

Angled decompositions of arborescent link complements

David Futer and François Guéritaud

ABSTRACT

This paper describes a way to subdivide a 3-manifold into *angled blocks*, namely polyhedral pieces that need not be simply connected. When the individual blocks carry dihedral angles that fit together in a consistent fashion, we prove that a manifold constructed from these blocks must be hyperbolic. The main application is a new proof of a classical, unpublished theorem of Bonahon and Siebenmann: that all arborescent links, except for three simple families of exceptions, have hyperbolic complements.

1. Introduction

In the 1990s, Andrew Casson introduced a powerful technique for constructing and studying cusped hyperbolic 3-manifolds. His idea was to subdivide a manifold M into *angled ideal tetrahedra*: that is, tetrahedra whose vertices are removed and whose edges carry prescribed dihedral angles. When the dihedral angles of the tetrahedra add up to 2π around each edge of M , the triangulation is called an *angled triangulation*. Casson proved that every orientable cusped 3-manifold that admits an angled triangulation must also admit a hyperbolic metric, and outlined a possible way to find the hyperbolic metric by studying the volumes of angled tetrahedra — an idea also developed by Rivin [17]. The power of Casson’s approach lies in the fact that the defining equations of an angled triangulation are both linear and local, making angled triangulations relatively easy to find and deform (much easier than to study an actual hyperbolic triangulation, as in [14, 20] or in some aspects of Thurston’s seminal approach [22]).

The aim of this paper is to extend this approach to larger and more complicated building blocks. These blocks can be ideal polyhedra instead of tetrahedra, but they may also have non-trivial topology. In general, an *angled block* will be a 3-manifold whose boundary is subdivided into faces looking locally like the faces of an ideal polyhedron (in a sense to be defined). The edges between adjacent faces carry prescribed dihedral angles. In Section 2, we will describe the precise combinatorial conditions that the dihedral angles must satisfy. These conditions will imply the following generalization of a result by Lackenby [11, Corollary 4.6].

THEOREM 1.1. *Let M be a compact orientable 3-manifold with non-empty boundary, subdivided into angled blocks in such a way that the dihedral angles at each edge of M sum to 2π . Then ∂M consists of tori, and the interior of M admits a complete hyperbolic metric.*

One can prove that a particular manifold with boundary is hyperbolic in a spectrum of practical ways, ranging from local to global. In some cases, a combinatorial description of M naturally guides a way to subdivide it into tetrahedra (see, for example, [9] or [24]). In these cases, angled triangulations are highly useful. On the other extreme, one can study the global

Received 18 December 2006; revised 24 January 2008; published online 18 July 2008.

2000 *Mathematics Subject Classification* 57M25, 57M50.

Futer is partially supported by NSF grant DMS-0353717 (RTG), and Guéritaud is partially supported by NSF grant DMS-0103511.

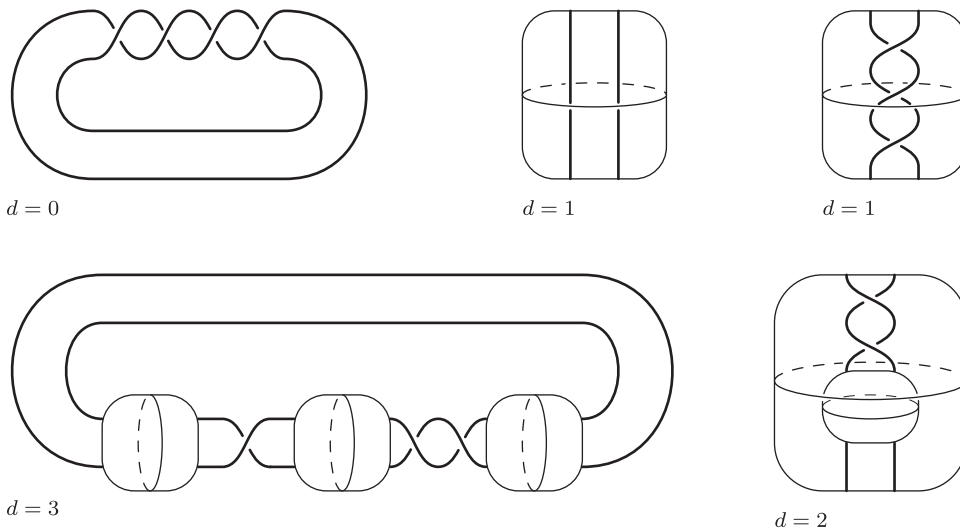


FIGURE 1.1. Examples of d -bracelets. The two 1-bracelets with different numbers of half-twists in their bands are homeomorphic.

topology of M and prove that it contains no essential spheres, disks, tori, or annuli; Thurston’s hyperbolization theorem then implies that $M \setminus \partial M$ is hyperbolic [23]. Theorem 1.1 provides a medium-range solution (still relying on Thurston’s theorem) for situations where M naturally decomposes into pieces that retain some topological complexity.

We will apply Theorem 1.1 to the complements of arborescent links, which are defined in terms of *bracelets*. We choose an orientation of \mathbb{S}^3 , to remain fixed throughout the paper.

DEFINITION 1.2. An *unknotted band* $A \subset \mathbb{S}^3$ is an annulus or Möbius band, whose core curve C is an unknotted circle. Such an A has a natural structure as an I -bundle over C , and we will refer to the fiber over a point of C as a *crossing segment* of the unknotted band A .

Consider the manifold M_d obtained by removing from \mathbb{S}^3 the open regular neighborhoods of d disjoint crossing segments of an unknotted band A , and let $K_d = M_d \cap \partial A$. Then a d -bracelet B_d is the pair (M_d, K_d) , as in Figure 1.1. We say that d is the *degree* of the bracelet.

Note that when $d > 0$, B_d is determined up to homeomorphism (of pairs) by the integer d . For example, when $d = 2$, B_2 is homeomorphic to the pair $(\mathbb{S}^2 \times I, \{4 \text{ points}\} \times I)$. When $d = 1$, M_1 is a 3-ball and K_1 is a pair of simultaneously boundary-parallel arcs; a 1-bracelet B_1 is commonly called a *trivial tangle*. When $d = 0$, B_0 is determined by the *number of half-twists* in the band: namely, the linking number of C with ∂A .

Let B_{d_1} and B_{d_2} be two bracelets with $d_i > 0$, and choose a boundary sphere S_i of each B_{d_i} . The S_i have natural orientations induced by the orientation of \mathbb{S}^3 , and we can glue S_1 to S_2 by any orientation-reversing homeomorphism sending the unordered 4-tuple of points $S_1 \cap K_1$ to the 4-tuple $S_2 \cap K_2$. The union of the K_{d_i} then defines a collection of arcs in a larger subset of \mathbb{S}^3 . More generally, if bracelets B_{d_1}, \dots, B_{d_n} are glued to form \mathbb{S}^3 (some of the d_i being 1), the arcs in these bracelets combine to form a link K in \mathbb{S}^3 , as in Figure 1.2.

DEFINITION 1.3. A link $K \subset \mathbb{S}^3$ is called *prime* if, for every 2-sphere S meeting K transversely in two points, at least one of the two balls cut off by S intersects K in a single

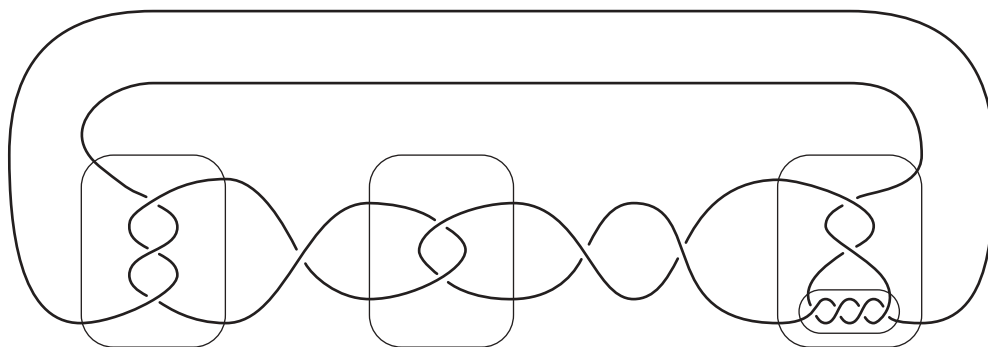


FIGURE 1.2. A generalized arborescent knot, obtained by gluing several bracelets.

boundary-parallel arc. If K is not prime, it is called *composite*. Note that with this convention, every split link (apart from the split link consisting of two unknots) is automatically composite.

DEFINITION 1.4. A knot or link $K = \cup_{i=1}^n K_{d_i}$ obtained when several bracelets are glued together to form \mathbb{S}^3 is called a *generalized arborescent link*. If, in addition, K is prime, then we say that it is an *arborescent link*.

The pattern of gluing bracelets to form a link can be represented by a tree T , in which a d -valent vertex corresponds to a d -bracelet and an edge corresponds to a gluing map of two neighboring bracelets. The term *arborescent*, from the Latin word *arbor* (tree), refers to this correspondence. Special cases of arborescent links include *two-bridge links*, which can be constructed by gluing two 1-bracelets, and *Montesinos links*, which can be constructed by gluing a single d -bracelet to d different 1-bracelets. Montesinos links are also known as *star links*, because the corresponding tree is a star.

The tree that represents an arborescent link carries a great deal of geometric and topological information. For example, Gabai has used trees to construct an algorithm that computes the genus of an arborescent link [7]. Bonahon and Siebenmann have used trees to completely classify arborescent links up to isotopy [3]. One geometric consequence of their work is the following result.

THEOREM 1.5 (Bonahon and Siebenmann). *The following three families, shown in Figure 1.3, form a complete list of non-hyperbolic arborescent links:*

- (I) K is the boundary of a single unknotted band;
- (II) K has two parallel components, each of which bounds a 2-punctured disk properly embedded in $\mathbb{S}^3 \setminus K$;
- (III) K or its reflection is the pretzel link $P(p, q, r; -1)$, where each of p, q, r is at least 2 and $1/p + 1/q + 1/r \geq 1$.

Furthermore, an effective algorithm decides whether a given generalized arborescent link K is prime, and whether it lies in one of the exceptional families.

