

Belted sum decompositions of fully augmented links

Porter Morgan, Brian Ransom, Dean Spyropoulos,
Rolland Trapp and Cameron Ziegler

ABSTRACT. Given two orientable, cusped hyperbolic 3-manifolds containing certain thrice-punctured spheres, Adams gave a diagrammatic definition for a third such manifold, their belted sum. Generalizing Adams' definition slightly, this work considers belted sum decompositions of fully augmented links (or FALs) in which all summands involved are also FALs. To do so, we provide explicit classifications of thrice-punctured spheres in FAL complements, making them easily recognizable. These classifications are used to characterize belted sum prime FALs geometrically, combinatorially and diagrammatically. Finally we prove that, in the context of belted sums, every FAL canonically decomposes into FALs which are either prime or two-fold covers of the Whitehead link.

CONTENTS

1. Introduction	1
2. Fully augmented links	5
3. Thrice-punctured spheres	11
4. Separating pairs	23
5. On b -prime FALs	31
6. Canonical FAL decompositions	37
References	41

1. Introduction

Within the context of orientable hyperbolic three-manifolds, Colin Adams in [1] uses embedded thrice-punctured spheres to form a sum of two cusped, orientable, hyperbolic three-manifolds as follows. First, slice along thrice-punctured spheres in each, resulting in manifolds whose boundaries contain two thrice-punctured spheres. Then glue the manifolds together along the boundary thrice-punctured spheres. Adams proves that the result is again hyperbolic and that volumes add under this operation (see [1, Theorem 4.5]). We

Received June 8, 2024.

2010 *Mathematics Subject Classification*. Primary: 57M25; Secondary: 57M27, 57M50.

Key words and phrases. hyperbolic fully augmented link, thrice-punctured spheres.

refer to the resulting manifold as a *punctured sphere sum* of the two original manifolds. A belted sum is a particular type of this operation (see [1, Corollary 5.2]), which Adams defines diagrammatically as in Figure 1. We extend this diagrammatic definition and allow belted sums to add a single crossing between the two strands (see Section 4 for details).




FIGURE 1. Forming the belted sum of M_1 and M_2

This paper studies the reverse of the belted sum operation on a particular class of hyperbolic links called fully augmented links, or FALs. Hyperbolic FALs are constructed from certain link diagrams by augmenting each twist region with an unknot by placing a trivial knot around each twist region and removing all full twists (see Figure 2, and Section 2 for more detail). FALs were introduced by Agol and Thurston in the appendix of [15], where they used FALs to improve Lackenby’s upper bound on volumes of hyperbolic alternating link complements. The geometry of an FAL complement can be described by gluing a right-angled ideal polyhedron to its reflection using products of reflections in its faces (see Section 2 or Purcell’s paper [19]). Since their introduction, the geometry of FAL complements has proven useful in a variety of contexts. Most Dehn fillings of the augmenting unknot components of an FAL yield hyperbolic links whose geometry can be related to that of the FAL. For example, FALs have been used to produce links that admit no exceptional surgeries, to bound the genera of surfaces in highly twisted link complements, to create diagrammatic bounds on volumes of sufficiently twisted links, and to bound crosscap numbers of alternating links (see [11], [5], [10], and [13], respectively). Thus FALs are a convenient and useful class of hyperbolic links to study.

In this work we focus on decomposing a given FAL complement into the belted sum of two simpler FALs, where the term “simpler” will develop to have meaning in topological, combinatorial, and geometric contexts. An FAL which cannot be so decomposed is called b -prime and, since several properties of manifolds are well-behaved under belted sums, a better understanding of b -prime FALs will shed light on FALs in general. We now highlight two potential applications.

First, since volumes are additive under belted sums, understanding volumes of b -prime FALs can be applied to the study of FAL volumes in general. Volumes of FALs, in turn, are of interest since they are related to (conjectured)

minimal volume n -cusped hyperbolic manifolds. Yoshida has shown that an FAL complement is the minimal volume orientable 4-cusped hyperbolic manifold (the fully twisted Borromean rings of Figure 20). Moreover, Agol [3] has conjectured that coverings of the Whitehead link complement are minimal volume for links with at least 11 components. Purcell proves (see [19, Proposition 3.6]) that $2(c-1)v_8$ is a sharp lower bound for the volume of an FAL with c crossing circles, where v_8 is the volume of a regular ideal octahedron, and there are octahedral FALs that satisfy Agol's conjectured minimal volume. These observations indicate that knowledge of FAL volumes, combined with the tractable geometry of FALs, could result in further applications.

Arithmetic invariants and, in particular, invariant trace fields, behave nicely under belted sum operations, providing a second application for b -prime FALs. Indeed, Section 5.6 of Maclachlan and Reid (see [16, Theorem 5.6.1]) shows that the invariant trace field of $M_1 \#_b M_2$ is the compositum of the fields of the summands M_i . Moreover, several authors have used the tractable geometry of FALs to study questions of arithmeticity and commensurability. Meyer-Millichap-Trapp computed invariant trace fields of minimally twisted chain links with an even number of components, and determined which such links are arithmetic and which are commensurable (see [18]). In addition, Rochy Flint ([9]) developed techniques for calculating invariant trace fields of FALs. Thus the explicit geometry of b -prime FALs combined with the belted sum operation can provide insight into invariant trace fields, arithmeticity, and commensurability questions for FALs.

It behooves us, then, to analyze belted sum decompositions of FALs. The first step in this analysis is to classify embedded, totally geodesic thrice-punctured spheres in FAL complements. Every FAL complement admits an orientation-reversing involution whose fixed point set is an embedded, totally geodesic *reflection surface*. We will see, in Lemma 4.7, that thrice-punctured sphere components of a reflection surface are never part of a belted sum decomposition. For this reason we focus on thrice-punctured spheres not contained in the reflection surface. Section 3 develops the necessary tools to classify such *non-reflection* thrice-punctured spheres, showing there are three types. The classification is given in Theorem 3.11, stated below. The three types of spheres are illustrated in Figures 15 and 17 (see also Definitions 3.7 and 3.10).

Theorem 3.11. *Let S be a non-reflection thrice-punctured sphere in an FAL complement $M = \mathbb{S}^3 \setminus \mathcal{A}$. If \mathcal{A} is not the fully twisted Borromean rings, then S is orthogonal to the reflection surface in M and is either a crossing, longitudinal or singly-separated disk.*

The three types of non-reflection thrice-punctured spheres—crossing disks, longitudinal disks, singly-separated disks—are classified by their punctures and how they intersect the reflection surface. Thus their geometry, and how it relates to the overall geometry of the FAL complement, is made quite explicit (see

Propositions 3.6 and 3.9). Moreover, they are easy to recognize using the tools developed in Section 2.

Section 4 determines which thrice-punctured spheres of Theorem 3.11 can be used in a belted sum decomposition. Section 5 applies results of Section 4 to give several characterizations of b -prime FALs. The background information of Section 2 is required to describe two of the characterizations given, so we focus on the third. There is a procedure to construct an FAL from a link diagram and FAL complements are (essentially) b -prime if and only if the diagram has only trivial flype orbits. The result is Theorem 5.4, stated below (technical definitions are postponed until Section 5).

Theorem 5.4. *Let D be a twist-reduced, non-split, prime link diagram with more than two twist regions, and let \mathcal{A} be a full augmentation of D . The fully augmented link \mathcal{A} is b -prime if and only if it induces only trivial flype orbits.*

We briefly mention, without detail, two other characterizations of b -primality, and refer the reader to Section 5 for precise statements. Theorem 5.2 is a combinatorial analogue of Theorem 5.4, and is stated in terms of perfect matchings on cubic planar graphs associated with FALs. Corollary 5.3 provides a geometric characterization of b -prime FALs stated in terms of the absence of certain thrice-punctured spheres. The combinatorial characterization provides an efficient method for tabulating b -prime FALs by number of crossing circles. All three characterizations of b -primality provide concrete approaches to the geometric and arithmetic questions described above.

An immediate application of the geometric characterization of Corollary 5.3 is the definition of a canonical belted sum decomposition in Section 6. Of the three types of non-reflection thrice-punctured spheres (crossing, longitudinal, and singly-separated disks), the canonical decomposition is obtained by decomposing only along crossing disks (see Definition 6.2). The result, stated here in Theorem 6.3, is that the summands of the canonical decomposition are either b -prime or 2-fold covers of the Whitehead link.

Theorem 6.3. *Let \mathcal{A} be an FAL with at least three crossing circles and let $M = \mathbb{S}^3 \setminus \mathcal{A}$ be its complement. Each summand in the canonical belted sum decomposition of M is either a b -prime FAL complement or the complement of the Borromean rings with at least one flat crossing circle.*

The paper is organized as follows. Section 2 reviews the definition of FALs, and the tools associated with them necessary for our study. Section 3 characterizes thrice-punctured spheres in FAL complements, and Section 4 classifies which can be used in belted sum decompositions. Section 5 gives the three characterizations of b -primality in FAL complements, and Section 6 introduces the canonical belted sum decomposition of an FAL.

2. Fully augmented links

Augmenting a link L consists of placing a trivial component around a twist region in a diagram of L . Adams, in [2], showed that augmenting an alternating hyperbolic link L results in another hyperbolic link. A fully augmented link \mathcal{A}_L (or FAL) is the result of augmenting every twist region and removing all full twists (see Figure 2(c)). Purcell relaxed the restriction that a diagram be alternating to show that an FAL created from any prime, twist reduced diagram with at least two twist regions is hyperbolic (see [20, Theorem 6.1]). In this paper we assume all FALs are hyperbolic. In addition, properties of a link and its complement will be referred to interchangeably, the meaning being clear from context.




FIGURE 2. Fully augmenting a link

This section reviews both geometric and combinatorial approaches toward studying FAL complements. The primary geometric tool will be a circle packing, unique up to Möbius transformation, which is the “footprint”, if you will, of a right-angled ideal polyhedron in hyperbolic space. The original FAL complement is made up of two such polyhedra. There are two graphs naturally associated with the circle packing—the nerve and its dual—which provide convenient combinatorial descriptions of FALs. Geometric and combinatorial tools will be used extensively throughout this paper, and we proceed with a more thorough description of them.

Let \mathcal{A}_L be the fully augmented link arising from L as described above. The additional trivial components are *crossing circles*, and the components that remain from the original link L will be called *knot circles*. A crossing circle with a half-twist is called *twisted*, while those without are *flat*. An FAL in which all crossing circles are flat will be called a *flat FAL*. Each crossing circle bounds a disk, called a *crossing disk*, which is punctured twice by knot circles. Crossing circles can bound more than one crossing disk, and when this occurs it always results in a possible belted sum decomposition (see Theorem 4.8). In what follows we assume a specific crossing disk is chosen for each crossing circle. Moreover, the complement $M = S^3 \setminus \mathcal{A}_L$ admits reflective symmetry across the projection plane, followed by a Dehn twist on twisted crossing circles. Purcell shows that the fixed point set of this reflection is an embedded, totally geodesic

surface $R \subset M$, called the *reflection surface* (see [19, Lemma 2.1]). If there is more than one reflection surface in an FAL complement we assume a specific choice of reflection surface has been made, just as a crossing disk is chosen for each crossing circle.

The complement $M = S^3 \setminus \mathcal{A}_L$ of every fully augmented link \mathcal{A}_L admits a *standard cell decomposition* \mathcal{C} . The standard cell decomposition is most easily seen in a flat FAL, so we begin with this case. There are two kinds of 2-cells in \mathcal{C} : the regions of the projection plane (e.g. U and V in Figure 3(a)) and the crossing disks (so \mathcal{C} depends on our initial choice of crossing disks). We refer to the projection plane 2-cells as *reflection 2-cells* since reflection across the plane of projection is a symmetry of the link \mathcal{A}_L . The 2-cells from crossing disks are *crossing 2-cells*. The projection plane cuts through the middle of each crossing disk, and the curves of intersection are the 1-cells of the standard cell decomposition (see the vertical and horizontal line segments of Figure 3(a)). Crossing disks are cut into two triangles by the 1-cells, one above and one below the plane of projection (e.g. A, A' in Figure 3(a)).




FIGURE 3. From cell decomposition to circle packing

There are no 0-cells in the standard cell decomposition as all 1-cells have endpoints on \mathcal{A}_L , which will become ideal points. The two 3-cells B_{\pm}^3 correspond to the regions of M above and below the projection plane.

