CONNECTING ESSENTIAL TRIANGULATIONS II: VIA 2-3 MOVES ONLY

TEJAS KALELKAR, SAUL SCHLEIMER, AND HENRY SEGERMAN

Abstract. In previous work we showed that for a manifold M, whose universal cover has infinitely many boundary components, the set of essential ideal triangulations of M is connected via 2-3, 3-2, 0-2, and 2-0 moves. Here we show that this set is also connected via 2-3 and 3-2 moves alone, if we ignore those triangulations for which no 2-3 move results in an essential triangulation. If we also allow V-moves and their inverses then the full set of essential ideal triangulations of M is once again connected. These results also hold if we replace essential triangulations with L–essential triangulations.

1. Introduction

1.1. **Paths of triangulations.** The various notions of equivalence of triangulations have been refined many times, notably by Pachner [6], who showed that bistellar moves connect any two triangulations of a given manifold. For a three-dimensional manifold M, these are the 2-3, 3-2, 1-4, and 4-1 moves. If we do not want to change the number of vertices, the latter two moves are almost always unnecessary, by the work of Matveev [5, Theorem 1.2.5], Piergallini [7, Theorem 1.2], and Amendola [1, Theorem 2.1]. To state their results, let $\mathbb{T}(M)$ be those one-vertex triangulations of M (if M is closed) *or* ideal triangulations (if M has boundary). Let $\mathbb{T}_2(M)$ be those triangulations of $\mathbb{T}(M)$ having at least two tetrahedra.

Theorem 1.2. *Suppose that M is a three-manifold. Then $\mathbb{T}_2(M)$ is connected via 2-3 and 3-2 moves.* □

An *essential* ideal triangulation has the property that none of its edges are homotopic, relative to their endpoints, into ∂M. (In this paper we will only consider ideal triangulations of three-manifolds with boundary.) Essentiality is required in many applications. For example, an ideal triangulation of a cusped hyperbolic three-manifold is essential if and only if it admits a solution to Thurston's gluing equations.

Date: July 23, 2024.

In the first paper in this series we prove a connectivity result similar to Theorem 1.2 but where the initial and terminal triangulations, as well as all intermediate triangulations, are essential. Also in that paper we generalise to *L–essential* triangulations. These involve a choice of *labelling* L of the boundary components Δ_M of the universal cover of the manifold M. See Definition 2.9 for the details. Let $\mathbb{T}(M, L)$ be those (necessarily ideal) triangulations in $\mathbb{T}(M)$ which are L–essential. The full result [3, Theorem 6.1] is as follows.

Theorem 1.3. *Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M with infinite image. Then $\mathbb{T}(M, L)$ is connected via 2-3, 3-2, 0-2, and 2-0 moves.* □

This result makes use of 0-2 and 2-0 moves to connect triangulations together. These are again local moves on the triangulation but are not bistellar moves. See Section 2.12.1. Let $\mathbb{T}^\circ(M, L)$ be the triangulations in $\mathbb{T}(M, L)$ which admit some 2-3 move preserving L–essentiality. The main goal of this paper is to prove the following.

Theorem 3.5. *Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M with infinite image. Then $\mathbb{T}^\circ(M, L)$ is connected via 2-3 and 3-2 moves.*

Thus we improve on Theorem 1.3 by removing the need for 0-2 and 2-0 moves, at the cost of being unable to connect to certain isolated L–essential triangulations. These isolated triangulations include those excluded in Theorem 1.2: that is, triangulations with only one tetrahedron. However, more subtle isolations occur in the context of L–essential triangulations. See Example 2.18.

On the other hand, each L–essential triangulation can be connected to $\mathbb{T}^\circ(M, L)$ by a single *V-move*, a more "local" kind of 0-2 move. See Definition 3.10 and Lemma 3.14. Combining Theorem 3.5 and Lemma 3.14, we obtain the following improvement of Theorem 1.3.

Corollary 1.4. *Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M with infinite image. Then $\mathbb{T}(M, L)$ is connected via 2-3 moves, V-moves, and their inverse moves.* □

The second and third authors will use Corollary 1.4 in their proof that the figure-eight knot complement, and certain other manifolds, have unique veering triangulations. In the proof, we carry the data of a circular order through a sequence of 2-3 moves, V-moves, and their inverses. Surprisingly, in this setting the 2-0 move is problematic. In particular, Theorem 1.3 does not suffice for our future work.

1.5. **Outline.** Given Theorem 1.3, it suffices to be able to implement a 0-2 move between L–essential triangulations via a sequence of 2-3 and 3-2 moves where all intermediate triangulations are also L–essential. In the absence of a requirement of L–essentiality, this can be done purely locally. See Lemma 1.2.11 and Proposition 1.2.8 of [5].

In Section 4, we lay out the hypotheses and tools needed to make this local construction go through with the additional requirement of L–essentiality. Our versions of the local construction are set out in Lemmas 4.13 and 4.16. However, as illustrated in Section 5.1, it may be that near the site of the 0-2 move there are no 3-2 moves, and any 2-3 move destroys L–essentiality. That is, any nearby 2-3 move introduces an edge between two vertices in the universal cover with the same label. Thus we must bring some "distant" vertex, with a different label, "close" to the site of the 0-2 move.

In Section 5 we give a collection of moves (the *augmented 2-3 move* and the *nature reserve moves*) that preserve L–essentiality and serve as tools to "transport" the distant vertex. In Section 6, we assemble these moves into two parallel sequences, appropriately commuting with the 0-2 move. These sequences move the distant vertex into contact with the site of the 0-2 move, at which point we apply Lemma 4.16.

Acknowledgements. The third author was supported in part by National Science Foundation grant DMS-2203993.

2. Background

Here we essentially follow the notation and definitions of the first paper in this series [3, Section 2].

2.1. **Triangulations and foams.** Our main theorem, Theorem 3.5, is stated in terms of triangulations. However, their dual *foams* are often easier to understand. (One possible explanation for this is that the complexity concentrated at the vertices of a triangulation is spread across the boundaries of the three-cells in the dual foam. For a simple example, see Figure 2.17.) Both triangulations and foams are complexes formed by taking a disjoint union of cells and gluing them together. We refer to the cells, before gluing, as *model* cells.

Definition 2.2. A *triangulation* $\mathcal{T}$ is a collection of model tetrahedra $\{t_k\}$ together with a collection of *face pairings* $\{\phi_{ij}\}$. Here ϕ_{ij} is an isomorphism from some model face f_i of some model tetrahedron $t_{k(i)}$ to some model face f_j of some model tetrahedron $t_{k(j)}$. The *realisation* of $\mathcal{T}$ is the topological space $|\mathcal{T}|$ obtained by taking the disjoint union of the t_i and forming the quotient by the ϕ_{ij}. The zero-skeleton of $|\mathcal{T}|$ is the image of the model vertices.

Suppose that M is a compact, connected three-manifold with boundary. A triangulation $\mathcal{T}$ is an *ideal triangulation* of M if $M - \partial M$ is homeomorphic to $|\mathcal{T}|$ minus its zero-skeleton. ◇

Definition 2.3. We denote the universal covering map by $\phi_M \colon \widetilde{M} \to M$. We use Δ_M to denote the set of boundary components of $\widetilde{M}$. ◇

Definition 2.4. Suppose that M is a compact, connected three-manifold, with boundary. Suppose that $\mathcal{T}$ is an ideal triangulation of M. We call $\mathcal{F}$, the dual two-complex to $\mathcal{T}$, a *foam* in M. We refer to the components of $M - \mathcal{F}$ as *(complementary) regions.* ◇

See Figures 2.5A, 2.5B, and 2.5C for small neighbourhoods of points of $\mathcal{F}$ in M.

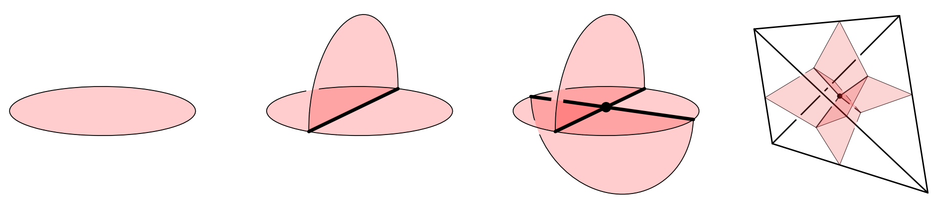

(A) (B) (C) (D)

FIGURE 2.5. Local pictures of foams.

Definition 2.6. An edge e of a foam is an *edge loop* if e has both ends at a single vertex. Any lift of e to a cover (including the trivial cover) is a *cyclic edge.* ◇

For examples of edge loops and cyclic edges see Figure 2.17.

2.7. Labellings and L–essentiality.

Definition 2.8. Suppose that $\mathcal{L}$ is a set of *labels* equipped with an action of $\pi_1(M)$. Suppose that $L \colon \Delta_M \to \mathcal{L}$ is a $\pi_1(M)$–equivariant function. Then we call L a *labelling* of Δ_M. ◇

Definition 2.9. Suppose that $\mathcal{T}$ is an ideal triangulation of M. Suppose that L is a labelling of Δ_M, as in Definition 2.8. Suppose that e is an edge of $\mathcal{T}$ with a lift $\widetilde{e}$ in $\widetilde{\mathcal{T}}$. Suppose that $\widetilde{u}$ and $\widetilde{v}$ are the endpoints of $\widetilde{e}$. If $L(\widetilde{u}) \neq L(\widetilde{v})$ then we say that e is *L–essential.* If all edges of $\mathcal{T}$ are L–essential then we say that $\mathcal{T}$ is *L–essential.* ◇

See [3, Section 2.9] for examples of labellings. The simplest labelling is the identity map on Δ_M. As noted in [3, Remark 2.18], with this labelling L–essential triangulations are essential triangulations in the sense of [2, Definitions 3.2 and 3.5] and [4, page 336].

Dually, our notions of L–essentiality apply to foams as follows.

Definition 2.10. Suppose that $\mathcal{F}$ is a foam in M. We extend the labelling function L to components of $\widetilde{M} - \widetilde{\mathcal{F}}$ as follows. Suppose that C is a component of $\widetilde{M} - \widetilde{\mathcal{F}}$ with boundary component $c \in \Delta_M$. Then we set $L(C) = L(c)$.

Now suppose that f is a face of a foam $\mathcal{F}$ with a lift $\widetilde{f}$ in $\widetilde{\mathcal{F}}$, with components U and V of $\widetilde{M} - \widetilde{\mathcal{F}}$ incident to $\widetilde{f}$. We say that f is *L–essential* if $L(U) \neq L(V)$. If all faces of $\mathcal{F}$ are L–essential then we say that $\mathcal{F}$ is *L–essential*. ◇

2.11. **Moves on foams.** The three-dimensional bistellar moves [6] are the 1-4, 2-3, 3-2, and 4-1 moves. These can be performed equally well on triangulations or their dual foams. Of these we only consider the 2-3 and 3-2 moves, applied to foams. The former is called the *T move* by Matveev [5, page 14]. It can be performed along any edge of $\mathcal{F}$ that is not an edge loop. See Figure 2.12. The 3-2 move can be performed on any triangular face whose closure is embedded in $\mathcal{F}$.

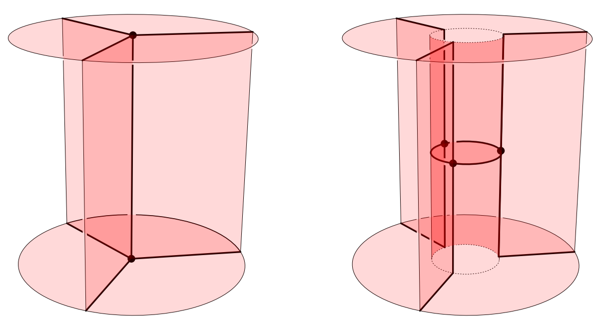

(A) (B)

FIGURE 2.12. The 2-3 move.

2.12.1. *The 0-2 move.* The *0-2 move*, defined by Figure 2.13, is called the *ambient lune* move by Matveev [5, page 17]. It is applied along an arc δ properly embedded in, and avoiding the vertices of, a face of $\mathcal{F}$. The 0-2 move creates two new vertices and a new bigon face. We denote the result by $\mathcal{F}[\delta]$.

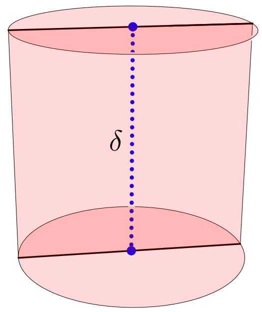

(A)

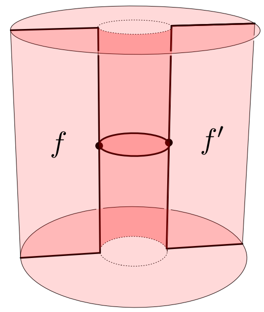

(B)

FIGURE 2.13. The 0-2 move. The vertical dotted arc δ in Figure 2.13A indicates the arc along which the 0-2 move acts.

Definition 2.14. Suppose that e is an edge of a foam $\mathcal{F}$. Let $\mathcal{N}(e)$ be a small regular neighbourhood of e. Suppose that $\mathcal{F}'$ is the result of applying a 2-3 move to $\mathcal{F}$ along e. We assume that the move is supported in $\mathcal{N}(e)$. Suppose that c is an open cell (or open complementary region) of $\mathcal{F}$. Suppose that c' is similarly obtained from $\mathcal{F}'$. If $c \cap c' - \mathcal{N}(e)$ is non-empty then we say that c the *ancestor* of c' and c' is the *descendant* of c.

We make similar definitions for the 3-2, 0-2, 2-0, and various other moves we define later in the paper. Finally we make the relation transitive through multiple moves. ◇

Remark 2.15. In a small abuse of notation we often use the same name for an ancestor and its descendants. ◇

2.16. **Edge loops cause issues.** Here we show that one cannot remove 0-2 and 2-0 moves from the statement of Theorem 1.3 without adding some other hypothesis.

Example 2.18. Figure 2.17A shows a foam in a solid torus. Mirroring this foam across the boundary torus produces a foam in $S^2 \times S^1$ with two vertices, four edges, three faces, and one complementary region. Let L be the identity labelling (as described in Section 2.7). All faces are essential (in the universal cover the regions to either side of each face are distinct). There are no triangular faces so no 3-2 move is possible. Two of the edges are edge loops so a 2-3 move cannot be applied along them. The other two edges bound a bigon face, which implies that a 2-3 move applied along them introduces an inessential face.

Applying various 0-2 moves produces other foams with all faces essential. Depending on which 0-2 move we apply, the resulting foam may or may not have a 2-3 move keeping the foam essential. For example,

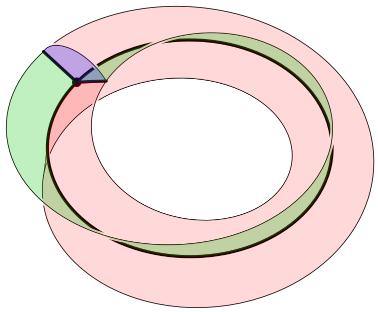

(A) A foam in $S^1 \times D^2$.

