also identifies the two edges with arrows. The result has an exterior consisting of two triangular faces, with two of the vertices of these triangles identified with each other.

find that the transverse coorientation also points from B into A' , then we find that $A = B$ and $B = A'$, so we reach the same contradiction.

Breaking symmetry, suppose that A is homeomorphic to a ball. Since A and A' are distinct, the result of connecting together A and A' is homeomorphic to A' . The topology of the other complementary regions is unchanged. \square

5.2. Building an arch-with-membrane. Using 2-3 and 3-2 moves, we can make an arch-with-membrane on any edge of a spine which has more than one vertex. The moves are applied in a neighbourhood of a vertex of the spine incident to the edge. This is illustrated in Figure 28.

5.3. Implementing a 1-4 move followed by introducing an arch using 2-3 and 3-2 moves. Similarly, we cannot perform a 1-4 move using only 2-3 and 3-2 moves. However, we can recreate the result of performing a 1-4 move followed by inserting an arch using 2-3 and 3-2 moves. This is achieved by constructing an arch-with-membrane, then applying two 2-3 moves. See Figure 29.

6. DETAILS IN THE MAIN PROOF

The point of time has come where we cannot further postpone diving into the intricacies of the proof of the main theorem. This section explains how we position and, if required, move arch marks, and how we connect waypoint triangulations with 2-3 and 3-2 moves. The overall aim is to preprocess the triangulations in such a way, that all that remains is a step that is easy to describe but requires work to justify in detail: sweeping a membrane across a ball, from one arch to another. This very last step is given in §7.

6.1. Arch mark positioning. First, we describe where the arch mark is on each of \mathcal{T}'_2 through \mathcal{T}'_{n-1} . Then we describe sequences of 2-3 and 3-2 moves to take \mathcal{T}_1 to $\widehat{\mathcal{T}}'_2$, to take each waypoint triangulation $\widehat{\mathcal{T}}'_i$ to the subsequent $\widehat{\mathcal{T}}'_{i+1}$, and to take $\widehat{\mathcal{T}}'_{n-1}$ to \mathcal{T}_n .

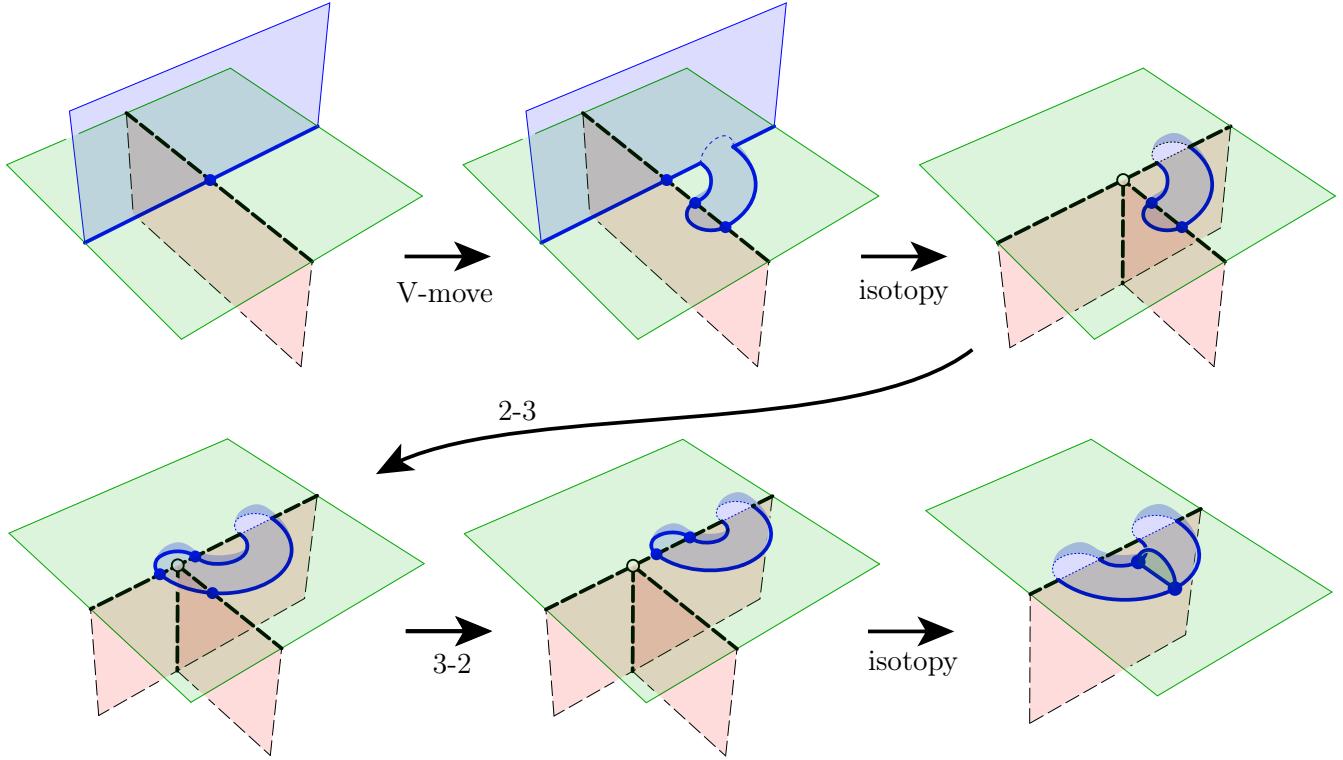


FIGURE 28. Building an arch-with-membrane. We perform a V-move, a 2-3 move and then a 3-2 move. An isotopy moves the bigon along the tube. (In the last diagram here we have zoomed in to the arch-with-membrane and no longer see the original vertex.) A further isotopy (not shown), pushing the bigon in the final diagram forward gives the arch-with-membrane as shown in Figure 26c.

In each triangulation \mathcal{T}'_i , for $i \in \{2, \dots, n-1\}$, the arch mark must connect a material vertex to some other vertex (either material or ideal).²

In \mathcal{T}'_2 , we choose the arch mark to connect the new vertex formed by the 1-4 move to one of the other vertices incident to it. See Figure 30. There are a number of cases under which we determine the new arch mark on \mathcal{T}'_{i+1} based on the arch mark on \mathcal{T}'_i . If an arch mark is on a triangle that is not deleted by the bistellar move relating \mathcal{T}_i and \mathcal{T}_{i+1} , then we leave it where it is, duplicating it on \mathcal{T}'_{i+1} . If an arch mark is on a triangle that is deleted by a 2-3 or 3-2 move, then we must first move it to another triangle of \mathcal{T}_i that is not deleted by the move. We denote the resulting marked triangulation by \mathcal{T}''_i . This done, we perform the bistellar move and get \mathcal{T}'_{i+1} . We will require that the arch marks on \mathcal{T}'_i and \mathcal{T}''_i share a material vertex. This will allow us to move the arch itself from $\widehat{\mathcal{T}}'_i$ to $\widehat{\mathcal{T}}''_i$ in Section 6.3. We deal with the following cases:

- (1) In the case of a 2-3 move, one triangle is removed from the interior of the six-sided polyhedron involved in the move. An arch mark on this triangle connects together two of the three vertices on the “equator” of the polyhedron. We move such a mark to either of the two triangles on the boundary of the polyhedron that share the same two vertices. This done, the arch mark connects the same two vertices together as before, so the arch marks on \mathcal{T}'_i and \mathcal{T}''_i share a material vertex. See Figure 31.

²Note that joining a material vertex to itself would produce an ideal vertex, and connecting two ideal vertices by an arch would act as a connect sum on their links.