Bonahon and Siebenmann's original proof of this theorem made strong use of the double branched covers of arborescent links. These covers are all *graph manifolds*, obtained by gluing Seifert fibered manifolds along incompressible tori that project to gluing spheres of d -bracelets. Their results and ideas were heavily quoted, but unfortunately the monograph containing the

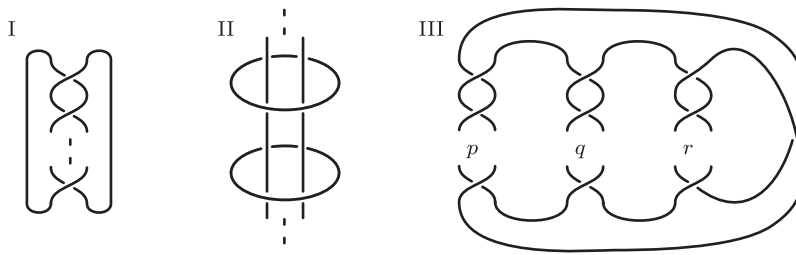


FIGURE 1.3. The three exceptional families of non-hyperbolic arborescent links. For family III, $p, q, r \geq 2$ and $1/p + 1/q + 1/r \geq 1$.

proof [3] remains unfinished. One of our primary motivations in this paper was to write down a version of the proof.

In the years since Bonahon and Siebenmann's monograph, several authors have re-proved parts of the theorem. Menasco [13] proved that a two-bridge link (more generally, a prime alternating link) is hyperbolic whenever it is not in family I. Oertel [15] proved that the complement of a Montesinos link contains an incompressible torus if and only if the link is in family III, with $1/p + 1/q + 1/r = 1$. Finally, it follows from Wu's work on Dehn surgery [25] that all non-Montesinos arborescent *knots* are hyperbolic.

It is fairly straightforward to check that the links listed in Theorem 1.5 are indeed non-hyperbolic. For families I and II, Figure 1.3 reveals an obvious annulus or Möbius band that forms an obstruction to the existence of a hyperbolic structure. Meanwhile, the pretzel link complements in family III contain (less obvious) incompressible tori when $1/p + 1/q + 1/r = 1$ (by Oertel's work [15]) and are Seifert-fibered when $1/p + 1/q + 1/r > 1$ by Sakuma's work [19] (in fact, such links are torus links unless (p, q, r) is a permutation of $(2, 2, n)$). In particular, all of these well-studied links are known to be prime. Thus we will focus our attention on proving that all the remaining arborescent links are indeed hyperbolic.

The proof is organized as follows. In Section 2, we will define angled blocks and prove Theorem 1.1. In Section 3, we will perform a detailed study of how d -bracelets can be glued along 4-punctured spheres. This will enable us to algorithmically simplify the bracelet presentation of any particular link and decide whether it is an exception. In Section 4, we will use the simplified bracelet structure to subdivide the link complement into tetrahedra and solid tori. The subdivision will work for all arborescent links except for families I and II. Finally, in Section 5, we will assign dihedral angles to edges on the boundary of the tetrahedra and solid tori. For links that are not in family III, these angles will satisfy the criteria of angled blocks, implying by Theorem 1.1 that the link complement is hyperbolic.

2. Angled blocks

In this section, we develop a theory of angled blocks that provides a practical way of proving that a given manifold is hyperbolic (Theorem 1.1). We lay out the necessary definitions in Subsection 2.1. In Subsection 2.2, we study the intersections between blocks and surfaces in a manifold, and prove that any surface can be placed into a sufficiently nice *normal form*. The angle structures on the blocks allow us to define a natural measure of complexity for the normal surfaces, called *combinatorial area*, which behaves like hyperbolic area. In Subsection 2.3, we will use combinatorial area considerations to show that M cannot contain any essential surfaces of non-negative Euler characteristic, so by Thurston's hyperbolization theorem M must admit a hyperbolic structure.

Our proof of Theorem 1.1 follows the same outline as Casson's proof that manifolds with an angled triangulation are hyperbolic, written down by Lackenby in [11, Section 4]. The credit for developing these ideas goes mainly to Casson and Lackenby.

2.1. From polyhedra to blocks

In studying a 3-manifold M , it is frequently useful to decompose M into pieces that are not contractible. This idea has been recently studied by other authors. Agol has described a way to cut a manifold into non-contractible *nanotubes* (personal communication), while Martelli and Petronio have cut a manifold into *bricks* [12]. Rieck and Sedgwick, among others, have investigated how a solid torus added during Dehn surgery can intersect a Heegaard surface [16]. Focusing on the individual pieces of the decomposition, Schlenker has studied manifolds with polyhedral boundary [21]. Our angled blocks fit into this theme.

DEFINITION 2.1. Let S be a closed oriented surface, and let $\Gamma \subset S$ be an embedded graph, each of whose vertices has degree at least 3. We say that Γ *fills* S if every component of $S \setminus \Gamma$ is an open disk, whose boundary consists of at least 3 edges of Γ . Given a graph Γ that fills a surface, we can construct a *dual graph* $\Gamma^* \subset S$, well defined up to isotopy, in the following fashion. Every disk of $S \setminus \Gamma$ defines a vertex of Γ^* . Every edge $e \subset \Gamma$ separates two faces of $S \setminus \Gamma$; we connect the corresponding vertices of Γ^* by a dual edge e^* . Finally, $S \setminus \Gamma^*$ is a union of disks, or faces, each corresponding to a vertex of Γ .

Note that this construction still makes sense if the surface S has several components. In this situation, both Γ and Γ^* will have as many components as S .

DEFINITION 2.2. A *block* \mathcal{P} is a compact, oriented, irreducible 3-manifold with boundary. We assume that \mathcal{P} does not contain any incompressible tori. (Even boundary-parallel incompressible tori in \mathcal{P} are prohibited, because they may cease to be boundary-parallel after several blocks are glued together.) As a consequence, if $\partial\mathcal{P}$ is a torus, then \mathcal{P} must be a solid torus.

Let Γ be a graph that fills $\partial\mathcal{P}$, whose edges are e_1, \dots, e_n . To every edge $e_i \subset \Gamma$ we assign an *internal angle* α_i and an *external angle* $\varepsilon_i = \pi - \alpha_i$. By duality, an edge $e_i^* \subset \Gamma^*$ receives the same angle as its dual edge $e_i \subset \Gamma$. We say that \mathcal{P} is an *angled block* if this assignment of angles satisfies the following properties:

- (1) $0 < \alpha_i < \pi$ for all i ,
- (2) $\sum_{\partial D} \varepsilon_i = 2\pi$ for every face D of $\partial\mathcal{P} \setminus \Gamma^*$, and
- (3) $\sum_{\gamma} \varepsilon_i > 2\pi$ for every simple closed curve $\gamma \subset \Gamma^*$ that bounds a disk in \mathcal{P} but is not the boundary of a face of $\partial\mathcal{P} \setminus \Gamma^*$.

Finally, we remove from \mathcal{P} all the vertices of Γ , making them *ideal vertices*. We will refer to the edges of Γ as the *edges of \mathcal{P}* , and to the faces of $\partial\mathcal{P} \setminus \Gamma$ as the *faces of \mathcal{P}* . Removing the vertices of Γ makes the faces of \mathcal{P} ideal polygons.

Property (1) says that \mathcal{P} is locally convex at every edge. Property (2) says that the link of every ideal vertex of \mathcal{P} has the angles of a convex Euclidean polygon. Property (3) is based on the following theorem of Rivin [18].

THEOREM 2.3 (Rivin). *Let \mathcal{P} be an angled polyhedron, that is, a contractible angled block. Then \mathcal{P} can be realized as a convex ideal polyhedron in \mathbb{H}^3 with the prescribed dihedral angles, uniquely up to isometry. Conversely, the dihedral angles of every convex ideal polyhedron in \mathbb{H}^3 satisfy (1)–(3).*

Such characterizations of polyhedra in \mathbb{H}^3 by their dihedral angles were first studied by Andreev [1]. We conjecture that an analogous result holds for non-contractible blocks as well.

CONJECTURE 2.4. *Let \mathcal{P} be an angled block. Then its universal cover $\tilde{\mathcal{P}}$ can be realized as a (possibly infinite) ideal polyhedron in \mathbb{H}^3 , with dihedral angles specified by \mathcal{P} , uniquely up to isometry.*

When \mathcal{P} is contractible, this conjecture is exactly Rivin's theorem. Schlenker [21, Theorem 8.15] has treated the case where \mathcal{P} has incompressible boundary. Finally, when all angles are at most $\pi/2$, the double of \mathcal{P} is a cone-manifold with cone angles at most π , and the conjecture follows from the cone-deformation proof of the hyperbolization theorem for orbifolds (see [2, 4]). We also note that the converse statement (that the ideal polyhedron $\tilde{\mathcal{P}}$ must satisfy (1)–(3)) is a fairly straightforward consequence of the Gauss–Bonnet theorem.

Our primary interest is in the manifolds that one may construct by gluing together angled blocks. To build a manifold with boundary, we first truncate all the ideal vertices of the blocks. As a result, a block \mathcal{P} has two kinds of faces: *interior faces* that are truncated copies of the original faces, and *boundary faces* that come from the truncated vertices. Similarly, \mathcal{P} has two kinds of edges: *interior edges* that are truncated edges of Γ , and *boundary edges* along the boundary faces. We note that a truncated block is a special case of a differentiable *manifold with corners* (modeled over \mathbb{R}_+^3 : see [5] for a general definition).

DEFINITION 2.5. Let $(M, \partial M)$ be a compact 3-manifold with boundary. An *angled decomposition* of M is a subdivision of M into truncated angled blocks, glued along their interior faces, such that $\sum \alpha_i = 2\pi$ around each interior edge of M . The boundary faces of the blocks fit together to tile ∂M .

Theorem 1.1 says that the interior of every orientable manifold with an angled decomposition must admit a hyperbolic structure. However, this is purely an existence result. An angled decomposition of a manifold is considerably weaker and more general than a hyperbolic structure, for two reasons. First, we do not know whether the blocks are actually geometric pieces — this is the content of Conjecture 2.4. Second, even when the blocks are known to be geometric, a geometrically consistent gluing must respect more than the dihedral angles. This means that faces of blocks that are paired in the gluing must be isometric — a condition that may not hold for faces larger than triangles. To obtain a complete hyperbolic structure, the truncated vertices of the blocks must also fit together to tile a horospherical torus, meaning that these Euclidean polygons must have consistent sidelengths as well as consistent angles.

There is an interesting contrast between the rigidity of a hyperbolic structure and the flexibility of angle structures. By Definitions 2.2 and 2.5, an angle structure on a block decomposition is a solution to a system of linear equations and (strict) linear inequalities. The solution set to this system, if non-empty, is an open convex polytope, so for every angled decomposition there is a continuum of deformations.

In fact, geometric angled blocks — for example, angled polyhedra — can serve as a stepping stone on the way to finding a complete hyperbolic structure. Every angled polyhedron has a well-defined volume determined by its dihedral angles, by Theorem 2.3. If the volume of an angled decomposition is critical in the polytope of deformations, we can exploit Schläfli's formula as in Rivin's theorem [17] and show that the polyhedra glue up to give a hyperbolic metric: this is carried out for some examples in [9] (where all blocks are tetrahedra). However, depending on the combinatorics of the decomposition, a critical point may or may not occur.