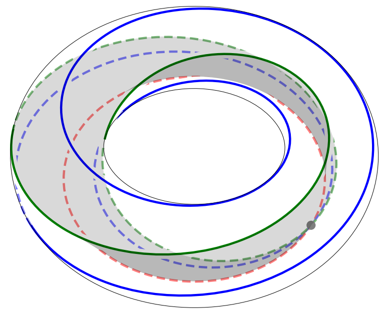

(B) The dual triangulation. The internal triangle is shaded.

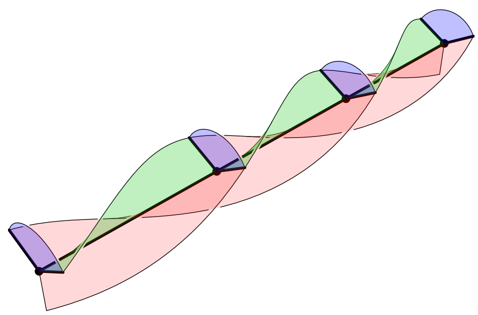

(C) Part of the universal cover for the foam in Figure 2.17A.

FIGURE 2.17. We obtain a foam in $S^1 \times S^2$ by doubling the foam in Figure 2.17A across the boundary of the solid torus. There is an edge loop in Figure 2.17A. Its lifts in Figure 2.17C are cyclic edges.

suppose that we perform a 0-2 move across the bigon face. Then in the resulting foam no 2-3 or 3-2 move is possible while maintaining essentiality.

Therefore there are at least some essential triangulations that are isolated if we can only use 2-3 and 3-2 moves. ◇

This is a similar situation to Matveev's, Piergallini's, and Amendola's original result (Theorem 1.2) in which triangulations with a single tetrahedron are isolated.

3. Implementing 0-2 moves

Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M.

Definition 3.1. Suppose that $\mathcal{T}$ is an L–essential ideal triangulation of M. Suppose that f is a face of $\mathcal{T}$ and suppose that performing a

2-3 move across f produces $\mathcal{T}'$, which is also L–essential. Then we say that f is L*–flippable*. We make the analogous definition for edges of foams. ◇

Remark 3.2. Suppose that $\mathcal{T}$ is an L–essential ideal triangulation. Suppose that $\mathcal{T}$ has a 3-2 move along the edge e that preserves L–essentiality. Then, for f any face adjacent to e, the 2-3 move across f also preserves L–essentiality. ◇

Definition 3.3. Let $\mathbb{T}^\circ(M, L)$ be the set of L–essential triangulations of M that have at least one L–flippable face. ◇

Remark 3.4. If L has infinite image then $\mathbb{T}(M, L)$ is non-empty by [3, Theorem 3.1]. Lemma 3.14 then implies that $\mathbb{T}^\circ(M, L)$ is non-empty. ◇

Our main result is the following.

Theorem 3.5. *Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M with infinite image. Then the set $\mathbb{T}^\circ(M, L)$ is connected via 2-3 and 3-2 moves.*

To prove this we require several tools.

3.6. **Finding L-flippable edges.**

Lemma 3.7. *Suppose that the labelling L has infinite image. Then there is an edge e of $\widetilde{\mathcal{F}}$ that is incident to complementary regions with five distinct labels.*

Proof. We prove the contrapositive. Suppose that e is an edge of $\widetilde{\mathcal{F}}$ with endpoints v and w. We assume that e is incident to at most four distinct labels. Let A, B, and C be the labels of the complementary regions meeting the interior of e. Let D and E be the labels of the complementary regions meeting v and w but not meeting the interior of e. The labels A, B, and C are distinct because $\mathcal{F}$ is L–essential. Similarly, the labels D and E are each distinct from A, B, and C. Thus we must have that $D = E$.

The one-skeleton of $\widetilde{\mathcal{F}}$ is connected, so propagating the above argument we find that every vertex of $\widetilde{\mathcal{F}}$ is incident to regions with labels A, B, C, and D. Thus $|L(\Delta_M)| = 4$. □

Lemma 3.8. *Suppose that the labelling L has infinite image. Suppose that $\mathcal{F}$ is an L–essential foam in M. Suppose that $\mathcal{F}$ has no cyclic edges. Then $\mathcal{F}$ contains an L–flippable edge.*

Proof. By Lemma 3.7 there is an edge e of $\widetilde{\mathcal{F}}$ that is incident to complementary regions with five distinct labels. By hypothesis, e is not cyclic. Therefore $\phi_M(e)$ is L–flippable. □

3.9. **V-moves.** We use some of the same tools as Matveev, beginning with the V-move [5, Definition 1.2.6].

Definition 3.10. Suppose that $\mathcal{F}$ is a foam. Suppose that v is a vertex of $\mathcal{F}$. Let $\mathcal{N}$ be a small regular neighbourhood of v. Suppose that δ_+ and δ_- are a pair of opposite edges of $\partial(\mathcal{F} \cap \mathcal{N})$. Then the *V-move* on v (in the direction of δ_+ and δ_-) is the 0-2 move along δ_+. (Equivalently, it is also the 0-2 move along δ_-). ◇

The V-move is shown for both a triangulation and the dual foam in Figure 3.11. Note that there is a reflection symmetry in the resulting triangulation and thus (combinatorially) in the foam.

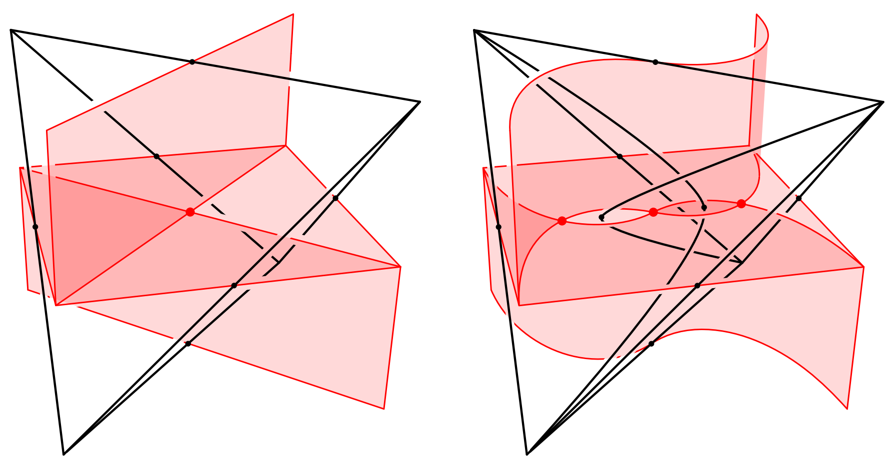

(A) (B)

FIGURE 3.11. The V-move, applied in the right/left direction. Small black dots indicate intersections between the edges of the triangulation and their dual faces in the foam.

Lemma 3.12. *Suppose that $\mathcal{F}$ is a foam in M. Suppose that applying a V-move to $\mathcal{F}$ produces $\mathcal{F}'$. Then we have the following.*

- *$\mathcal{F}$ is L–essential if and only if $\mathcal{F}'$ is L–essential.*
- *If $\mathcal{F}$ has an L–flippable edge then so does $\mathcal{F}'$.*
- *If e is a cyclic edge of $\mathcal{F}'$ then it has a cyclic ancestor in $\mathcal{F}$.*

Proof. The first conclusion follows since no new pair of regions comes into contact as a result of the V-move or its inverse. For the second conclusion, note that an edge of $\mathcal{F}$ and its descendant in $\mathcal{F}'$ (Definition 2.14) are incident to the same complementary regions. For the third conclusion, note that there are four edges in $\mathcal{F}'$ without ancestors;

none of these are cyclic. If some other edge in $\mathcal{F}'$ is cyclic then its ancestor in $\mathcal{F}$ is also cyclic. □

3.13. **Avoiding cyclic edges.** As usual, we assume that M is a compact, connected three-manifold with boundary.

Lemma 3.14. *Suppose that $L\colon \Delta_M \to \mathcal{L}$ has infinite image. Suppose that $\mathcal{F}$ is an L–essential foam in M. Suppose that $\mathcal{F}$ contains no L–flippable edges. Then there is a V-move on $\mathcal{F}$ that results in an L–essential foam $\mathcal{F}'$ which contains an L–flippable edge.*

Proof. By Lemma 3.7, there is an edge e of $\widetilde{\mathcal{F}}$ that is incident to complementary regions with five distinct labels. Since $\mathcal{F}$ contains no L–flippable edges, $\phi_M(e)$ must be cyclic. Let $\mathcal{N}$ be a small regular neighbourhood of v, the unique endpoint of $\phi_M(e)$. Let δ_+ be as given in Definition 3.10. We obtain $\mathcal{F}'$ by performing a 0-2 move along δ_+. By Lemma 3.12, the foam $\mathcal{F}'$ is L–essential. Also, the descendant of $\phi_M(e)$ not contained in $\mathcal{N}$ is L–flippable. □

Also using notation as in Definition 3.10, we have the following.

Lemma 3.15. *Suppose that e is a cyclic edge of $\mathcal{F}$ with both endpoints at v. Suppose that exactly one endpoint of δ_+ lies on e. Then the V-move at v obtained by applying a 0-2 move along δ_+ gives a foam $\mathcal{F}'$ with fewer cyclic edges than $\mathcal{F}$.*

Proof. The cyclic edge e is destroyed by the vertices added by the 0-2 move. By Lemma 3.12, no new cyclic edges are created. □

Lemma 3.16. *Suppose that $\mathcal{F}$ is an L–essential foam in M. Suppose that $\mathcal{F}$ contains an L–flippable edge. Then there is a sequence of L–essential foams $\mathcal{F} = \mathcal{F}_0, \ldots, \mathcal{F}_n$ where*

- *each foam is related to the next by a 0-2 move,*
- *each foam contains an L–flippable edge, and*
- *$\mathcal{F}_n$ has no cyclic edges.*

Proof. We repeatedly apply Lemma 3.15. By Lemma 3.12, each resulting foam is L–essential and has an L–flippable edge. □

The following is a refinement of Theorem 1.3.

Lemma 3.17. *Suppose that the labelling L has infinite image. Suppose that $\mathcal{F}$ and $\mathcal{F}'$ are L–essential foams in M, each with no cyclic edges. Then there is a path $\mathcal{F} = \mathcal{F}_0, \ldots, \mathcal{F}_n = \mathcal{F}'$ of L–essential foams without cyclic edges where each foam is related to the next by a 2-3, 3-2, 0-2, or 2-0 move.*

Proof. Let $\mathcal{F} = \mathcal{G}_0, \dots, \mathcal{G}_n = \mathcal{F}'$ be the path of L–essential foams (dual to triangulations) given to us by Theorem 1.3. To be precise in our counting, we say that a move destroys a cyclic edge when the move alters all neighbourhoods of the edge. A move creates a cyclic edge when the reverse move destroys it. Note that a single move can destroy one cyclic edge and create another. Some finite number of cyclic edges are each created and then destroyed in the sequence, since $\mathcal{F}$ and $\mathcal{F}'$ have no cyclic edges. We modify this sequence recursively, so that after each modification one fewer cyclic edge is created. After a finite number of these modifications, we have the desired sequence.

Suppose that a cyclic edge e, with both ends at a vertex v, is created by the move m transforming $\mathcal{G}_p$ into $\mathcal{G}_{p+1}$. There are four cases to consider, as the move m is a 2-3, 3-2, 0-2, or 2-0 move. In each case, we apply a 0-2 move along an arc δ before the move m to avoid making the cyclic edge. In some cases, we then apply a second 0-2 move along an arc δ' followed by a 2-0 move to undo the first 0-2 move. The foam resulting from this process also results from performing m and then destroying the cyclic edge by performing a V-move at v. (We choose δ and δ' so that all intermediate foams remain L–essential.) As the V-move does not alter the foam outside of a small neighbourhood of v, it does not affect later moves in our sequence, until we reach a move that destroys the cyclic edge e. Suppose that this occurs between $\mathcal{G}_{q-1}$ and $\mathcal{G}_q$. Viewing the sequence in reverse, between $\mathcal{G}_q$ and $\mathcal{G}_{q-1}$ we create the cyclic edge. We can therefore get from $\mathcal{G}_{q-1}$ to $\mathcal{G}_q$ using the reverse of one of the same four constructions we use for the forward direction.

Suppose that m is a 2-3 move. See Figure 3.18A. Breaking symmetry, in $\mathcal{G}_{p+1}$ the two edge-ends labelled e are connected together at the vertex v after the 2-3 move, as shown in Figure 3.18B. Before m we apply a 0-2 move along the arc marked δ in Figure 3.18A. This move does not itself introduce another cyclic edge. Next we apply m. Note that m no longer creates a cyclic edge because of the two extra vertices formed by the 0-2 move along δ. Next, we apply a 0-2 move along the arc marked δ' in Figure 3.18B. Last, we undo the 0-2 move that occurred along δ.

Now suppose that m is a 3-2 move. See Figure 3.18C. Breaking symmetry, in $\mathcal{G}_{p+1}$ the two edge-ends labelled e are connected together at the vertex v after the 3-2 move, as shown in Figure 3.18D. Before m we apply a 0-2 move along the arc marked δ in Figure 3.18C. Again, this move does not itself introduce another cyclic edge. We then apply m. In this case we do not need to do a second 0-2 move because we already have the result of a V-move at v.