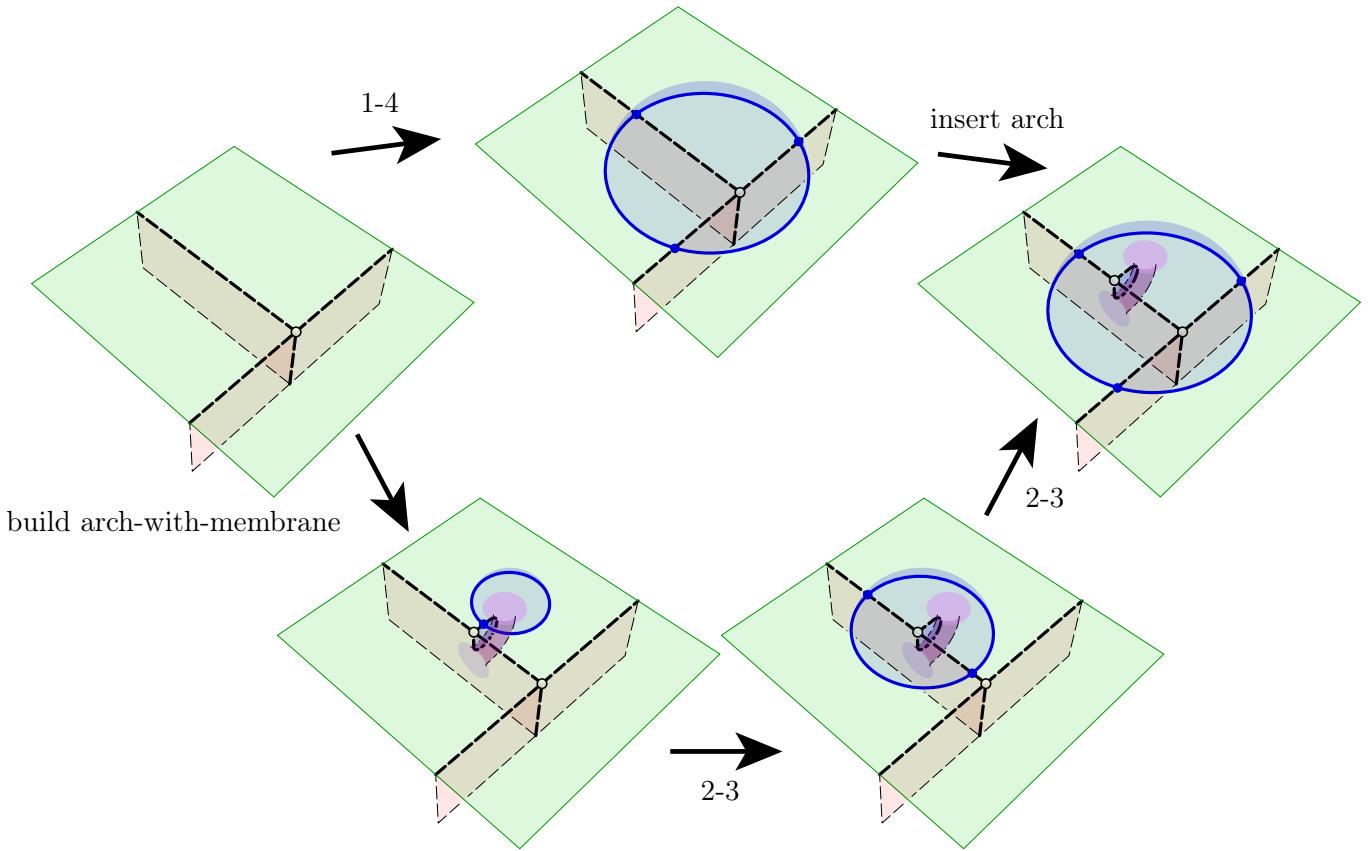


FIGURE 29. By building an arch-with-membrane and performing two 2-3 moves, we implement the composition of a 1-4 move with inserting an arch while keeping the number of material vertices constant.

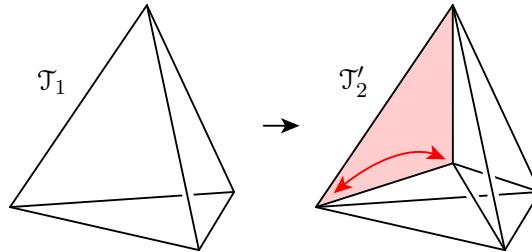


FIGURE 30. We draw an arch mark to connect the vertex created by the 1-4 move to one of the existing vertices.

(2) In the case of a 3-2 move, three triangles are removed from the interior of the six-sided polyhedron involved in the move. An arch mark on one of these triangles is in one of two cases: either it connects a vertex on the equator of the polyhedron to one of the two polar vertices of the polyhedron, or it connects the two polar vertices together. In the first case, we again move such a mark to one of the two triangles of the polyhedron that share the same two vertices. Again, the new arch mark connects the same two vertices so the arch marks on \mathcal{T}'_i and \mathcal{T}''_i share a material vertex. See Figure 32a. In the second case, label the two polar vertices N and S , and let A be one of the equatorial vertices. Without loss of generality, assume that N is a material vertex. If A is distinct from N then we connect A

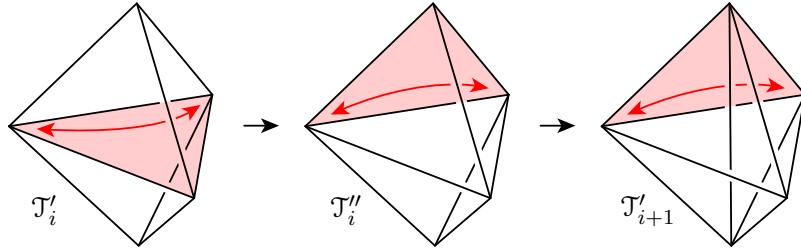
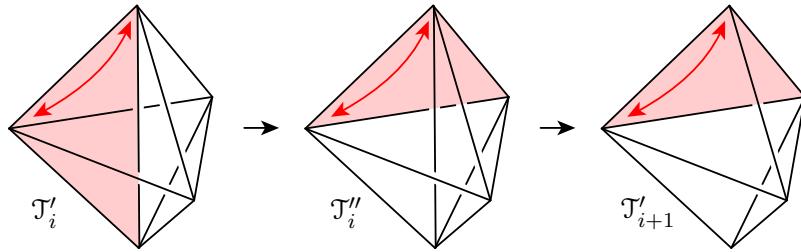
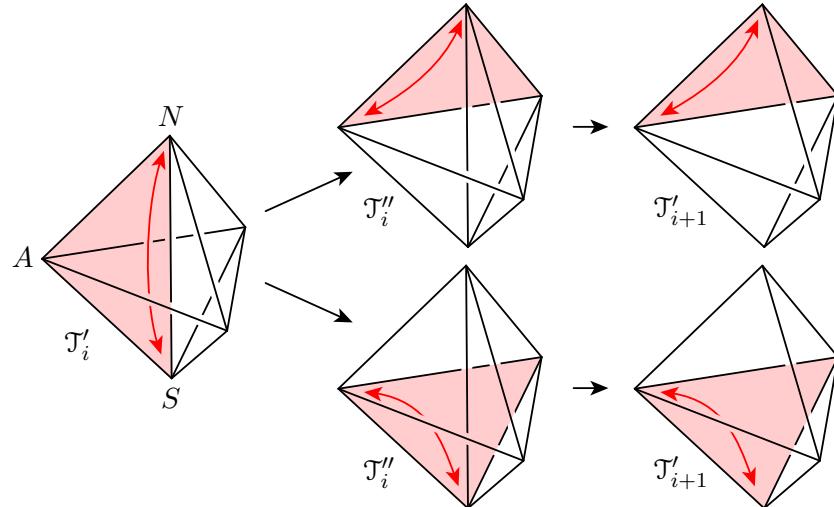


FIGURE 31. If a 2-3 move deletes the face that the arch mark is on, we move the arch mark so as to connect the same pair of vertices before performing the 2-3 move.



(A) Case 1. Here, we move the arch mark so as to connect the same pair of vertices before performing the 3-2 move.



(B) Case 2. We take the upper path if A is distinct from N , and the lower path if not. Here, the new arch mark may or may not connect the same pair of vertices together as the old arch mark.

FIGURE 32. If a 3-2 move deletes the face that the arch mark is on, we move the arch mark before performing the 3-2 move.

to N , again by adding an arch mark on a boundary triangle of the polyhedron. Otherwise we connect A to S , again by adding an arch mark on a boundary triangle of the polyhedron. By hypothesis, $A = N$ is distinct from S . Here, the arch marks on T'_i and T''_i share the material vertex N . See Figure 32b.

With the above moves, we can add arch marks to all triangulations $\mathcal{T}_2, \dots, \mathcal{T}_{n-1}$ in the sequence, and produce all of the waypoint triangulations $\widehat{\mathcal{T}}'_i$, all of which have the same number of material vertices as \mathcal{T}_A and \mathcal{T}_B by construction.

6.2. Moving from \mathcal{T}_1 to $\widehat{\mathcal{T}}'_2$. The triangulation \mathcal{T}_1 is related to \mathcal{T}_2 by a 1-4 move. we need a sequence of 2-3 and 3-2 moves that produces the result of a 1-4 move followed by adding an arch, converting \mathcal{T}_1 to $\widehat{\mathcal{T}}'_2$. We gave details of this in Section 5.3.

6.3. Moving between waypoint triangulations using 2-3 and 3-2 moves. If the arch mark does not move between \mathcal{T}'_i and \mathcal{T}'_{i+1} (since the 2-3 or 3-2 move does not destroy its triangle), then $\widehat{\mathcal{T}}'_i$ and $\widehat{\mathcal{T}}'_{i+1}$ are related by the same 2-3 or 3-2 move, so we connect them using that move.

If however we had to move the arch mark first, via a marked triangulation \mathcal{T}''_i , then we need a way to move the arch via 2-3 and 3-2 moves, connecting the triangulations $\widehat{\mathcal{T}}'_i$ and $\widehat{\mathcal{T}}''_i$. Having moved the arch, we then connect $\widehat{\mathcal{T}}''_i$ to $\widehat{\mathcal{T}}'_{i+1}$ by the same 2-3 or 3-2 move that connects \mathcal{T}_i to \mathcal{T}_{i+1} .

We have arranged matters so that the arch marks on \mathcal{T}'_i and \mathcal{T}''_i share a material vertex, v say. In the dual spine picture for \mathcal{T}_i , there is a three-ball B dual to v since v is a material vertex. In $\widehat{\mathcal{T}}'_i$, B is connected via a cylinder in the neck of the arch to the dual region R to the vertex on the other end of the arch mark on \mathcal{T}'_i . See Figure 33a. Here the arch is to the right of the figure, and in the middle we have drawn part of the boundary of B , as seen from within B .

First we build a *arch-with-membrane* (see Figure 26c) in the location of the new arch mark. This can be made using 2-3 and 3-2 moves – see Section 5.2. The result is shown in Figure 33b. This done, we see B with an arch glued to one end and an arch-with-membrane glued to the other.

Next, we sweep the membrane from the arch-with-membrane, through B , and into the other arch. This turns the arch-with-membrane into an arch, and the arch into an arch-with-membrane. See Figure 33c. Then, we deconstruct the new arch-with-membrane, following the reverse of the sequence of moves used to build such an arch-with-membrane. Having performed these moves, we have moved the arch. See Figure 33d. We give details of this sweep move in Section 7.