Slicing along the reflection 2-cells separates M into two pieces which are reflections of each other, and crossing disks contribute a triangle in each half. Further slicing along the crossing disks yields the two three balls B_{\pm}^3 , and the boundary of B_+ (viewed from “inside” B_+) is pictured in Figure 3(b). Shrinking the arcs of the original link to vertices gives a cell decomposition on the boundary of B_+ . This cell decomposition can be checkerboard colored so that the unshaded faces correspond to reflection 2-cells, while shaded faces are all triangular and come from crossing disks.

To transition from topology to geometry, consider the boundary two-sphere of B_+ to be the one-point compactification of the plane which is also the boundary of the upper half space model \mathbb{H}^3 of hyperbolic space. Then an application

of Andreev’s Theorem (see [4]) implies the unshaded faces of the cell decomposition can be isotoped to form a circle packing. Note that, in this circle packing, two circles are tangent if the corresponding regions in the plane share a knot circle or are both punctured by the same crossing circle. For example, the circles corresponding to regions U and W are tangent at the vertex q corresponding to the crossing circle bounding A (see Figure 3(c)).

This circle packing determines a right-angled ideal polyhedron P_+ , providing the geometric structure on B_+ , which can be described as the intersection of half-spaces in \mathbb{H}^3 associated with 2-cells of ∂B_+ . Each reflection 2-cell U is bounded by a circle of the circle packing which, in turn, is the boundary of a plane $\mathcal{P}_U \subset \mathbb{H}^3$. Let H_U denote the half-space of $\mathbb{H}^3 \setminus \mathcal{P}_U$ whose boundary *does not* contain U . In Figure 3(c), for example, H_U is the half-space above \mathcal{P}_U while H_V is the half-space below \mathcal{P}_V . Each crossing 2-cell A of ∂B_+ is represented by a shaded interstice in the circle packing. The three vertices of ∂A determine a unique circle in the plane which bounds a hyperbolic plane \mathcal{P}_A . Let H_A denote the half-space of $\mathbb{H}^3 \setminus \mathcal{P}_A$ whose boundary *does not* contain A . The intersection of these half-spaces, taken over all 2-cells of B_+ , forms P_+ . Letting P_- be the corresponding polyhedron for B_- , we refer to the polyhedra P_{\pm} as *standard polyhedra* for the FAL \mathcal{A}_L .

Since the 3-cells B_{\pm}^3 are reflections of each other across the projection plane, the polyhedra P_{\pm} are reflections of each other in \mathbb{H}^3 across an unshaded face. Letting P_- be the reflection of P_+ across an unshaded face G , the polyhedron $\mathcal{F} = P_+ \cup P_- \subset \mathbb{H}^3$ is a fundamental domain for M . Figure 4(a) illustrates the fundamental domain \mathcal{F} for the complement of the Borromean rings in Figure 3(a) after transforming the vertex p to infinity.




FIGURE 4. Gluing patterns viewed from infinity in \mathbb{H}^3

The gluing maps on the faces of \mathcal{F} are quite explicit. To describe the gluings between unshaded faces we refer to faces, edges and vertices of \mathcal{F} that are reflections of each other as *corresponding* faces, edges or vertices, respectively. If F and F' are corresponding unshaded faces, then they are glued by

identifying corresponding edges and vertices. This is realized by the isometry $r_G \circ r_F : F \rightarrow F'$, where r_F and r_G represent reflections across the respective faces (recall that G is the shared face of P_{\pm}). Indeed, reflection in F fixes F point-wise while reflection in G identifies corresponding components of F and F' .

For a flat FAL, each shaded triangle on P_+ will be glued to an adjacent shaded triangle by the parabolic isometry that fixes their shared point (for example, the point q between the triangles labeled A in Figure 3). Similarly for shaded triangles on P_- .

The universal cover \tilde{M} of M is tessellated by the images of \mathcal{F} under the group generated by the gluing maps just described. The fundamental domain \mathcal{F} is an example of what we'll call a standard domain for M , since it was formed using its standard cell decomposition. A *standard domain* for M is the union of P_+ with any copy of P_- sharing an unshaded face.

To see how twisted crossing circles change the above description, note that a twisted crossing circle can be obtained from a flat one by slicing along the crossing disk and regluing with a half-twist. This glues the top and bottom of one crossing disk to the bottom and top of the other. Carrying the projection plane along in the process yields the reflection 2-cells which form a reflection surface for the twisted disk. The reflection is the homeomorphism defined by reflecting in the plane of projection followed by full twists inside each twisted crossing circle.

FALs that differ by half-twists will be called *half-twist partners*. Note that the polyhedra P_{\pm} are the same for the half-twist partners, but shaded triangles in opposite polyhedra are identified. Figure 4(a), for example, illustrates a fundamental domain for the Borromean rings in which the vertex p of Figure 3(c) is placed at infinity in the upper half-space model. Figure 4 parts (b) and (c) illustrate the gluing patterns for the Borromean rings' half-twist partners.

The standard polyhedra P_{\pm} give a nice geometric description of M and, following [19], we describe combinatorial descriptions that also arise from the circle packing. The *nerve* of a circle packing is the planar graph obtained by placing a vertex at the center of each circle and an edge through points of tangency. Since the shaded regions are triangular, the nerve of an FAL circle packing is always a triangulation of \mathbb{S}^2 . Figure 5 illustrates a FAL with reflection 2-cells labeled, the circle packing for P_+ , and its corresponding nerve.

The nerve Γ can be painted, as described in [19], to record which shaded triangles are glued or, equivalently, which ideal vertices correspond to crossing circle cusps. Glued shaded triangles must share an ideal vertex that corresponds to a crossing circle cusp, so simply paint edges in the nerve through these shared vertices as in Figure 5(c). Painting the nerve in this way results in a triangulation of \mathbb{S}^2 in which one edge of each triangle is painted.

Conversely, given a triangulation of \mathbb{S}^2 for which each triangle has one painted edge, one can construct a unique flat FAL from it. Each painted edge is replaced with a crossing circle that punctures the plane near its vertices. Knot circles lie





FIGURE 5. Painted nerve of an FAL

in the plane and connect midpoints of non-painted edges through the crossing circles just constructed. The result is a flat FAL. Of course, arbitrarily twisting the crossing circles leads to other FALs that give rise to the same painted nerve.


 FIGURE 6. Painted crushtacean Γ^*

The planar dual Γ^* of the nerve Γ provides a second combinatorial structure for studying FALs. Following Chesebro, DeBlois and Wilton, we call the planar dual Γ^* the *crushtacean* (see [8] and Figure 6). Vertices of Γ^* correspond to shaded triangles in the circle packing and edges to points of tangency between them.

Since the nerve Γ is a triangulation of \mathbb{S}^2 , the crushtacean Γ^* is a trivalent planar graph. A painting on Γ determines one on the crushtacean simply by painting edges that are dual to painted edges of Γ . Note that each vertex of Γ^* has one painted edge since the same is true of each face of Γ , so a painted crushtacean is simply a perfect matching on a trivalent planar graph. Moreover, vertices sharing a painted edge correspond to shaded triangles in P_+ that are glued together.

A painted nerve and crushtacean, then, correspond to each FAL. Conversely, Purcell shows in [19, Lemma 2.4] how to create an FAL from a painted nerve Γ . Although [19, Lemma 2.4] is stated in terms of painted nerves, the proof's construction resorts to crushtaceans. Given the painted crushtacean Γ^* , simply replace each painted edge with a crossing circle that links it once and smooth the adjacent unpainted edges as in Figure 7 (which is Figure 5 of [19]). The FAL resulting from this process determines the same polyhedra with appropriate

gluing pattern. Of course, one alters this process appropriately to accommodate half-twists partners with the same painted crushtacean Γ^* .



FIGURE 7. Creating a FAL from a painted crushtacean

Features of the circle packing and crushtacean will be useful in characterizing thrice-punctured spheres and belted sums of FAL complements. We introduce two that will be particularly useful.

A shaded triangular interstice in the circle packing corresponds to a shaded triangular face of a standard polyhedron, say P_+ . The shaded face is an ideal triangle which is orthogonal to the unshaded faces of P_+ that it intersects, and which projects to half a crossing disk in the FAL complement. Another way to describe this is that three mutually tangent circles that bound an interstice correspond to an ideal triangular face of P_+ that projects to (a portion of) a crossing disk. More generally, we now observe that three mutually tangent circles in the circle packing correspond to an ideal triangle in P_{\pm} that intersects unshaded faces orthogonally. Later, Proposition 3.6 will demonstrate that these always project to thrice-punctured spheres in the FAL complement.

Let C_1, C_2, C_3 be three mutually tangent circles in the circle packing, let p, q, r denote their points of tangency, and let C^* the circle determined by p, q, r . We first show that the hyperbolic plane bounded by C^* is orthogonal to those determined by the C_i . Apply a Möbius transformation to the four circles that takes p to infinity and use the same labels to refer to the transformed figures. Then C_1 and C_2 are parallel lines and C_3 is a circle tangent to each at q, r , so that segment qr is a diameter of C_3 . The transformed C^* , then, is the line \overrightarrow{qr} , which is orthogonal to the C_i and the original C^* is orthogonal to the original C_i . Thus C^* bounds a plane $\mathcal{P} \subset \mathbb{H}^3$ and $\mathcal{P} \cap P_{\pm}$ is an ideal triangle orthogonal to unshaded faces of P_{\pm} . If the C_i bound a shaded triangular interstice in the circle packing, then $\mathcal{P} \cap P_{\pm}$ is a shaded face of P_{\pm} , and we say the C_i form a standard triple of mutually tangent circles. Otherwise, the circles C_1, C_2, C_3 are a *non-standard triple* of mutually tangent circles (see Figure 8 parts (a) and (b) for two views).

Non-standard triples admit a convenient combinatorial description as well. The three points of tangency in any set of three mutually tangent circles correspond to three edges e_1, e_2, e_3 in the crushtacean Γ^* which form a 3-edge cut. If the set is a non-standard triple, then the three edges are not adjacent, and neither component of $\Gamma^* \setminus \{e_1, e_2, e_3\}$ is a single vertex. A *non-trivial 3-edge cut* of Γ^* will be one in which neither component of $\Gamma^* \setminus \{e_1, e_2, e_3\}$ is a single vertex (see Figure 8(c)). Thus every non-standard triple in the circle packing corresponds to a non-trivial 3-edge cut of the crushtacean, and observe that the converse is also true.




FIGURE 8. A non-standard triple and corresponding non-trivial 3-edge cut

Two upcoming results are easily stated in terms of non-trivial 3-edge cuts. First, Corollary 3.8 shows that non-trivial 3-edge cuts in Γ^* correspond to one class of thrice-punctured spheres in FAL complements. Moreover, Theorem 5.2 shows that an FAL complement is belted sum prime if and only if every non-trivial 3-edge cut of the cruschacean is fully painted (i.e. all edges in the cut are painted).

3. Thrice-punctured spheres

This section classifies embedded, totally geodesic, thrice-punctured spheres in FAL complements. Adams proves that every essential, embedded thrice-punctured sphere in an orientable hyperbolic 3-manifold is totally geodesic (see [1]), so assume all thrice-punctured spheres in this work are such. We begin by specifying two closely related classes of them, non-standard and non-reflection disks.


Recall that, by convention, an FAL comes with a specific choice of reflection surface and a choice of crossing disk for each crossing circle relative to it. Such choices lead to the standard cell decomposition \mathcal{C} and standard polyhedra P_{\pm} of Section 2. Thrice-punctured spheres in \mathcal{C} will be called *standard disks*. The chosen crossing disks are standard, as are any thrice-punctured sphere components of the reflection surface. These are special since they form faces on P_{\pm} . *Non-standard disks* are thrice-punctured spheres not in \mathcal{C} . The second class of disks, called *Non-reflection disks*, are thrice-punctured spheres that are not part of the reflection surface of an FAL complement. Thus non-reflection disks are the non-standard disks together with the chosen crossing disks.

The main result of this section is Theorem 3.11, which combines results of this section into a classification of non-reflection disks in an FAL complement. The classification is based on two properties: the intersection of non-reflection disks with the reflection surface and the slopes of their punctures along components of the FAL. We begin by describing possible intersections with the reflection surface.

Both the reflection surface and a non-reflection disk are embedded, totally geodesic surfaces. If two such surfaces intersect non-trivially, they do so in a union of pairwise disjoint, simple geodesics on each. Thrice-punctured spheres have only the six simple geodesics pictured in Figure 9(a). Those labeled a, b, c are *non-separating* geodesics since slicing along one does not separate the sphere. Slicing along x, y or z separates S , so they are *separating* geodesics. We shall see that, if \mathcal{A}_L is any FAL other than the fully-twisted Borromean rings, then a non-reflection disk S is always orthogonal to the reflection surface R . The geodesics in $S \cap R$ help distinguish non-reflection disks. Propositions 3.4, 3.6, and 3.9 will show that, if \mathcal{A}_L has at least three crossing circles, then $S \cap R$ will either consist of all the non-separating geodesics on S (termed n -disks), or exactly one separating geodesic of S (called s -disks).