In fact, numerical experiments show that some of the decompositions that we will define for arborescent link complements in Section 4 admit angle structures, but have no critical point, hence no geometric realization in the hyperbolic metric.

2.2. Normal surface theory in angled blocks

To prove Theorem 1.1, we study the intersections between blocks and (smooth) essential surfaces.

DEFINITION 2.6. A surface $(S, \partial S) \subset (M, \partial M)$ is called *essential* if S is incompressible, boundary-incompressible, and not boundary-parallel, or if S is a sphere that does not bound a ball.

Our aim is to move any essential surface into a form where its intersections with the individual blocks are particularly nice.

DEFINITION 2.7. Let \mathcal{P} be a truncated block, and let $(S, \partial S) \subset (\mathcal{P}, \partial \mathcal{P})$ be a surface. We say that S is *normal* if it satisfies the following properties:

- (1) every closed component of S is incompressible in \mathcal{P} ;
- (2) S and ∂S are transverse to all faces and edges of \mathcal{P} ;
- (3) no component of ∂S lies entirely in a face of $\partial \mathcal{P}$;
- (4) no arc of ∂S in a face of \mathcal{P} runs from an edge of \mathcal{P} back to the same edge;
- (5) no arc of ∂S in an interior face of \mathcal{P} runs from a boundary edge to an adjacent interior edge.

Given a decomposition of M into blocks, a surface $(S, \partial S) \subset (M, \partial M)$ is called *normal* if for every block \mathcal{P} , the intersection $S \cap \mathcal{P}$ is a (possibly non-connected) normal surface in \mathcal{P} .

THEOREM 2.8. Let $(M, \partial M)$ be a manifold endowed with a fixed block decomposition.

- (a) If M is reducible, then M contains a normal 2-sphere.
- (b) If M is irreducible and ∂M is compressible, then M contains a normal disk.
- (c) If M is irreducible and ∂M is incompressible, then any essential surface can be moved by isotopy into normal form.

Proof. The following argument is the standard procedure for placing surfaces in normal form with respect to a triangulation or polyhedral decomposition [10]. As long as all faces of all blocks are disks, the topology of the blocks never becomes an issue. We will handle part (c) first, followed by (b) and (a).

For (c), assume that M is irreducible and ∂M is incompressible. Let $(S, \partial S)$ be an essential surface in $(M, \partial M)$. To move S into normal form, we need to check the conditions of Definition 2.7. Since S is essential in M , it automatically satisfies (1). Furthermore, a small isotopy of S ensures the transversality conditions of (2).

Consider the intersections between S and the open faces of the blocks, and let γ be one component of intersection. Note that by Definition 2.1, the face F containing γ is contractible. We want to make sure that γ satisfies (3), (4), and (5).

(3) Suppose that γ is a closed curve, violating (3). Without loss of generality, we may assume that γ is innermost on the face F . Then γ bounds a disk $D \subset F$, whose interior is disjoint from S . However, since S is incompressible, γ also bounds a disk $D' \subset S$. Furthermore, since we have assumed that M is irreducible, the sphere $D \cup_{\gamma} D'$ must bound a ball. Thus we may isotope S through this ball, moving D' past D . This isotopy removes the curve γ from the intersection between S and F .

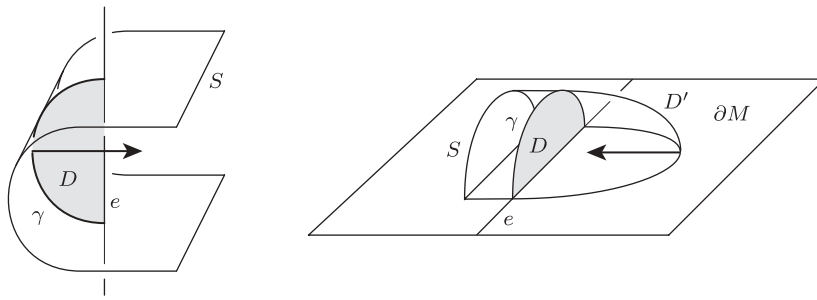


FIGURE 2.1. When a surface violates condition (4) of normality, an isotopy in the direction of the arrow removes intersections between S and the faces of M .

(4) Suppose that γ runs from an edge e back to e , violating (4). Then γ and e co-bound a disk $D \subset F$, and we can assume that γ is innermost (that is, S does not meet D again). If e is an interior edge, we can use this disk D to guide an isotopy of S past the edge e , as in the left panel of Figure 2.1. This isotopy removes γ from the intersection between S and F (some intersection components between S and the interiors of faces other than F may merge, but their total number always decreases).

If γ lies in a boundary face, then the situation is very similar to the previous paragraph. This time, the disk D guides an isotopy of S along ∂M , simplifying the intersection between S and the faces of the blocks.

Finally, if e is a boundary edge and F is an interior face, then D is a boundary compression disk for S . Since S is boundary-incompressible, γ must also cut off a disk $D' \subset S$, as in the right panel of Figure 2.1. Since M is irreducible and ∂M is incompressible, it follows that the disk $D \cup_\gamma D'$ is boundary-parallel: $D \cup D' \cup \Delta$ bounds a ball B , for some disk $\Delta \subset \partial M$. We must ask on which side of $D \cup_\gamma D'$ the ball B lies: if a neighborhood of the arc γ in the surface S meets the interior of B , then S is a disk of B and is boundary-parallel (recall that S does not meet D again, because γ is innermost among the arcs running from e back to e). Therefore, S does not meet the interior of B . In particular, D' is isotopic to D by an isotopy sweeping out B and missing $S \setminus D'$. This defines an isotopy of S that can be extended slightly to move D' past D , thus removing the curve γ from the intersection between S and F .

(5) Suppose that γ runs from a boundary edge to an adjacent interior edge, violating (5). Then γ once again cuts off a disk D . By isotoping S along this disk, as in Figure 2.2, we remove γ from the intersection.

It is immediate to check that each of the last three moves reduces the number of components of $S \cap Z$, where Z is the union of the interiors of the faces of M . Thus, after a finite number of isotopy moves, S becomes normal.

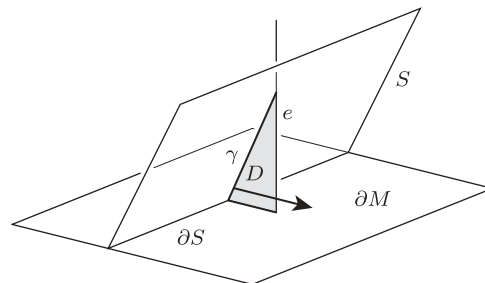


FIGURE 2.2. When a surface violates condition (5) of normality, a ∂M -preserving isotopy of S along the disk D , in the direction of the arrow, removes intersections between S and the faces of M .

For part (b), assume that M is irreducible and ∂M is compressible. Let S be an essential disk in M ; under our assumptions, S must be a compression disk for ∂M . To move S into normal form, we follow a procedure very similar to that in part (c). In particular, condition (1) of Definition 2.7 is vacuous because S has no closed components. Furthermore, a small isotopy of S ensures the transversality conditions of (2). Focusing our attention on conditions (3)–(5), let γ be one component of intersection between S and a face F of a block.

If γ is a simple closed curve, violating (3), the argument is exactly the same as above. We find that γ bounds a disk $D \subset F$ and an isotopic disk $D' \subset S$, because S is incompressible and M is irreducible. Thus we may isotope S past D .

If γ runs from an edge e back to e , violating (4), the argument is mostly the same as above. If e is an interior edge, or γ lies in a boundary face, then the exact isotopies described in part (c-4) will guide S past e . If e is a boundary edge and γ lies in an interior face F , then γ and e co-bound a disk $D \subset F$; up to replacing γ with an outermost arc of $D \cap S$ on F , we may assume that $D \cap S = \gamma$ so that the disk D realizes a boundary compression of S . The situation is similar to the right panel of Figure 2.1, except that now γ splits S into disks D_1 and D_2 (since S itself is a disk). At least one $D_i \cup_\gamma D$ must be essential in M , because if they were both boundary-parallel, S would be boundary-parallel too. If by S we now denote this essential disk $D_i \cup_\gamma D$, then S can be pushed away from the face F .

Finally, if γ runs from a boundary edge to an adjacent interior edge e , violating (5), an isotopy of S as in Figure 2.2 will remove γ from the intersections between S and the faces of the blocks. Each of the last three moves simplifies the intersections between S and the faces, so a repeated application will place S in normal form.

For part (a), assume that the manifold M is reducible, and let $S \subset M$ be a sphere that does not bound a ball. We will move S into normal form by checking the conditions of Definition 2.7. Note that by Definition 2.2, an essential sphere can never be contained in a single block, so condition (1) is vacuous. A small isotopy of S ensures the transversality conditions of (2). Note as well that condition (5) is vacuous, because S is closed. To satisfy conditions (3) and (4), let γ be one arc of intersection between S and a face F of a block.

If γ is a simple closed curve, violating (3), we may assume as before that γ is innermost in F . Thus γ bounds a disk $D \subset F$ whose interior is disjoint from S . Because S is a sphere, we may write $S = D_1 \cup_\gamma D_2$ for disks D_1 and D_2 . Suppose that each $D_i \cup_\gamma D$ bounds a ball B_i . Because the boundaries of B_1 and B_2 intersect exactly along a single disk D , either one ball contains the other or they have disjoint interiors. In either scenario, it follows that $S = D_1 \cup D_2$ must bound a ball, which is a contradiction. Thus, since at least one $D_i \cup_\gamma D$ must fail to bound a ball, we can replace D by one of the D_i . The resulting sphere, which we continue to call S , can be pushed away from the face F .

If γ runs from an edge e back to e , violating (4), then γ and e co-bound a disk D . As before, we can use D to guide an isotopy of S past e . (See Figure 2.1, left.) Note that since S is closed, γ must be an interior edge.

By repeating these moves, we eventually obtain a sphere in normal form. \square

2.3. Combinatorial area

So far, we have not used the dihedral angles of the blocks. Their use comes in estimating the complexity of normal surfaces.