Now suppose that m is a 0-2 move. See Figure 3.18E. Breaking symmetry, in $\mathcal{G}_{p+1}$ the two edge-ends labelled e are connected together

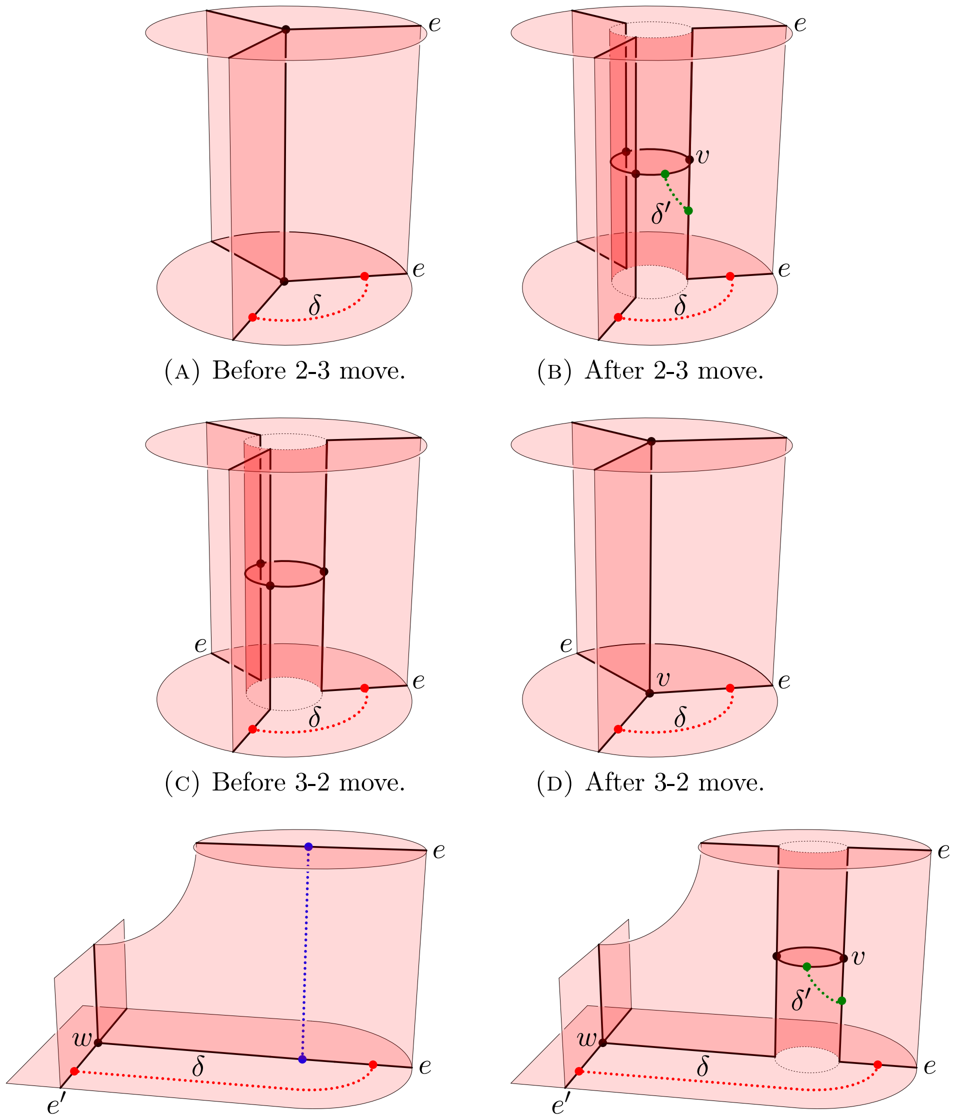

(A) Before 2-3 move. (B) After 2-3 move.

(C) Before 3-2 move. (D) After 3-2 move.

(E) Before 0-2 move. (F) After 0-2 move.

FIGURE 3.18. Steps to avoid creating a cyclic edge.

at the vertex v after the 0-2 move, as shown in Figure 3.18F. Following the edge in the opposite direction, we find a vertex w. Breaking symmetry again, we find an edge e' incident to w which is not e. Before m we apply a 0-2 move along the arc marked δ in Figure 3.18E. We then apply m. Next, we apply a 0-2 move along the arc marked δ' in Figure 3.18F. Last, we undo the 0-2 move that occurred along δ.

Finally, suppose that m is a 2-0 move. (There are multiple ways in which such a 2-0 move can create a cyclic edge depending on how the four edge ends exiting Figure 2.13B are connected to each other.) In this case, the vertex v exists before m is applied. Before applying m, we apply a 0-2 move to implement a V-move at v, chosen as in Lemma 3.15. Again we do not need to do a second 0-2 move because we already have the result of a V-move at v. □

The following is the final tool we use to prove Theorem 3.5.

Proposition 3.19. *Suppose that the labelling L has infinite image. Suppose that $\mathcal{F}$ is an L–essential foam in M. Suppose that applying a 0-2 move along an arc δ in a face of $\mathcal{F}$ produces $\mathcal{F}[\delta]$, which is also L–essential. Suppose that each of $\mathcal{F}$ and $\mathcal{F}[\delta]$ has an L-flippable edge. Then there is a path $\mathcal{F} = \mathcal{F}_0, \ldots, \mathcal{F}_n = \mathcal{F}[\delta]$ of L–essential foams where each foam is related to the next by a 2-3 or 3-2 move.*

Proof of Theorem 3.5. Let $\mathcal{F}$ and $\mathcal{F}'$ be foams dual to the triangulations $\mathcal{T}$ and $\mathcal{T}'$ given in the statement of Theorem 3.5. We apply Lemma 3.16 to $\mathcal{F}$ and to $\mathcal{F}'$ to produce sequences of L–essential foams ending at $\mathcal{G}$ and $\mathcal{G}'$ say (respectively). Thus $\mathcal{G}$ and $\mathcal{G}'$ have no cyclic edges. Lemma 3.16 also tells us that every foam in these sequences has an L–flippable edge.

We apply Lemma 3.17 to produce a sequence of L–essential foams without cyclic edges connecting $\mathcal{G}$ to $\mathcal{G}'$. By Lemma 3.8, every foam in the sequence connecting $\mathcal{G}$ to $\mathcal{G}'$ has an L–flippable edge.

Concatenating the three sequences together, we obtain a sequence of L–essential foams connecting $\mathcal{F}$ to $\mathcal{F}'$. Consecutive foams are related by a 2-3, 3-2, 0-2, or 2-0 move, and each foam contains an L–flippable edge. We then use Proposition 3.19 to replace each 0-2 or 2-0 move in the sequence with a sequence of 2-3 and 3-2 moves. □

The proof of Proposition 3.19 is quite difficult and takes up the remainder of the paper.

4. Following Matveev

Without the L–essentiality condition, Proposition 3.19 is proved by Matveev (combining Lemma 1.2.11 and Proposition 1.2.8 of [5]).

We will give a series of constructions that prove Proposition 3.19 under increasingly general circumstances. Following [5, Proposition 1.2.8], we begin with the construction of a V-move in Lemma 4.1, where the arc δ that the 0-2 move is to be applied along cuts off a single vertex v of its model face.

Lemma 4.1 requires that some edge incident to the vertex v is L–flippable. In the completely general case there may be no L–flippable edge anywhere near δ, so the general construction must start work at such an edge and work towards δ. Thus we require a "pre-processing" stage (see Sections 6.1 and 6.9) where we generate good circumstances around δ that allows us to perform the 0-2 move. Lemma 4.16 gives an implementation of a 0-2 move under these good circumstances. That lemma relies on Lemma 4.13, which implements a 0-2 move under even more stringent circumstances.

Lemma 4.1. *Suppose that applying a V-move to a vertex v of an L–essential foam $\mathcal{F}$ produces $\mathcal{F}'$. Suppose that some edge e incident to v is L–flippable. Then there is a sequence of 2-3 and 3-2 moves from $\mathcal{F}$ to $\mathcal{F}'$ such that all foams are L–essential.*

Proof. Figure 4.2 shows how to implement a V-move using three 2-3 moves followed by a 3-2 move. (Note that if we rotate our picture of the foam around e before applying these moves we also obtain the other two V-moves on v.) By hypothesis, the first move along e is possible and does not introduce an L–inessential face. Each of the subsequent moves takes place on a collection of distinct vertices because those vertices were produced by the previous moves. Thus each move can be applied. Moreover, at each step, faces without ancestors are only created between regions that already have a face in common. Thus the moves do not produce an L–inessential face. □

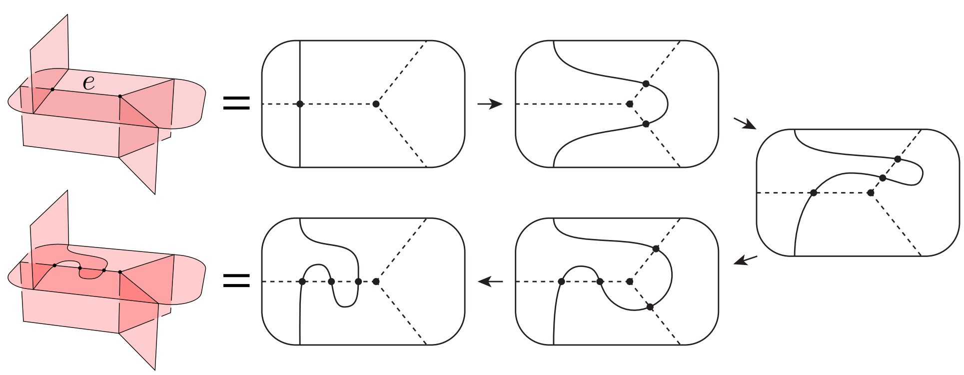

FIGURE 4.2. A V-move can be implemented by three 2-3 moves followed by a 3-2 move. Adapted from [5, Figure 1.15].

It will be convenient for drawing pictures and describing the combinatorics to talk about 0-2 moves in the following way.

Definition 4.3. Suppose that $\mathcal{F}$ is a foam in a manifold M. Suppose that f is a co-oriented model face of $\mathcal{F}$. Suppose that δ is an oriented arc properly embedded in f which is disjoint from the vertices of f. (It will be convenient to conflate δ with its image in M.) Let $\mathcal{N}$ be a regular neighbourhood of δ in M. Let S be the component of $\mathcal{N} - \mathcal{F}$ meeting the interior of δ and pointed at by the co-orientation of $\mathcal{F}$. We say that S is the *snakelet* generated by δ.

After taking the closure, the boundary of S consists of two bigons b and b' and two rectangles r and r'. Suppose that b is the bigon pointed at by δ and suppose that r is the rectangle contained in f. We call b and b' the *head* and *tail* of S respectively. We call r and r' the *belly* and *back* of S respectively. See Figure 4.4. ◇

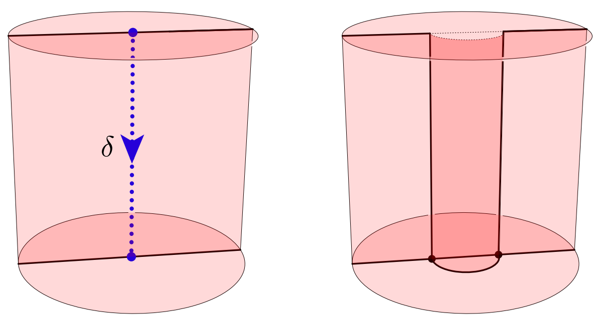

(A) (B)

FIGURE 4.4. A 0-2 move generated by a snakelet along the arc δ. Since δ points down, the head of the snakelet is at the bottom of Figure 4.4B while the tail is at the top. Compare with Figure 2.13.

Note that the foam $\mathcal{F}' = (\mathcal{F} \cup r') - \text{interior}(b')$ is combinatorially identical to the foam obtained by performing a 0-2 move along δ.

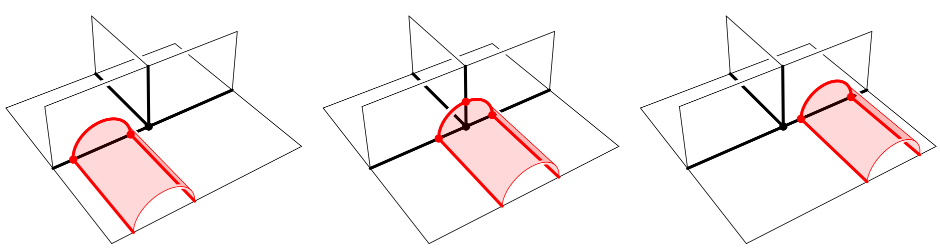

(A) (B) (C)

FIGURE 4.5. Moving the head of a snakelet past a vertex.

We now discuss how to move the head of a snakelet across a vertex.

Lemma 4.6. *Suppose that δ and δ' are arcs in f that have the same start point but differ by an isotopy moving the terminal point from one edge of f to the next. Suppose that $\mathcal{F}[\delta]$ and $\mathcal{F}[\delta']$ are L–essential. Then we can connect $\mathcal{F}[\delta]$ to $\mathcal{F}[\delta']$ by a 2-3 move followed by a 3-2 move without introducing an L–inessential face.*

Proof. See Figure 4.5. Since the arcs move by isotopy, their terminal points do not cross their initial points. Note that the only new regions that come into contact are also in contact in $\mathcal{F}[\delta']$. Therefore all faces are L–essential throughout. Also note that the edge along which we apply the 2-3 move cannot be cyclic since its endpoints are distinct: one endpoint is a vertex of f and the other is part of the snakelet. □

4.7. **Moving snakelets to realise some 0-2 moves.** Throughout this subsection we make the following assumptions.

Hypotheses 4.8. Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M. Suppose that $\mathcal{F}$ is an L–essential foam in M. Suppose that f is a model face of $\widetilde{\mathcal{F}}$. Suppose that δ is an arc properly embedded in f and disjoint from the model vertices of f. Suppose that $\mathcal{F}$ and $\mathcal{F}[\phi_M(\delta)]$ are L–essential. ◇

To simplify our notation, for the remainder of the paper we will conflate δ with $\phi_M(\delta)$ and f with $\phi_M(f)$.

Definition 4.9. The two components of $f - \delta$ are the *sides* of δ. ◇

To aid our exposition, we choose a co-orientation on f and an orientation for δ. This allows us to realise the 0-2 move along δ as generating a snakelet S, as in Definition 4.3.

Definition 4.10. Suppose that s is a side of δ. Recall that r is the belly of the snakelet: a small regular neighbourhood, in f, of δ. Choose a point x in $s - r$. We denote by f_s the descendant (under the 0-2 move) of f in $\mathcal{F}[\delta]$ containing x. ◇

Remark 4.11. If f has no self-gluings then f_s is a subset of f. However, when a model edge of $s \cap \partial f$ other than the first or last meets δ then f_s is more complicated. To see this, consider Figure 4.4A. In that figure, the dotted line indicating δ is on the face f. However, suppose that one of the other two disks at the top of the figure lies in (a translate of) s. In this case f_s extends along either the belly or the back of the snakelet to meet the head of the snakelet. If instead one of the two disks at the bottom of the figure lies in (a translate of) s then f_s meets the head of the snakelet directly. ◇

Suppose that s is one of the sides of δ. We orient $s \cap \partial f$ from the initial point to the terminal point of δ. We name the model edges of f, that meet $s \cap \partial f$ as $e_0, \ldots e_N$, with index increasing in the direction of the orientation. Note that s meets at least two model edges; if it did not then the foam $\mathcal{F}[\delta]$ would have an L–inessential face. We also name the model vertex of f where e_i meets e_{i+1} as v_i.

Definition 4.12. Suppose that e is a model edge of f. Let x be a point in the interior of e; let $\eta(x)$ be a small regular neighbourhood of x, taken in $\widetilde{M}$. Note that $\eta(x)$ is a three-ball. Let y be a point in the intersection of $\eta(x)$ with an open collar of the model edge e taken in f. Let η' be the component of $\eta(x) - \widetilde{\mathcal{F}}$ whose closure does not contain y. The *outer region* for e is the component E of $\widetilde{M} - \widetilde{\mathcal{F}}$ that contains η'. We also say that E is an *outer region for* f. ◇

Let E_i be the outer region for e_i. Recall that $\phi\colon \widetilde{M} \to M$ is the universal covering map.