6.4. Moving from $\widehat{\mathcal{T}}'_{n-1}$ to \mathcal{T}_n . If our arch mark is connected to the material vertex w we are about to remove in the 4-1 move then we perform the reverse procedure to Section 6.2 and we are done. If however the arch mark is connected to some other material vertex v , then we must do some further work first. If v and w share an edge e , then we move the arch mark to a triangle containing e so that it connects v to w . As in the previous section, moving the corresponding arch can be performed by sweeping across a three-ball, here it is the three-ball dual to v . If v and w do not share an edge, then we move w until they do. After the 4-1 move our triangulation is \mathcal{T}_n , and we may think of \mathcal{T}_{n-1} as \mathcal{T}_n with one of its tetrahedra, σ say, subdivided into four. See Figure 34a. By performing a 3-2 move on one of the four edges incident to w , we convert \mathcal{T}_{n-1} into the result of inserting a triangular pillow into a face of \mathcal{T}_n incident to σ . See Figure 34b. Note that we can perform this 3-2 move because the three tetrahedra incident to this edge are distinct, since we were able to perform a 4-1 move on them and one other tetrahedron. The 3-2 move creates a new face that forms one of the two outer faces of the triangular pillow. By performing a 2-3 move on the other outer face, we then get the result of subdividing a tetrahedron σ' of \mathcal{T}_n incident to σ . See Figure 34c. Note that we can perform this 2-3 move because the two tetrahedra incident to the face are distinct: one is part of the triangular pillow, while the other is not.

These two bistellar moves transport the vertex w from σ to σ' , and once again we are a 4-1 move away from \mathcal{T}_n . By repeatedly moving w in this way, we can arrange for v and w to share an edge.

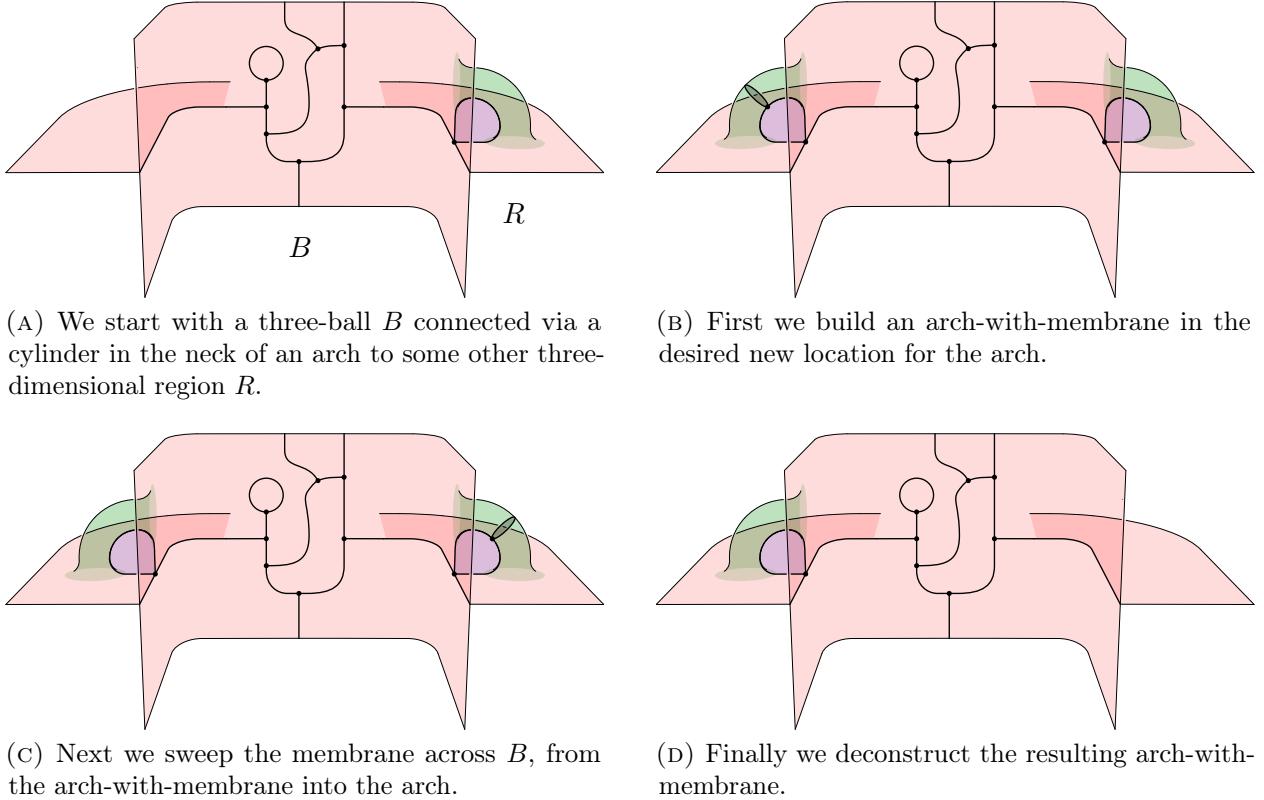


FIGURE 33. Moving an arch across a material vertex, drawn in the dual spine picture.

Also note that in moving w in this way, we do not destroy the triangle containing the arch mark, because in order to do so, w and v would already share an edge.

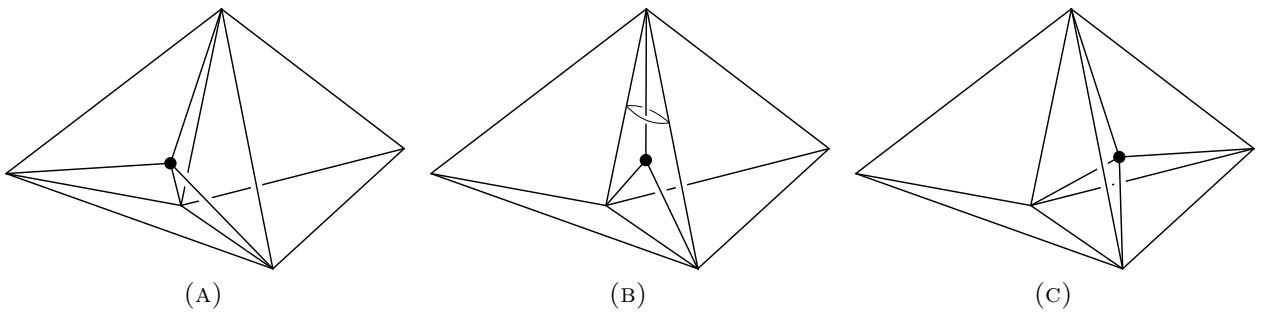


FIGURE 34. The vertex w can be moved from one tetrahedron of \mathcal{T}_n to a neighbouring tetrahedron by a 3-2 move followed by a 2-3 move.

7. MOVING A MEMBRANE ACROSS A BALL, FROM ONE ARCH TO ANOTHER

There are two key steps at which we depart from the proofs of Matveev and Amendola in how we sweep a membrane across a ball. We point these out in the discussion below.

We must connect together triangulations \mathcal{T}_α and \mathcal{T}_β , say, shown in dual spine form in Figures 33b and 33c respectively. We note that this is the setting of Lemma 1.2.16 (2-cell replacement) of [48]. Matveev's method is to slide the membrane across the ball from the arch-with-membrane into the arch. The combinatorial changes as one slides the membrane are 2-3 and 3-2 moves, and what we call *quadrilateral 0-2 moves*. In the literature this move is often called simply a “0-2 move”, but we add the adjective “quadrilateral” to distinguish them from our triangular 0-2 move.³ Matveev refers to a quadrilateral 0-2 move as a *lune move*. The implementation of the lune move in terms of 2-3 and 3-2 moves is quite subtle in full generality, due to self-gluings in the vicinity of the move. In our setting, we will be able to avoid these subtleties (although at the cost of adding other subtleties!) The use of 2-3, 3-2 and lune moves to sweep the membrane across the ball is also somewhat tricky. The situation is (again) complicated by the fact that the ball may have self gluings, so when we isotope the membrane across a self-glued face of the ball, it potentially interacts with parts of itself meeting that face from the other side. Thus, Matveev's proof that the moves can be done to achieve the sweep involves a PL topology argument.

In contrast, our method is entirely combinatorial. Our strategy is to first apply moves to each of the triangulations \mathcal{T}_α and \mathcal{T}_β to convert them to triangulations \mathcal{T}'_α and \mathcal{T}'_β , which again only differ by the position of the membrane, and for which the ball has no self-gluings. This means that we can then organise the isotopy of the membrane from one arch to the other in a combinatorial way, and convert \mathcal{T}'_α into \mathcal{T}'_β .