DEFINITION 2.9. Let \mathcal{P} be an angled block, and denote by ε_δ the exterior dihedral angle at the edge δ . Truncate the ideal vertices of \mathcal{P} , and label every boundary edge δ with a dihedral angle of $\varepsilon_\delta = \pi/2$. Let S be a normal surface in \mathcal{P} , and let $\delta_1, \dots, \delta_n$ be the edges of the truncated block \mathcal{P} met by ∂S (each edge may be counted several times). We define the

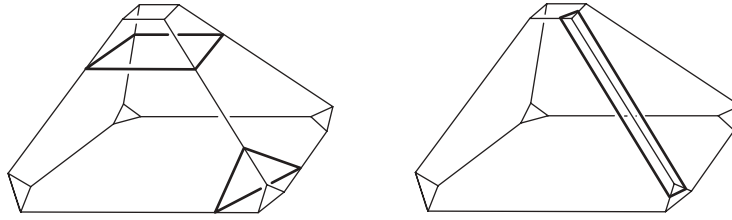


FIGURE 2.3. In any angled block, vertex links (left) and boundary bigons (right) are the only connected normal surfaces of area 0.

combinatorial area of S to be

$$a(S) = \sum_{i=1}^n \varepsilon_{\delta_i} - 2\pi\chi(S).$$

For the sake of brevity, we will refer to the above sum of dihedral angles ($\sum_{i=1}^n \varepsilon_{\delta_i}$) as $\sum_{\partial S} \varepsilon_i$.

Note that by the Gauss–Bonnet theorem, $a(S)$ is just the area of a hyperbolic surface with piecewise geodesic boundary, with exterior angles ε_i along the boundary and Euler characteristic $\chi(S)$.

PROPOSITION 2.10. *Let S be a normal surface in a truncated angled block \mathcal{P} . Then $a(S) \geq 0$. Furthermore, if $a(S) = 0$, then every component of S is a vertex link (boundary of a regular neighborhood of a boundary face) or a boundary bigon (boundary of a regular neighborhood of an interior edge), as in Figure 2.3.*

Proof. Because combinatorial area is additive over multiple components of S , it suffices to consider the case when S is connected. Furthermore, when $\chi(S) < 0$, $a(S) > 0$, so it suffices to consider the case when $\chi(S) \geq 0$. By Definition 2.2, \mathcal{P} is irreducible and atoroidal, so S cannot be a sphere or torus. If S is an annulus, then $a(S) = \sum_{\partial S} \varepsilon_i > 0$, because ∂S must intersect some edges and the dihedral angle on each edge is positive. Thus the only remaining case is when S is a disk.

For the rest of the proof, let $D \subset \mathcal{P}$ be a normal disk. We consider three cases, conditioned on n , the number of intersections between D and the boundary faces.

Case 0: $n = 0$. Recall, from Definitions 2.1 and 2.2, that every interior face of \mathcal{P} corresponds to a complementary region of the graph Γ and to a vertex of the dual graph Γ^* . Thus ∂D defines a closed path γ through the edges of Γ^* ; this is a non-backtracking path because no arc of ∂D runs from an edge back to itself. The path γ may pass through an edge multiple times, but it contains a simple closed curve in Γ^* . Thus, by Definition 2.2, $\sum_{\partial D} \varepsilon_i \geq 2\pi$. Equality can occur only when ∂D encircles an ideal vertex, in other words when D is a vertex link.

Case 1: $n = 1$. The two boundary edges crossed by ∂D contribute π to the external angle sum of ∂D . Thus we may isotope ∂D off the boundary face F without increasing the angle sum, since by Definition 2.2 the interior edges meeting F have a total angle of 2π . Let D' be the resulting disk, and a_1, \dots, a_k be the intersection points, numbered consecutively, of $\partial D'$ with interior edges of the block near F (in Figure 2.4, $k = 2$).

We claim that D' is normal. Since $n = 1$, the only way that D' can fail Definition 2.7 is if an arc of $\partial D'$ violates condition (4) and runs from an interior edge e back to itself. This cannot happen between a_i and a_{i+1} , otherwise the block would have a monogon face, in contradiction with Definitions 2.1 and 2.2. Hence condition (4) is violated by an arc starting from a_1 in the

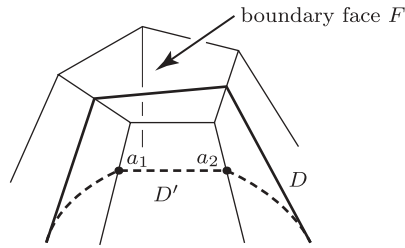


FIGURE 2.4. We may isotope ∂D off a boundary face of \mathcal{P} , producing a normal disk D' with $a(D') \leq a(D)$.

direction opposite a_2 to end on the same interior edge e (or by an analogous arc from a_k). But then the corresponding arc of D must connect e to an adjacent boundary edge, contradicting condition (5). Thus, by case (0), $a(D) \geq a(D') \geq 0$.

It remains to show that the area of D is strictly positive. If $a(D') > 0$, this is automatic. If $a(D') = 0$, then D' must be a vertex link, parallel to a boundary face F_1 . This F_1 cannot be the boundary face F that intersects D , for otherwise all of D is parallel to F , forcing some arc of ∂D to violate condition (5). As a result, the interior edges containing points a_1, \dots, a_k must run in parallel between F_1 and F . However, by Definition 2.1, the boundary of a block cannot have parallel edges, so in fact $k = 1$. Thus, because the (exterior) dihedral angle at a_1 is strictly less than π , we have $a(D) > a(D') \geq 0$.

Case 2: $n \geq 2$. Since ∂D crosses at least 4 boundary edges, $a(D) \geq 0$, with equality only if $n = 2$ and ∂D is disjoint from the interior edges. We restrict our attention to this case, and claim that D is a boundary bigon.

Push ∂D off the two boundary faces F_1 and F_2 , in a way that minimizes the angle sum of the new disk D' . Denote by a_1, \dots, a_k and b_1, \dots, b_l the points where $\partial D'$ crosses interior edges near F_1 and F_2 , respectively. Orient the edge germs containing the a_i and b_j away from the faces F_1 and F_2 . If A and B are the sums of the angles of D' at the a_i and b_i , respectively, then $A, B \leq \pi$.

Suppose that D' is normal. Since we know that $a(D') \leq a(D) = 0$, it follows by case (0) that D' must be the vertex link associated to a boundary face F' . Moreover, we have $A = B = \pi$, and hence $k, l \geq 2$. Let us isotope D' into $\partial \mathcal{P}$ while keeping its boundary fixed, so that after the isotopy, D' contains the boundary face F' as well as initial segments of all interior edges starting at F' : these initial segments end at $a_1, \dots, a_k, b_1, \dots, b_l$, in that cyclic order around F' . Suppose that the orientations on the interior edges through the a_i are inward for D' (in particular, this will happen whenever $F_1 \neq F'$). Then, since $k \geq 2$, it follows that $\partial \mathcal{P}$ contains an ideal bigon, which is impossible. Therefore the orientations point outward, which implies notably that $F_1 = F'$. Similarly, $F_2 = F'$. As a result, the boundary of the original disk D violates condition (4), for example, at the boundary edge situated between a_k and b_1 (Figure 2.5, left), which is a contradiction.

Therefore D' is not normal: define $\Delta = D'$ (we are going to modify Δ , but not D'). Then the loop $\partial \Delta$ must violate (4), running in a U-turn from an interior edge e back to e : we can isotope the disk Δ so as to erase this U-turn. The angle sum of Δ decreases to a value less than 2π , so Δ is even less normal now (by case (0)). If $\partial \Delta$ still crosses any (interior) edges, then we can repeat the operation, until $\partial \Delta$ violates (3), and Δ can be isotoped into an interior face (recall that the block \mathcal{P} is irreducible). Therefore D' can be isotoped, with fixed boundary, to a disk in the union of all (open) interior faces and interior edges. Interior edges must connect across D' the points $a_1, \dots, a_k, b_1, \dots, b_l$ of $\partial D'$, which are still cyclically ordered (Figure 2.5, right.) If some edge goes from a_i to a_j (where $i < j$) then there must be an edge from a_s to a_{s+1} for some $i \leq s < j$, and therefore $\partial \mathcal{P}$ contains an ideal monogon, which is impossible. Therefore every edge across D' runs from an a_i to a b_j , in fact to b_{k+1-i} (and we have $k = l$). If $k \geq 2$,

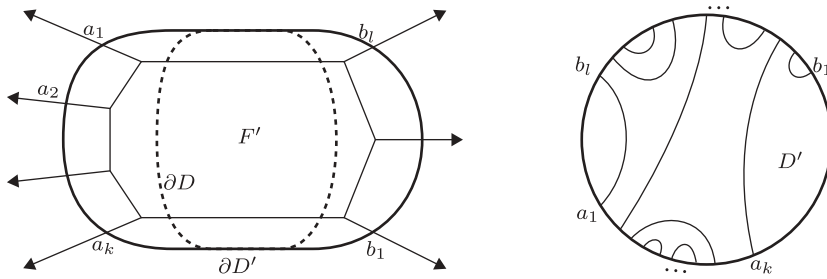


FIGURE 2.5. When $n = 2$ and $a(D) = 0$, we have a contradiction for D' normal (left), as well as for D' non-normal (right), unless $k = l = 1$.

then $\partial\mathcal{P}$ contains an ideal bigon. Therefore $k = 1$, so D' is traversed by a single edge e , and the original disk D is the boundary bigon associated to e . \square

For an essential surface $(S, \partial S) \subset (M, \partial M)$, we can define the combinatorial area $a(S)$ by adding up the areas of its intersections with the blocks. This definition of combinatorial area was designed to satisfy a Gauss–Bonnet relationship.

PROPOSITION 2.11. *Let $(S, \partial S) \subset (M, \partial M)$ be a surface in normal form. Then*

$$a(S) = -2\pi\chi(S).$$

Proof. Consider the decomposition of S into S_1, \dots, S_n , namely its components of intersection with the various blocks. Let $S' = S_{i_1} \cup \dots \cup S_{i_k}$ be a union of some S_{i_i} glued along some (not necessarily all) of their edges: S' is a manifold with polygonal boundary. Define the interior angle of S' at a boundary vertex to be the sum of the interior angles of the adjacent S_{i_α} , and the exterior angle as the complement to π of the interior angle. It is enough to prove that

$$\sum_{\alpha=1}^k a(S_{i_\alpha}) = \sum_{\partial S'} \varepsilon_i - 2\pi\chi(S'), \tag{2.1}$$

where the ε_i are the exterior angles of S' : the result will follow by taking $S' = S$ (the union of all S_i glued along all their edges), because all ε_i are then equal to $\pi - (\pi/2 + \pi/2) = 0$. Since M is orientable, up to replacing S with the boundary of its regular neighborhood, we can restrict to the case where S is orientable.

We prove (2.1) by induction on the number of gluing edges, where the set of involved components S_{i_α} is chosen once and for all. When no edges are glued, (2.1) follows from Definition 2.9. It remains to check that the right-hand side of (2.1) is unchanged when two edges are glued together. In what follows,

- ν is the number of boundary vertices of S' ,
- θ is the sum of all interior angles along $\partial S'$, and
- χ is the Euler characteristic of S' .

Thus the right-hand side of (2.1) is $\nu\pi - \theta - 2\pi\chi$.