Lemma 4.13. *Assuming Hypothesis 4.8, suppose that s is a side of the arc δ. Suppose that s meets only three model edges, e_0, e_1, and e_2. Let $E = E_1$. Suppose that $L(E)$ appears precisely once as a label of an outer region for an edge of s, and precisely once as a label for an outer region for an edge of f_s. Suppose that e_1 is not cyclic. Suppose that $\phi_M(e_1)$ is distinct from $\phi_M(e_0)$ and $\phi_M(e_2)$. Then there is a path $\mathcal{F} = \mathcal{F}_0, \ldots, \mathcal{F}_n = \mathcal{F}[\delta]$ of L–essential foams of M where each foam is related to the next by a 2-3 or 3-2 move.*

Proof. Let $A = E_0$ and $B = E_2$. See Figure 4.14A. Let N and S be the regions of $\widetilde{M} - \widetilde{\mathcal{F}}$ incident to the interior of f. The labels on A, B, E, N, and S are all distinct because $\mathcal{F}[\delta]$ is L–essential. Because $L(A) \neq L(B)$ and e_1 is not cyclic, we have that e_1 is L–flippable. By Lemma 4.1 we may perform a V-move at the vertex v_0 where e_0 meets e_1. See Figure 4.14B. This produces the *red snakelet*. Its head is on e_1 and its tail is on e_0. (Here we think of e_1 and e_0 as subsets of M, since after the V-move, they are no longer edges of the foam.) Since $\phi_M(e_1)$ is distinct from $\phi_M(e_0)$, we may isotope the head of the red snakelet to the other end of e_1. Applying Lemma 4.6, we move the head of the red snakelet onto e_2. At this stage we are done unless we are in the case that $\phi_M(e_0) = \phi_M(e_2)$, and the head and the tail of the red snakelet are in the wrong order along the edge e_2. (Matveev swaps the ends of a snakelet in [5, Figure 1.19]; maintaining L–essentiality requires additional work as follows.)

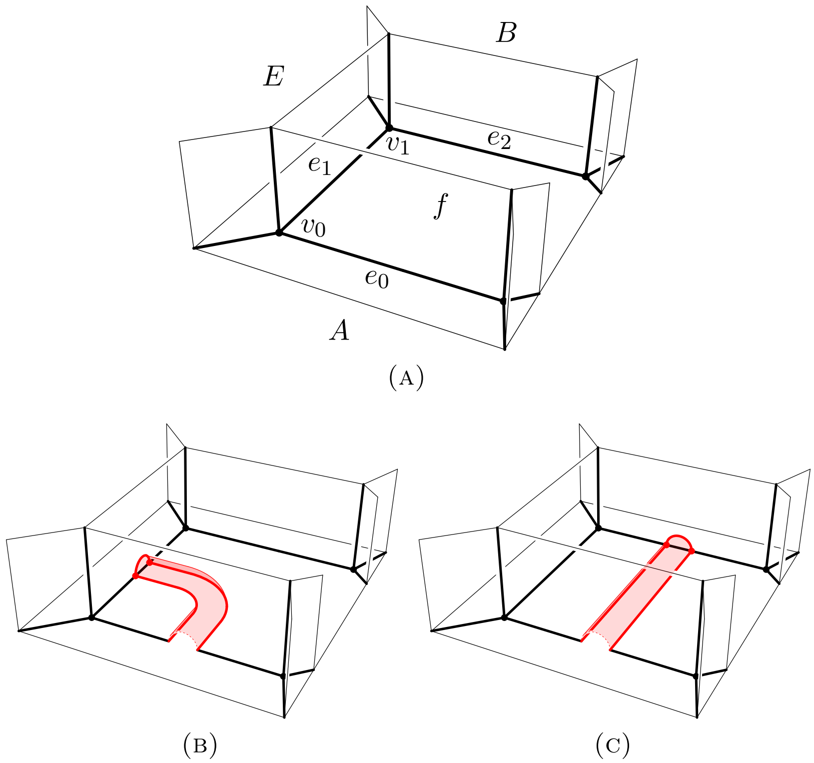

FIGURE 4.14

So, suppose that $\phi_M(e_0) = \phi_M(e_2)$. Figure 4.15A shows the foam $\mathcal{F}'$ obtained after performing the V-move in this case. (In Figure 4.15 we no longer draw the edge e_0 so that we may continue to illustrate all cases without additional figures. We would see different configurations of snakelets along e_0 depending on how the orientations of e_0 and e_2 relate to the orientation of $\phi_M(e_0) = \phi_M(e_2)$.) The stabiliser of e_0 in $\pi_1(M)$ acting on $\widetilde{M}$ is trivial. Thus there is a unique $\gamma \in \pi_1(M)$ taking e_0 to e_2.

Let e_1' and e_2' be the edges of $\widetilde{\mathcal{F}'}$ leaving v_1 in the direction of e_1 and e_2 (in the original foam $\widetilde{\mathcal{F}}$). See Figure 4.15A. Note that neither e_1' nor e_2' is cyclic because the vertex at one end, v_1, is part of the original foam $\widetilde{\mathcal{F}}$, while the vertex at the other end is one of the two vertices at the head of the red snakelet. The regions E and $\gamma(E)$ meet the endpoints of e_2' but not the interior of e_2'. There are two cases depending on whether or not $L(E) = L(\gamma(E))$.

Case $(L(E) = L(\gamma(E)))$. In this case we perform a V-move at v_1 as shown in Figure 4.15B. This is possible by Lemma 4.1 using the fact

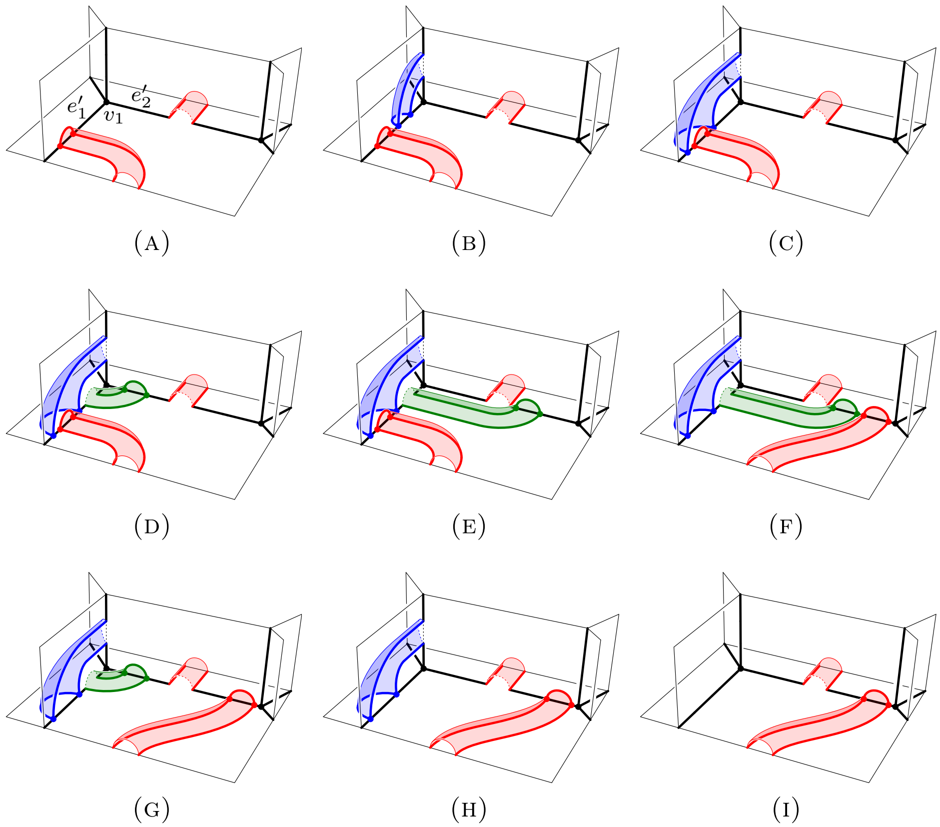

FIGURE 4.15

that e'_1 is not cyclic and that the regions at its endpoints, A and B, have distinct labels. This builds the *blue snakelet*. We then apply Lemma 4.6 to move the head of the red snakelet past one vertex of the blue snakelet. See Figure 4.15C. The hypotheses of Lemma 4.6 are satisfied because $L(A) \neq L(B)$.

Now the regions at the ends of e'_2 have labels $L(E)$ and $L(\gamma(B))$. (The latter is because B is at the end of δ, and therefore $\gamma(B)$ is at the end of $\gamma(\delta)$.) Since $\mathcal{F}$ is L–essential, we have that $L(E)$ is not equal to $L(B)$. By equivariance, $L(\gamma(E))$ is not equal to $L(\gamma(B))$. By assumption, $L(E) = L(\gamma(E))$. Thus the labels $L(E)$ and $L(\gamma(B))$ (of the regions at the ends of e'_2) are distinct. Since e'_2 remains non-cyclic it is L–flippable. We apply Lemma 4.1 to perform a V-move at v_1 as shown in Figure 4.15D. This builds the *green snakelet*.

Next, we slide the head of the green snakelet along the edge e'_2, into the red snakelet, over the two vertices of the red snakelet (using Lemma 4.6), and out of the red snakelet to reach Figure 4.15E. Again we are able to do this without creating any L–inessential faces because $L(E) \neq L(\gamma(B))$.

Next, we slide the head of the red snakelet off of the belly of the blue snakelet and along the back of the green snakelet to reach Figure 4.15F. In this process, the head of the red snakelet meets regions B, E, and then B again. The labels on these are distinct from the label on A. We then slide the green snakelet back through the inside of the red snakelet and back to its original position, as shown in Figure 4.15G. In this process, the head of the green snakelet meets regions B, $\gamma(B)$, and then B again. The labels on these are distinct from the label on E.

Finally we perform reverse V-moves (Lemma 4.1) to remove the green snakelet (as shown in Figure 4.15H), and the blue snakelet, to reach Figure 4.15I. To remove the green snakelet we again use the fact that $L(E) \neq L(\gamma(B))$. To remove the blue snakelet we use the fact that $L(A) \neq L(B)$.

Case $(L(E) \neq L(\gamma(E)))$. In this case we do not add the blue snakelet and we omit all steps mentioning it. The various applications of Lemma 4.1 and Lemma 4.6 are justified similarly. The one subtle step is when we slide the green snakelet back through the inside of the red snakelet and back to its original position, as shown in Figure 4.15G (again the blue snakelet is not present). Here we require that $L(E)$ is distinct from $L(\gamma(B))$. This holds by the hypothesis that $L(E)$ appears precisely once as the label of an outer region for f_s. □

Lemma 4.16. *Assuming Hypothesis 4.8, suppose that the side s meets $N \geq 3$ model edges. Suppose there is an index k with the following properties. Set $E = E_k$. Suppose that $L(E)$ appears precisely once as a label of an outer region for an edge of s, and precisely once as a label of an outer region for an edge of f_s. Suppose that e_{k-1}, e_k, and e_{k+1} are not cyclic. Suppose that the image $\phi_M(e_k)$ under the covering map is distinct from $\phi_M(e_i)$ for $i \neq k$. Then there is a path $\mathcal{F} = \mathcal{F}_0, \ldots, \mathcal{F}_n = \mathcal{F}[\delta]$ of L–essential foams of M where each foam is related to the next by a 2-3 or 3-2 move.*

Proof. Our goal is to provide a set of moves connecting the given foam $\mathcal{F}$ to a foam with the hypotheses of Lemma 4.13, apply that lemma, and then provide a set of moves connecting the output of that lemma to the desired foam $\mathcal{F}[\delta]$.

If $N = 3$ then we apply Lemma 4.13 and we are done. Now suppose that $N > 3$. Let $A = E_0$ and $B = E_N$. See Figure 4.17A. Assume for now that k is not equal to either 1 or $N - 1$. Using the hypotheses that e_{k+1} is not cyclic, and that $L(E)$ appears once as the label of an

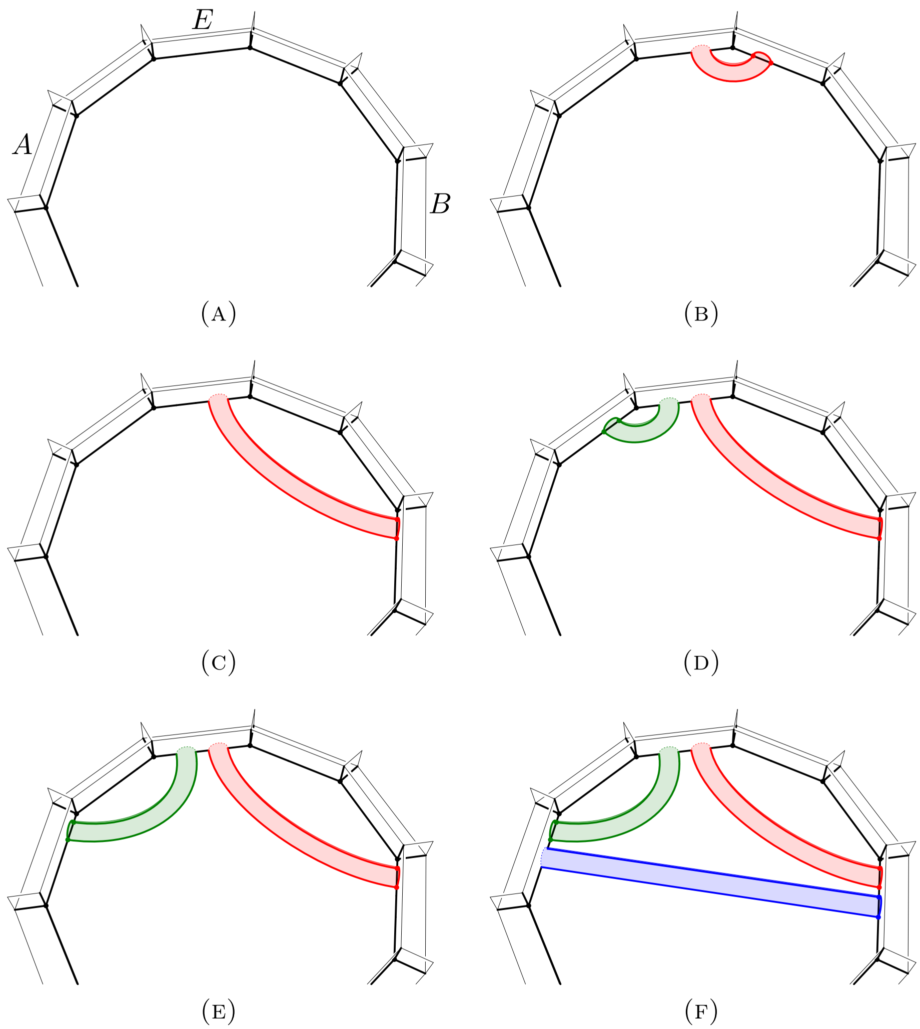

FIGURE 4.17. Creating red and green snakelets in the proof of Lemma 4.16.

outer region for s, by Lemma 4.1 we may apply a V-move, creating a *red snakelet* as shown in Figure 4.17B. Repeatedly applying Lemma 4.6 (and again using the hypothesis that $L(E)$ appears once as the label of an outer region for s), we move the head of the red snakelet just past v_{N-1} and onto e_N. See Figure 4.17C. Note that since $\phi_M(e_k)$ is distinct from $\phi_M(e_i)$ for $i \neq k$, when moving the head of the red snakelet we never need it to cross its own tail.