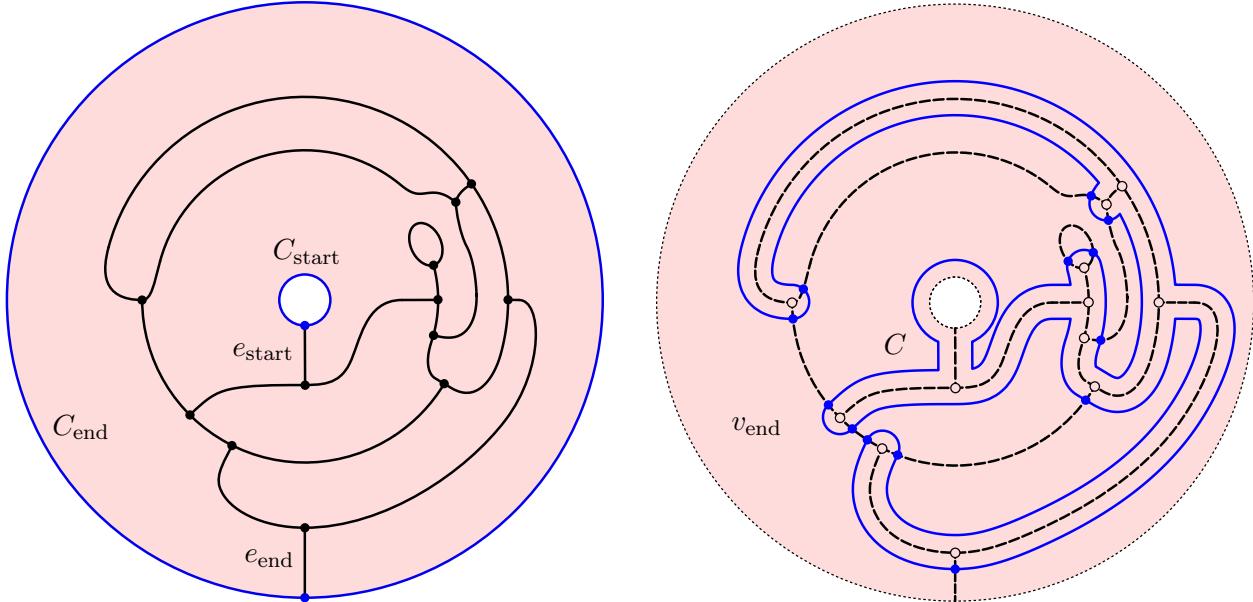
First of all, we see what to do in the case that the ball has no self-gluings.

7.1. The no self-gluings case. Again, we must get from Figure 33b to Figure 33c by sweeping the membrane disk across the three-ball B , which we assume for the moment has no self-gluings. Equivalently, we have to slide the boundary of the disk, a circle C , across the annulus A formed from the boundary of the B minus the initial and final positions of the membrane. Let C_{start} and C_{end} be the start and end positions of C , which are therefore also the two boundary components of A . The key point is that the pattern on the annulus A formed from the intersection of the boundary of B with parts of the spine outside of B is a trivalent graph Γ . Let e_{start} and e_{end} be the edges of Γ that intersect C_{start} and C_{end} respectively. See Figure 35a.

We slide C across A in two stages. First, we push C off C_{start} slightly into A , then choose a spanning tree T for Γ (not including e_{end}), and push C over T . This is achieved by a sequence of 2-3 moves, as shown in Figure 36a. Figure 35b shows the result in our example. In the second stage, we slide C over the 2-cells of A , as follows. These 2-cells themselves form a tree T' , dual to T . There is one special vertex v_{end} of T' , corresponding to the 2-cell incident to C_{end} . We may collapse T' onto v_{end} by performing a sequence of quadrilateral 2-0 moves (Figure 36b shows the reverse move). We perform a quadrilateral 2-0 move on a leaf of T' to remove it from the tree as we push C over the corresponding 2-cell. Since T' is a tree, we can always find a leaf vertex other than v_{end} to perform a quadrilateral 2-0 move on.

7.2. Performing the quadrilateral 2-0 moves. In Lemma 1.2.11 of [48], Matveev shows that the quadrilateral 0-2 move can be performed using 2-3 and 3-2 moves in a very general setting. In this paper however, we only prove that it can be performed in our more restrictive setting.

³The name comes from the dual triangulation picture: the quadrilateral 0-2 move inserts a quadrilateral pillow consisting of two tetrahedra, instead of the triangular pillow inserted in the triangular 0-2 move.



(A) The annulus A . The boundary of the membrane is shown at its initial position C_{start} and final position C_{end} , which are also the boundary curves of A .

(B) After pushing C over the spanning tree T .

FIGURE 35. The setting in which we sweep a membrane across a three-ball, from an arch-with-membrane to another arch, assuming that the ball has no self-gluing.

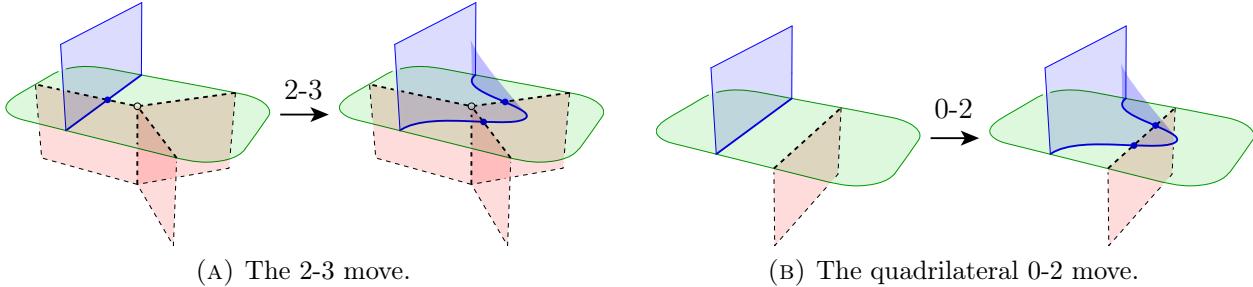


FIGURE 36. Sliding the upper disk past parts of the spine below the plane is achieved by the 2-3 move and (the inverse of the) quadrilateral 0-2 move.

Here we state the moves purely in terms of the circle intersecting the planar graph. The meaning of “0-2”, “2-3” and “3-2” moves is as for the boundary of the upper disk moving against the boundary of the spine below the plane, as shown in Figure 36.

See Figure 37. We have a “bump” in the curve C , which protrudes into the face f . We wish to push the bump into g . There are always two components A_{start} and A_{end} of $A - C$, which contain C_{start} and C_{end} respectively, and we want to push the bump from the A_{start} side towards the A_{end} side.

The bump intersects ∂g in two points, say p_1 and p_2 . The circle C also intersects ∂g at some other point on ∂g . This follows since if not, after the 2-0 move C would lie in the interior of the

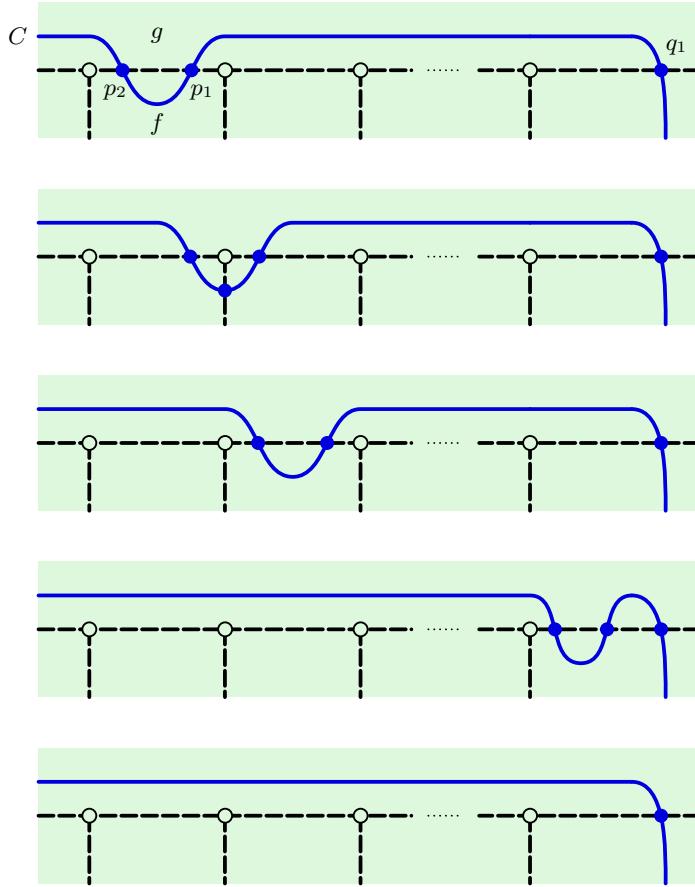


FIGURE 37. Moves to perform a quadrilateral 2-0 move.

face g of the spine. But C cannot contract to either side since it is essential in A . Therefore g cannot be a disk, which is a contradiction. Let q_1 and q_2 be the first such intersection points moving away from the bump in the directions of p_1 and p_2 respectively. Let a_1 and a_2 be the subarcs of ∂g between p_1 and q_1 , and between p_2 and q_2 respectively. At least one of a_1 and a_2 does not contain e_{start} . Relabelling if necessary, we assume that a_1 does not contain e_{start} .

In our setting, a_1 is in A_{start} . Therefore there are no intersections between C and a_1 , and by our previous choice a_1 does not contain e_{start} . So, we can move the bump along a_1 via a sequence of 2-3 and 3-2 moves, staying inside of A , as shown in Figure 37, until it is next to q_1 . Finally, the inverse of a V-move removes the bump.

After performing all of these moves, the position of C is identical to its original position, except that the original bump has been removed. This then achieves the quadrilateral 2-0 move.

In this process, we should be concerned that some moves we make might be blocked by self-identifications in our graph Γ . It is possible that some of the edges of a_1 are identified with each other, corresponding to identifications among the sides of the face g . Note that any identification between these edges must be orientation reversing since the graph is planar. Moreover, there can be no identification between adjacent edges since the graph is trivalent. Each such step of the bump from sitting in one edge to sitting in the next involves only a small neighbourhood of the vertex between the two edges. By the above comments, there can be no identification between parts of Γ involved in the step.