Each puncture of S induces a slope on a torus neighborhood of one component of \mathcal{A}_L , which also helps classify non-reflection disks. Recall that a *slope* is the isotopy class of an unoriented simple closed curve on a torus. Given a component J of \mathcal{A}_L , let $T_J \subset \mathbb{S}^3$ denote the torus boundary of a tubular neighborhood $V(J)$ of J . Since J is a knot in \mathbb{S}^3 , there is a natural meridional slope m and a longitudinal slope ℓ on T_J , determined by the following properties: m links J once and intersects ℓ once, while ℓ is homologous to J in $V(J)$ and null-homologous in $\mathbb{S}^3 \setminus V(J)$ (see [6, Theorem 3.1]). We refer to any slope of the form $\ell \pm n \cdot m$ as a *generalized longitude*. Geometrically, J corresponds to a cusp of M , and we think of T_J as the boundary of a neighborhood of this cusp. If an embedded totally geodesic surface F is punctured by J , we frequently refer to the puncture by the slope $F \cap T_J$. For example, a crossing disk of M has two meridional knot circle punctures and one longitudinal crossing circle puncture. Alternatively, consider a component R_1 of the reflection surface. It has (generalized) longitudinal punctures along knot circles, and two types of punctures along a crossing circles, depending on whether the crossing circle is flat or twisted. If a flat crossing circle punctures R_1 , it does so in meridian(s) of the crossing circle. On the other hand, placing a half-twist in a crossing circle connects the two meridians, forming a single puncture of R_1 consisting of one longitude plus two meridians. These observations will be useful in proving Lemma 4.7. Types of punctures, then, also help classify non-reflection disks.

The main tool used for studying non-reflection disks will be their pre-images in a standard domain \mathcal{F} . We begin by describing the preimages of chosen crossing disks, the simplest non-reflection disks. Let D be a crossing disk in an FAL complement, and consider how geodesics in a standard domain \mathcal{F} project to the simple geodesics on D . Observe that the non-separating geodesics on D are the one-cells of \mathcal{C} since they join distinct punctures. Thus the edges of \mathcal{F} project to non-separating geodesics on crossing disks. Separating geodesics of D intersect non-separating geodesics orthogonally, so altitudes of shaded triangles on $\partial\mathcal{F}$ project to form separating geodesics in D (here an altitude of an ideal triangle is the geodesic ray orthogonal to one side and toward the opposite vertex). The


 FIGURE 9. Neighborhood of C in P_+

altitudes labeled \tilde{y} in Figure 9(b), for example, project to the separating geodesic y of Figure 9(a). The preimage of a crossing disk in \mathcal{F} , then, lies in the boundary of a standard domain and is well understood.

For other non-reflection disks we apply techniques of Knafl and Trapp in [14], who studied embedded, totally geodesic surfaces in FAL complements by analyzing their pre-images in a standard domain \mathcal{F} . We review some relevant results from [14], and begin with their definition of a geodesic disk.

Definition 3.1. Let $S \subset M$ be an embedded totally geodesic surface in the FAL complement M , and let \tilde{S} denote its preimage in the universal cover. A *geodesic disk* is a connected component of $\tilde{S} \cap P_+$, or of $\tilde{S} \cap P_-$.

Further, faces of the polyhedra P_{\pm} are *standard geodesic disks* since they project to the standard cell decomposition \mathcal{C} . A *non-standard geodesic disk* is one that is not a face of P_{\pm} .

Since S is an embedded totally geodesic surface, its preimage \tilde{S} in the universal cover is a union of disjoint planes in \mathbb{H}^3 . Thus a geodesic disk is the intersection of a plane \mathcal{P} in \mathbb{H}^3 with either P_+ or P_- . The boundary γ of a geodesic disk S'_+ in P_+ , say, is a planar polygonal curve given by $\gamma = \mathcal{P} \cap \partial P_+$. For example, Figure 10(a) depicts a non-reflection disk in M , with some planes of its preimage \tilde{S} in the universal cover found in Figure 10(b). The geodesic disks S'_{\pm} are highlighted in Figure 10(c), both quadrilaterals with two ideal and two material vertices.

The boundary of a non-standard geodesic disk can intersect a shaded triangle T of ∂P_{\pm} only along an edge or altitude. To see this, let D be a crossing disk and recall that its preimage in ∂P_{\pm} consists of two pairs of shaded triangles (see Figure 9(b)). Edges of ∂P_{\pm} project to the non-separating, and altitudes of shaded triangles project to the separating, geodesics of D . Now a non-reflection sphere S is an embedded, totally geodesic surface, as are components of \mathcal{C} . Thus S must intersect each component of \mathcal{C} in a collection of disjoint, simple geodesics




FIGURE 10. Non-reflection sphere S , its \tilde{S} , and geodesic disks

on each surface. In particular, if $S \cap D$ is non-empty, then it is a collection of disjoint separating and non-separating geodesics on each surface. But then $\mathcal{P} \cap T$ must be an edge or altitude of T .


With these initial observations in hand, we classify the types of geodesic disks that result from non-reflection disks. Area considerations limit the types of geodesic disks that can arise from a non-standard thrice-punctured sphere. Recall that the area of a hyperbolic polygon with external angles ϵ_i is $\sum_{i=1}^n \epsilon_i - 2\pi$, with the convention that $\epsilon = \pi$ for ideal vertices. We begin by demonstrating the general observation that the area of a geodesic disk in an FAL fundamental domain is always an integer multiple of π .

Lemma 3.2. *Let $S' \subset P_{\pm}$ be a non-standard geodesic disk for the embedded totally geodesic surface S in an FAL complement. Then*

- i. *Material vertices of $\gamma = \partial S'$ are right-angled,*
- ii. *Material vertices can be paired so that their adjacent altitudes share an ideal vertex, and*
- iii. *The area of S' is an integer multiple of π .*

Proof. Let $S' \subset P_{\pm}$ be a non-standard geodesic disk for the embedded totally geodesic surface $S \subset M$, with boundary curve $\gamma = S' \cap \partial P_{\pm}$.

Every edge of P_{\pm} bounds one unshaded face and one shaded triangle. If q is a material vertex of γ , then one edge $\gamma_0 \subset \gamma$ intersects the interior of a shaded triangular face $T_0 \subset \partial P_{\pm}$ (see Figure 11). As above, the geodesic ray γ_0 is an altitude of T_0 and projects to (half of) a separating geodesic on some crossing disk $D \subset M$. This is precisely the situation of [14, Lemma 4.2.i.], which shows that S' is orthogonal to both faces of ∂P_{\pm} incident with q (labeled T_0 and R in Figure 11). Then q is the intersection of three pairwise orthogonal planes in \mathbb{H}^3 , one of which contains S' , so the angle of S' at q is a right angle.


 FIGURE 11. Geodesic disk S' with material vertex q

To see that material vertices come in pairs, let p be the ideal “endpoint” of γ_0 , and let T_1 be the shaded triangle in ∂P_{\pm} that shares p with T_0 . Now apply a Möbius transformation to \mathcal{F} so that p is at infinity. The geodesic disk S' is orthogonal to T_0 and P_{\pm} is right-angled, so S' is orthogonal to the face T_1 opposite to T_0 at p . Thus $S' \cap T_1$ is an altitude of T_1 with one material endpoint q' , and there are an even number of material endpoints on γ . Further, the altitudes with base material vertices q, q' share the ideal vertex p on γ .

Since each ideal vertex of γ contributes π to the area of S' and there are an even number of material vertices, each contributing $\pi/2$, the area of S' is an integer multiple of π . \square

Lemma 3.2 has the following immediate corollary.

Corollary 3.3. *If S is an embedded, totally geodesic thrice-punctured sphere in an FAL complement, then S has either two geodesic disks of area π or one of area 2π .*

Proof. Let S is an embedded, totally geodesic thrice-punctured sphere in an FAL complement. Then S has area 2π , so Lemma 3.2 (iii) implies it either has two geodesic disks of area π or one of area 2π . \square

The first proposition shows that there is a unique FAL complement that contains a non-reflection disk with a single area 2π geodesic disk.

Proposition 3.4. *If S is a non-standard disk in the FAL complement M with a single geodesic disk of area 2π , then M is the fully twisted Borromean rings complement.*

Proof. Suppose that S' is a non-standard geodesic disk with area 2π that projects to a non-reflection disk S . Up to a reflection of M we can suppose $S' \subset P_+$. If q were a material vertex of $\partial S'$, the unshaded face adjacent to q contains an edge of $\partial S'$ which glues S' to another geodesic disk. Since S' is the only geodesic

disk for S , this implies all of its boundary vertices must be ideal. Thus $\gamma = \partial S'$ is an ideal square (see Figure 12(a)).




FIGURE 12. Area 2π geodesic disk

The edges of γ lie on distinct faces since the plane containing S' intersects a face of ∂P_+ in at most one geodesic arc. In addition, vertices of P_+ are 4-valent and checkerboard colored, so the faces of ∂P_+ on one side of γ alternate between shaded triangles and unshaded faces as γ is traversed. Let A and B be the shaded triangles of P_+ adjacent to one side of γ , and let U, V denote the unshaded faces adjacent to the same side (see Figure 12(b)). Both U and V share an edge with the shaded triangle A , so they share the vertex of A not on γ . Similarly, they share the third vertex of B . Now U and V share at most one ideal vertex (since their corresponding circles in the circle packing are tangent at most once), so the third vertices of shaded triangles A and B coincide. Thus one side of γ consists solely of the four ideal triangles A, U, B, V .

The same argument applies to the other side of γ , making P_{\pm} regular ideal octahedra. Then \mathcal{F} is the standard domain for the Borromean rings or its half-twist partners. Moreover, \mathcal{F} can be chosen so that S' is on the vertical half plane that intersects P_+ along a diagonal as in Figure 12(c).

To see that the gluing pattern on \mathcal{F} is that of the fully twisted Borromean rings (see Figure 4(c)), apply a Möbius transformation to \mathcal{F} so that the two vertical edges of $S' \subset P_+$ are identified as in Figure 12(a). The vertical shaded triangles on the left of \mathcal{F} are glued to ones on the right by parabolic isometries that preserve infinity. Since the gluing maps also identify the vertical edges of γ , the faces labeled B in Figure 12(c) must be identified, resulting in a twisted crossing disk. A similar argument shows the gluing pattern on A faces results in a twisted crossing disk as well, finishing the proof. \square

Ideal square geodesic disks occur for more general embedded, totally geodesic surfaces in FAL complements. The point of Proposition 3.4, however, is that in most cases they are glued to other geodesic disks, giving the resulting

surfaces too much area to be a thrice-punctured sphere. See [14, Example 5.2] for examples of surfaces with ideal square geodesic disks in flat, 2-bridge FAL complements. The proof of Lemma 4.6 contains a more thorough discussion of thrice-punctured spheres in the fully-twisted Borromean rings complement.

Now consider area π geodesic disks. An area π geodesic disk must have exterior angle sum 3π , and it either has material vertices or not. Lemma 3.2 showed that material vertices are right-angled and come in pairs, with an ideal vertex between them, thus contributing 2π to the angle sum. The area of a geodesic disk with more than two material vertices, then, is at least 3π and cannot project to a thrice-punctured sphere. Thus an area π geodesic disk that projects to a non-standard disk will either be an ideal triangle (with no material vertices), or a quadrilateral whose boundary contains two material vertices separated by two ideal vertices. The term *material-vertex disk* refers to an area π geodesic disk with two material and two ideal vertices (see Figure 10(c) for an example of material-vertex disks), while a *non-standard ideal triangle* is an ideal triangle geodesic disk that is not a face of P_{\pm} .




FIGURE 13. Area π geodesic disks in the circle packing

A non-standard geodesic disk results from intersecting P_{\pm} with a plane \mathcal{P} in \mathbb{H}^3 , where \mathcal{P} projects to a non-standard disk in an FAL complement. We describe area π disks in terms of the circle packing for P_{\pm} and the boundary circle for \mathcal{P} (see Figure 13). The circle packing for P_{\pm} must have three mutually tangent circles for a non-standard ideal triangle to exist. Moreover, the interstices (regions U and v in Figure 13(a)) between the circles must contain non-trivial portions of the circle packing, otherwise the ideal triangle $\mathcal{P} \cap P_{\pm}$ is a face of P_{\pm} . Thus non-standard ideal triangles correspond to the non-standard triples of Section 2. The four circles involved in a material-vertex disk (two at each ideal point of its boundary) are mutually tangent in the circle packing, forming a Descartes quadruple. The interstices U and v of Figure 13(b) may contain non-trivial portions of the circle packing, while T and T' are shaded triangular faces of P_{\pm} containing the altitudes in the boundary of the disk.