Next, using the hypothesis that e_{k-1} is not cyclic, we apply a V-move to create a *green snakelet*, as shown in Figure 4.17D. We then move the head of the green snakelet just past v_0 and onto e_0, as shown in Figure 4.17E. These moves on the green snakelet are possible unless the

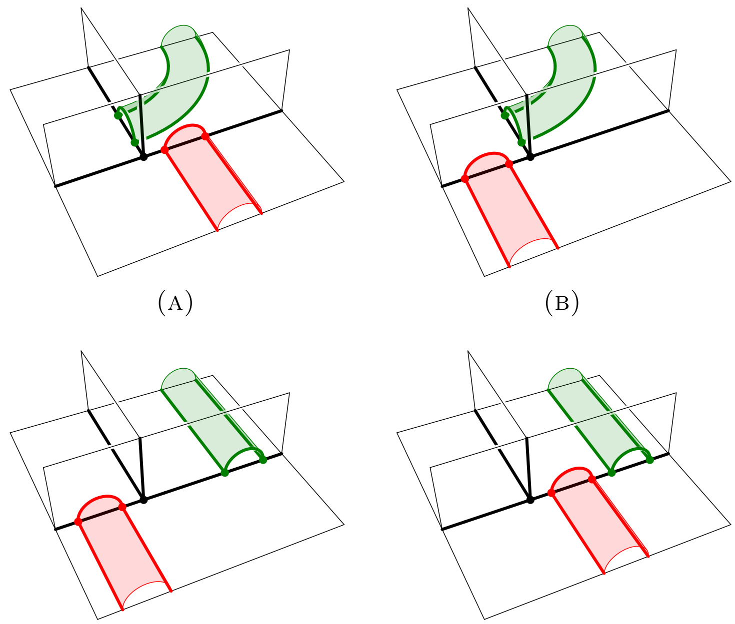

(A) (B) (C) (D)

FIGURE 4.18. A regular neighbourhood of the vertex v_{N-1}. The red snakelet moves back one step to let the green snakelet past, then moves forward into place again.

head of the red snakelet blocks us. If it does, we move the red snakelet back one step onto e_{N-1}, apply the desired move to the green snakelet, and then move the red snakelet forward to e_N again. This is illustrated in Figure 4.18 for the case of moving the head of the green snakelet past a vertex. We deal with the case of creating the green snakelet with a V-move (when the head of the red snakelet is in the way) in a similar fashion.

With the red and green snakelets in place, the side s of δ has been reduced to three edges, with the middle edge being made up of a small segment of the original edge e_k, together with parts of the one-skeletons of the red and green snakelets. See Figure 4.17E. This edge is not cyclic because its endpoints are distinct, being vertices on different snakelets. One can check that the remaining hypotheses of Lemma 4.13 also hold, and thus we can do the 0-2 move, building the *blue snakelet*. The result is illustrated in Figure 4.17F in the case that $\phi_M(e_0) \neq \phi_M(e_N)$.

With the blue snakelet in place, we must deconstruct the red and green snakelets. This is essentially the reverse of the process we used to create the red and green snakelets and move them into place. The only difference is that our snakelet heads move around the face f_s instead of s. Note that it is possible for a 0-2 move to introduce a cyclic edge when $\phi_M(e_0) = \phi_M(e_N)$. However, such a cyclic edge is based at one of the new vertices (on the blue snakelet) created by the 0-2 move, not any of the v_i. In particular, the edges next to E around f_s are not cyclic, so Lemma 4.1 can be used in reverse to deconstruct the red and green snakelets.

The case that $k = 1$ or $k = N - 1$ is similar but simpler. We need build only one of the "helper" snakelets (red or green) in order to reach the hypotheses of Lemma 4.13. □

5. Distant labels

We now turn to the problem of implementing a 0-2 move in general.

5.1. **Locally frozen configurations.** As an example, suppose we want to perform a 0-2 move along an arc δ that connects opposite sides of a hexagonal face f. Suppose that the outer regions around f alternate between two labels, a and b say. Then performing a 0-2 move along δ creates a new face between regions with labels a and b, so does not create an L–inessential face. However, none of the edges of ∂f are L–flippable. Thus, there is no local 2-3 move that we can use in Lemma 4.1 to start building snakelets. Depending on the combinatorics around f, there may be no 2-3 or 3-2 move that we can perform anywhere near f. So we may need to start work very far away. This is why Proposition 3.19 includes the assumption that the image of L is infinite.

5.2. **Strategy.** In the remainder of the paper we show how to use a (possibly distant) L–flippable edge to perform the 0-2 move by reducing to the case of Lemma 4.16. In order to do this, we will grow a complementary region E, having a "distant" label, through the foam to bring it into contact with the target face f. We will use *augmented 2-3 moves* to do this. See Section 5.4. This done, we apply Lemma 4.16. Then we must undo all of the augmented 2-3 moves. However, there is a potential obstacle to performing exactly the inverse moves in reverse order; the 0-2 move along the arc δ alters the foam. Said another way, our sequence of moves growing E must commute with the 0-2 move along δ.

To achieve this goal, we introduce two variants of a new *nature reserve move*, described in Section 5.11. These "protect" the endpoints of the arc δ from being disturbed by augmented 2-3 moves. Furthermore, they also protect the snakelet formed by performing a 0-2 move along δ from being disturbed by the same augmented 2-3 moves.

5.3. **Handle structures for foams.** We use the following notion of *handle structures* to organise the argument.

Suppose that M is a compact, connected three-manifold with boundary. Suppose that $\mathcal{F}$ is a foam in M. For each vertex v of $\mathcal{F}$ we choose a regular neighbourhood $\eta(v) \subset M$ of v which we call the *zero-handle* for v. We make the neighbourhoods small enough that the zero-handles are disjoint. To agree with the notation below we also set $\overline{\eta}(v) = \eta(v)$.

For each edge e of $\mathcal{F}$ we choose a regular neighbourhood $\eta(e)$ of

$$e - \bigcup_{v \in \partial e} \overline{\eta}(v) \qquad \text{in} \qquad M - \bigcup_{v \in \partial e} \overline{\eta}(v)$$

which we call the *one-handle* for e. We make the neighbourhoods small enough that the one-handles are disjoint. We define

$$\overline{\eta}(e) = \eta(e) \cup \bigcup_{v \in \partial e} \overline{\eta}(v)$$

For each face f of $\mathcal{F}$ we choose a regular neighbourhood $\eta(f)$ of

$$f - \bigcup_{e \in \partial f} \overline{\eta}(e) \qquad \text{in} \qquad M - \bigcup_{e \in \partial f} \overline{\eta}(e)$$

which we call the *two-handle* for f. We make the neighbourhoods small enough that the two-handles are disjoint. We define

$$\overline{\eta}(f) = \eta(f) \cup \bigcup_{e \in \partial f} \overline{\eta}(e)$$

We say that the collection of zero-, one-, and two-handles is a *handle structure* $\eta(\mathcal{F})$ for $\mathcal{F}$. We say that a foam $\mathcal{G}$ of M is *carried* by $\mathcal{F}$ if there is a *carrying function* C from the vertices, edges, and faces of $\mathcal{G}$ to the vertices, edges, and faces of $\mathcal{F}$ so that for all cells c of $\mathcal{G}$ we have:

(1) c lies in $\overline{\eta}(C(c))$,
(2) c does not lie in $\overline{\eta}(d)$ for any model facet d of $C(c)$, and
(3) the dimension of $C(c)$ is at most the dimension of c.

Note that when $\mathcal{G}$ is carried by $\mathcal{F}$, the carrying function is unique.

5.4. **Augmented 2-3 moves.** The main move we use is the following.

Definition 5.6. Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M. Suppose that $\mathcal{G}$ is an L–essential foam in M. Suppose that e is an L–flippable edge of $\mathcal{G}$. Suppose that e is equipped with an orientation, pointing from the vertex u to the vertex v. See Figure 5.5A. Note that $u \neq v$ because e is not cyclic. Fix small regular open neighbourhoods $\mathcal{N}(e)$, $\mathcal{N}(u)$, and $\mathcal{N}(v)$ of e, u, and v such that $\mathcal{N}(u)$ and $\mathcal{N}(v)$ lie within $\mathcal{N}(e)$. We

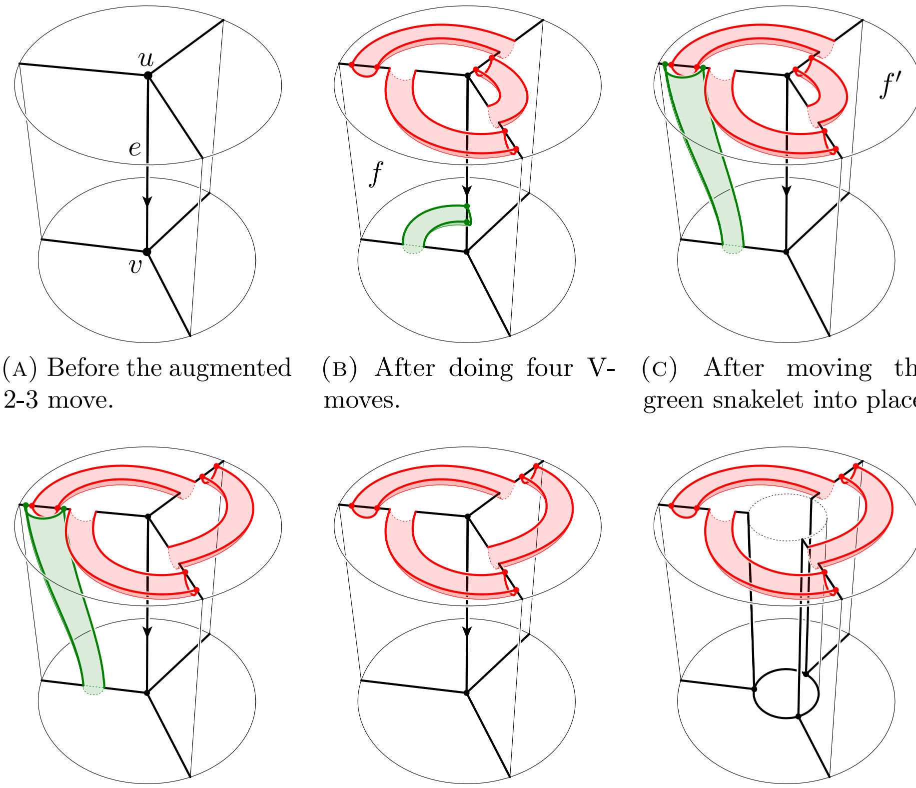

(A) Before the augmented 2-3 move.

(B) After doing four V-moves.

(C) After moving the green snakelet into place.

(D) After sliding the innermost red snakelet out.

(E) After deconstructing the green snakelet.

(F) After the 2-3 move.

FIGURE 5.5. An augmented 2-3 move.

perform an *augmented 2-3 move* along e to produce a foam $\mathcal{G}_e$ as shown in Figure 5.5F.

We require the following.

(1) The augmented 2-3 move is supported within $\mathcal{N}(e)$.
(2) The three red snakelets shown in Figure 5.5F are in $\mathcal{N}(u)$.
(3) All cells generated by the 2-3 move that do not have ancestors are in $\mathcal{N}(v)$.
(4) Suppose additionally that $\mathcal{G}$ is carried by some foam $\mathcal{F}$ with carrying function C. In this case we also require that $\mathcal{N}(c)$ lies within $\overline{\eta}(C(c))$ as c ranges over e, u, and v. ◇

Lemma 5.7. *The augmented 2-3 move along e can be realised by a sequence of 2-3 and 3-2 moves from $\mathcal{G}$ to $\mathcal{G}_e$ that pass through L–essential foams.*

Proof. First, we apply Lemma 4.1 four times, creating three *red snakelets* and one *green snakelet*. See Figure 5.5B. Next we slide the head of the green snakelet around the boundary of the face f, using Lemma 4.6 three times. After a 2-3 and then a 3-2 move (similar to the moves between Figure 4.15B and Figure 4.15C), we reach Figure 5.5C, in which the belly of the green snakelet covers the head of one of the red snakelets. This done, using Lemma 4.6 we slide the innermost red snakelet around the face f' to get to Figure 5.5D. The green snakelet ensures that we do not introduce an L–inessential face as we do this. Next, we reverse the movement of the green snakelet by sliding it back around the boundary of f. We undo the V-move, deconstructing the green snakelet. This takes us to Figure 5.5E. Finally we apply the 2-3 move along e, giving the foam $\mathcal{G}_e$ shown in Figure 5.5F. □

We collect several useful properties of the augmented 2-3 move. In particular, augmented 2-3 moves do not create cyclic edges, and so can safely be used to destroy them. The third property below refers to Nat, a collection of *nature reserve edges*. See Section 6.

Lemma 5.8. *Suppose that $\mathcal{G}$ is carried by $\mathcal{F}$. Suppose that $\mathcal{G}$ is L–essential. Suppose that e is an L–flippable oriented edge of $\mathcal{G}$. Suppose that $\mathcal{G}_e$ is the result of applying an augmented 2-3 move along e. Then we have the following.*

1. *The foam $\mathcal{G}_e$ is L–essential.*
2. *The foam $\mathcal{G}_e$ has an L–flippable edge.*
3. *No cyclic edges of $\mathcal{G}_e$ meet $\mathcal{N}(e)$.*
4. *The foam $\mathcal{G}_e$ is carried by $\mathcal{F}$.*
5. *Suppose that w is a vertex of $\mathcal{F}$. Suppose that Nat is a collection of edges in $\mathcal{G}^{(1)}$. Suppose that e is not in Nat. Let Nat_e be the collection of descendants of the edges of Nat under the augmented 2-3 move. Suppose that $X = \mathcal{G}^{(1)} \cap \eta(w) - \mathsf{Nat}$ is connected. Then $X_e = \mathcal{G}_e^{(1)} \cap \eta(w) - \mathsf{Nat}_e$ is connected.*

Proof. Property (1) follows from Lemma 3.12 (or Lemma 5.7) and the fact that e is L-flippable. Each of the three edges created by the final 2-3 move along e is L–flippable. This gives property (2).

Since the augmented 2-3 move occurs within a ball, all nine vertices it generates are distinct. By consulting Figure 5.5F we see that there are no cyclic edges entirely contained within the figure, and no vertex has more than one edge-end leaving the figure. This gives property (3).