7.3. Removing self-gluings. Now that we know how to sweep a membrane across a ball when the ball has no self-gluings, we now have to deal with balls with self-gluings.

Fixing some notation, we have membranes m_α and m_β in \mathcal{T}_α and \mathcal{T}_β respectively, and the difference in the position of the two membranes is the only difference between the two spines. Let $\mathcal{T}_{\alpha\beta} = (\mathcal{T}_\alpha \cup m_\beta) = (\mathcal{T}_\beta \cup m_\alpha)$. See Figure 38.

The ball B is one of the connected regions of M minus the spine $\mathcal{T}_{\alpha\beta}$. Let R_α be the other incident region to m_α , and similarly for R_β and m_β . We have $R_\alpha \neq B \neq R_\beta$, although R_α may or may not be the same region as R_β . The 1-cell ∂m_α is incident to two 2-cells other than m_α , namely f_α and g_α . The 2-cell f_α separates B from another region C_α , while g_α separates R_α from C_α . The region C_α meets itself along a 2-cell d_α associated to the arch-with-membrane. We have similar notation for the β side of the picture. The regions C_α and C_β may each be the same as any of the other regions listed, including B if f_α or f_β is a self-gluing of B . Moreover, we may have $f_\alpha = f_\beta$ if the intersection of the 1-skeleton of the spine with ∂B does not separate the two arches, and we may have $g_\alpha = g_\beta$, connecting around outside of B .

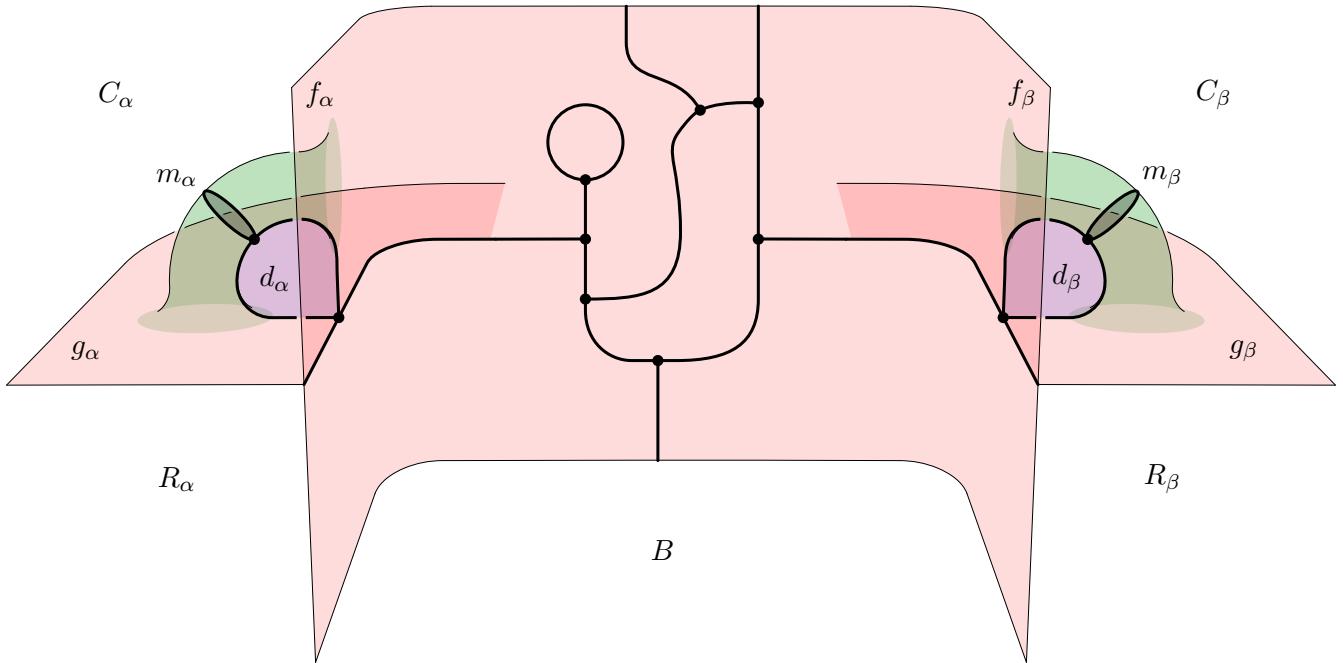


FIGURE 38. The spine $\mathcal{T}_{\alpha\beta}$.

Note that ∂B (before we glue up the faces of B) consists of $(m_\alpha \cup m_\beta)$, together with an annulus A . After gluing, the ball B may be glued to itself along some faces, which means that the annulus has some faces identified.

The pattern on $\partial B \cup g_\alpha \cup g_\beta$ formed from the intersection of the boundary of B with parts of the spine outside of B (and parts of ∂d_α and ∂d_β intersecting g_α and g_β) is a trivalent graph Γ .

We will remove self-gluings of B , using moves analogous to the barycentric subdivision moves from Section 2.5. As we modify our triangulations to remove self-gluings, we abuse notation and continue to refer to the results of our various moves as \mathcal{T}_α , \mathcal{T}_β and $\mathcal{T}_{\alpha\beta}$. Similarly, we will maintain the same names for the various cells of our spines, even as they are modified or moved. Having removed all self-gluings, we move the membrane from its position in \mathcal{T}_α to its position in \mathcal{T}_β ,

similarly to as in Section 7.1 (with a few modifications we will give after this section). This connects the two triangulations.

In this process, we may do moves to \mathcal{T}_α in the vicinity of m_β . We will alter the corresponding moves applied to \mathcal{T}_β (and $\mathcal{T}_{\alpha\beta}$), so as to give the same result. To be more precise, we will apply various *composite moves* to \mathcal{T}_α , \mathcal{T}_β and $\mathcal{T}_{\alpha\beta}$. The same composite moves on these triangulations may consist of a different sequence of bistellar moves, to deal with the fact that membranes may or may not be present in different positions for the three spines. However, after each such composite move, we require that the identity $\mathcal{T}_{\alpha\beta} = (\mathcal{T}_\alpha \cup m_\beta) = (\mathcal{T}_\beta \cup m_\alpha)$ holds. Our composite moves are given in the following sections.

7.3.1. Ensure that there are no BBB spine edges. First, we will require that there are no edges of $\mathcal{T}_{\alpha\beta}$ (thought of as a spine) whose only incident region is the ball B (“*BBB*-edges”). We can remove *BBB*-edges as follows. First, if there is a *BBB*-edge then there must be a vertex v of the spine $\mathcal{T}_{\alpha\beta}$ incident to B three times and some other region, R say, distinct from B . If no such vertex can be found, then the *BBB*-edges and their incident regions form a connected component of the manifold, which therefore contains only B . But we know that B is incident to a vertex of the spine inside the arch-with-membrane, which is incident to $R_\alpha \neq B$, a contradiction. Having found such a vertex v , we perform a 2-3 move along its incident *BBB*-edge e , noting that the 2-3 move is possible: e cannot connect v to itself since no other edge incident to v is a *BBB*-edge. This 2-3 move reduces the number of *BBB*-edges by one. By induction we can remove all *BBB*-edges.

If ∂d_β is a *BBB*-edge in \mathcal{T}_α , then this procedure will perform a 2-3 move along it. We must apply appropriate bistellar moves to \mathcal{T}_β so that \mathcal{T}_β is identical to $\mathcal{T}_\alpha \cup m_\beta$ in the vicinity of m_β . This is done by using two 2-3 moves along ∂d_β instead of just one. See Figure 39. We apply a similar technique if ∂d_α is a *BBB*-edge in \mathcal{T}_β .

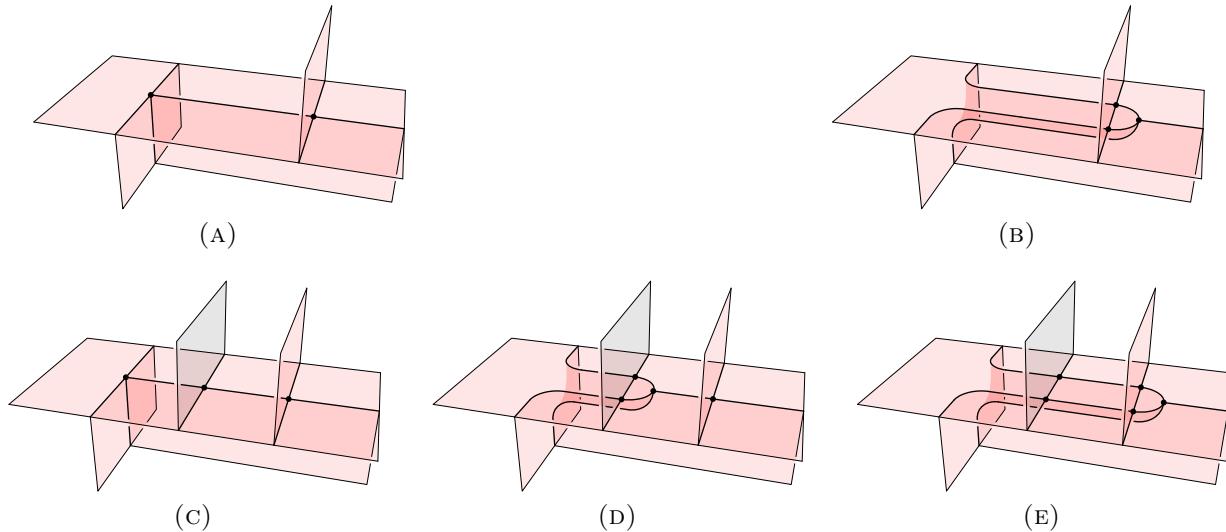


FIGURE 39. Top row: a 2-3 move applied to an edge. Bottom row: the same result is achieved despite a membrane disk (shaded gray) being in the way by applying two 2-3 moves.