The above discussion is summarized in the following characterization of area π non-standard geodesic disks.

Lemma 3.5. *Let M be an FAL complement with standard polyhedra P_{\pm} , and let S' be an area π geodesic disk for a non-standard disk $S \subset M$.*

Then S' is orthogonal to ∂P_{\pm} and is either a non-standard ideal triangle or a material-vertex disk.

Proof. The preceding discussion restricts area π disks to non-standard ideal triangles or material-vertex disks, so it must be shown that both types are orthogonal to the standard cell decomposition. The discussion leading up to the definition of non-standard triples shows that a non-standard ideal triangle meets \mathcal{C} orthogonally. Further, the boundary of a material-vertex disk contains two altitudes of shaded triangles, and [14, Lemma 4.2.i.] states that any geodesic disk whose boundary contains an altitude meets \mathcal{C} orthogonally. \square

Area π disks can also be recognized in the nerve and crushtacean of an FAL. Indeed, we have already seen that a non-standard triple corresponds to a non-trivial 3-edge cut in the crushtacean. A material-vertex disk corresponds to a K_4 subgraph in the nerve, two of whose regions are faces of the triangulation while the other two regions may contain additional faces in the triangulation. This corresponds to two 3-edge cuts in the crushtacean that share an edge (the edge between U and v in Figure 14(b)).




FIGURE 14. Recognizing disks from graphs

The following proposition characterizes non-standard disks whose geodesic disks are ideal triangles. It is interesting to note that every non-standard ideal triangle and its reflection glue to form a non-standard disk. In contrast, whether or not material-vertex disks project to non-standard disks depends on how faces of standard polyhedra are identified.

Proposition 3.6. *Let $M = \mathbb{S}^3 \setminus \mathcal{A}$ be an FAL complement with reflection surface R and standard domain $\mathcal{F} = P_+ \cup P_-$.*

The surface $S \subset M$ is a non-standard disk disjoint from chosen crossing disks if and only if it is the projection of a non-standard ideal triangle and its reflection in P_{\pm} .

Moreover, in this case S satisfies:

- i. S intersects R orthogonally in its three non-separating geodesics.*

- ii. *The punctures of S are either one crossing-circle longitude and two knot-circle meridians, or three longitudes on distinct crossing circles.*

Proof. Let M be an FAL complement and S a non-standard disk in M that is disjoint from the crossing disks. Then S , by definition, is not in the standard cell decomposition \mathcal{C} and geodesic disks for S must be non-standard as well. Moreover, since S is disjoint from crossing disks, its geodesic disks will be disjoint from shaded faces of P_{\pm} . By Proposition 3.4 and the definition of material-vertex disks, area 2π geodesic disks and material-vertex disks always intersect shaded faces of P_{\pm} ; therefore, the geodesic disks for S must be two non-standard ideal triangles S_{\pm} .

Conversely we must show that each non-standard ideal triangle and its reflection, denoted S'_{\pm} , project to a non-standard disk $S \subset M$. Recall that the interstices U and v of Figure 13(a) contain additional circles of the circle packing since, otherwise, S'_{\pm} would be standard. Thus the boundary edges of S'_{\pm} are interior to unshaded faces of P_{\pm} and the disks S'_{\pm} are disjoint from shaded faces of P_{\pm} . Then boundary edges of S_{\pm} are identified with their reflections on S_{\mp} when projecting to M . By Lemma 3.5 the disks S'_{\pm} are orthogonal to ∂P_{\pm} , so they project to an embedded, totally geodesic surface $S \subset M$ which is disjoint from chosen crossing disks. Moreover, the surface S is two ideal triangles with edges glued to their reflections, resulting in a thrice-punctured sphere and proving the first claim of the proposition.

We now verify the stated properties of S . Note that slicing the projection of $S'_+ \cup S'_-$ along one geodesic of $\partial S'_{\pm}$ does not separate it; therefore S intersects the reflection surface along its non-separating geodesics. Moreover, we've already noted that S'_{\pm} are orthogonal to ∂P_{\pm} , so S is orthogonal to the reflection surface.

We now address the punctures of S . Since S is disjoint from all crossing disks, and every cusp is intersected by some crossing disk, it will have the same meridional and longitudinal punctures as a crossing disk. Hence knot circle punctures of S will be meridians while crossing circle punctures are longitudinal.

To see that crossing circle punctures of S must be on distinct cusps note that exactly one vertex of P_+ projects to each crossing circle cusp of M (and similarly for P_-). Thus vertices of S'_{\pm} that project to crossing circle cusps correspond to distinct crossing circles as desired.

It remains to show that S has an even number of knot circle punctures. If all punctures of S were meridians of knot circles, then S would be a sphere in S^3 punctured three times by a link, which is impossible. Similarly, if S had exactly one knot circle puncture, then the other two punctures would be longitudes of crossing circles. Gluing the two crossing disks to S would result in a sphere in S^3 punctured five times by a link, arriving at the same contradiction. Thus S has an even number of knot circle punctures, completing the proof. \square

Proposition 3.6 gives a very precise description of non-standard disks which are disjoint from crossing disks. Notice that, like crossing disks, they intersect the reflection surface orthogonally in their non-separating geodesics. For this

reason we refer to them as n -disks, the “ n ” emphasizing this observation. There are two kinds of n -disks though, depending on their punctures, and we make the following formal definition.

Definition 3.7. Let M be an FAL complement with chosen reflection surface R . An n -disk is a non-reflection disk $S \subset M$ is one that intersects R orthogonally along its three non-separating geodesics. There are two types of n -disks:

- i. A *longitudinal disk* is an n -disk that has three longitudinal punctures on distinct crossing circles (see Figure 15(a)).
- ii. A *crossing disk* is an n -disk that has one crossing-circle longitudinal and two knot-circle meridional punctures (see Figure 15(b)).

Note that Definition 3.7(ii) generalizes the term “crossing disk”. In fact, generalized crossing disks are potential crossing disks for the crossing circle. Generalized crossing disks will play a crucial role in classifying belted sum decompositions of FALs, as well as in defining a canonical one.




FIGURE 15. Two types of n -disks

The fact that non-standard ideal triangles can be recognized in circle packings and crushtaceans leads immediately to the following corollary (see Figures 13(a) and 14(a)).

Corollary 3.8. *Let $M = \mathbb{S}^3 \setminus \mathcal{A}$ be an FAL complement. Non-standard disks in M which are disjoint from crossing disks are in one-to-one correspondence both with non-standard triples in the circle packing and with non-trivial 3-edge cuts of the crushtacean.*

Furthermore, (non-standard) crossing disks correspond to non-trivial 3-edge cuts with one painted edge while longitudinal disks correspond to those with all edges painted.

Proof. The follows from simply interpreting Proposition 3.6 in terms of the circle packing and crushtacean. Do recall that painted edges in the crushtacean represent vertices of P_{\pm} that correspond to crossing circles, so the number of

painted edges in the 3-edge cut will correspond to the number of crossing circle punctures on the disk. \square

One slices along thrice-punctured spheres when performing belted sum decompositions in FAL complements. In a standard domain this corresponds to slicing along geodesic disks. Therefore, it will be helpful to analyze how non-standard ideal triangles and their reflections separate standard domains. Figure 16 illustrates a standard domain $\mathcal{F} = P_+ \cup P_-$ containing the union S' of a non-standard ideal triangle and its reflection. The shaded regions labeled U and V contain additional circles in the circle packing. We can assume that the vertex at infinity corresponds to a crossing circle cusp, since either one or all vertices of S' correspond to crossing circle cusps by Proposition 3.6.




FIGURE 16. Slicing a standard domain along an n -disk

Slicing \mathcal{F} along S' yields two components as in Figure 16, and we denote the two components of $\mathcal{F} \setminus S'$ by $\mathcal{F}_1, \mathcal{F}_2$. The faces of the components inherit gluing instructions from those of \mathcal{F} , which we now analyze. Unshaded faces of P_{\pm} glue to their reflections, and boundary edges of S' are identified under these gluing maps, so an unshaded face of one component \mathcal{F}_i glues to another on the same component. Shaded triangles are paired with adjacent ones to determine gluings. Shaded triangles which share a vertex on $\partial S'$, then, are the only ones which can identify faces between components. In the case of a longitudinal disk, all shaded triangles sharing a vertex of S' are paired, providing three identifications between faces of \mathcal{F}_1 and \mathcal{F}_2 . On the other hand, if S' projects to a generalized crossing disk, only infinity corresponds to a crossing circle and the vertical shaded faces of \mathcal{F} are the only shaded faces identified between the two components.

Having characterized non-reflection thrice-punctured spheres with ideal triangle geodesic disks, we turn our attention to material-vertex disks. Material-vertex disks have two altitudes in their boundaries which share an ideal vertex. The next proposition shows that a material-vertex disk (together with its reflection) projects to a non-standard disk if and only if the ideal vertex shared by these altitudes corresponds to a flat crossing circle. In this case, the remaining ideal vertex corresponds to a crossing circle as well.

Proposition 3.9. *Let $M = \mathbb{S}^3 \setminus \mathcal{A}$ be an FAL complement with standard domain $\mathcal{F} = P_+ \cup P_-$.*

The surface $S \subset M$ is a non-reflection thrice-punctured sphere that separates a crossing disk D if and only if S is the projection of a material-vertex disk and its

reflection in P_{\pm} whose altitude edges project to a separating geodesic on the flat crossing disk D .

In this case, S intersects the reflection surface orthogonally along one of its separating geodesics. Moreover, the punctures of S are a longitude of one crossing-circle and two meridians of a different crossing-circle.

Proof. Let $S \subset M$ be a non-reflection thrice-punctured sphere that intersects the crossing disk D along one of its separating geodesics γ_D . Since the preimage in P_{\pm} of γ_D consists of altitudes of shaded triangles, at least one geodesic disk has material vertices, call it S'_+ . Then S'_+ is a material-vertex disk, and must glue to its reflection S'_- along the rays interior to unshaded faces in P_{\pm} . Area considerations imply these are the only two geodesic disks for S , so the altitudes must identify as well. If D is a twisted crossing disk, the identification of S'_{\pm} results in a non-orientable surface. Therefore D is a flat crossing disk.

Conversely, suppose the material vertex disks S'_{\pm} are reflections of each other, and their altitudinal boundary rays project to a separating geodesic on a flat crossing disk D . This implies the altitudes of $\partial S'_+$ are identified, and the ideal vertex they share corresponds to the crossing circle C which bounds D (and similarly for the altitudes of $\partial S'_-$). The punctures of S'_{\pm} shared by altitudes, then, project to two distinct meridional punctures of the crossing circle C . As in the proof of Lemma 3.2, applying [14, Lemma 4.2.i.] we see that S'_{\pm} is orthogonal to all faces of ∂P_{\pm} that they meet. Identifying edges, then, results in an embedded, totally geodesic surface S in M . Moreover, the gluing pattern (altitudes glue within and other rays glue between geodesic disks) is that of a thrice-punctured sphere. Note that first gluing altitudes within each S'_{\pm} does not connect the disks. Thus the unshaded boundary rays project to a separating geodesic of S , which is its intersection with the reflection surface of M .

It remains to show that the final puncture is a longitude of a crossing circle. The remaining ideal vertex of S'_+ joins two geodesic rays interior to unshaded faces. Therefore it and its reflection identify to either a meridian of a knot circle or a longitude of a crossing circle. If it were the meridian of a knot circle, then S would be a sphere in \mathbb{S}^3 punctured three times by the link \mathcal{A} (twice by C and once by the knot circle), which is impossible. Therefore, the remaining puncture is the longitude of a crossing circle. \square

The situation of Proposition 3.9 is illustrated both in Figures 10(a) and 17(a) and, as they appear frequently in the remainder of the paper, we make the following definition.

Definition 3.10. Let M be an FAL complement with chosen reflection surface R . A *singly-separated disk*, or *s-disk*, is a non-reflection disk $S \subset M$ satisfying:

- i. S intersects R orthogonally,
- ii. $S \cap R$ is one separating geodesic on S , and
- iii. S has two meridional punctures along one crossing circle and one longitudinal puncture along a different crossing circle.