Suppose that C is the carrying function for $\mathcal{G}$ in $\mathcal{F}$. Suppose that e was oriented from vertex u to vertex v. Suppose that d is a cell of $\mathcal{G}_e$. There are three cases. First suppose that d is disjoint from

$\mathcal{N}(e)$. Then $c = d$ is the ancestor of d in $\mathcal{G}$. Thus d is carried by the cell $C(c)$ in $\mathcal{F}$. Second, suppose that d meets but is not contained in $\mathcal{N}(e)$. (There are six edge ends and nine pieces of faces of this type.) Consulting Figure 5.5F we see that d is strictly contained in a cell c of $\mathcal{G}$. Furthermore, d is carried by $C(c)$. Finally, suppose that d is contained in $\mathcal{N}(e)$. (There are nine vertices, fifteen edges, and seven faces of this type.) By Definition 5.6(4), each of these is carried by one of $C(u)$, $C(v)$, or $C(e)$. Thus we have established (4).

Let c be any one of u, v, or e. Consulting Figure 5.5A (respectively Figure 5.5F) we see that Nat (Nat_e) meets $\mathcal{G}^{(1)} \cap \mathcal{N}(c)$ ($\mathcal{G}_e^{(1)} \cap \mathcal{N}(c)$) in at most only edge ends. Again consulting the figures, we see that $\mathcal{G}^{(1)} \cap \mathcal{N}(c) - \mathsf{Nat}$ and $\mathcal{G}_e^{(1)} \cap \mathcal{N}(c) - \mathsf{Nat}_e$ are each connected. We now recall the following.

Fact 5.9. Suppose that $A \cup B$ and $A \cup B'$ are graphs with A being a subgraph of both $A \cup B$ and $A \cup B'$, while B, and B' are subgraphs of $A \cup B$ and $A \cup B'$ respectively. Suppose that $A \cap B = A \cap B'$. Suppose that B and B' are both connected. Suppose that $A \cup B$ is connected. Then $A \cup B'$ is connected. ◇

The three-ball $\eta(w)$ either:

- is disjoint from $\mathcal{N}(e)$,
- contains $\mathcal{N}(e)$,
- contains $\mathcal{N}(u)$ but not $\mathcal{N}(v)$,
- contains $\mathcal{N}(v)$ but not $\mathcal{N}(u)$, or
- contains $\mathcal{N}(u) \sqcup \mathcal{N}(v)$ but not $\mathcal{N}(e)$.

Thus the number of components of $\eta(w) \cap \mathcal{N}(e)$ is either zero, one, or two. If $\eta(w) \cap \mathcal{N}(e)$ is empty then $X = X_e$ and we are done. Suppose instead that $\eta(w)$ contains $\mathcal{N}(e)$. We are given that X is connected. By taking closures, X and X_e become graphs. Take $A = X - \mathcal{N}(e) = X_e - \mathcal{N}(e)$, take $B = X \cap \mathcal{N}(e) = \mathcal{G}^{(1)} \cap \mathcal{N}(e) - \mathsf{Nat}$, and take $B' = X_e \cap \mathcal{N}(e) = \mathcal{G}_e^{(1)} \cap \mathcal{N}(e) - \mathsf{Nat}_e$. By the above argument, B and B' are both connected. Then $A \cup B = X$, $A \cup B' = X_e$, and $A \cap B = A \cap B'$. Applying Fact 5.9, we get that X_e is connected.

The argument is similar in the other cases, replacing $\mathcal{N}(e)$ with $\mathcal{N}(u)$, $\mathcal{N}(v)$, or their disjoint union. (In the latter case, we apply Fact 5.9 twice.) Thus we obtain (5). □

Remark 5.10. Carrying only places restrictions on the zero-, one-, and two-dimensional cells. The three-dimensional complementary regions (components of $M - \mathcal{G}_e$) are not similarly restricted. Indeed the main purpose of the augmented 2-3 move is to grow complementary regions throughout the universal cover. ◇

5.11. **Nature reserve moves.** In order to define a nature reserve move, we make the following assumptions: Suppose that M is a compact, connected three-manifold with boundary. Suppose that L is a labelling of Δ_M. Suppose that $\mathcal{G}$ is an L–essential foam in M. Suppose that e is an L–flippable edge of $\mathcal{G}$. Suppose that p is one of the endpoints of e (since e is not cyclic it has distinct endpoints). Let $\mathcal{N}(p)$ be a small regular neighbourhood of p. If $\mathcal{G}$ is carried by some foam $\mathcal{F}$ with carrying function C then we further require that $\mathcal{N}(p)$ lies within $\eta(C(p))$.

Suppose that δ is an arc in a face of $\mathcal{G}$ with at least one endpoint on e. Suppose that $\mathcal{G}[\delta]$ is L–essential and has an L–flippable edge. There are two possibilities. Either the arc δ has *exactly one* endpoint d on e, or it has *both* endpoints, d and d', on e.

5.11.1. *Singleton nature reserve move.* Suppose that the arc δ has precisely one endpoint d on e.

Definition 5.13. We perform *singleton nature reserve move* at p along e to produce a foam $\mathcal{G}_p$. This foam is shown in Figure 5.12B. We properly isotope δ within f to move d along e so that d lies within $\mathcal{N}(p)$. We introduce three *red snakelets* and one *green snakelet*, all contained within $\mathcal{N}(p)$. The green snakelet lies on a (descendant of a) face g of $\mathcal{G} \cap \mathcal{N}(p)$ so that

- g is incident to the interior of e,
- g does not intersect the interior of δ, and
- along e, the endpoint d lies between p and the vertices on the green snakelet.

(There are two possibilities for g.) See Figure 5.12D. In $\mathcal{G}_p$, the edge e has been split into three segments. The first of these is the *nature reserve edge*, e', which "protects" the endpoint d of δ. The second, with interior incident to the green snakelet, is e'' say. The third and last is the *remainder*, e''' say. See Figure 5.12D. ◇

Lemma 5.14. *For either choice of vertex p, the singleton nature reserve move at p along e can be realised by a sequence of 2-3 and 3-2 moves from $\mathcal{G}$ to $\mathcal{G}_p$ with all foams being L–essential.*

Proof. Since e is L–flippable, we follow the steps of Section 5.4, within $\mathcal{N}(p)$, to reproduce the construction shown in Figures 5.5A to 5.5E, (with e oriented away from p). We call the resulting foam $\mathcal{G}'$. See Figure 5.12A.

We perform a V-move at the vertex p (shown in closeup in Figure 5.12C) to produce the foam $\mathcal{G}_p$; the hypotheses of Lemma 4.1 are satisfied because e is L–flippable. □

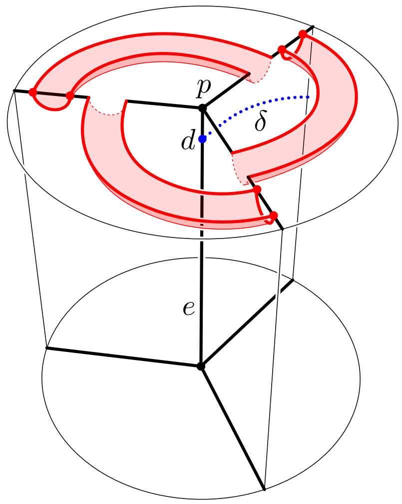

(A) Part of the foam $\mathcal{G}'$. The arc δ has an endpoint d on e.

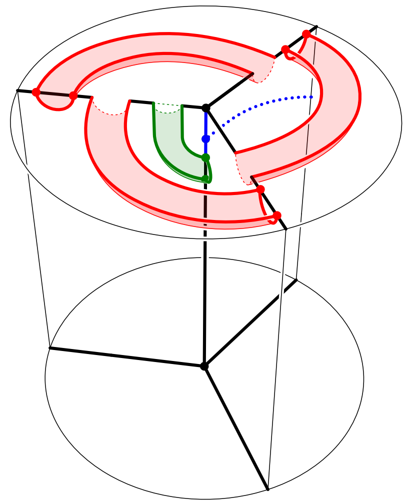

(B) Part of the foam $\mathcal{G}_p$. A nature reserve edge protects d.

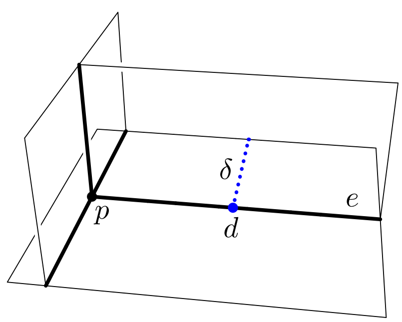

(C) Close up of Figure 5.12A.

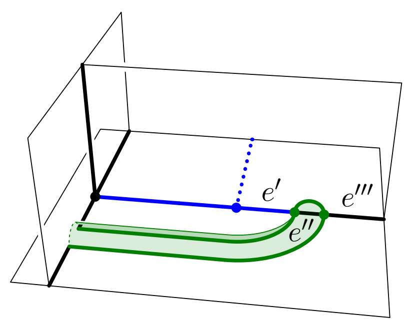

(D) Close up of Figure 5.12B.

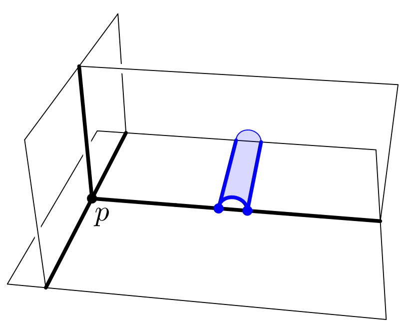

(E) Part of the foam $\mathcal{G}[\delta]$.

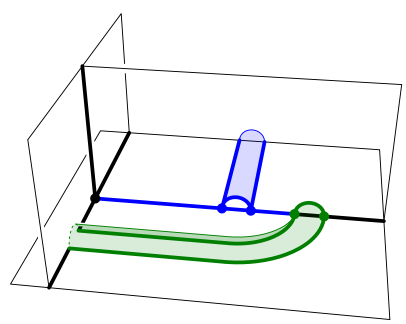

(F) Part of the foam $\mathcal{G}_p[\delta]$.

FIGURE 5.12. Introducing a nature reserve edge e'. In Figure 5.12E, we are so close to p that this is also a correct picture of $\mathcal{G}'[\delta]$.

Lemma 5.15. *There is a choice of vertex p so that there is a sequence of 2-3 and 3-2 moves from $\mathcal{G}[\delta]$ to $\mathcal{G}_p[\delta]$ with all foams being L–essential.*

Proof. Suppose that $\widetilde{e}$ is a lift of e. Let P and Q be the regions of $\widetilde{M} - \widetilde{\mathcal{G}}$ that meet the endpoints of $\widetilde{e}$ but not the interior. Since e is L–flippable, $L(P) \neq L(Q)$. After performing the 0-2 move along δ, Figure 5.12C becomes Figure 5.12E. Here we see an additional region R of $\widetilde{M} - \widetilde{\mathcal{G}}[\delta]$ that meets the interior of $\widetilde{e}$. Within the figure, R is inside the snakelet. Since $L(P) \neq L(Q)$, the label $L(R)$ is different from at least one of $L(P)$ and $L(Q)$. Let $\widetilde{p}$ be an endpoint of $\widetilde{e}$ incident to a region (either P or Q) which does not have label $L(R)$. Let $p = \phi_M(\widetilde{p})$ be the image under the covering map.

We deduce that the edge between p and the head of the blue snakelet is L–flippable. Thus following the steps of Section 5.4 as before, we can get from $\mathcal{G}[\delta]$ to $\mathcal{G}'[\delta]$. Again by our choice of p we can apply Lemma 4.1 and Lemma 4.6 to build and then slide the head of the green snakelet past the head of the blue snakelet built along δ. This done, we have reached $\mathcal{G}_p[\delta]$. Again, see Figure 5.12F. □

5.15.2. *Pair nature reserve move.* Suppose that the arc δ has both endpoints d and d' on e.

Definition 5.16. We perform a *pair nature reserve move* at p along e to produce a foam $\mathcal{G}_p$. This foam is shown in Figure 5.18B. We properly isotope δ within f to move d and d' along e so that both lie within $\mathcal{N}(p)$. Exactly as in Definition 5.13, we introduce three *red snakelets* and one *green snakelet*, all contained within $\mathcal{N}(p)$. We arrange matters as in Definition 5.13, except that now

- along e, the endpoints d and d' lie between p and the vertices on the green snakelet.

(In Definition 5.13 there were two possibilities for g; now there is only one.) See Figure 5.18D. As in Definition 5.13, in $\mathcal{G}_p$ the edge e has been split into three segments. These are the *nature reserve edge* e' (containing d and d'), the edge e'', whose interior is incident to the green snakelet, and the *remainder* e'''. See Figure 5.18D. ◇

The proof of the following lemma is identical to that of Lemma 5.14 and we omit it.

Lemma 5.17. *For either choice of vertex p, the pair nature reserve move at p along e can be realised by a sequence of 2-3 and 3-2 moves from $\mathcal{G}$ to $\mathcal{G}_p$ with all foams being L–essential.* □

We do not yet have the tools to prove the following lemma. We defer its proof to Section 6.15.

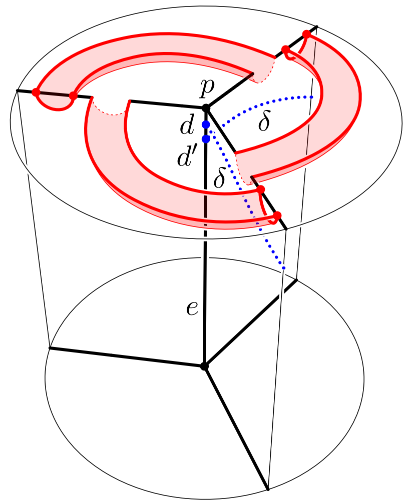

(A) Part of the foam $\mathcal{G}'$. The arc δ has both endpoints, d and d' on e.

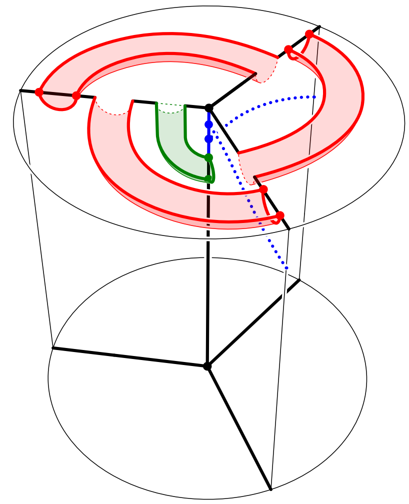

(B) Part of the foam $\mathcal{G}_p$. A nature reserve edge protects d and d'.

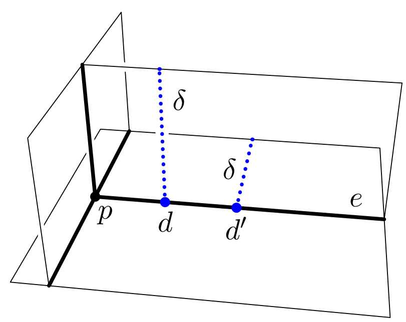

(C) Close up of Figure 5.18A.