Having completed this step, no triangle of $\mathcal{T}_{\alpha\beta}$ (thought of as a triangulation) has three B vertices.

7.3.2. Thickening the 1-skeleton of B . Next, let K be the 1-skeleton of B (in $\mathcal{T}_{\alpha\beta}$, thought of as a spine), union the edges in ∂g_α , and ∂g_β , and those edges that intersect ∂m_α and ∂m_β , but minus ∂m_α and ∂m_β themselves. We wish to “thicken up” K , inspired by the implementation of barycentric subdivision given in Section 2.5, which uses 1-4 moves. Of course we cannot perform any 1-4 moves here, but we can perform the composition of a 1-4 move with adding an arch, as shown in Figure 29. Moreover, since every vertex of K is incident to the region B at most twice, we can arrange matters so that the arch is on an edge of the resulting spine which is not incident to B . In the dual triangulation picture, we apply a 1-4 move and can indicate the arch with an arch mark. The picture is the same as shown in Figure 30. Here, we put the arch mark on the internal triangle incident to two non- B vertices. Following this convention, after these moves our triangulation \mathcal{T}_α is the result of adding arches at arch marks drawn on a triangulation \mathcal{T}'_α say, where \mathcal{T}'_α has an extra material vertex for every 1-4 move. On \mathcal{T}'_α , we label the new vertices of the triangulation “3”, as in the first stage of barycentric subdivision, given in Section 2.5. This thickens up the vertices of K . To thicken up the edges of K , we continue with the barycentric subdivision procedure, adding vertices labelled “2” to \mathcal{T}'_α , again with arches connecting the new vertices back to non- B vertices, with arch marks placed on triangles with vertices labelled 2, 3 and a non- B 0 vertex. Again this is always possible to do, since no face of the triangulation has three B vertices. In the second part of this stage of barycentric subdivision, we perform a 2-3 move on a face of the triangulation with three 0 vertices. We are able to do this without destroying one of our arch marks because all of our arch marks are on triangles with at least one non-0 vertex. This done, we have thickened K .

When we thicken an edge that intersects ∂m_α , on \mathcal{T}_α we apply two 2-3 moves instead of one, just as in Section 7.3.1, in order to keep up with \mathcal{T}_β . The same applies with α and β swapped.

Having completed this step, each face of the spine $\mathcal{T}_{\alpha\beta}$ (and also $\mathcal{T}'_{\alpha\beta}$) incident to B has all of its incident vertices and edges distinct. Moreover, in $\mathcal{T}'_{\alpha\beta}$, no arch mark is on a triangle with a B -vertex.

7.3.3. Thickening self-gluings of B . We now wish to continue the barycentric subdivision theme, thickening up the faces of B that are self-gluings. For such a face, we are precisely in the setting of adding a vertex labelled 1, as in Figure 14. Again here we implement the 1-4 and arch move on \mathcal{T}_α with 2-3 and 3-2 moves. The corresponding arch mark is placed on a triangle of \mathcal{T}'_α not incident to a B vertex. In fact we may place it on a triangle with vertices labelled 1, 2 and 3. Next we follow the subsequent steps, as shown in Figure 15. These are all 2-3 moves, deleting triangles with two B vertices, which therefore do not have arch marks on them, and a final 3-2 move. Each of the three triangles deleted in the final 3-2 move has a B vertex, and so again there are no arches in the way of this move.

In this step, special handling for m_α (and m_β respectively) is needed in the case that $B = C_\alpha$, so f_α is a self-gluing of B . In this case, for \mathcal{T}_β , the membrane m_α is not present, so $f_\alpha = g_\alpha$, and the entire face is thickened. So, we must do the same for \mathcal{T}_α , despite m_α being in the way. To achieve this, first of all let’s break the operation of thickening up a face, as shown in Figure 20, into two stages.

- (1) First, from Figure 20a to 20b we create a new three-cell region (here, connected via an arch to one of the neighbouring regions).
- (2) Second, from Figure 20b onwards, we expand the new region outwards, *collapsing* the face away.

The data determining a face collapse move consists of

- (1) the face f to be collapsed, together with
- (2) the *collapsing edge* of f which expands into f as f collapses.

In Figure 20b the collapsing edge is drawn dashed. In order to perform a face collapse, we require that all vertices of the face be distinct, and that there are at least three vertices.

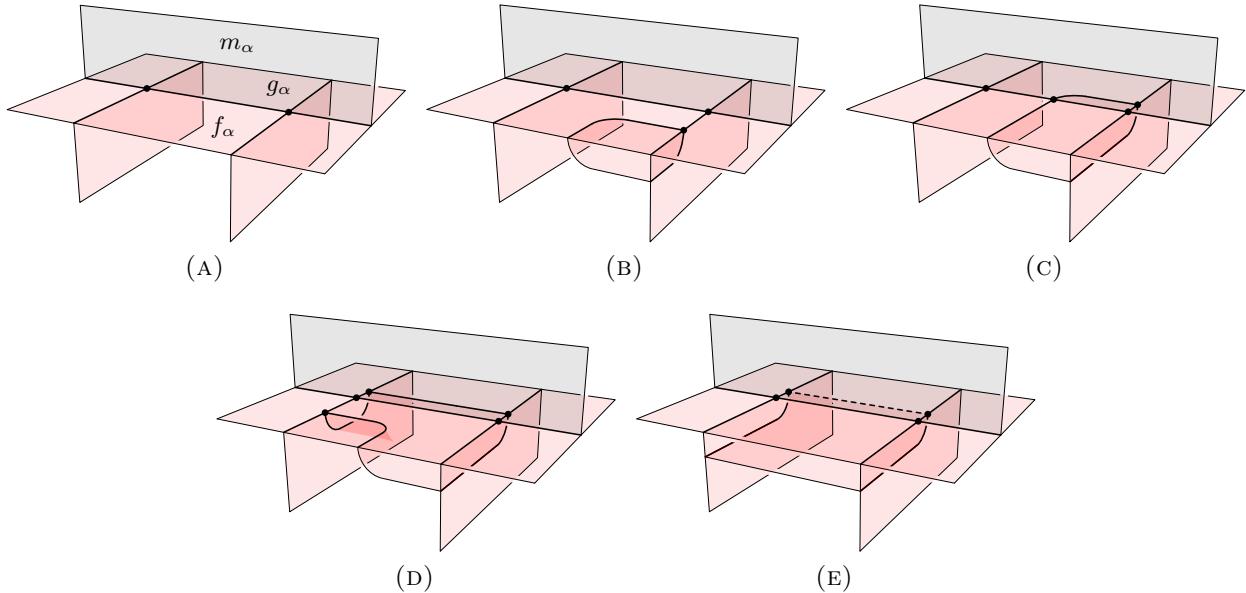


FIGURE 40. Collapsing a face past a membrane (shaded gray).

Figure 40 shows the effect of thickening a face f_α incident to a membrane m_α . (Note that this is not a completely general picture: the degree of the remains of f_α could reach three earlier than shown, in which case we perform a 3-2 move rather than further 2-3 moves.) Now, to continue the thickening of f_α into g_α , we perform a second collapsing step, this time collapsing g_α , without first creating a new region. Instead we use the region already created to thicken f_α , collapsing g_α with the collapsing edge shown in Figure 40e. Note that by construction, at the start of the second collapsing step the degree of g_α is at least three, so we are able to perform the collapse: If the degree of g_α were two, then before the thickening move we would have had an arc of Γ cross ∂m_α and return, forming a bigon. But our previous moves cannot produce such a configuration.

Also note that if $g_\alpha = g_\beta$, then it is possible that we collapse g_α first, which then requires us to collapse into f_α . Finally, it is also possible that we have that $f_\alpha = f_\beta$, since the graph Γ does not separate them, and moreover that $g_\alpha = g_\beta$, connecting around outside of B . In this case, the membrane (in either position) cannot be ignored - without it there is one face $f_\alpha = f_\beta = g_\beta = g_\alpha$, which is therefore not a disk. When we collapse this face, we continue on from Figure 40e to Figure 41a. Completing the collapse in the rest of Figure 41, we gain an extra rectangular face, parallel to the membrane, which separates the region that expanded into $f_\alpha = f_\beta = g_\beta = g_\alpha$ from itself. Here, we need to do a little more work to maintain $(\mathcal{T}_\alpha \cup m_\beta) = (\mathcal{T}_\beta \cup m_\alpha)$. In particular, after thickening $f_\alpha = f_\beta = g_\beta = g_\alpha$, we have a rectangle near m_α in \mathcal{T}_α , but one near m_β in \mathcal{T}_β .