If we glue edges ab and cd , where a, b, c, d are distinct vertices of S' , then θ is unchanged, but ν goes down by 2 and χ goes down by 1: (2.1) is preserved.

If we glue edges ab and bc by identifying a and c , where a, b, c are distinct vertices, then θ goes down by 2π , because b becomes an interior vertex, while ν goes down by 2 and χ is unchanged: (2.1) is preserved.

If we glue two different edges of the form ab , closing off a bigon boundary component of S' , then θ goes down by 4π because both a and b become interior vertices. Since ν goes down by 2, and χ goes up by 1, again (2.1) is preserved.

If we glue an edge ab to a monogon boundary component cc , where a, b, c are distinct vertices, then θ is unchanged, while ν goes down by 2 and χ goes down by 1. If we glue two boundary monogons aa and bb together (where $a \neq b$), then θ goes down by 2π , while ν goes down by 2 and $\chi(S')$ is unchanged. In all cases, (2.1) is preserved. \square

We are now ready to complete the proof of Theorem 1.1.

THEOREM 1.1. *Let $(M, \partial M)$ be an orientable 3-manifold with an angled decomposition. Then ∂M consists of tori, and the interior of M is hyperbolic.*

Proof. Each component of ∂M is tiled by boundary faces of the blocks. Just inside each boundary face, a block has a normal disk of area 0. These vertex links glue up to form a closed, boundary-parallel normal surface S of area 0. By Proposition 2.11, $\chi(S) = 0$, and since M is orientable, the boundary-parallel surface S must be a torus. Thus ∂M consists of tori.

By Thurston's hyperbolization theorem [23], the interior of M carries a complete, finite-volume hyperbolic metric if and only if M contains no essential spheres, disks, annuli, or tori. By Theorem 2.8, if M has such an essential surface, then it has one in normal form. A normal sphere or disk has a positive Euler characteristic, and hence a negative area; thus it cannot occur.

A normal torus $T \subset M$ has area 0 and thus, by Proposition 2.10, must be composed of normal disks of area 0. Since T has no boundary, these must all be vertex links, which glue up to form a boundary-parallel torus. Similarly, a normal annulus $A \subset M$ must be composed entirely of bigons, since a bigon cannot be glued to a vertex link. However, a chain of bigons forms a tube around an edge of M , which is certainly not essential. Thus we can conclude that M is hyperbolic. \square

3. A simplification algorithm for arborescent links

Recall, from the introduction, that a generalized arborescent link is constructed by gluing together a number of d -bracelets. In this section, we describe an algorithm that takes a particular link and simplifies its bracelet presentation into a reduced form. This algorithm, directly inspired by Bonahon and Siebenmann's work [3], has several uses. First, if a given generalized arborescent link is composite, the algorithm will decompose it into its prime arborescent pieces. Second, the simplified bracelet description will allow us to rapidly identify the non-hyperbolic arborescent links listed in Theorem 1.5. In particular, the algorithm recognizes the unknot from among the family of generalized arborescent links. Finally, the simplified bracelet form of an arborescent link turns out to be the right description for the block decomposition of the link complement that we undertake in Section 4.

3.1. Slopes on a Conway sphere

Whenever two bracelets are glued together, they are joined along a 2-sphere that intersects the link K in 4 points. This type of sphere, called a *Conway sphere*, defines a 4-punctured sphere in the link complement. Our simplification algorithm is guided by the way in which gluing maps act on arcs in Conway spheres.

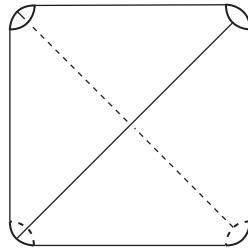


FIGURE 3.1. Arcs of slope 0, 1, and ∞ give an ideal triangulation of a 4-punctured sphere S .

DEFINITION 3.1. Let S be a 4-punctured sphere. An *arc pair* $\gamma \subset S$ consists of two disjoint, properly embedded arcs γ_1 and γ_2 , such that γ_1 connects two punctures of S and γ_2 connects the remaining two punctures of S . A *slope* on S is an isotopy class of arc pairs, and is determined by any one of the two arcs.

To visualize slopes, it helps to picture S as a pillowcase in \mathbb{R}^3 surrounding the unit square of \mathbb{R}^2 , with punctures at the corners. (See Figure 3.1.) Any arc pair on the pillowcase can then be straightened so that its intersections with the front of the pillow have a well-defined Euclidean slope. A *marking* of S (that is, a fixed homeomorphism between S and the pillowcase of Figure 3.1) induces a bijection between slopes on S and elements of $\mathbb{P}^1\mathbb{Q} = \mathbb{Q} \cup \{\infty\}$.

Slopes on 4-punctured spheres can be neatly represented by the *Farey complex* \mathcal{F} , shown in Figure 3.2. Vertices of \mathcal{F} correspond to slopes (arc pairs), edges of \mathcal{F} to disjoint slopes, and triangles to triples of disjoint slopes. Observe that a choice of three disjoint arc pairs of different slopes gives an ideal triangulation of S . Figure 3.2 also illustrates that \mathcal{F} (with its vertices removed) is homeomorphic to the Poincaré disk and can be endowed with a hyperbolic metric, making the triangles of \mathcal{F} straight ideal triangles. The dual of \mathcal{F} is an infinite trivalent planar tree.

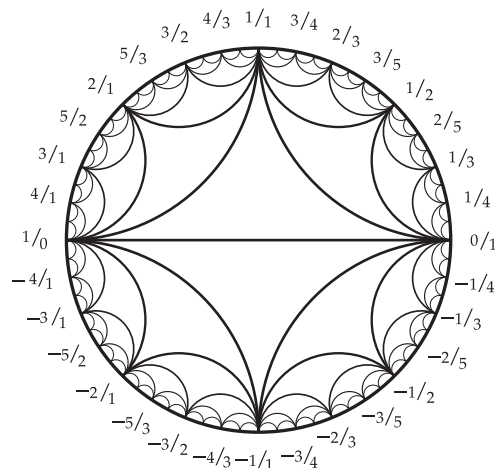


FIGURE 3.2. The Farey complex \mathcal{F} of a 4-punctured sphere (graphic by Allen Hatcher).

DEFINITION 3.2. Let S be a boundary sphere of a d -bracelet. This Conway sphere will be assigned a *preferred slope*, as follows. When $d > 1$, choose a crossing segment on each side of S (see Definition 1.2). If we isotope these two segments into S , we get an arc pair whose slope is the preferred slope of S . When $d = 1$, K_1 consists of two arcs that can be isotoped into S ; their slope is then the preferred slope of S . Note that the two definitions (for $d = 1$ and $d > 1$) are truly different. The case $d = 0$ is empty (a 0-bracelet has no boundary spheres).

3.2. *The algorithm*

We will perform the following sequence of steps to simplify the bracelet presentation of a generalized arborescent link.

(1) *Remove all 2-bracelets.* As Figure 1.1 illustrates, the two boundary components of a 2-bracelet B_2 are isotopic, and moreover B_2 is homeomorphic to the pair

$$(\mathbb{S}^2 \times I, \{4 \text{ points}\} \times I).$$

Thus, whenever a 2-bracelet sits between two other bracelets, those other bracelets can be glued directly to one another, with the gluing map adjusted accordingly.

(2) *Remove needless 1-bracelets.* Suppose that a 1-bracelet B_1 is glued to a d -bracelet B_d (with $d > 1$), and that their preferred slopes at the gluing Conway sphere are Farey neighbors. Then the two arcs of K_1 can be isotoped to lie on the Conway sphere ∂B_1 , without intersecting the crossing segments of B_d . As a result, the arcs $K_d \cup K_1$ combine to form the band of a $(d-1)$ -bracelet, as in Figure 3.3. Thus we may remove B_1 and replace B_d by a bracelet B_{d-1} , with one fewer boundary component.

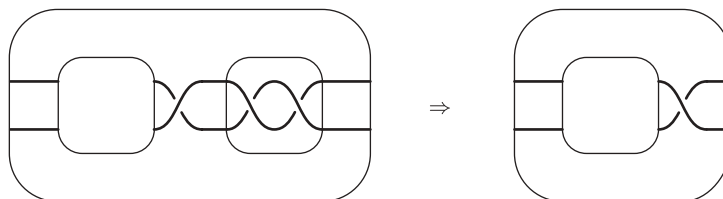


FIGURE 3.3. *Removing a needless 1-bracelet.*

(3) *Undo connected sums.* Suppose that a 1-bracelet B_1 is glued to a d -bracelet B_d (with $d > 1$), and that their preferred slopes are equal. Then there are several different 2-spheres that pass through the trivial tangle of B_1 and intersect K in a pair of points connected by a crossing segment. In this situation, we cut K along the crossing segments of B_d , decomposing it as a (possibly trivial) connected sum of $d - 1$ other links, as in Figure 3.4. On the level of bracelets, each piece of K that was glued to a Conway sphere of B_d will instead be glued to its own 1-bracelet, whose slope on the gluing sphere is given by the crossing segments of B_d . After this cutting operation, we will work separately with each of the $d - 1$ new links. That is, we will apply the reduction algorithm to simplify the bracelet presentation of each of these links. The algorithm may reveal that one or more of the new links is actually the unknot (see Theorem 3.9), and thus that we have undone a trivial connected sum. In this case, we may simply throw away the trivial pieces, having still gained the benefit of a simpler bracelet presentation of K .

(4) *Repeat steps (1)–(3), as necessary.* Note that removing a needless 1-bracelet can create a new 2-bracelet (as in Figure 3.3), and removing a 2-bracelet can change the gluing map of a 1-bracelet. However, since each of the above steps reduces the total number of Conway spheres in the construction of K , eventually we reach a point where none of these reductions is possible.

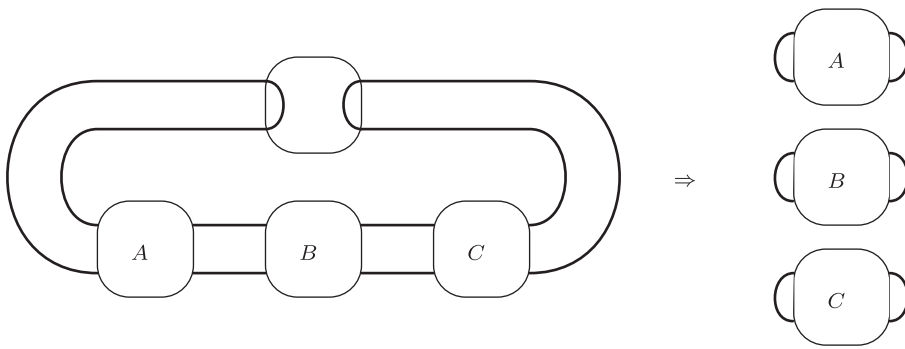


FIGURE 3.4. Special 1-bracelets decompose a link as a connected sum.