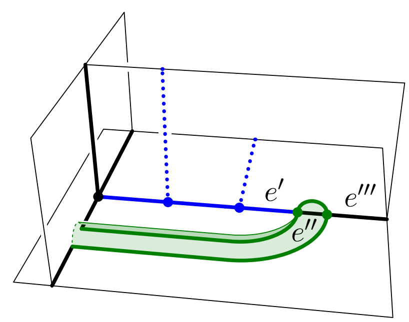

(D) Close up of Figure 5.18B.

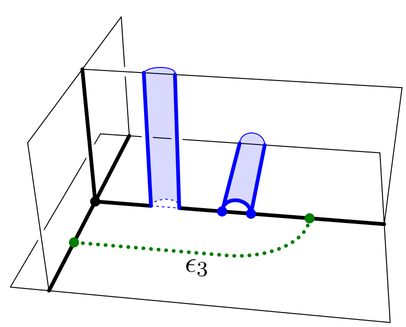

(E) Figure 5.18C with the snakelet generated by the 0-2 move. This is part of the foam $\mathcal{G}'[\delta]$.

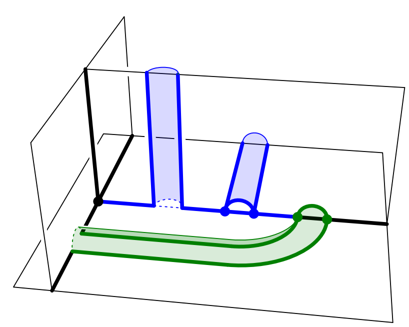

(F) Figure 5.18D with the snakelet generated by the 0-2 move. This is part of the foam $\mathcal{G}_p[\delta]$.

Figure 5.18. Introducing a nature reserve (the segment of e drawn in blue) in the case that both ends of the arc δ lie on the same edge e.

Lemma 5.19. *For either choice of p the following holds. Suppose that $\mathcal{G}_p$ is the result of applying the pair nature reserve move at p along e. Then there is a sequence of 2-3 and 3-2 moves from $\mathcal{G}[\delta]$ to $\mathcal{G}_p[\delta]$ with all foams being L–essential.*

5.19.3. *Properties of nature reserve moves.* We collect several useful properties of the singleton and pair nature reserve moves.

Remark 5.20. The five regions meeting the closure of e''' in either nature reserve move are the same as the five regions meeting the closure of e. Also, one endpoint of e''' lies on the head of the green snakelet while the other lies outside of $\mathcal{N}(p)$ (because e was not cyclic). Thus e''' is not cyclic. Thus in both cases the remainder edge e''' is L–flippable. $\diamond$

The proof of the following lemma is very similar to the proof of Lemma 5.8 and we omit it.

Lemma 5.21. *Suppose that $\mathcal{G}$ is carried by $\mathcal{F}$. Suppose that $\mathcal{G}$ is L–essential. Suppose that e is an L–flippable edge of $\mathcal{G}$. Fix an end p of e. Suppose that $\mathcal{G}_p$ is the result of applying a (singleton or pair) nature reserve move at p along e, producing a nature reserve edge e'. Then we have the following.*

1. *The foam $\mathcal{G}_p$ is L–essential.*
2. *The foam $\mathcal{G}_p$ has an L–flippable edge.*
3. *No cyclic edges of $\mathcal{G}_p$ meet $\mathcal{N}(p)$.*
4. *The foam $\mathcal{G}_p$ is carried by $\mathcal{F}$.*
5. *Suppose that w is a vertex of $\mathcal{F}$. Suppose that Nat is a collection of edges in $\mathcal{G}^{(1)}$. Suppose that e is not in Nat. Let Nat_p be the collection of descendants of the edges of Nat under the singleton nature reserve move, union with e'. Suppose that $\mathcal{G}^{(1)} \cap \eta(w) - \mathsf{Nat}$ is connected. Then $\mathcal{G}_p^{(1)} \cap \eta(w) - \mathsf{Nat}_p$ is connected.* $\square$

6. Parallel sequences

Suppose that $\mathcal{F}$ is the given L–essential foam. Suppose that δ is the given arc. Performing a 0-2 move along δ produces the foam $\mathcal{F}[\delta]$. In this section, we give the proof of Proposition 3.19. That is, we give a sequence of 2-3 and 3-2 moves from $\mathcal{F}$ to $\mathcal{F}[\delta]$ with all foams being L–essential.

We deal with $\mathcal{F}$ and $\mathcal{F}[\delta]$ in parallel, working in from both ends of our eventual sequence of 2-3 and 3-2 moves. That is, we produce a sequence $\mathcal{F} = \mathcal{F}_0, \mathcal{F}_1, \ldots, \mathcal{F}_K$, where $\mathcal{F}_i$ and $\mathcal{F}_{i+1}$ (respectively $\mathcal{F}_i[\delta]$ and $\mathcal{F}_{i+1}[\delta]$) are related by a sequence of 2-3 and 3-2 moves through L–essential foams. Furthermore, F_K and $F_K[\delta]$ are the input and output foams of Lemma 4.16.

All foams $\mathcal{F}_i$ are carried by $\mathcal{F}$ with carrying function C_i, say. To each foam $\mathcal{F}_i$ we associate a list Vis_i of *visited* vertices of $\mathcal{F}$. We also maintain a list Nat_i of zero, one, or two edges of $\mathcal{F}_i$; these are the *nature reserve edges*. We abuse notation and refer to the descendant of f in $\mathcal{F}_i$ as f.

We break the sequence of moves into two stages.

- In the *loosening stage*, taking us from $\mathcal{F}_0$ to $\mathcal{F}_J$, we apply loosening moves along all edges of $\mathcal{F}_0$, destroying any and all cyclic edges. Furthermore, after the loosening stage the set Nat_J contains either one or two nature reserve edges, one for each edge containing an endpoint of δ.
- In the *contacting stage*, taking us from $\mathcal{F}_J$ to $\mathcal{F}_K$, we bring a "distant" region into contact with the face f.

6.1. **Loosening move.**

Definition 6.2. Suppose that $\mathcal{F}_i$, C_i, Vis_i, and Nat_i are the given foam, with its carrying function, list of visited vertices, and nature reserve edges. Suppose that e_i is an oriented and L–flippable edge of $\mathcal{F}_i$ which is not in Nat_i. Then the *loosening move* along e_i produces $\mathcal{F}_{i+1}$, C_{i+1}, Vis_{i+1}, and Nat_{i+1}, as follows.

- Suppose that e_i does not meet an endpoint of δ. Then we perform an augmented 2-3 move along e_i.
- Suppose instead that e_i meets one or both endpoints of δ. First, we apply a nature reserve move (singleton or pair, as appropriate) to $\mathcal{F}_i$ along e_i. Let e''' be the remainder edge produced by the nature reserve move. We orient e''' in the same direction as e_i. Second, we apply an augmented 2-3 move along the remainder edge e'''. (Note that the augmented 2-3 move along e''' is possible because e''' is L–flippable by Remark 5.20.)

The carrying function C_{i+1} is given by Lemma 5.8(4) or Lemma 5.21(4).

Let v_i be the vertex of $\mathcal{F}_i$ at the terminal point of e_i. We set $\mathsf{Vis}_{i+1} = \mathsf{Vis}_i \cup \{C_i(v_i)\}$.

- Suppose that e_i does not meet an endpoint of δ. Since e_i is not in Nat_i, every edge of Nat_i has a unique descendant under the augmented 2-3 move. We place this unique descendant in Nat_{i+1}.
- Suppose that e_i meets one or both endpoints of δ. Since e_i is not in Nat_i, every edge of Nat_i has a unique descendant under both the nature reserve move and the augmented 2-3 move. We place this unique descendant in Nat_{i+1}, together with the new nature reserve edge e'. ◇

To either $\mathcal{F}_i$ or $\mathcal{F}_{i+1}$ we may apply the 0-2 move along δ to obtain $\mathcal{F}_i[\delta]$ or $\mathcal{F}_{i+1}[\delta]$ respectively. We record this information in Diagram 6.4, together with a dashed arrow connecting $\mathcal{F}_i[\delta]$ to $\mathcal{F}_{i+1}[\delta]$. Depending on the number of endpoints of δ on e_i, a subcollection of Lemmas 5.7, 5.14, 5.15, 5.17, and 5.19 proves the following.

Corollary 6.3. *For each horizontal arrow in Diagram 6.4 there is a sequence of 2-3 and 3-2 moves with all foams being L–essential.* □

$$\begin{array}{ccc} \mathcal{F}_i & \longmapsto & \mathcal{F}_{i+1} \\ \big\downarrow{\scriptstyle 0\text{-}2} & & \big\downarrow{\scriptstyle 0\text{-}2} \\ \mathcal{F}_i[\delta] & \dashrightarrow & \mathcal{F}_{i+1}[\delta] \end{array} \tag{6.4}$$

Thus the loosening move commutes with the 0-2 move along δ. Note that we have not yet established Lemma 5.19 (the case in which δ has both endpoints on e_i). The proof of Lemma 5.19 in fact requires Corollary 6.3 in the case in which δ has zero or one endpoint on e_i.

6.5. **Loosening stage.** From here and until Section 6.15, we will assume the following.

Hypothesis 6.6. The endpoints of δ lie on distinct edges of $\mathcal{F}$. ◇

(Under this hypothesis, as discussed in Section 5.11, we do not use the pair nature reserve move. Hypothesis 6.6 is used only when a loosening move is performed along an edge containing an endpoint of δ.) We now recursively choose the sequence of edges along which we apply loosening moves.

We first deal with the base case. We set $\mathcal{F}_0 = \mathcal{F}$, and we take both Vis_0 and Nat_0 to be empty. We now choose e_0 to be any L–flippable edge of $\mathcal{F}_0$. Such an edge exists by the hypotheses of Proposition 3.19. We arbitrarily orient e_0. Let u_0 and v_0 be the initial and terminal vertices of e_0. We apply a loosening move along e_0. This takes us from $\mathcal{F}_0$ to $\mathcal{F}_1$. The set of visited vertices becomes $\mathsf{Vis}_1 = \{v_0\}$. If e_0 does not meet an endpoint of δ then $\mathsf{Nat}_1 = \mathsf{Nat}_0$ is empty. If e_0 contains an endpoint of δ then $\mathsf{Nat}_1 = \{e_0'\}$, where e_0' is the nature reserve edge generated by the nature reserve move.

Our next step is to choose e_1 to be one of the three edges of $\mathcal{F}_1$ that intersect $\eta(e_0)$. One endpoint of e_1 lies in $\eta(u_0)$ while the other lies in $\eta(v_0)$. We orient e_1 from the latter to the former; we now apply a loosening move (necessarily an augmented 2-3 move) along e_1. This takes us from $\mathcal{F}_1$ to $\mathcal{F}_2$, and $\mathsf{Vis}_2 = \{u_0, v_0\}$. Here it is not possible that e_1 contains an endpoint of δ so we set $\mathsf{Nat}_2 = \mathsf{Nat}_1$.

We now deal with the inductive step. We are given a foam $\mathcal{F}_i$ with its list of visited vertices Vis_i and nature reserve edges Nat_i. The induction hypothesis tells us that $\mathcal{F}_i$ is L–essential. There are now two possibilities. Either there is an edge of $\mathcal{F}_i$ which is a descendant of an edge of $\mathcal{F}_0$, or not. In the latter case we set $J = i$ and we are done with the loosening stage.

Suppose instead that we have such an edge. We choose this edge to be e_i. (Thus $C_i(e_i)$ is an edge rather than a vertex.) Since the one-skeleton of $\mathcal{F}$ is connected, we may assume (possibly by choosing a different such edge e_i) that e_i has at least one endpoint, u_i say, in a zero-handle $\eta(C_i(u_i))$, where $C_i(u_i) \in \mathsf{Vis}_i$. We will now grow a nearby region so that e_i becomes L–flippable. We will then perform a loosening move along e_i.

Set $w = C_i(u_i)$. Pick a lift $\widetilde{w}$ of w; let $\eta(\widetilde{w})$ be the corresponding lift of $\eta(w)$.

Claim 6.7. The zero-handle $\eta(\widetilde{w})$ intersects regions of $\widetilde{M} - \widetilde{\mathcal{F}}_i$ having at least five labels. Furthermore, for all vertices u in $\mathcal{F}_i \cap \eta(w)$, we have that u is not incident to a cyclic edge in $\mathcal{F}_i$.

Proof. The loosening move that initially added w to the visited list Vis_j ($j \leq i$) added a fifth region to $\eta(\widetilde{w}) - \widetilde{\mathcal{F}}_j$. See Figure 5.5F. These regions have distinct labels because $\mathcal{F}_j$ is L–essential and the regions are pairwise adjacent. The second statement follows by Lemma 5.8(3) and Lemma 5.21(3). □

Thus e_i is not cyclic. Let v_i be the other endpoint of e_i. (It is possible that $C_i(u_i) = C_i(v_i)$.) Let $\widetilde{e_i}$ be a lift of e_i meeting $\eta(\widetilde{w})$. Let $\widetilde{u_i}$ and $\widetilde{v_i}$ be the corresponding lifts of u_i and v_i.

6.8. **Tunnelling through a zero-handle.** Since $w = C_i(u_i)$ lies in Vis_i, the above Claim 6.7 tells us that the zero-handle $\eta(C_i(\widetilde{u_i}))$ of $\widetilde{\mathcal{F}}$ intersects regions having at least five labels. In particular, it contains a region E with label different from the labels of the four regions incident to $\widetilde{v_i}$. By Lemma 5.8(5), Lemma 5.21(5), and induction, there is a path in the one-skeleton of $\widetilde{\mathcal{F}}_i \cap \eta(C_i(\widetilde{u_i}))$ from E to $\widetilde{u_i}$ that does not pass through any edge of Nat_i. Take a minimal such path γ and orient it towards $\widetilde{u_i}$. Orient e_i from u_i to v_i and form γ' by adding $\widetilde{e_i}$ to the end of γ. Note that the edges containing endpoints of δ are either in Nat_i or they are descendants of edges of $\mathcal{F}$, in which case they are carried by edges of $\mathcal{F}$. In either case, no endpoint of δ lies on γ.

We repeatedly apply loosening moves along γ', creating foams $\mathcal{F}_{i+1}$, $\mathcal{F}_{i+2}$, and so on, until we perform a loosening move along (a descendant of) e_i. Strictly speaking, this requires another recursive construction with the following hypotheses.

- Each remaining edge of γ' has a descendant in each foam (γ lies in the ball $\eta(C_i(\widetilde{u}))$ and by the choice of e_i).
- The edges of γ' are not cyclic (Lemmas 5.8(3) and 5.21(3)).
- The label $L(E)$ appears as exactly one of the five labels incident to each edge (by the minimality of γ and the definition of the loosening move).