We can move the rectangle past a membrane using a 2-3 followed by a 3-2 move, and we can move it around the boundary of the thickened face using similar moves. See Figure 42. Note that although we are sliding a disk (the rectangle) through a ball (the thickened face), this is a much simpler setting than our original problem: there are no self-identifications and the thickened face

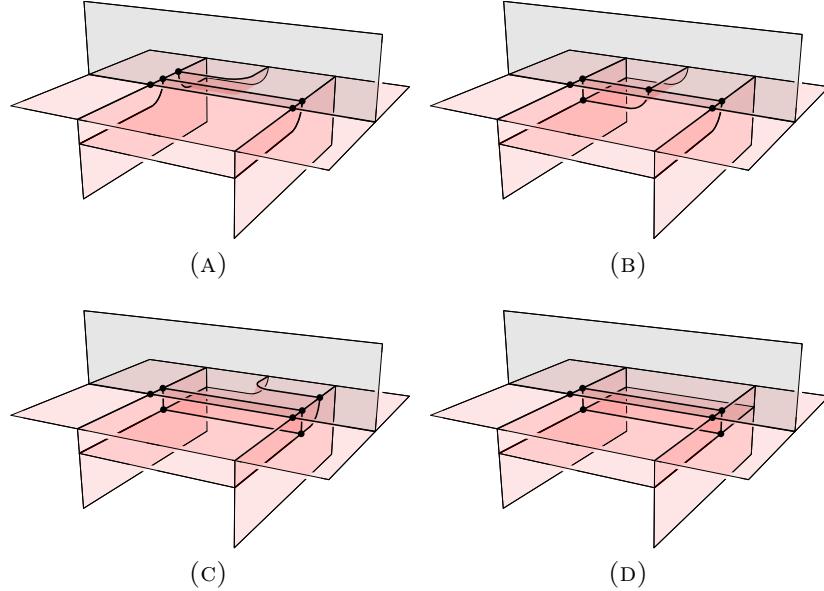


FIGURE 41. Collapsing a face past a membrane (shaded gray) when the membrane is a self-gluing.

has a much simpler graph on its boundary. These two rectangle moves suffice to transport it from its location near m_α to its location near m_β , as long as we don't attempt to move it past the arch connecting the thickened face to the outside. Thankfully, we can either move the rectangle through B or around the outside of B , and one of these two paths does pass by the arch.

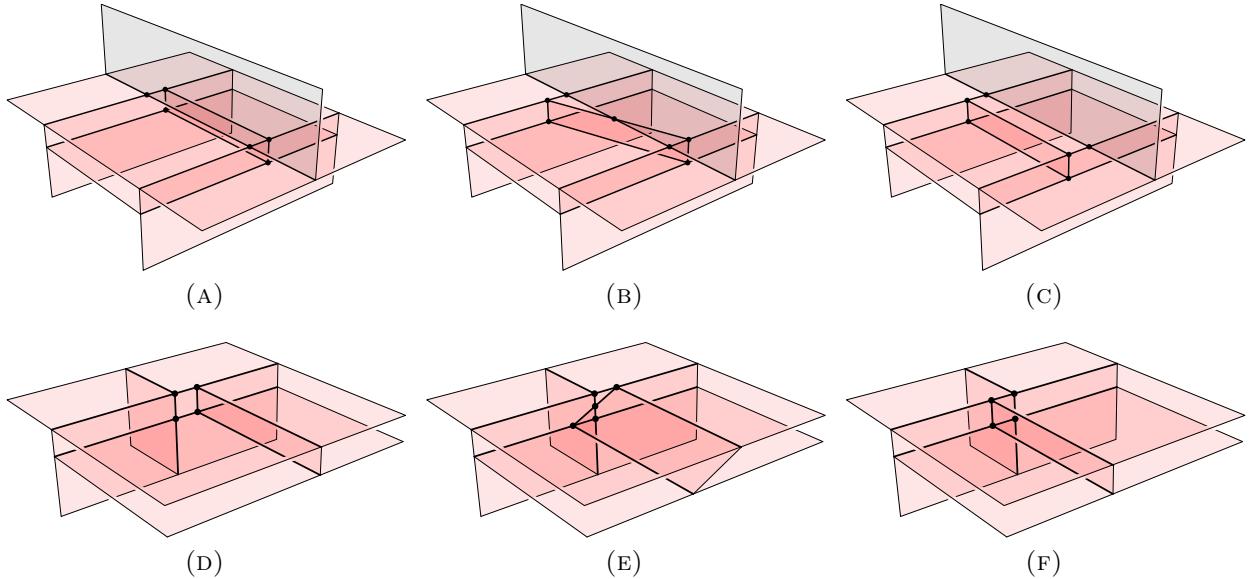


FIGURE 42. Sliding the rectangle past a membrane (first row) or a corner of the thickened face (second row). In both cases, this is achieved by applying a 2-3 move followed by a 3-2 move.

7.4. Modifications to moving a membrane across a ball. The only difference in the setup here in comparison to that in Section 7.1 is that there may be more than one edge of Γ exiting

the annulus A through each of its boundary curves C_{start} and C_{end} : If ∂d_α is a BBB -edge then in Section 7.3.1 we remove it. The resulting spine has two edges intersecting ∂m_α . When we thicken up edges in Section 7.3.2, this again doubles the number of edges intersecting ∂m_α . The same is true for edges intersecting ∂m_β . Our algorithm is however almost identical to the one given in Section 7.1. All we need to do is choose one of the edges intersecting C_{start} to be e_{start} , and have the spanning tree include that e_{start} but no other edges intersecting C_{start} or C_{end} . The rest of the algorithm then goes through identically.

Acknowledgements. Research of the first and third authors is supported in part by the Australian Research Council under the Discovery Projects funding scheme (DP160104502). The second author was supported in part by National Science Foundation grant DMS-1708239. The authors thank Robert Löwe for helpful comments on a previous draft. We are most grateful to the referee for their careful reading and helpful remarks.

REFERENCES

- [1] Colin C. Adams, *The knot book*, American Mathematical Society, Providence, RI, 2004, An elementary introduction to the mathematical theory of knots, Revised reprint of the 1994 original.
- [2] Ian Agol, Joel Hass, and William Thurston, *The computational complexity of knot genus and spanning area*, Trans. Amer. Math. Soc. **358** (2006), no. 9, 3821–3850.
- [3] Selman Akbulut and John D. McCarthy, *Casson’s invariant for oriented homology 3-spheres*, Mathematical Notes, vol. 36, Princeton University Press, Princeton, NJ, 1990, An exposition.
- [4] James W. Alexander, *The combinatorial theory of complexes*, Ann. of Math. (2) **31** (1930), no. 2, 292–320.
- [5] Gennaro Amendola, *A calculus for ideal triangulations of three-manifolds with embedded arcs*, Math. Nachr. **278** (2005), no. 9, 975–994.
- [6] Markus Banagl and Greg Friedman, *Triangulations of 3-dimensional pseudomanifolds with an application to state-sum integrals*, Algebraic & Geometric Topology **4** (2004), 521–542.
- [7] Stéphane Baseilhac and Riccardo Benedetti, *Analytic families of quantum hyperbolic invariants*, Algebr. Geom. Topol. **15** (2015), no. 4, 1983–2063.
- [8] Riccardo Benedetti and Carlo Petronio, *A finite graphic calculus for 3-manifolds*, Manuscripta Math. **88** (1995), no. 3, 291–310.
- [9] ———, *Branched standard spines of 3-manifolds*, Lecture Notes in Mathematics, vol. 1653, Springer-Verlag, Berlin, 1997.
- [10] Edward M. Brown, *The Hauptvermutung for 3-complexes*, Trans. Amer. Math. Soc. **144** (1969), 173–196.
- [11] Benjamin A. Burton, Ryan Budney, William Pettersson, et al., *Regina: Software for low-dimensional topology*, <http://regina-normal.github.io/>, 1999–2017.
- [12] Benjamin A. Burton and Rodney G. Downey, *Courcelle’s theorem for triangulations*, J. Combin. Theory Ser. A **146** (2017), 264–294.
- [13] Gaspare Carbone, Mauro Carfora, and Annalisa Marzuoli, *Wigner symbols and combinatorial invariants of three-manifolds with boundary*, Comm. Math. Phys. **212** (2000), no. 3, 571–590.
- [14] David Coulson, Oliver A. Goodman, Craig D. Hodgson, and Walter D. Neumann, *Computing arithmetic invariants of 3-manifolds*, Experiment. Math. **9** (2000), no. 1, 127–152.
- [15] Alexander Coward and Marc Lackenby, *An upper bound on Reidemeister moves*, Amer. J. Math. **136** (2014), no. 4, 1023–1066.
- [16] Marc Culler, Nathan Dunfield, and Jeffrey R. Weeks, *SnapPy*, a computer program for studying the geometry and topology of 3-manifolds, <http://snappy.computop.org>.
- [17] Tudor Dimofte and Sergei Gukov, *Quantum field theory and the volume conjecture*, Chern-Simons gauge theory: 20 years after, AMS/IP Stud. Adv. Math., vol. 50, Amer. Math. Soc., Providence, RI, 2011, pp. 19–42.
- [18] Nathan M. Dunfield and Anil N. Hirani, *The least spanning area of a knot and the optimal bounding chain problem*, Computational geometry (SCG’11), ACM, New York, 2011, pp. 135–144.
- [19] D. B. A. Epstein, *Curves on 2-manifolds and isotopies*, Acta Math. **115** (1966), 83–107.