DEFINITION 3.3. Let $A \subset \mathbb{S}^3$ be an unknotted band, and let $T = \mathbb{S}^3 \setminus A$ be the open solid torus equal to the complement of A . Let L_n be a link consisting of n parallel, unlinked copies of the core of T .

Recall from Definition 1.2 that a d -bracelet B_d is the pair (M_d, K_d) , where M_d is the complement of a regular neighborhood of d disjoint crossing segments of A , and $K_d = M_d \cap \partial A$. We define an n -augmented d -bracelet to be $B_{d,n} = (M_d, K_d \cup L_n)$. Thus a traditional d -bracelet B_d corresponds to taking $n = 0$. When $n \geq 1$, for all positive d (including $d = 1$) the preferred slope of $B_{d,n}$ at a boundary (Conway) sphere is the slope of the crossing segments of A at this sphere.

Augmented bracelets naturally arise from certain configurations of d -bracelets. We continue our algorithm as follows.

(5) *Create augmented bracelets.* Let B_3 be a 3-bracelet glued to 1-bracelets B_1 and B'_1 . Suppose that there is a marking of the boundary spheres of B_3 such that the preferred slope of B_3 is ∞ and the preferred slopes of B_1 and B'_1 are in $\mathbb{Z} + 1/2$. (For example, in the two trivial tangles of Figure 3.5, these slopes are $-1/2$ and $1/2$; but other values would just correspond to adding half-twists in the “horizontal” band of B_3 . An intrinsic criterion is: the slopes of B_3 and B_1 (respectively B'_1) are not Farey neighbors, but they share exactly two common Farey neighbors.) In this situation, we will replace $B_3 \cup B_1 \cup B'_1$ by a once-augmented 1-bracelet $B_{1,1}$, as in Figure 3.5. Note that the closed loop in $B_{1,1}$ can be isotoped to lie on the boundary sphere S . Up to isotopy, there is exactly one arc pair on S that is disjoint from this loop; its slope is the preferred slope of $B_{1,1}$.

Remark: If B_3 is glued to three different 1-bracelets, each with slope in $\mathbb{Z} + 1/2$ (so the link contains no other bracelets than these four), we break the symmetry by choosing two of the 1-bracelets for augmentation.

DEFINITION 3.4. A (possibly augmented) bracelet $B_{d,n}$ is *large* if $d \geq 3$ or $n \geq 1$.

(6) *Combine large bracelets when possible.* Suppose that large bracelets $B_{d,n}$ and $B_{d',n'}$ are glued together along a Conway sphere, with their preferred slopes equal. Then we will combine them into a single $(d + d' - 2)$ -bracelet, augmented $(n + n')$ times. Note that at the beginning of this step, the only augmented bracelets are of the form $B_{1,1}$, created in step (5). However, under certain gluing maps, several bracelets of this form may combine with other large bracelets to form n -augmented d -bracelets, with d and n arbitrarily large.

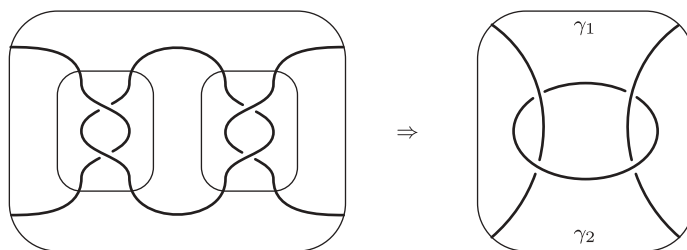


FIGURE 3.5. Creation of an augmented 1-bracelet. The arc pair $\gamma_1 \cup \gamma_2$ defines the preferred slope of $B_{1,1}$.

(7) *Form 0-bracelets and augmented 0-bracelets.* Consider a 1-bracelet B_1 , with preferred slope s . For any arc pair $\gamma \subset \partial B_1$ whose slope is a Farey neighbor of s , we can construct a rectangular strip in B_1 with boundary $K_1 \cup \gamma$. Therefore, when bracelets B_1 and B'_1 are glued together and their preferred slopes share a common neighbor in \mathcal{F} , we can glue these two rectangular strips to form an annulus or Möbius band whose boundary is $K_1 \cup K'_1$. In this situation, we replace $B_1 \cup B'_1$ by a single 0-bracelet.

In a similar fashion, an augmented 1-bracelet $B_{1,n}$ contains a rectangular strip whose intersection with the boundary sphere defines the preferred slope of $B_{1,n}$. Therefore, when $B_{1,n}$ is glued to a 1-bracelet B_1 and their preferred slopes are Farey neighbors, we once again have an annulus or Möbius band. In this situation (similar to step (2)), we replace $B_1 \cup B_{1,n}$ by a single augmented bracelet $B_{0,n}$, as in Figure 3.6.

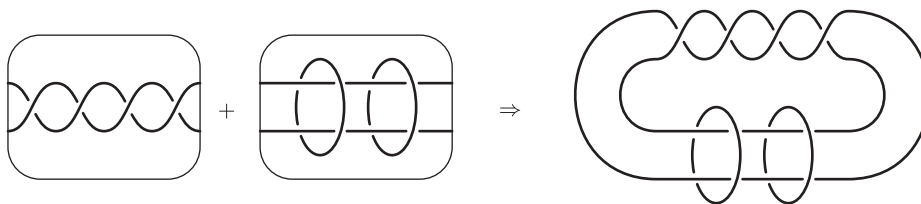


FIGURE 3.6. Creating an augmented 0-bracelet.

REMARK 3.5. No further instances of steps (1)–(3) occur after the creation of augmented bracelets in step (5). This is clear for step (1): no (unaugmented) 2-bracelets appear, not even in step (6) because the bracelets that merge in (6) are already large. Next, observe that the preferred slopes on Conway spheres of large bracelets are never changed after step (5) (not even when Conway spheres are cancelled in (6)). An easy discussion then implies that steps (2)–(3), or their analogues for augmented bracelets, never occur.

The following result summarizes the output of the simplification algorithm.

PROPOSITION 3.6. *For every generalized arborescent link given as input, the algorithm above produces several ‘output links’, of which the input is a connected sum. Let K be such an output link. Then K is expressed as a gluing of (possibly augmented) bracelets, in which all 2-bracelets are augmented.*

Furthermore, suppose that bracelets B and B' are glued along a Conway sphere. Any path through the 1-skeleton of the Farey complex connecting the preferred slope of B to the preferred slope of B' must contain at least the following number of edges (Table 1):

TABLE 1.

	$B'_{1,0}$	$B'_{d,n}$ large
$B_{1,0}$	3	2
$B_{d,n}$ large	2	1

Proof. Observe that the reduction algorithm only changes the topological type of K in step (3), where it cuts K into (possibly trivial) connected summands. Thus the output links do in fact sum to K .

Top-left entry of the table: if the preferred slopes of bracelets $B_{1,0}$ and $B'_{1,0}$ are at distance 2 (or less) in the Farey graph, they share a Farey neighbor and thus step (7) has reduced them to a single 0-bracelet. Similarly, the reduction of step (6) accounts for the bottom-right entry, and step (2) for the non-diagonal entries. □

3.3. Analysing the output

We are now ready to recognize the non-hyperbolic arborescent links listed in Theorem 1.5. After the simplification algorithm, they can appear in any one of the following four ways.

(1) *Bracelets augmented more than once.* A bracelet $B_{d,n}$, where $n \geq 2$, will contain two isotopic link components, as in Figure 3.6. Each of these parallel components bounds a disk that is punctured twice by the strands of K_d . Thus any link containing such a bracelet falls in the exceptional family II.

(2) *0-bracelets.* By Definition 1.2, the link contained in a 0-bracelet is the boundary of an unknotted band. These links fall in the exceptional family I.

(3) *Once-augmented 0-bracelets.* Let K be the link contained in an augmented 0-bracelet with r half-twists. By reflecting K if necessary, we may assume that $r \geq 0$. Now, we consider the following three cases.

(a) $r = 0$. Then K is the link depicted in Figure 3.7(a). We note that K is composite, and thus not arborescent by Definition 1.4.

In fact, we claim that this case (a) is void because the reduction algorithm will have cut this link into its prime components (two copies of the Hopf link). The augmented 0-bracelet was necessarily created in step (7) from a 1-bracelet B_1 and an augmented bracelet $B_{1,1}$, which in turn was necessarily created from a 3-bracelet B_3 in step (5). However, B_3 must have been glued to B_1 with their preferred slopes equal, so in step (3) the algorithm will have recognized K as a connected sum.

(b) $r = 1$. Then, as Figure 3.7(b) shows, K is the boundary of an unknotted band with 4 half-twists, which falls in the exceptional family I.

(c) $r \geq 2$. Then, as Figure 3.7(c) shows, K is the pretzel link $P(r, 2, 2; -1)$. Because $r \geq 2$ and $1/2 + 1/2 + 1/r > 1$, K falls in the exceptional family III.

(4) *Exceptional Montesinos links.* Recall, from the introduction, that a Montesinos link can be constructed by gluing a bracelet B_d to d different 1-bracelets. Consider such a link K , with $d \geq 3$. By Proposition 3.6, in any Montesinos output link, the preferred slope of B_d is not a Farey neighbor of the preferred slope of any of the 1-bracelets. Thus there is a marking of each Conway sphere, in which the preferred slope of B_d is ∞ and the preferred slope of the 1-bracelet glued to that Conway sphere is not in $\mathbb{Z} \cup \{\infty\}$.

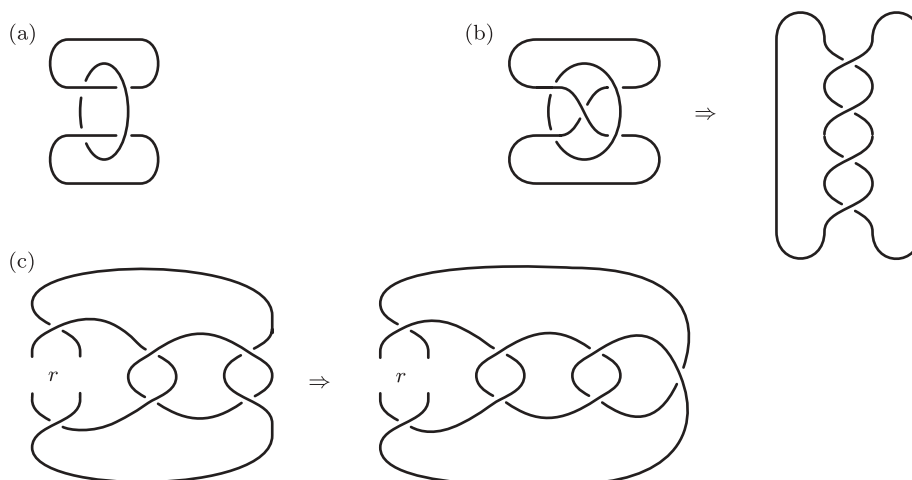


FIGURE 3.7. Augmented 0-bracelets form exceptional links in three different ways.