This completes the recursive step and so reduces the number of edges which are descendants of edges in $\mathcal{F}$.

When the loosening stage completes, the resulting foam $\mathcal{F}_J$ contains (by Hypothesis 6.6) two nature reserve edges. Furthermore, no edge of $\mathcal{F}_J$ is cyclic. This follows since the induction halts when no edge of $\mathcal{F}_J$ is a descendant of an edge of $\mathcal{F}_0$. Thus every such edge at some point is destroyed by either an augmented 2-3 move or a nature reserve move. By Lemma 5.8(3) and Lemma 5.21(3), these moves destroy all cyclic edges.

6.9. **Contacting stage.** The face f (that is, its descendants) survives all loosening moves. The same holds for the arc δ. Let s and s' be the two sides of δ (see Definition 4.9). Consider the set of complementary regions in $\widetilde{\mathcal{F}}_J$ that are incident to f, unioned with the sets of complementary regions in $\widetilde{\mathcal{F}}_J[\delta]$ that are incident to f_s and $f_{s'}$ (see Definition 4.10). The resulting union is finite. Since L has infinite image, there is a label ℓ different from the labels of all regions in that union. We choose, in the one-skeleton of $\widetilde{\mathcal{F}}_J$ minus $\widetilde{\mathsf{Nat}}_J$, a minimal path γ connecting a vertex of a region E, with $L(E) = \ell$, to a vertex of f. Our goal is to bring E into contact with f. To do this we extend our sequence of foams by performing further augmented 2-3 moves.

Fix a handle structure η' for $\mathcal{F}_J$. The foams $\mathcal{F}_{J+1}, \ldots, \mathcal{F}_K$ will be carried by $\mathcal{F}_J$. Let $v_0, v_1, \ldots, v_m$ be the vertices of γ in $\widetilde{\mathcal{F}}_J$. Let $e_1, e_2, \ldots, e_m$ be the edges of γ in $\widetilde{\mathcal{F}}_J$. We begin by performing an augmented 2-3 move along $\phi_M(e_1)$, from $\phi_M(v_0)$ to $\phi_M(v_1)$. The edge e_1 is L–flippable by the minimality of γ and because in the loosening stage we destroyed all cyclic edges.

We now proceed recursively. Assume that the previous augmented 2-3 move, producing the foam $\mathcal{G}$ (one of the $\mathcal{F}_i$), was the first to introduce the region E (and the label $L(E)$) to the zero-handle $\eta'(v_j)$. We now tunnel through $\eta'(v_j)$ in a process very similar to that described in Section 6.8. Here are some of the details. Let β_j be a minimal path in the one-skeleton of $\eta'(v_j) \cap \widetilde{\mathcal{G}}$ connecting a vertex incident to E with an endpoint of an edge e'_{j+1} that meets $\eta'(e_{j+1})$. Form β'_j by adding e'_{j+1} to the end of β_j. We perform augmented 2-3 moves along (the image under the covering map ϕ of) each edge of β'_j in turn. The last augmented 2-3 move in our sequence, along $\phi_M(e'_{j+1})$, is the first to introduce the region E (and the label $L(E)$) to the zero-handle $\eta'(v_{j+1})$.

This completes the recursive step. Once we have processed all edges of γ, we have produced a foam $\mathcal{H}$. The last augmented 2-3 move introduced, for the first time, the region E (and the label $L(E)$) to the zero-handle $\eta'(v_m)$, which also contains at least one vertex of f. Let β_m be a minimal path in $\eta'(v_m) \cap \mathcal{H}$ connecting a vertex incident to E with f. We perform augmented 2-3 moves along each edge of β_m in turn. Again this process is very similar to that described in Section 6.8.

6.10. **Unique contact.** Having processed through all of β_m we reach a foam in which the region E is now in contact with f. There are three possible positions for f in Figure 5.5F. These are the three sectors at the bottom of the figure. Suppose that more than one of these is f. In this case, before performing the last augmented 2-3 move we perform some extra V-moves to block off all but one of the sectors from E. That is, we perform V-moves at the terminal end of the edge and then perform the augmented 2-3 move to produce the final foam $\mathcal{F}_K$. Figure 6.11 shows the result when all three sectors are f. The effect is that the region E now meets f along precisely one edge. We denote this edge by e_E.

6.12. **Verifying hypotheses.** Following the loosening and contacting stages builds us a sequence of 2-3 and 3-2 moves taking us from $\mathcal{F}_0$ to $\mathcal{F}_K$ via L–essential foams. Applying Corollary 6.3 at each step builds us a sequence of 2-3 and 3-2 moves taking us from $\mathcal{F}_0[\delta]$ to $\mathcal{F}_K[\delta]$ via L–essential foams. What remains is to check that we can apply Lemma 4.16 to connect $\mathcal{F}_K$ to $\mathcal{F}_K[\delta]$.

Lemma 6.13. *The foam $\mathcal{F}_K$ satisfies the hypotheses of Lemma 4.16, taking E_k to be E.*

Proof. Recall that s is the side of δ containing e_E. Thus $\phi_M(s)$ contains $\phi_M(e_E)$ and the two nature reserve edges containing the ends of δ. Since $\phi_M(e_E)$ is not a nature reserve edge there are at least three edges on the side s.

By Lemma 5.8(4) and Lemma 5.21(4), the face f is always carried by its ancestor in $\widetilde{\mathcal{F}}_J$. This and the minimality of γ (Section 6.9) ensure that f does not become incident to E (or any other region with the same label as E) before the last move of our sequence. Thus the label $L(E)$ appears exactly once as the label of an outer region for f.

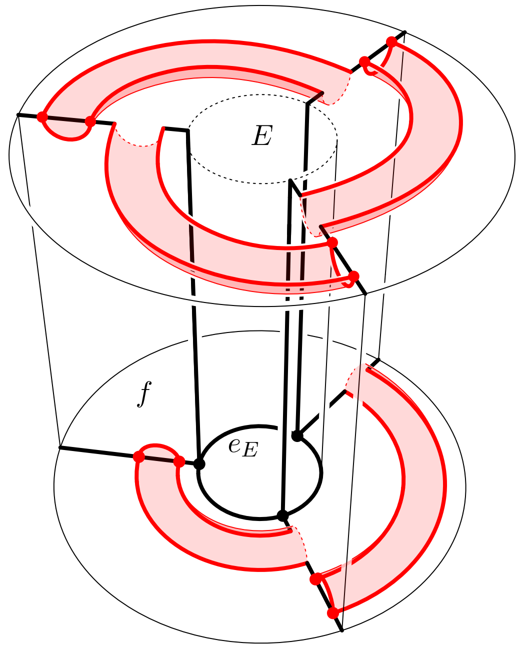

FIGURE 6.11. We apply extra V-moves to limit the number of edges where E meets f to one, namely the edge e_E.

Suppose that D is an outer region for the face f_s but not an outer region for f. Let e_D be the edge of f_s in $\widetilde{\mathcal{F}}_K[\delta]$ along which D meets f_s. Thus $\phi_M(e_D)$ is contained in a nature reserve edge e' in $\mathcal{F}_K$. (See Figures 5.12F and 4.15I as well as Remark 4.11.) Since e' is a nature reserve edge, augmented 2-3 moves taking us from $\mathcal{F}_J$ to $\mathcal{F}_K$ do not destroy it. Thus D is (a descendant of) one of the regions incident to the face f_s as it sat in $\widetilde{\mathcal{F}}_J[\delta]$. Our choice of E (made at the start of Section 6.9) then implies that $L(D) \neq L(E)$. Thus the label $L(E)$ appears exactly once as the label of an outer region for f_s.

The edge e_E and the two adjacent edges around f are not cyclic edges since we destroyed all cyclic edges in the loosening step and augmented 2-3 moves do not introduce them. (Also, the extra V-moves of Section 6.10 do not introduce any cyclic edges.) Finally, $\phi_M(e_E)$ is distinct from all other edges of $\phi_M(\partial f)$ because the other two faces of $\mathcal{F}_K$ incident to $\phi_M(e_E)$ are not equal to $\phi_M(f)$. This is because these two faces were created in the final augmented 2-3 move; thus they have no ancestors. □

We now apply Lemma 4.16 to $\mathcal{F}_K$. This gives the desired sequence of moves connecting $\mathcal{F}_K$ to $\mathcal{F}_K[\delta]$.

Remark 6.14. This completes the proof of Proposition 3.19 assuming Hypothesis 6.6: that is, that the endpoints of δ lie on distinct edges of $\mathcal{F}$. ◇

6.15. **Both endpoints of δ lie on one edge.** We now replace Hypothesis 6.6 by the following.

Hypothesis 6.16. The endpoints of δ lie on the same edge of $\mathcal{F}$. ◇

The proof of Proposition 3.19 proceeds as before except that when a loosening move is applied along the edge containing both endpoints of δ in the loosening stage we use the pair nature reserve move. The only piece of the proof remaining is the following.

Proof of Lemma 5.19. We now describe moves to get from $\mathcal{G}[\delta]$ to $\mathcal{G}_p[\delta]$ with all foams being L–essential. Here we do not use V-moves to create the three red snakelets and the green snakelet (as we did in the proof of Lemma 5.15). Instead we apply the special case of Proposition 3.19 that we proved under Hypothesis 6.6; see Remark 6.14. That is, we produce a sequence of five L–essential foams

$$\mathcal{G}[\delta] = \mathcal{G}_0, \quad \mathcal{G}_1, \quad \mathcal{G}_2, \quad \mathcal{G}_3 = \mathcal{G}'[\delta], \quad \mathcal{G}_4 = \mathcal{G}_p[\delta]$$

with the following properties.

- $\mathcal{G}'$ is as in Figure 5.18A.
- Each foam has an L–flippable edge.
- $\mathcal{G}_j[\epsilon_j] = \mathcal{G}_{j+1}$. That is, $\mathcal{G}_{j+1}$ is obtained from $\mathcal{G}_j$ by performing a 0-2 move along an arc ϵ_j.
- The endpoints of ϵ_j lie on distinct edges of $\mathcal{G}_j$. (This ensures that Hypothesis 6.6 is satisfied.)

See Figure 5.18B. The arcs ϵ_0 through ϵ_2 are the cores of the red snakelets of $\mathcal{G}_p$, while ϵ_3 is the core of the green snakelet. Let a, b, and c be the three edge-ends in $\mathcal{G}$ incident to p that are not part of e. We assume that a is not part of the same edge as either b or c. (It is possible that b and c are part of the same cyclic edge.) We choose the order of ϵ_0 through ϵ_2 as follows.

(1) Let ϵ_0 be the arc connecting a to b.
(2) Let ϵ_1 be the arc connecting a to c (with the end on a between p and the end of ϵ_0).
(3) Finally, let ϵ_2 be the arc connecting b to c (with the end on c between p and the end of ϵ_1, and the end on b not in the segment between p and the end of ϵ_0).

We now check the hypotheses of Proposition 3.19 (including Hypothesis 6.6) for each arc ϵ_j and for each foam $\mathcal{G}_j$.

Claim 6.17. $\mathcal{G}_0 = \mathcal{G}[\delta]$ is L–essential and has an L–flippable edge.

Proof. Recall that $\mathcal{G} = \mathcal{F}_i$ where the latter is a foam created during the loosening stage (Section 6.5). Under Hypothesis 6.16 we apply only one pair nature reserve move, which is part of a loosening move taking us from $\mathcal{F}_i$ to $\mathcal{F}_{i+1}$. There are two cases.

- Suppose that $i = 0$. That is, this is the very first loosening move. Then by the hypotheses of Proposition 3.19, the foam $\mathcal{F}_0[\delta] = \mathcal{F}[\delta]$ is L–essential and has an L–flippable edge.
- Suppose that $i > 0$. Then all previous loosening moves were augmented 2-3 moves (since there is only one pair nature reserve move). Thus $\mathcal{F}_i[\delta]$ is L–essential and has an L–flippable edge by Lemma 5.8(1) and (2).

This proves Claim 6.17. □

It follows from Claim 6.17 that $\mathcal{G}_j$ is L–essential for $j > 0$; each snakelet creates a new bigon face connecting two regions that were already in contact in $\mathcal{G}_0$. The 0-2 moves along ϵ_0 and ϵ_1 are in fact V-moves, so by two applications of Lemma 3.12, the foams $\mathcal{G}_1$ and $\mathcal{G}_2$ each have an L-flippable edge. A similar analysis shows that for each of the remaining two 0-2 moves, if an edge not entirely contained in $\mathcal{N}(p)$ is L-flippable then its descendant (that is also not entirely contained in $\mathcal{N}(p)$) is also L-flippable. (The existence of an L–flippable edge on $\mathcal{G}_3 = \mathcal{G}_p[\delta]$ also follows from Remark 5.20.)

Hypothesis 6.6 holds for ϵ_0 by construction (see (1) above). For both ϵ_1 and ϵ_2, one endpoint is on an edge contained within $\mathcal{N}(p)$ while the other is not contained within $\mathcal{N}(p)$. For ϵ_3, one endpoint is on an edge incident to p, while the other is not.

Thus the hypotheses of Proposition 3.19 and Hypothesis 6.6 hold for each arc. □

This completes the proof of Proposition 3.19. □

References

[1] Gennaro Amendola. A calculus for ideal triangulations of three-manifolds with embedded arcs. *Math. Nachr.*, 278(9):975–994, 2005. doi:10.1002/mana.200310285. [1]

[2] Craig D. Hodgson, J. Hyam Rubinstein, Henry Segerman, and Stephan Tillmann. Triangulations of 3-manifolds with essential edges. *Ann. Fac. Sci. Toulouse Math. (6)*, 24(5):1103–1145, 2015. doi:10.5802/afst.1477. [5]

[3] Tejas Kalelkar, Saul Schleimer, and Henry Segerman. Connecting essential triangulations I: via 2-3 and 0-2 moves, 2024. arXiv:2405.03539. [2, 3, 5, 8]

[4] Feng Luo, Stephan Tillmann, and Tian Yang. Thurston's spinning construction and solutions to the hyperbolic gluing equations for closed hyperbolic 3-manifolds. *Proc. Amer. Math. Soc.*, 141(1):335–350, 2013. doi:10.1090/S0002-9939-2012-11220-1. [5]

[5] S. Matveev. *Algorithmic Topology and Classification of 3-Manifolds.* Algorithms and Computation in Mathematics. Springer Berlin Heidelberg, 2007. https://books.google.com/books?id=vFLgAyeVSqAC. [1, 3, 5, 9, 13, 14, 17]

[6] Udo Pachner. Bistellare Äquivalenz kombinatorischer Mannigfaltigkeiten. *Arch. Math. (Basel)*, 30(1):89–98, 1978. doi:10.1007/BF01226024. [1, 5]

[7] Riccardo Piergallini. Standard moves for standard polyhedra and spines. Number 18, pages 391–414. 1988. Third National Conference on Topology (Italian) (Trieste, 1986). [1]