Figures 13(b) and 14(b) illustrated recognizing material-vertex disks from the circle packing and crushtacean. The proof of Proposition 3.9 justifies the edge-paintings of Figure 14(b), showing that the vertices corresponding to painted edges do correspond to crossing circles. Transferring vertex p of Figure 13(b) to infinity via a Möbius transformation yields Figure 17(b), from which it is easier to see that slicing along geodesic disks for an s -disk separates P_{\pm} into two pieces. Moreover, the only gluing of faces between the two pieces of $P_{+} \setminus S'_{+}$ is along the vertical shaded triangular faces. This is because the reflection between P_{\pm} preserves the pieces of $P_{\pm} \setminus S'_{\pm}$ and unshaded faces glue to their reflections. Further, shaded triangular faces are identified in adjacent pairs and, since the shaded faces intersected by S'_{\pm} project to a flat crossing disk, the only pair separated by slicing along S'_{\pm} is the vertical triangular faces.




FIGURE 17. A singly-separated disk with circle-packing view

Combining Propositions 3.4, 3.6, and 3.9 classifies non-reflection thrice-punctured spheres in FAL complements, proving the following theorem.

Theorem 3.11. *Let S be a non-reflection thrice-punctured sphere in an FAL complement $M = \mathbb{S}^3 \setminus \mathcal{A}$. If \mathcal{A} is not the fully twisted Borromean rings, then S is orthogonal to the reflection surface in M and is either a crossing, longitudinal or singly-separated disk.*

4. Separating pairs

Now that we can recognize thrice-punctured spheres in FAL complements, we consider which spheres can be used to decompose FALs. Adams introduced a cut-and-paste method for constructing a hyperbolic manifold from two others containing incompressible, embedded thrice-punctured spheres (see [1, Theorem 4.5]). We review his construction here. Let $S_i \subset M_i$ be an incompressible, embedded thrice-punctured sphere in an orientable hyperbolic three manifold M_i , for $i = 1, 2$. Let $M'_i = M_i \setminus N(S_i)$ be the manifolds M_i with a neighborhood

of the spheres S_i removed. Thus the boundary of M'_i consists of two thrice-punctured spheres which we label S_i^0 and S_i^1 . For $j \in \{0, 1\}$, choose homeomorphisms $\lambda_j : S_1^j \rightarrow S_2^j$ which either both preserve or both reverse orientation. Let M be the manifold obtained by identifying the boundaries of M'_i using the homeomorphisms λ_j . Theorem 4.5 of [1] proves that the manifold M constructed in this way is hyperbolic, and the volume of M is the sum of the volumes of M_1 and M_2 . We define M to be a *punctured sphere sum* of M_1 and M_2 , and note that different choices of homeomorphisms λ_1, λ_2 generally result in distinct punctured sphere sums.

As an application of this process, Adams introduces the notion of a belted sum of two links in S^3 . His definition of a belted sum is diagrammatic, saying that the link L is the belted sum of links L_1, L_2 if they admit diagrams as in Figure 1. Clearly belted sums are special cases of punctured sphere sums in which the S_i look like crossing disks in FALs, containing one longitudinal and two meridional punctures. The homeomorphisms that yield the link L identify punctures of the same type (meridional or longitudinal) without introducing additional crossings.

When considering the process in reverse, a twisted crossing circle can lead to two distinct belted sum decompositions depending on which side of the crossing disk the additional crossing is placed (see Figure 18).




FIGURE 18. Two belted sum decompositions of twisted crossing circles

To eliminate this ambiguity when defining a canonical belted sum decomposition of an FAL, we extend Adams' definition of a belted sum to allow for the addition of a single crossing in the process. In the reverse process, then, a twisted crossing circle can be replaced by flat crossing circles in both summands (as in Definition 6.2(2)). We formalize this in the following definition.

Definition 4.1. Let L_1, L_2 be links with diagrams as in the left side of Figure 19. The link L is a belted sum of L_1, L_2 if it has either of the diagrams after the arrow in Figure 19.

In addition, if $M_i = \mathbb{S}^3 \setminus L_i$ and $M = \mathbb{S}^3 \setminus L$ we say that M is the belted sum of the M_i and let $M = M_1 \#_b M_2$ denote a belted sum decomposition.




FIGURE 19. Extended belted sum definition

In what follows we will use the term *belted sum* interchangeably to refer both to a diagrammatic relationship between links as well as a geometric one between their complements. Note that, using this more general definition, a belted sum is still a special case of a punctured sphere sum, with slightly more flexibility than Figure 1 in the choice of gluing homeomorphisms.

We are concerned with the belted sum process in reverse and, in particular, applying it to fully augmented link complements. The reverse processes are referred to as *punctured sphere decompositions* and *belted sum decompositions* of a manifold (or link), respectively. One consequence of the close relationship between diagrammatic and geometric descriptions of FALs, is that every punctured sphere decomposition of an FAL complement can be realized as a belted sum decomposition (see Corollary 4.9). This justifies focusing on belted sum decompositions in the last two sections of this paper.

By definition, if M is a punctured sphere sum of two manifolds then it contains a pair of disjoint thrice-punctured spheres that separate it. To search for such decompositions, then, we introduce the notion of a separating pair of thrice-punctured spheres.

Definition 4.2. Let M be an FAL complement. The pair $\{S_1, S_2\}$ of disjoint, essential, thrice-punctured spheres in M is a *separating pair* if $M \setminus (S_1 \cup S_2)$ is disconnected.

We also formally define primality in the context of belted sums, and introduce notation.

Definition 4.3. An FAL is *b-prime* if its complement is not the belted sum of two manifolds.

Section 5 will use the absence of separating pairs to provide several characterizations of *b-prime* FALs, and Section 6 introduces a canonical belted sum decomposition of FALs. First, a clear understanding of separating pairs in FAL complements is required.

The remainder of this section is devoted to determining when two thrice-punctured spheres in an FAL complement form a separating pair. To do so, recall our analysis from the previous section. There are several types of thrice-punctured spheres in FAL complements: standard thrice-punctured spheres,

the non-standard ones of Proposition 3.4 that exist only in the fully twisted Borromean rings, as well as non-standard crossing disks, longitudinal disks, and singly-separated disks (see Theorem 3.11). The following lemma will prove a useful observation.

Lemma 4.4. *Let \mathcal{A} be an FAL and let $\{S_1, S_2\}$ be a separating pair in $M = S^3 \setminus \mathcal{A}$. If the component $J \in \mathcal{A}$ punctures the pair $\{S_1, S_2\}$, then J punctures the pair more than once.*

Proof. Suppose that $\{S_1, S_2\}$ is a separating pair and that the component $J \in \mathcal{A}$ punctures $\{S_1, S_2\}$ exactly once, say on S_1 . Then the cusp of M corresponding to J has a neighborhood which is disjoint from S_2 , and we let T_J be the torus boundary of this cusp neighborhood. The curve $S_1 \cap T_J$ is a torus knot on T_J whose complement on T_J is an annulus. This annulus is disjoint from S_2 and connects one side of S_1 to the other, implying that both copies of S_1 are in the boundary of the same connected component of $M \setminus (S_1 \cup S_2)$. Re-identifying the boundary copies of S_1 , then, does not change the connectivity of $M \setminus (S_1 \cup S_2)$ and yields the manifold $M \setminus S_2$. But a single thrice-punctured sphere cannot separate an orientable three-manifold with torus boundary components, so $M \setminus S_2$ is connected. Since $M \setminus (S_1 \cup S_2)$ and $M \setminus S_2$ have the same connectivity, this contradicts the fact that $\{S_1, S_2\}$ is a separating pair. Thus a component of \mathcal{A} that punctures the spheres $\{S_1, S_2\}$ must do so more than once. \square

Lemma 4.4 will be used to show that longitudinal disks, the thrice-punctured sphere of Proposition 3.4, and reflection thrice-punctured spheres (except in the Borromean rings) do not contribute to separating pairs. We begin with the case of longitudinal disks.

Lemma 4.5. *If C_1, C_2 are distinct crossing circles of the FAL \mathcal{A} , then there is at most one longitudinal disk with punctures C_1, C_2 . In particular, a longitudinal disk is never part of a separating pair.*

Proof. Let C_1, C_2 be distinct crossing circles of the FAL \mathcal{A} , and let p_1, p_2 be the ideal vertices of P_+ projecting to cusps C_1, C_2 . Corollary 3.8 implies that, if a longitudinal disk S_ℓ has punctures C_1, C_2 , then p_1, p_2 are two of the three points of tangency in a non-standard triple of the circle packing for P_+ . Given vertices p_1, p_2 of P_+ , however, there is at most one non-standard triple containing them. Indeed, if p_1, p_2 are vertices such a set, they both lie on one of the circles. The other two circles in the set must be those through p_1, p_2 , so the vertices determine one such set. Hence, there is at most one longitudinal disk with punctures C_1, C_2 .

To see that separating pairs do not contain longitudinal disks, let S_ℓ be a longitudinal disk. Since S_ℓ has punctures on distinct crossing circles, Lemma 4.4 implies any thrice-punctured sphere that forms a separating pair with S_ℓ has the same punctures. The previous paragraph, however, shows there is at most one such disk, so separating pairs do not contain longitudinal disks. \square

A second consequence of Lemma 4.4 is that the fully twisted Borromean rings complement does not contain a separating pair. Before proving this, note that the fully twisted Borromean rings is also the minimally twisted chain of four components (Figure 20 parts (a) and (b) are isotopic).




FIGURE 20. The belted sum of L_1 and L_2

Lemma 4.6. *The fully twisted Borromean rings contains no separating pairs and is b -prime. In particular, the thrice-punctured sphere with a single geodesic disk of area 2π is not part of a separating pair.*

Proof. The fully twisted Borromean rings is the link depicted in Figure 20 (see also [23, Figure 6]). Yoshida proves, [23, Lemma 3.13], its complement has exactly eight thrice-punctured spheres. The spheres are the “obvious” twice-punctured disks of Figure 20(b), together with those of Figure 20(c) containing twisted bands. Note that each has boundary slopes along distinct cusps, and that each set of three distinct cusps contributes exactly two thrice-punctured spheres—those of Figure 20. Two spheres with the same cusps intersect and do not form a separating pair. Further, two thrice-punctured spheres that do not share the same cusps must have an isolated boundary slope and cannot separate by Lemma 4.4. Thus the twisted Borromean rings do not contain a separating pair. \square

We now show that reflection thrice-punctured spheres do not contribute to separating pairs in an FAL other than the Borromean rings.

Lemma 4.7. *Let \mathcal{A} be an FAL other than the Borromean rings, and let $M = S^3 \setminus \mathcal{A}$ be its complement. No separating pair in M contains a reflection thrice-punctured sphere. Moreover, two standard thrice-punctured spheres never form a separating pair for M .*

Proof. We can assume \mathcal{A} is not the fully twisted Borromean rings, since Lemma 4.6 implies the result in this case. We first argue that a reflection thrice-punctured sphere cannot form a separating pair with a non-reflection disk. If S is a reflection thrice-punctured sphere, the punctures of S along knot circles are

(generalized) longitudes. Additionally, let m, ℓ denote the meridian and longitude slopes of a crossing circle C that punctures S . If C is a flat crossing circle it punctures S in two meridians, and if C is twisted it punctures S in the $\ell + 2m$ slope. Theorem 3.11 implies that punctures of a non-reflection disk are either meridians of knot circles, longitudes of crossing circles, or meridians of crossing circles, all of which intersect the punctures of S . Thus any non-reflection disk that shares a puncture with S intersects it, and cannot form a separating pair with S .

Now assume $\{S, T\}$ are both reflection thrice-punctured spheres that form a separating pair. Since $M \setminus (S \cup T)$ is disconnected, and no proper subset of the reflection surface separates M , the reflection surface must be $S \cup T$. We show that the reflection surface equaling $S \cup T$ implies that \mathcal{A} is the Borromean rings, whose reflection surface does form a separating pair. Now the nerve of the circle packing for \mathcal{A} is a triangulation of \mathbb{S}^2 , so it has at least four vertices. Vertices of the nerve correspond to unshaded faces so P_+ has at least four unshaded faces, each of which is an ideal polygon. Since they glue up to two thrice-punctured spheres, area considerations imply the unshaded faces of P_+ must be four ideal triangles. Therefore, P_+ is a regular ideal octahedron and the only FALs resulting from gluing two regular ideal octahedra are the Borromean rings, possibly with twisted crossing circles. Twisting one crossing circle in the Borromean rings makes one of the reflection surface components non-orientable, so it cannot be a thrice-punctured sphere, and we've already assumed \mathcal{A} is not the fully twisted Borromean rings. Thus the only FAL with two thrice-punctured spheres for a reflection surface is the Borromean rings themselves.