Once these markings are chosen, there is a unique unknotted band consisting of the arcs of K_d and arcs of slope 0 along the Conway spheres. We define the *number of half-twists* in the band of B_d to be the number of half-twists in this band. If we modify the marking on some sphere by $k/2$ Dehn twists about the preferred slope of B_d , the slope of the 1-bracelet glued to that sphere goes up by k , while the number of half-twists in the band goes down by k . Thus, by employing Dehn twists of this sort, we can choose markings in which the preferred slope of B_d is still ∞ and the preferred slope of each 1-bracelet is in the interval $(0,1)$.

We perform two final normalizations. The edge pairs of slopes $0, 1, \infty$ decompose each Conway sphere into four triangles. The orientation of the bracelet B_d induces an orientation for each of the d Conway spheres of ∂B_d . We stipulate that the edges of any triangle of any Conway sphere S have slope $0, 1, \infty$, in that clockwise order around the triangle, for the induced orientation of S . (We recommend that the reader disregard all orientation issues at a first reading. The choices in this paragraph are unimportant, as long as they are consistent from one Conway sphere to the next.) This can be assumed, up to changing the marking of some of the Conway spheres by an orientation-reversing homeomorphism, exchanging the slopes $0, 1$ and fixing ∞ . Finally, up to reflecting the link K , we can ensure that at least one of the 1-bracelet slopes falls in the interval $(0, 1/2]$. Now, we can recognize the exceptional links:

- (a) If $d = 4$, and the slope of every 1-bracelet is $1/2$, and there are 2 half-twists in the band of B_d , then K is the pretzel link $P(2, -2, 2, -2)$, as in Figure 3.8. Thus K falls in exceptional family II.
- (b) If $d = 3$, and the slope of every 1-bracelet is of the form $1/n$, and there is 1 half-twist in the band of B_d , then K is the pretzel link $P(p, q, r; -1)$. If $1/p + 1/q + 1/r \geq 1$, then K falls in the exceptional family III.

DEFINITION 3.7. An arborescent link K is called a *candidate link* if it can be constructed from (possibly augmented) bracelets in the following fashion:

- The link K is not a 0-bracelet or augmented 0-bracelet.
- All bracelets are either unaugmented or augmented once.
- All 2-bracelets are augmented once.
- The gluing maps of bracelets along Conway spheres satisfy the minimum-distance table (Table 1) of Proposition 3.6.

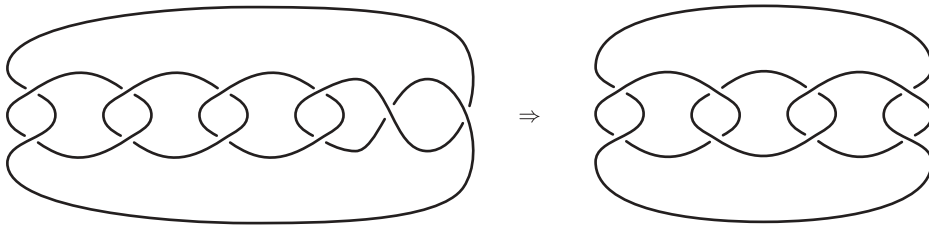


FIGURE 3.8. The pretzel link $P(2, -2, 2, -2)$ falls in the exceptional family II.

– If K has exactly one large bracelet $B_{a,0}$, the normalization process of (4a, b) above does not reveal K as an exceptional Montesinos link.

So far, we have proved that every link K contained in the output of the algorithm (and thus, in particular, every arborescent link) either falls into one of the three exceptional families, or else is a candidate link. To complete the proof of Theorem 1.5, it remains to show the following theorem.

THEOREM 3.8. *Let K be a candidate link. Then $\mathbb{S}^3 \setminus K$ admits a decomposition into angled blocks. Thus, by Theorem 1.1, every candidate link is hyperbolic.*

The proof of Theorem 3.8 occupies Sections 4 and 5. In the meantime, we record the following consequence of Theorem 3.8, which shows that our reduction algorithm does more than merely sort generalized arborescent links into hyperbolic and non-hyperbolic bins.

THEOREM 3.9. *Let K be a generalized arborescent link, given in any presentation as a union of unaugmented bracelets. If we apply the simplification algorithm to K , it will output several links K_1, \dots, K_n , with the following properties:*

- (1) $K = K_1 \# \dots \# K_n$, a connected sum of the output links;
- (2) every K_i is prime; and
- (3) if some K_i is the unknot, it appears as a 0-bracelet with ± 1 half-twist.

Therefore, this algorithm recognizes the unknot and factors a link into its prime summands.

Proof. Statement (1) was part of Proposition 3.6. To prove (2), let K_i be a link produced in the output of the algorithm. By Proposition 3.6 and the discussion that follows it, K_i is either a candidate link or a known exception. By Theorem 3.8 (which we have yet to prove), the candidate links are hyperbolic, and thus prime. As we mentioned in the introduction, the exceptional links in families I and III are also prime. When K_i is in the exceptional family II, we can identify every group of isotopic components to a single circle, producing a new link K'_i that is either hyperbolic or in families I or III. By the cases already discussed, K'_i is prime, and thus K_i is also prime.

To prove (3), suppose that K_i is the unknot. Again, it is easy to check that apart from the 0-bracelet with ± 1 half-twist, all the exceptional knots discussed above are non-trivial. By Theorem 3.8, the candidate links are hyperbolic, and hence also non-trivial. Thus the unknot K_i can in fact be recognized as a 0-bracelet with ± 1 half-twist. \square

REMARK 3.10. To solve the general link isotopy problem for arborescent links, Bonahon and Siebenmann [3] developed a special type of calculus on weighted trees representing links. Our algorithm for simplifying bracelet presentations is directly inspired by their algorithm for simplifying trees (the latter only keeps more careful track of the order of Conway spheres along each band). The main result of [3], which uses the machinery of double branched covers and equivariant JSJ decompositions, can be paraphrased as follows: for every generalized arborescent link (with a handful of exceptions), any two bracelet presentations produce the same output under the algorithm above, up to the number of trivial components in the connected sum. The few exceptions include, for instance, the pretzel link $P(2, -2, 2, -2)$ depicted in Figure 3.8, which can appear either in pretzel form (as a 4-bracelet glued to four 1-bracelets), or as a twice-augmented 0-bracelet (in two different ways). Thus, by the results of [3], the algorithm can be said to essentially classify all arborescent links.

4. Block decomposition of the link complement

In this section, we consider a candidate link K (see Definition 3.7) and devise a block decomposition of $\mathbb{S}^3 \setminus K$ (up to applying the algorithm of Section 3, we may and will assume that the bracelet presentation of K satisfies the conditions of Definition 3.7). In what follows, we will use one ‘large’ block for each large bracelet $B_{d,n}$ ($d \geq 3$ or $n = 1$). The Farey combinatorics involved in gluing these blocks to one another and to 1-bracelets will be encoded in a certain number of ideal tetrahedra in the decomposition.

Consider a large d -bracelet $B_{d,n}$ (where $n = 0$ or 1), and recall that its underlying space M_d is \mathbb{S}^3 minus d open balls. Denote by $K_{d,n}$ the tangle (union of arcs) contained in $B_{d,n}$. Between any pair of consecutive Conway spheres, M_d contains a rectangular strip consisting of crossing segments; we call this strip a *crossing rectangle*. The boundary of each crossing rectangle consists of two arcs of K_d and two arcs on Conway spheres, which define the preferred slopes of these two Conway spheres.

The large block $\mathcal{P}_{d,n}$ corresponding to the large bracelet $B_{d,n}$ will be constructed in Subsection 4.3: $\mathcal{P}_{d,n}$ will be a solid torus whose boundary is decomposed into ideal polygons, and some pairs of edges (or even faces) of $\mathcal{P}_{d,n}$ will be identified in the block decomposition of $\mathbb{S}^3 \setminus K$ (in addition, the core curve of the solid torus $\mathcal{P}_{d,n}$ will be removed if $n = 1$). More precisely, each crossing rectangle in $\mathbb{S}^3 \setminus K$ will eventually be collapsed to a crossing edge, which will be realized as an ideal edge of $\mathcal{P}_{d,n}$ (in fact, as a pair of identified edges of $\mathcal{P}_{d,n}$; see Figure 4.3 for a preview). Because of this collapsing operation, we need to be very careful that we do not change the nature of the manifold $\mathbb{S}^3 \setminus K$ up to homeomorphism. This will be proved in due course (Subsection 4.2) before actually constructing $\mathcal{P}_{d,n}$ (Subsection 4.3). Meanwhile, we simply insist that the arc pairs of preferred slope (defined, for example, by crossing rectangles) in a Conway sphere of a large bracelet will be realized as ideal edges of the corresponding large blocks. This fact motivates the whole construction of Subsection 4.1, where we describe all the small blocks (ideal tetrahedra).

4.1. Gluings through tetrahedra

In this section, we use the Farey complex \mathcal{F} in order to encode the gluing homeomorphisms between the boundary spheres of two bracelets into sequences (layers) of ideal tetrahedra, glued at the interface between large bracelets. This follows [9, Section 5]. We also realize trivial tangles as gluings of ideal tetrahedra, following [9, Appendix].

4.1.1. *Gluing of two large bracelets.* Consider two large bracelets $B_{d,n}$ and $B_{d',n'}$ glued to one another along a Conway sphere. If $K_{d,n}$ denotes the union of arcs contained in $B_{d,n}$, then define $\mathcal{C} := M_d \setminus K_{d,n}$ and $\mathcal{C}' := M_{d'} \setminus K_{d',n'}$. For the later purpose of realizing (a retract of)

\mathcal{C} as a polyhedral block, we assume that all the arc pairs of the preferred slope are marked on the Conway spheres of $\partial\mathcal{C}$ (and similarly for \mathcal{C}'). We will now enhance these graphs on $\partial\mathcal{C}, \partial\mathcal{C}'$ to filling graphs (Definition 2.1), then glue \mathcal{C} to \mathcal{C}' via a union of ideal tetrahedra attached to the Conway spheres of $\mathcal{C}, \mathcal{C}'$ (thus realizing the gluing homeomorphism between $B_{d,n}$ and $B_{d',n'}$).

Let s and s' be the preferred slopes of \mathcal{C} and \mathcal{C}' , respectively (see Definition 3.2). By Proposition 3.6, we know that $s \neq s'$. The obstruction to gluing \mathcal{C} to \mathcal{C}' directly is that arcs of slopes s and s' may have high intersection number. We regard $s, s' \in \mathbb{P}^1\mathbb{Q}$ as vertices in the Farey diagram.