- [20] D. B. A. Epstein and R. C. Penner, *Euclidean decompositions of noncompact hyperbolic manifolds*, J. Differential Geom. **27** (1988), no. 1, 67–80.
- [21] George K. Francis and Jeffrey R. Weeks, *Conway’s ZIP proof*, Amer. Math. Monthly **106** (1999), no. 5, 393–399.
- [22] Roberto Frigerio and Carlo Petronio, *Construction and recognition of hyperbolic 3-manifolds with geodesic boundary*, Trans. Amer. Math. Soc. **356** (2004), no. 8, 3243–3282.
- [23] Wolfgang Haken, *Ein Verfahren zur Aufspaltung einer 3-Mannigfaltigkeit in irreduzible 3-Mannigfaltigkeiten*, Math. Z. **76** (1961), 427–467.
- [24] ———, *Theorie der Normalflächen*, Acta Math. **105** (1961), 245–375.
- [25] ———, *über das Homöomorphieproblem der 3-Mannigfaltigkeiten. I*, Math. Z. **80** (1962), 89–120.
- [26] ———, *Some results on surfaces in 3-manifolds*, Studies in Modern Topology, Math. Assoc. Amer. (distributed by Prentice-Hall, Englewood Cliffs, N.J.), 1968, pp. 39–98.
- [27] A. J. S. Hamilton, *The triangulation of 3-manifolds*, Quart. J. Math. Oxford Ser. (2) **27** (1976), no. 105, 63–70.
- [28] Joel Hass, Jeffrey C. Lagarias, and Nicholas Pippenger, *The computational complexity of knot and link problems*, J. ACM **46** (1999), no. 2, 185–211.
- [29] A. Hatcher, *The Kirby torus trick for surfaces*, ArXiv e-prints (2013), 1312.3518.
- [30] Allen Hatcher, *On triangulations of surfaces*, Topology Appl. **40** (1991), no. 2, 189–194.
- [31] Neil Hoffman, Kazuhiro Ichihara, Masahide Kashiwagi, Hidetoshi Masai, Shin’ichi Oishi, and Akitoshi Takayasu, *Verified computations for hyperbolic 3-manifolds*, Exp. Math. **25** (2016), no. 1, 66–78.
- [32] William Jaco and Ulrich Oertel, *An algorithm to decide if a 3-manifold is a Haken manifold*, Topology **23** (1984), no. 2, 195–209.
- [33] William Jaco, Hyam Rubinstein, and Stephan Tillmann, *Minimal triangulations for an infinite family of lens spaces*, J. Topol. **2** (2009), no. 1, 157–180.
- [34] William Jaco and J. Hyam Rubinstein, *PL minimal surfaces in 3-manifolds*, J. Differential Geom. **27** (1988), no. 3, 493–524.
- [35] ———, *PL equivariant surgery and invariant decompositions of 3-manifolds*, Adv. in Math. **73** (1989), no. 2, 149–191.
- [36] ———, *0-efficient triangulations of 3-manifolds*, J. Differential Geom. **65** (2003), no. 1, 61–168.
- [37] William Jaco and J. Hyam Rubinstein, *Layered-triangulations of 3-manifolds*, ArXiv Mathematics e-prints (2006), math/0603601.
- [38] Yuichi Kabaya, *Pre-Bloch invariants of 3-manifolds with boundary*, Topology Appl. **154** (2007), no. 14, 2656–2671.
- [39] R. M. Kashaev, *R-matrix knot invariants and triangulations*, Interactions between hyperbolic geometry, quantum topology and number theory, Contemp. Math., vol. 541, Amer. Math. Soc., Providence, RI, 2011, pp. 69–81.
- [40] Hellmuth Kneser, *Geschlossene flächen in dreidimensionalen mannigfaltigkeiten*, Jahresber. Dtsch. Math.-Ver. **38** (1929), 248–260.
- [41] Tao Li, *An algorithm to determine the Heegaard genus of a 3-manifold*, Geom. Topol. **15** (2011), no. 2, 1029–1106.
- [42] W B Raymond Lickorish, *Simplicial moves on complexes and manifolds*, Proceedings of the Kirbyfest, Geometry & Topology Monographs, vol. 2, 1999, pp. 299–320.
- [43] Monika Ludwig and Matthias Reitzner, *Elementary moves on triangulations*, Discrete Comput. Geom. **35** (2006), no. 4, 527–536.
- [44] Ciprian Manolescu, *Triangulations of manifolds*, ICCM Not. **2** (2014), no. 2, 21–23.
- [45] ———, *Lectures on the triangulation conjecture*, Proceedings of the Gökova Geometry-Topology Conference 2015, Gökova Geometry/Topology Conference (GGT), Gökova, 2016, pp. 1–38.
- [46] ———, *Pin(2)-equivariant Seiberg-Witten Floer homology and the triangulation conjecture*, J. Amer. Math. Soc. **29** (2016), no. 1, 147–176.
- [47] Sergei V. Matveev, *Transformations of special spines, and the Zeeman conjecture*, Izv. Akad. Nauk SSSR Ser. Mat. **51** (1987), no. 5, 1104–1116, 1119.
- [48] ———, *Algorithmic topology and classification of 3-manifolds*, second ed., Springer, 2007.
- [49] Edwin E. Moise, *Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung*, Ann. of Math. (2) **56** (1952), 96–114.
- [50] Lee Mosher, *Tiling the projective foliation space of a punctured surface*, Trans. Amer. Math. Soc. **306** (1988), no. 1, 1–70.
- [51] Hitoshi Murakami, *Current status of the volume conjecture*, Sugaku Expositions **26** (2013), no. 2, 181–203.

- [52] Walter D. Neumann, *Realizing arithmetic invariants of hyperbolic 3-manifolds*, Interactions between hyperbolic geometry, quantum topology and number theory, Contemp. Math., vol. 541, Amer. Math. Soc., Providence, RI, 2011, pp. 233–246.
- [53] Walter D. Neumann and Alan W. Reid, *Arithmetic of hyperbolic manifolds*, Topology '90 (Columbus, OH, 1990), Ohio State Univ. Math. Res. Inst. Publ., vol. 1, de Gruyter, Berlin, 1992, pp. 273–310.
- [54] Walter D. Neumann and Jun Yang, *Bloch invariants of hyperbolic 3-manifolds*, Duke Math. J. **96** (1999), no. 1, 29–59.
- [55] Walter D. Neumann and Don Zagier, *Volumes of hyperbolic three-manifolds*, Topology **24** (1985), no. 3, 307–332.
- [56] M. H. A. Newman, *On the foundations of combinatorial analysis situs*, Proc. Royal Acad. Amsterdam **29** (1926), 610–641.
- [57] Udo Pachner, *Bstellare äquivalenz kombinatorischer Mannigfaltigkeiten*, Arch. Math. (Basel) **30** (1978), no. 1, 89–98.
- [58] R. C. Penner, *The decorated Teichmüller space of punctured surfaces*, Comm. Math. Phys. **113** (1987), no. 2, 299–339.
- [59] Robert C. Penner, *Decorated Teichmüller theory*, QGM Master Class Series, European Mathematical Society (EMS), Zürich, 2012, With a foreword by Yuri I. Manin.
- [60] Riccardo Piergallini, *Standard moves for standard polyhedra and spines*, Rend. Circ. Mat. Palermo (2) Suppl. (1988), no. 18, 391–414, Third National Conference on Topology (Italian) (Trieste, 1986).
- [61] Joachim H. Rubinstein, *An algorithm to recognize the 3-sphere*, Proceedings of the International Congress of Mathematicians, Vol. 1, 2 (Zürich, 1994), Birkhäuser, Basel, 1995, pp. 601–611.
- [62] Saul Schleimer, *Sphere recognition lies in NP*, Low-dimensional and symplectic topology, Proc. Sympos. Pure Math., vol. 82, Amer. Math. Soc., Providence, RI, 2011, pp. 183–213.
- [63] Horst Schubert, *Bestimmung der Primfaktorzerlegung von Verkettungen*, Math. Z. **76** (1961), 116–148.
- [64] Boris Springborn, *The hyperbolic geometry of Markov's theorem on Diophantine approximation and quadratic forms*, Enseign. Math. **63** (2017), no. 3-4, 333–373.
- [65] Abigail Thompson, *Thin position and the recognition problem for S^3* , Math. Res. Lett. **1** (1994), no. 5, 613–630.
- [66] William P. Thurston, *Three-dimensional geometry and topology. Vol. 1*, Princeton Mathematical Series, vol. 35, Princeton University Press, Princeton, NJ, 1997, Edited by Silvio Levy.
- [67] Stephan Tillmann and Sampson Wong, *An algorithm for the Euclidean cell decomposition of a cusped strictly convex projective surface*, J. Comput. Geom. **7** (2016), no. 1, 237–255.
- [68] V. G. Turaev, *Quantum invariants of knots and 3-manifolds*, De Gruyter Studies in Mathematics, vol. 18, Walter de Gruyter & Co., Berlin, 1994.
- [69] V. G. Turaev and O. Ya. Viro, *State sum invariants of 3-manifolds and quantum 6j-symbols*, Topology **31** (1992), no. 4, 865–902.

SCHOOL OF MATHEMATICS AND STATISTICS, THE UNIVERSITY OF MELBOURNE, VIC 3010, AUSTRALIA

DEPARTMENT OF MATHEMATICS, OKLAHOMA STATE UNIVERSITY, STILLWATER, OKLAHOMA, OK 74078, USA

SCHOOL OF MATHEMATICS AND STATISTICS, THE UNIVERSITY OF SYDNEY, NSW 2006, AUSTRALIA

E-mail address: joachim@unimelb.edu.au

E-mail address: segerman@math.okstate.edu

E-mail address: stephan.tillmann@sydney.edu.au