To finish the proof it remains to show that two standard crossing disks do not form a separating pair. However, distinct standard crossing disks are punctured by distinct crossing circle components, and Lemma 4.4 shows they do not form a separating pair. \square

The following theorem characterizes separating pairs in FAL complements, summarizing the results of this section.

Theorem 4.8. *Let \mathcal{A} be an FAL other than the Borromean rings with complement $M = S^3 \setminus \mathcal{A}$. Suppose that S_1, S_2 are disjoint thrice-punctured spheres in M .*

The pair $\{S_1, S_2\}$ is a separating pair if and only if each is either a crossing disk or a singly-separated disk, and their longitudinal slopes coincide.

Proof. Suppose $\{S_1, S_2\}$ is a separating pair. Proposition 3.4 and Lemma 4.6 imply that neither S_1 nor S_2 has an area 2π geodesic disk. Lemmas 4.5 and 4.7 show that neither S_1 nor S_2 are longitudinal disks or reflection thrice-punctured spheres. After eliminating these possibilities, crossing disks (standard or non-standard) and singly-separated disks are the only remaining types for S_1 and S_2 . Now both crossing and singly-separated disks have a unique longitudinal slope (which is along a crossing circle). If the longitudinal slopes of S_1 and S_2 are along distinct crossing circle cusps of M , then either $S_1 \cap S_2 \neq \emptyset$ showing

$\{S_1, S_2\}$ is not a separating pair or Lemma 4.4 implies $\{S_1, S_2\}$ is not a separating pair. Thus the longitudinal slopes of S_1 and S_2 must coincide.

Conversely, let S_1 and S_2 be two crossing or singly-separated disks that share a longitudinal puncture along the same crossing circle C . One obtains an embedded two sphere \mathbb{S}^2 in \mathbb{S}^3 from $S_1 \cup S_2$ by including the crossing circle C , together with the points of \mathcal{A} corresponding to meridional punctures of $S_1 \cup S_2$. Since \mathbb{S}^2 separates \mathbb{S}^3 , $M \setminus (S_1 \cup S_2)$ is also disconnected, and $\{S_1, S_2\}$ is a separating pair. \square

As observed earlier, Theorem 4.8 shows that every separating pair in an FAL complement consists of a shared crossing-circle longitudinal puncture together with meridional punctures. Thus every punctured sphere decomposition in an FAL complement can be realized as a belted sum decomposition. We highlight this as a corollary.

Corollary 4.9. *If M is an FAL complement, then every separating pair $\{S_1, S_2\} \subset M$ yields a belted sum decomposition of M .*

Proof. Suppose L is an FAL and $\{S_1, S_2\}$ form a separating pair in $M = \mathbb{S}^3 \setminus L$. We must show that L has a diagram as in the right-hand side of Figure 19. By Theorem 4.8 there are three cases: both S_1, S_2 are generalized crossing disks, both are singly-separated disks, or there is one of each. Figure 23 illustrates the situation when both are crossing disks. In fact, this figure applies to all cases. If an S_i is singly-separated, its two meridional punctures in Figure 23 are from a crossing circle, whose crossings with knot circles are subsumed in tangles A, B, C . See Figure 24 for an example of the case when the pair $\{S_1, S_2\}$ contains one of each type of disk. \square

Separating pairs consisting of two crossing disks will play an important roll in what follows. Corollary 5.3, for example, shows that an FAL is b -prime if and only if there are no separating pairs with two crossing disks. Moreover, such separating pairs are exactly those used to define canonical belted sum decompositions of FALs (see Definition 6.2). For these reasons we now prove that every separating pair consisting of two crossing disks decomposes an FAL into the belted sum of two FALs.

Lemma 4.10. *Let $\{S_1, S_2\}$ be a separating pair consisting of two crossing disks in the FAL complement $M = \mathbb{S}^3 \setminus \mathcal{A}$, and let $M = M_1 \#_b M_2$ denote a belted sum decomposition they determine. Then both M_1 and M_2 are FAL complements.*

Proof. Let $\{S_1, S_2\}$ be a separating pair consisting of two crossing disks and $M = M_1 \#_b M_2$ their corresponding belted sum decomposition. Consider the case where neither S_i is standard (the case where one of the S_i is the standard crossing disk is similar). Then S_1 and S_2 share a longitudinal puncture along the same crossing circle C of \mathcal{A} .

We begin by slicing a fundamental polyhedron for M into two polyhedra and determining face pairings in the individual pieces. Then we show these face pairings result in FALs. This is the reverse of Adams proof of [1, Theorem

4.5], where he glues fundamental polyhedra for the individual manifolds along faces corresponding to chosen thrice-punctured spheres. The standard polyhedral decomposition for FALs facilitates our process for us and, since P_- is the reflection of P_+ , our description focuses on P_+ .

The crossing circle C shared by $\{S_1, S_2\}$ corresponds to exactly one ideal vertex of P_+ and we assume, possibly after a Möbius transformation, that it corresponds to infinity (using the upper-half space model for \mathbb{H}^3). Figure 21(a) illustrates P_+ with geodesic disks S'_i for the non-standard crossing disks S_i . When viewed from above, P_+ is rectangular and the vertical shaded faces (labeled T') project to the standard crossing disk for C . The geodesic disks S'_i are ideal triangles parallel to the T' since all have longitudinal punctures along C . The shaded interstices of Figure 21(a) may or may not contain additional circles of the circle packing.




FIGURE 21. belted sum decomposition if both $\{S_1, S_2\}$ are non-standard crossing disks

Now slice P_+ along the S'_i to obtain three pieces, then glue the outside pieces along the shaded triangles T' . The result is two polyhedra, and let P_1^+ denote the one between the S'_i while P_2^+ denotes that formed from the outsides (see Figure 21(b)). A similar slicing along P_- yields the reflections P_1^- and P_2^- . Note that in the case where T' projects to a twisted crossing disk, the polyhedron P_2^+ will consist one outside piece from each of P_\pm .

Now consider the gluing maps on faces of P_i^\pm . The gluing maps of P_\pm identify unshaded faces of P_i^\pm to each other since the disks S'_i are orthogonal to the reflection surface. Moreover, since the S_i are crossing disks, exactly one vertex of the ideal triangles S'_i corresponds to a crossing circle. This means the triangular faces T' are the only ones of P_\pm adjacent to S'_i that are identified on P_\pm , and they are glued to form P_2^\pm . The face pairings on the remaining shaded triangles of ∂P_\pm do not connect P_1^\pm to P_2^\pm . Thus all faces of P_i^\pm inherit gluing maps from P_\pm except the shaded triangles corresponding to the S'_i , and those maps do not connect the P_1^\pm to P_2^\pm . Pairing the copies of S'_i of P_i^+ along their shared

vertex completes the gluing instructions so that unshaded faces glue to their reflections and shaded faces glue to their mate or its reflection.

Identifying the S'_i either way yields a so-called *admissible* gluing pattern from [12], and [12, Theorem 3.2] guarantees that P_i^\pm with these face-pairings yield complete hyperbolic manifolds M_i . One can then form the circle packing and nerve corresponding to these polyhedral decompositions, and apply [19, Lemma 2.4] to conclude the manifolds are FAL complements. Let M_i denote the FAL resulting from gluing P_i^\pm . The only ambiguity in this process is whether the S'_i project to flat or twisted crossing disks.

Beginning with M_1 and M_2 and forming the belted sum $M_1 \#_b M_2$ along the disks D_i , as in [1], reverses this process on the polyhedral level, thereby yielding M by construction.




FIGURE 22. $\{S_1, S_2\}$ with one standard and one non-standard crossing disk

The case where the disk S_1 of the pair $\{S_1, S_2\}$ is a standard crossing disk is similar. The difference is that one only slices P_+ along S'_2 (respectively P_- along the reflection of S'_2), yielding only two pieces which are the P_1^\pm and P_2^\pm described (see Figure 22). \square

The case just argued, where both S_i are non-standard crossing disks, can be illustrated in the crushtacean as well (see Figure 23). The tangles A, B, C of Figure 23 correspond to the interstices of the same label in Figure 21.

Thus every separating pair consisting of two crossing disks decomposes an FAL complement into the belted sum of two simpler FAL complements. We turn our attention to the questions of b -primality and canonical decompositions.

5. On b -prime FALs

This section provides three characterizations of b -prime FALs. The Borromean rings and its half-twist partners are a special case, and considered separately. For FALs with at least three crossing circles, Theorem 5.2 shows that




FIGURE 23. FAL belted sum and crushtaceans

an FAL is b -prime if and only if each non-trivial 3-edge cut in its crushtacean is fully painted. Interpreting this in terms of crossing disks shows that an FAL is b -prime if and only if each crossing circle bounds a unique crossing disk (see Corollary 5.3). These results follow readily from those of Section 4. After introducing diagrammatic terminology, Theorem 5.4 shows that b -prime FALs come from fully augmenting diagrams with trivial flype orbits.

First consider the special case of FALs with two crossing circles—the Borromean rings and its half-twist partners.

Lemma 5.1. *The fully twisted Borromean rings complement is b -prime. The complement of the Borromean rings with both, or exactly one, flat crossing circle are b -composite.*

In fact, both are the belted sum of two Whitehead link complements and both have a separating pair consisting of one crossing and one singly-separated disk.

Proof. We've already seen in Lemma 4.6 that the fully twisted Borromean rings is b -prime. Moreover, Adams observed (see [1]) that the Borromean rings complement is the belted sum of two copies of the Whitehead link complement. In our terminology, the singly-separated disk of Figure 24 forms a separating pair with the top crossing disk. This pair, with the appropriate choice of gluing homeomorphisms, decomposes the Borromean rings into two copies of the Whitehead link.

Finally, the separating pair of the original Borromean rings is preserved when adding a half-twist to the top crossing disk; therefore, the complement of the Borromean rings with one twisted crossing circle is b -composite. As in the case of the Borromean rings complement, it is the belted sum of two Whitehead link complements, with appropriate choice of gluing homeomorphisms. \square

We now treat the general case, characterizing b -prime FALs in terms of their painted crushtaceans.

Theorem 5.2. *Let \mathcal{A} be a fully augmented link with at least three crossing circles, and let Γ^* be its crushtacean. Then the complement $M = \mathbb{S}^3 \setminus \mathcal{A}$ is b -prime if and only if every 3-edge cut of Γ^* is fully painted.*




FIGURE 24. Borromean rings as belted sum of two Whitehead links

Proof. We begin with the observation that every non-trivial 3-edge cut in Γ^* has either one or three painted edges. To see this, note that trivalent graphs have an even number of vertices since the sum of the degrees is both $3V$ and $2E$ (where V, E represent the number of vertices and edges, respectively). A non-trivial 3-edge cut separates the graph into two components, say Γ_1 and Γ_2 , and collapsing Γ_2 to a single vertex results in a new trivalent graph Γ_3 . Then Γ_3 has an even number of vertices as well as having one more vertex than Γ_1 (the one corresponding to the collapsed Γ_2). Hence Γ_1 has an odd number of vertices. Now painted edges with both vertices in Γ_1 use up an even number of vertices, leaving an odd number whose painted edges belong to the 3-edge cut.

To prove that b -prime implies fully painted non-trivial 3-edge cuts, we show the contrapositive. If Γ^* has a once-painted, non-trivial 3-edge cut, then M has a non-standard crossing disk by Corollary 3.8. This forms a separating pair with the standard disk sharing the crossing circle puncture, by Theorem 4.8, and M is not b -prime.

Conversely, suppose every 3-edge cut of Γ^* is fully painted, so that every non-standard n -disk is longitudinal. Theorem 4.8, then, implies that every separating pair must contain a singly-separated disk. We now argue that if M contains a singly-separated disk, it contains a non-standard crossing disk, contradicting the fact that n -disks in M are assumed longitudinal.

Let S be a singly-separated disk in an FAL complement M , and S' one of its material vertex disks. Proposition 3.9 shows that circle packing near S' is as illustrated in Figure 25(a). If both interstices U and v are trivial, this is the circle packing for an FAL with two crossing circles. Since M has at least three crossing circles, at least one of the interstices, say U , has additional circles in the circle packing. Now S' has two neighboring ideal triangles labeled S'_i in Figure 25(b), and we focus on S'_2 . Since q is the only vertex of the triangles T, T' corresponding to a crossing circle, the ideal vertices at the base of S'_2 correspond to knot circles in \mathcal{A} . Proposition 3.6 implies that S'_2 projects to a crossing disk and, since interstice U is non-trivial, the disk S'_2 is not a face of P_+ . Thus S'_2 projects




FIGURE 25. Ideal triangular neighbors of material-vertex disks

to a non-standard crossing disk, contradicting the assumption that every 3-edge cut is fully painted.

If every 3-edge cut of Γ^* is fully painted then, among non-reflection thrice-punctured spheres, M contains only standard crossing disks and is void of singly-separated disks. Lemma 4.7 and Theorem 4.8 then imply M is void of separating pairs, and M is b -prime. \square

The correspondence between once-painted, 3-edge cuts of the crushtacean and non-standard crossing disks immediately leads to the following characterization of b -prime FALs.

Corollary 5.3. *Let \mathcal{A} be a fully augmented link with at least three crossing circles. The complement $M = \mathbb{S}^3 \setminus \mathcal{A}$ is b -prime if and only if every crossing circle of \mathcal{A} bounds a unique crossing disk.*

Proof. Every crossing circle of \mathcal{A} bounds a standard crossing disk. Corollary 3.8 implies that non-standard crossing disks are in one-to-one correspondence with once-painted 3-edge cuts of the crushtacean. Thus every crossing circle of \mathcal{A} bounds a unique crossing disk (i.e. the standard one) if and only if there are no once-painted 3-edge cuts of the crushtacean. The result follows from Theorem 5.2. \square

Observe that the criteria of Corollary 5.3 can be rephrased to say that $M = \mathbb{S}^3 \setminus \mathcal{A}$ contains no non-standard crossing disks.

We now interpret FAL b -primality in terms of the link diagram that generated the FAL. The main result is that an FAL \mathcal{A}_D obtained by augmenting a diagram D is b -prime if and only if it generates only trivial flype orbits. We review the necessary diagrammatic properties involved before proving the theorem, and refer the reader to [7] for further diagrammatic details.

An m -tangle in a regular link diagram D is a disk \mathcal{D} in the projection plane whose boundary intersects D in m points. An m -tangle is trivial if it contains no crossings of D . A flype changes one diagram of a link L into another by rotating

a 4-tangle 180° , moving a crossing over the tangle as in Figure 26. A flype is non-trivial if neither of the 4-tangles involved are trivial. Flypes were introduced by Tait during his original efforts to enumerate links in the late 1800's, when he conjectured that any two reduced alternating diagrams of a link are related by a sequence of flypes. This conjecture was finally proven by Menasco and Thistlethwaite in [17].




FIGURE 26. Flype on a link diagram

A *flype tangle* for a crossing is a 4-tangle that cannot be decomposed into smaller flype-admitting 4-tangles for that crossing. A flype orbit for a crossing is a collection of flype tangles for the crossing, and a flype orbit is trivial if it consists of a single flype tangle. Each crossing generates two flype orbits, and Calvo shows that at most one of them is non-trivial (see [7, Lemma 4]).

Figure 27 illustrates these definitions. Knot 9_{28} of Rolfsen's tables has six twist regions (see Figure 27(a)). The twists with nontrivial flype orbits consist of single crossings, while the 2-crossing twists have only trivial flype orbits. Parts (b) and (c) of Figure 27 depict flype tangles and orbits for crossings a and b , respectively (the flype orbit for crossing b is more easily seen after an isotopy of the projection sphere). By symmetry, rotating the diagram of Figure 27(c) about a horizontal axis yields the flype tangles and orbit of crossing c .




FIGURE 27. A knot and its flype orbits

Let D be a reduced alternating diagram for a link L , and recall that fully augmenting L from D amounts to choosing a crossing circle around each twist in D . There is a unique choice of crossing circle associated with each twist of D containing more than one crossing, and the corresponding flype orbit may or may not be trivial. Single-crossing twists in D admit two choices for crossing circles, and at most one of the choices corresponds to a non-trivial flype orbit by [7, Lemma 4]. Given an FAL \mathcal{A}_L arising from D , the *flype orbits induced by \mathcal{A}_L* are those determined by the crossing circles.




FIGURE 28. b -composite and b -prime augmentations of knot 9_{28}

For example, the augmentation \mathcal{A}' of Figure 28(b) is obtained by changing the choice of crossing circle for every single-crossing twist in Figure 28(a). The FAL \mathcal{A}' induces all trivial flype orbits to augment D , while \mathcal{A} induces non-trivial flype orbits for all single-crossing twists. The crushtaceans of Figure 28 illustrate the fact that the same diagram can have a b -prime augmentation as well as b -composite ones. Note that every painted edge corresponding to a twisted crossing circle in Figure 28(a) is contained in at least one non-trivial, once-painted 3-edge cut; whereas the crushtacean for \mathcal{A}' contains no non-trivial, 3-edge cuts.

We now have the terminology needed to give the following diagrammatic characterization of b -prime FALs. A technical note in the following theorem: the assumption that D has more than two twist regions prevents \mathcal{A} from being the Borromean rings or one of its half-twist partners.

Theorem 5.4. *Let D be a twist-reduced, non-split, prime link diagram with more than two twist regions, and let \mathcal{A} be a full augmentation of D . The fully augmented link \mathcal{A} is b -prime if and only if it induces only trivial flype orbits.*

Proof. The assumptions on the diagram D ensure that \mathcal{A} is a hyperbolic FAL other than the Borromean rings or its half-twist partners (See [19, Theorem 2.5]). The key observation to prove is that the manifold $M = \mathbb{S}^3 \setminus \mathcal{A}$ contains non-standard crossing disks if and only if \mathcal{A} induces a non-trivial flype orbit.

The fact that the changes made in passing from the diagram D to the FAL \mathcal{A} are local allows us to argue diagrammatically. The result then follows from Corollary 5.3.

Suppose the augmentation \mathcal{A} of D induces a non-trivial flype orbit for a twist region τ with the chosen crossing circle C . Let $\mathcal{T}_1, \dots, \mathcal{T}_n$, with $n \geq 2$, denote the flype tangles in the order encountered traversing the orbit for τ . One can choose pairwise disjoint three-ball neighborhoods $\mathbb{B}_i \subset \mathbb{S}^3$ of the \mathcal{T}_i such that $\cup_{i=1}^n \mathbb{B}_i$ contains all components of \mathcal{A} except for C and the strands of D connecting the \mathcal{T}_i . In addition, choose the \mathbb{B}_i so that they are disjoint from the crossing circle C and its standard crossing disk S . The crossing circle C bounds a disk S_i disjoint from $\cup_{i=1}^n \mathbb{B}_i$ that intersects the arcs of D joining \mathcal{T}_i and \mathcal{T}_{i+1} . Each S_i is a thrice-punctured sphere with a longitudinal puncture along C and two meridional knot circle punctures, so S_i is a non-standard crossing disk. Moreover, the S_i are distinct (not isotopic) since the \mathbb{B}_i contain components of \mathcal{A} which separate them. Hence if \mathcal{A} induces a non-trivial flype orbit, M contains non-standard crossing disks and is b -composite.

Conversely, suppose C is a crossing circle of \mathcal{A} that bounds a non-standard crossing disk S' in M , and let S denote the standard crossing disk for C . Let τ be the twist region of the diagram D augmented by C , we wish to show that τ has a non-trivial flype orbit. By Theorem 4.8, the pair $\{S, S'\}$ is a separating pair. The non-separating geodesics of $\{S, S'\}$ are contained in the projection plane of D and form, together with points of their punctures, a simple closed loop γ in D . This is true even when C is a twisted crossing circle since the half-twist changes the reflection surface and not the projection plane containing D .

Let $\gamma \subset D$ be the simple closed curve described in the previous paragraph. The boundary curves of an appropriately chosen product neighborhood $\gamma \times (-\epsilon, \epsilon) \subset D$, then, form 4-tangles containing non-trivial portions of D , since $\{S, S'\}$ is a separating pair. These tangles allow for a non-trivial flype, showing the flype orbit of τ is non-trivial.

Since \mathcal{A} induces non-trivial flype orbits if and only if M contains non-standard crossing disks, the proof is complete by Corollary 5.3. \square

We now have three characterizations of b -prime FALs: one combinatorial, another geometric, and the third diagrammatic. The next section applies these results to create canonical belted sum decompositions of FALs in which each summand is either a b -prime FAL or the Borromean rings (possibly with one twisted crossing circle).

6. Canonical FAL decompositions

There are two desirable properties for belted sum decompositions of an FAL complement: first, summands ought to be b -prime and, second, summands should be FALs to maintain the tractable geometry of FALs in the process. These two properties can nearly be achieved simultaneously. More precisely, this section introduces a canonical method for decomposing an FAL in which

every summand is an FAL complement and, if a summand is not b -prime, then it is a copy of the Borromean rings with at least one flat crossing circle. If desired, each summand can be made b -prime by decomposing each copy of the Borromean rings into two Whitehead links by Lemma 5.1. The result can then be recast as follows: there is a canonical belted sum decomposition of an FAL complement in which each summand is either a b -prime FAL or a copy of the Whitehead link. The decomposition is canonical in the sense that the prescribed choices in the process produce a unique decomposition.

A technical lemma is required before defining the canonical belted sum decomposition, which is that distinct crossing disks in an FAL complement are disjoint. This follows from the more general result that n -disks are disjoint in an FAL complement.

Lemma 6.1. *If $M = \mathbb{S}^3 \setminus \mathcal{A}$ is an FAL complement, then distinct n -disks in M are disjoint.*

Proof. Any two n -disks S_i and S_j are both orthogonal to the reflection surface R of M , and both intersect R in their non-separating geodesics (Proposition 3.6). If $S_i \cap S_j$ contains a non-separating geodesic γ , then $S_i = S_j$ since they both are orthogonal to R along γ . Moreover, distinct n -disks cannot intersect in a geodesic that is separating on one and non-separating on the other, since separating geodesics on n -disks are not contained in R while non-separating geodesics are. Thus $S_i \cap S_j$ can only contain a geodesic that is separating on both. Yoshida, however, has classified intersections of two thrice-punctured spheres in an orientable hyperbolic three manifolds (see [23, Proposition 3.1]), showing that they cannot intersect in a geodesic that separates both. Hence n -disks are pairwise disjoint. \square

In particular, Lemma 6.1 shows that crossing disks in an FAL complement M are disjoint, and we can slice M along any or all of them simultaneously. Recall that two crossing disks $\{S_1, S_2\}$ form a separating pair if and only if they share the same longitudinal crossing circle puncture. Also, the proof of Lemma 4.10 shows that such a pair decomposes M into two FAL summands, and the only ambiguity is whether the corresponding crossing circles in the summands are twisted or flat. In defining the canonical decomposition, we make the convention to identify copies of the S_i within P_+ , resulting in flat crossing circles. We are now prepared to define the canonical belted sum decomposition of an FAL.

Definition 6.2. Let \mathcal{A} be an FAL with complement $M = \mathbb{S}^3 \setminus \mathcal{A}$. The **canonical belted sum decomposition** of M is obtained by the following procedure:

- (i) First slice M along all crossing disks that belong to a separating pair, then
- (ii) Identify boundary spheres in the resulting components which correspond to separating pairs of M so as to create flat crossing disks in the summands.

Complements of FALs with two crossing circles are their own canonical belted sum decompositions, since each crossing circle bounds a unique crossing disk. More generally, Corollary 5.3 shows that if M is already b -prime then each crossing circle bounds a unique crossing disk, so a b -prime FAL is its own canonical belted sum decomposition.

The canonical belted sum decomposition is well-defined. Indeed, since all crossing disks are disjoint by Lemma 6.1, the result of Step (1) is a collection of manifolds with thrice-punctured spheres in their boundaries. Moreover, boundary spheres are paired along their common longitudinal crossing circle punctures, uniquely determining which boundary spheres are to be identified. Finally, the restriction that the crossing circles resulting from these identifications be flat uniquely determines the summands of the canonical decomposition. Thus there is a unique canonical decomposition for each given FAL.

We now prove that summands in canonical decompositions are either b -prime FALs or copies of the Borromean rings with zero or one twisted crossing circles. To prove this, we must show that decomposing an FAL complement along crossing disks can be done simultaneously, that the summands are always FALs, and that each summand that is not b -prime is a copy of the Borromean rings with at most one twisted crossing circle.

Theorem 6.3. *Let A be an FAL with at least three crossing circles and let $M = \mathbb{S}^3 \setminus A$ be its complement.*

Each summand in the canonical belted sum decomposition of M is either a b -prime FAL complement or the complement of the Borromean rings with at least one flat crossing circle.

Proof. Let $M = M_1 \#_b \dots \#_b M_m$ be the canonical belted sum decomposition of M and let S_1, \dots, S_n denote all crossing disks (both standard and non-standard) involved in a separating pair of M . As noted above, the slicing operation $M \setminus (\cup_{i=1}^n S_i)$ can be done simultaneously since distinct crossing disks in M are disjoint.

Now suppose C bounds the crossing disks S_1^C, \dots, S_k^C , with $k \geq 2$. Further, suppose they are labeled in the order encountered along a meridian of C , with S_1^C denoting the standard crossing disk. Any two of them form a separating pair, so the two copies of S_i^C are on distinct components of $M \setminus (\cup_{i=1}^n S_i)$. In addition, (at least) one component of $M \setminus (\cup_{i=1}^n S_i)$ has one copy of each consecutive S_i^C, S_{i+1}^C in its boundary, and they share a longitudinal puncture along C .

Let M'_j denote the component of $M \setminus (\cup_{i=1}^n S_i)$ that forms the summand M_j when boundary spheres are identified. The previous paragraph implies the boundary of M'_j consists of an even number of thrice-punctured spheres which are paired according to their longitudinal punctures in the original manifold M . Preserving longitudinal and meridional punctures while identifying paired thrice-punctured spheres in M'_j yields a FAL M_i by Lemma 4.10. Our initial goal is to show that each M_j with at least three crossing circles is b -prime.

Our first task is to show that crossing disks of M contribute a unique crossing disk for each crossing circle of M_j , which we choose to be standard. First note that if C is crossing circle in M that bounds a unique crossing disk S_C , then the belted sum operations in the canonical decomposition of M are disjoint from a neighborhood of $C \cup S_C$. Thus C bounds the disk S_C in M_j . Now suppose that C bounds at least one non-standard crossing disk in M . Then copies of consecutive crossing disks in M along C are glued to form a disk S_C for C in M_j . In either case, the crossing disks of M contribute a unique crossing disk S_C for each C in M_j , and we choose these to be the standard disks in M_j .

We now argue that the process of canonical decompositions do not create additional crossing disks in the summands. Lemma 4.10 implies that each summand M_j is an FAL, as each can be formed by consecutive belted sums along separating pairs of crossing disks. If M_j is b -composite, then Corollary 5.3 implies there is a crossing circle C of M_j that bounds a non-standard crossing disk S . Note that S is not part of a belted sum operation needed to build M , since only chosen standard crossing disks of M_j are used to reconstruct M . By Lemma 6.1 applied to the FAL M_j , S is disjoint from all standard crossing disks of M_j , and so is preserved under the belted sum operations used to reconstruct M . Thus S is a thrice-punctured sphere in M . By Proposition 3.6 applied to M_j , S is orthogonal to the reflection surface in M_j and has two meridional knot circle punctures and one longitudinal crossing circle puncture. Since the belted sum operation on crossing disks preserves both the reflection surface and component types, S must have these properties in M as well. The classification of non-standard disks in FALs implies that S is a crossing disk in M as well. But then M contributes two crossing disks to M_j , which is a contradiction. Therefore each summand with at least three crossing circles must be b -prime.

Now suppose a summand M_i has two crossing circles. At least one crossing disk D in M_i is used in the belted sum reconstruction of M , which implies that it is the result of slicing along and gluing consecutive crossing disks in M . By definition, the gluing results in a flat crossing circle, so that D must be flat in M_i . Since M_i has two crossing circles, at least one of which is flat, it cannot be the fully twisted Borromean rings, completing the proof. \square

It may seem unnatural to ignore singly-separated disks in defining the canonical belted sum decomposition. Indeed, Lemma 6.1 can be generalized to show that all disks involved in any separating pair are disjoint. One can then slice along all of them and glue corresponding boundary spheres to construct an alternative decomposition. The components with a singly-separated disk will glue to Whitehead links, and the rest to either b -prime FALs with at least three crossing circles or the Borromean rings with at least one flat crossing circle. Further decomposing Borromean rings summands into two Whitehead link complements in both this decomposition and the canonical one results in the same decomposition. For simplicity of exposition, we presented the crossing disk method.

Neither approach, that of excluding or of including singly-separated disks, guarantees uniqueness over all belted sum decompositions. The difficulty is that a non-FAL summand may be produced in the process of decomposing an FAL. The characterizations of thrice-punctured spheres in Section 3 and of separating pairs in Section 4 cannot be applied to non-FALs, and it is unclear at this writing how to proceed. For now, then, we content ourselves with the level of uniqueness provided in Theorem 6.3.

Acknowledgements. This research was supported in part by NSF-REU Grants DMS-1461286 and DMS-1758020, as well as California State University, San Bernardino. We are very grateful to Christian Millichap for many (many!) conversations and suggestions that have significantly improved this work. Thanks also to Jeff Meyer for helpful conversations. Finally, thanks to the referee for their careful reading and helpful suggestions that improved this work.

References

- [1] ADAMS, C. Thrice-punctured spheres in hyperbolic 3-manifolds. *Trans. Amer. Math. Soc.* **287** (1985), no. 2, 645–656. [MR768730](#) (86k:57008), [Zbl 0527.57002](#), doi: [10.2307/1999666](#). [1](#), [2](#), [11](#), [23](#), [24](#), [30](#), [31](#), [32](#)
- [2] ADAMS, C. Augmented alternating link complements are hyperbolic. “*Low-dimensional topology and Kleinian groups*” (Coventry/Durham, 1984), *London Math. Soc. Lecture Note Ser.* **112** (1986), 115–130. [MR903861](#) (89f:57003), [Zbl 0632.57008](#). [5](#)
- [3] AGOL, I. The minimal volume orientable hyperbolic 2-cusped 3-manifolds. *Proc. Amer. Math. Soc.* **138:10** (2010), 3723–3732. [MR2661571](#) (2011k:57023), [Zbl 1203.57006](#), doi: [10.1090/S0002-9939-10-10364-5](#). [3](#)
- [4] ANDREEV, E. M. Convex polyhedra in Lobacevskii spaces. *Mat. Sb. (N.S.)* **81(123)** (1970), 445–478. [MR259734](#). [7](#)
- [5] BLAIR, R.; FUTER, D.; TOMOVA, M. Essential surfaces in highly twisted link complements. *Algebraic & Geometric Topology.* **15** (2015), no. 3, 1501–1523. [MR3361143](#), [Zbl 1339.57007](#), doi: [10.2140/agt.2015.15.1501](#). [2](#)
- [6] BURDE, G.; ZIESCHANG, H. Knots, 2nd rev. and extended ed. *De Gruyter studies in mathematics*, 5, Berlin, 2003. xii+559 pp. [MR1959408](#), [Zbl 1009.57003](#). [12](#)
- [7] CALVO, J. Knot enumeration through flypes and twisted splices. *J. Knot Theory and Ramifications.* **6** (1997), no. 6, 785–798. [MR1483786](#), [Zbl 0920.57004](#), doi: [10.1142/S021821659700042X](#). [34](#), [35](#), [36](#)
- [8] CHESEBRO, E.; DEBLOIS, J.; WILTON, H. Some virtually special hyperbolic 3-manifold groups. *Comment. Math. Helv.* **87** (2012), no. 3, 727–787. [MR2980525](#), [Zbl 1283.57007](#), [arXiv:0903.5288](#), doi: [10.4171/CMH/267](#). [9](#)
- [9] FLINT, R. Intercusp Geodesics and Cusp Shapes of Fully Augmented Links. *ProQuest LLC*. Ann Arbor, MI, 2017, Thesis (Ph.D.) - City University of New York. [MR3664936](#), [arXiv:1811.07397](#). [3](#)
- [10] FUTER, D.; KALFAGIANNI, E.; PURCELL, J. Dehn filling, volume and the Jones polynomial. *J. Differential Geom.* **78** (2008), no. 3, 429–464. [MR2396249](#), [Zbl 1144.57014](#). [2](#)
- [11] FUTER, D.; PURCELL, J. Links with no exceptional surgeries. *Comment. Math. Helv.* **82** (2007), no. 3, 629–664. [MR2314056](#) (2008k:57008), [Zbl 1134.57003](#), [arXiv:math/0412307](#), doi: [10.4171/CMH/105](#). [2](#)
- [12] HARNOIS, J.; OLSON, H.; TRAPP, R. Hyperbolic tangle surgeries and nested links. *Algebraic & Geometric Topology.* **18** (2018), no. 3, 1573–1602. [MR3784013](#), [Zbl 1396.57008](#), doi: [10.2140/agt.2018.18.1573](#). [31](#)

- [13] KALFAGIANI, E.; LEE, C. R. S. Crosscap numbers and the Jones Polynomial. *Adv. Math.* **286** (2016), 308–337. [MR3415687](#), [Zbl 1329.57008](#), doi: [10.1016/j.aim.2015.09.017](#). [2](#)
- [14] KNAVEL, S.; TRAPP, R. Embedded totally geodesic surfaces in fully augmented links. *Comm. Anal. Geom.* **31** (2023), 563–593. [MR4685031](#), [Zbl 1545.57005](#). [13](#), [14](#), [17](#), [18](#), [22](#)
- [15] LACKENBY, M. The volume of hyperbolic alternating link complements. *Proc. London Math. Soc. (3)* **88** (2004), no. 1, 204–224, With an appendix by Ian Agol and Dylan Thurston. [MR2018964](#) (2004i:57008), [Zbl 1041.57002](#), [arXiv:math/0012185](#), doi: [10.1112/S0024611503014291](#). [2](#)
- [16] MACLACHLAN C.; REID, A.W. The Arithmetic of Hyperbolic 3–Manifolds. *Graduate Texts in Mathematics*, vol. 219, Springer-Verlag, New York, (2003). [MR1937957](#) (2004i:57021), [Zbl 1025.57001](#). [3](#)
- [17] MENASCO, W.; THISTLETHWAITE, M. The classification of alternating links. *Annals of Math.* **138**, (1993), 113–171. [MR1230928](#), [Zbl 0809.57002](#), doi: [10.2307/2946636](#). [35](#)
- [18] MEYER, J. S.; MILLICHAP, C.; TRAPP, R. Arithmeticity and hidden symmetries of fully augmented pretzel link complements. *New York J. Math.* **26**, (2020), 149–183. [MR4063955](#), [Zbl 1435.57005](#). [3](#)
- [19] PURCELL, J. An introduction to fully augmented links. *Interactions between hyperbolic geometry, quantum topology and number theory*. 205–220, *Contemp. Math., Amer. Math. Soc., Providence, RI* **541**, (2011). [MR2796634](#) (2012c:57019), [Zbl 1236.57006](#), doi: [10.1090/conm/541.2](#), [3](#), [6](#), [8](#), [9](#), [31](#), [36](#)
- [20] PURCELL, J. Cusp shapes under cone deformations. *J. Differential Geom.* **80**, (2008), no. 3, 453–500. [MR2472480](#), [Zbl 1182.57015](#). [5](#)
- [21] THURSTON, W. The Geometry and Topology of 3-Manifolds. *Lecture Notes, Princeton University Math. Dept.* (1978).
- [22] YOSHIDA, K. The minimal volume orientable hyperbolic 3-manifold with 4 cusps. *Pacific J. Math.* **266**, (2013), no. 2, 457–476. [MR3130632](#), [Zbl 1294.57010](#), doi: [10.2140/pjm.2013.266.457](#).
- [23] YOSHIDA, K. Unions of 3-punctured spheres in hyperbolic 3-manifolds. *Comm. Anal. Geom.* **29**, (2021), 1643–1689. [MR4429241](#), [Zbl 1498.57019](#). [27](#), [38](#)

(P. Morgan) DEPARTMENT OF MATHEMATICS, UNIVERSITY OF MASSACHUSETTS AMHERST, AMHERST, MA 01003, USA
pamorgan@umass.edu

(B. Ransom) DEPARTMENT OF MATHEMATICS, UNIVERSITY OF CALIFORNIA IRVINE, IRVINE, CA 92697, USA
bransom@uci.edu

(D. Spyropoulos) DEPARTMENT OF MATHEMATICS, MICHIGAN STATE UNIVERSITY, EAST LANSING, MI 48824, USA
spyropou@msu.edu

(R. Trapp) DEPARTMENT OF MATHEMATICS, CALIFORNIA STATE UNIVERSITY, SAN BERNARDINO, CA 92407, USA
rtrapp@csusb.edu

(C. Ziegler) DEPARTMENT OF MATHEMATICS, UNIVERSITY AT BUFFALO, BUFFALO, NY 14260, USA
cz22@buffalo.edu

This paper is available via <http://nyjm.albany.edu/j/2025/31-1.html>.