Consider the simplest case, where s and s' are Farey neighbors. Then the four arcs of slope s and s' define a subdivision of a Conway sphere of $\partial\mathcal{C}$ into two ideal squares. Similarly the arcs of slope s' and s subdivide a Conway sphere of $\partial\mathcal{C}'$ into two ideal squares. In this special case, we may glue $\partial\mathcal{C}$ directly to $\partial\mathcal{C}'$ along these two squares.

If s, s' are not Farey neighbors, we need to consider the sequence of Farey triangles (T_0, \dots, T_m) crossed by the geodesic line from s to s' (where $m \geq 1$, and s and s' are vertices of T_0 and T_m , respectively). For each $0 \leq i \leq m$, the vertices of T_i define three slopes, and the corresponding arc pairs provide an ideal triangulation τ_i of the Conway sphere S . Moreover, if x, y are the ends of the Farey edge $T_i \cap T_{i+1}$, then the two arc pairs whose slopes are x and y define a subdivision σ of the 4-punctured sphere S into two ideal squares: both triangulations τ_i and τ_{i+1} are refinements of σ . In fact, τ_{i+1} is obtained from τ_i by a pair of *diagonal exchanges*: remove two opposite edges of τ_i (thus liberating two square cells, the cells of σ , of which the removed edges were diagonals); then insert back the other diagonals. Each of these two diagonal exchanges defines (up to isotopy) a topological *ideal tetrahedron* in the space $S \times I$: more precisely, the union of these two ideal tetrahedra is bounded by two topological *pleated surfaces* S_i, S_{i+1} isotopic to $S \times \{*\}$ in $S \times I$, which are pleated along the edges of τ_i and τ_{i+1} , and intersect each other precisely along the edges of σ (see Figure 4.1). Doing the same construction for all $0 \leq i < m$, we thus obtain an ideal triangulation of a strong deformation retract of $S \times I$, whose bottom and top are pleated along an ideal triangulation of S containing the arc pair of slopes s and s' , respectively.

Finally, the remaining two pairs of edges in the triangulation τ_0 (in addition to the pair of slope s) define a subdivision of the boundary of the space \mathcal{C} . The same occurs for \mathcal{C}' . We have completed our aim of gluing \mathcal{C} to \mathcal{C}' , with boundaries suitably triangulated, using a sequence of (pairs of) ideal tetrahedra as an interface. Note that the choice of ‘suitable’ boundary triangulations of $\mathcal{C}, \mathcal{C}'$ is forced by the gluing homeomorphism itself.

DEFINITION 4.1. The family of ideal tetrahedra between \mathcal{C} and \mathcal{C}' is called a *product region*. The same term also refers to the union of that family. Topologically, the product region is a strong deformation retract of $S \times I$; when s and s' have no common Farey neighbors, the product region is homeomorphic to $S \times I$.

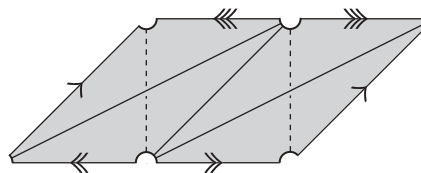


FIGURE 4.1. A layer of two tetrahedra, caught between two topological pleated surfaces S_i, S_{i+1} . Edges with identical arrows are identified.

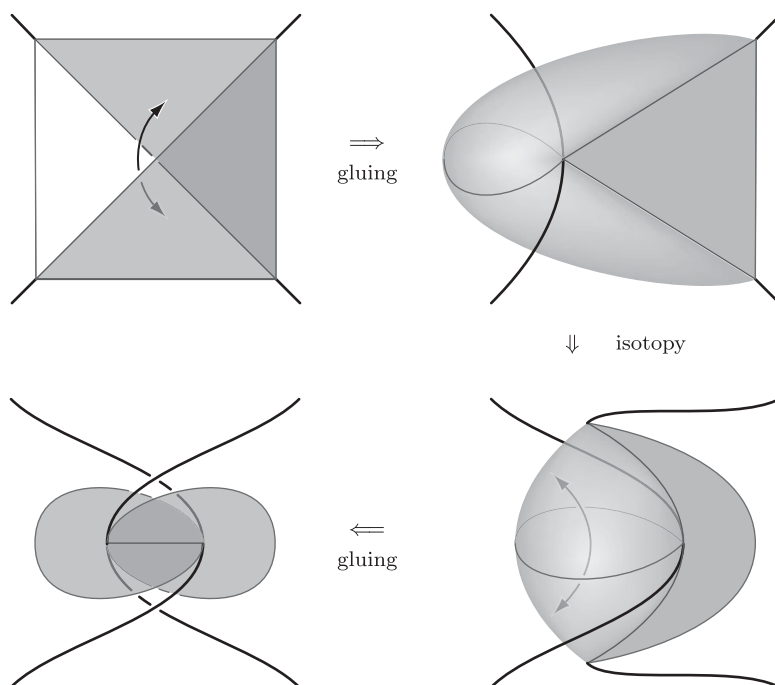


FIGURE 4.2. The surface S_{m-1} is glued to itself, realizing a 1-bracelet (trivial tangle).

4.1.2. *Gluing a large bracelet to a trivial tangle.* Consider the gluing of a large bracelet $B_{d,n}$ to a trivial tangle $B_{1,0}$ along a Conway sphere. As before, define $\mathcal{C} := M_d \setminus K_{d,n}$ and $\mathcal{C}' := M_1 \setminus K_{1,0}$. We will triangulate the Conway sphere S of \mathcal{C} and attach ideal tetrahedra to S to realize a space homeomorphic to $\mathcal{C} \cup_S \mathcal{C}'$. (Note: we will not need to attach a copy of \mathcal{C}' itself, only ideal tetrahedra.)

Let s and s' be the preferred slopes of \mathcal{C} and \mathcal{C}' , respectively. By the minimum-distance table, Table 1, of Proposition 3.6, we know that s and s' are not equal and are not Farey neighbors. We can thus consider the sequence of Farey triangles (T_0, \dots, T_m) from s to s' , where $m \geq 1$. We perform the same construction as in Paragraph 4.1.1 above, using topological pleated surfaces S_0, \dots, S_{m-1} whose triangulations are given by T_0, \dots, T_{m-1} (note the omission of T_m). To realize the trivial tangle complement \mathcal{C}' , we will now glue the faces of the pleated surface S_{m-1} together in pairs, following Sakuma and Weeks's construction in [20].

Without loss of generality, we may assume that the vertices of T_{m-1} and T_m are $0, 1, \infty$ and $0, 1, \frac{1}{2}$, respectively. Each face (ideal triangle) f of S_{m-1} has exactly one edge e of slope ∞ , shared with another face f' . We simply identify f and f' by a homeomorphism respecting e . The result is shown in Figure 4.2: it is straightforward to check that the simple closed curve in S_{m-1} of slope $s' = \frac{1}{2}$ becomes contractible. The gluing thus realizes a 1-bracelet of slope $s' = \frac{1}{2}$.

REMARK 4.2. If $m = 1$, note that the Conway sphere of \mathcal{C} , made of four ideal triangles, is S_{m-1} and has been directly collapsed to two ideal triangles, without gluing any tetrahedra. More generally, for any m , all four edges whose slopes are Farey neighbors of s' are collapsed to just one edge (the horizontal edge in the last panel of Figure 4.2): though none of these four edges can have slope s (because s, s' are not Farey neighbors), some of them may belong to the Conway sphere in $\partial\mathcal{C}$. In spite of these collapsings, for any candidate link K , we can realize

the space $\mathbb{S}^3 \setminus K$ with a well-defined *manifold* structure by gluings of the type above, because the Conway spheres of \mathcal{C} are pairwise disjoint.

4.1.3. *Gluing two trivial tangles together.* Finally, when two trivial tangles are glued to one another, we obtain a 2-bridge link K . The strands in each bracelet can be isotoped to proper pairs of arcs in the Conway sphere, of slopes s and s' . If s, s' are sufficiently far apart in the Farey diagram, we can perform the gluing operation above both near \mathcal{C} and near \mathcal{C}' (both \mathcal{C} and \mathcal{C}' being homeomorphic to 1-bracelets $M_1 \setminus K_{1,0}$). The resulting decomposition into tetrahedra was constructed by Sakuma and Weeks [20], and also described in the appendix to [9]. For completeness, we include the following result.

PROPOSITION 4.3. *If s, s' have no common Farey neighbors (that is, satisfy the minimum-distance table, Table 1, of Proposition 3.6), the union of the tetrahedra defined by the construction above is a triangulated manifold homeomorphic to $\mathbb{S}^3 \setminus K$, where K is a 2-bridge link.*

Proof. First, the path of Farey triangles (T_0, \dots, T_m) from s to s' satisfies $m \geq 3$: indeed, if $m = 2$, then two vertices of T_1 are Farey neighbors of s , and two are Farey neighbors of s' ; so s and s' have a common neighbor. Therefore the first and last pleated surfaces S_1 and S_{m-1} are distinct, and there is at least one layer of tetrahedra. Consider the union of all tetrahedra *before* folding S_1 and S_{m-1} onto themselves. Denote by x, y the ends of the Farey edge $T_1 \cap T_2$, and thicken the corresponding tetrahedron layer between S_1 and S_2 by replacing each edge whose slope is x or y with a bigon. The resulting space is homeomorphic to $S_1 \times [0, 1]$: therefore, after folding S_1 and S_{m-1} , we do obtain the manifold $\mathbb{S}^3 \setminus K$. It remains to collapse the four bigons back to ordinary edges, without turning the space into a non-manifold.

Recall (Remark 4.2) that the foldings of S_1 and S_{m-1} identified all four edges whose slopes are ends of $T_0 \cap T_1$ (respectively $T_m \cap T_{m-1}$) to one edge, and caused no other edge identifications. At most one of x, y belongs to the Farey edge $T_0 \cap T_1$; at most one of x, y belongs to $T_m \cap T_{m-1}$; and none of x, y belongs to both Farey edges simultaneously (s, s' have no common neighbors). Thus under the folding of S_1 and S_{m-1} , the two bigons of slope x may become glued along one edge; the two bigons of slope y may become glued along one edge, and no further identifications occur between points of the four bigons. When two bigons are identified along one edge, consider their union as just one bigon. All (closed) bigons are now disjointly embedded in $\mathbb{S}^3 \setminus K$, so we can collapse each of them to an ideal segment without changing the space $\mathbb{S}^3 \setminus K$ up to homeomorphism. \square

Thus, by the results of [9, Theorem A.1], or more broadly by Menasco's theorem [13], we already know that 2-bridge links admit angle structures (and are hyperbolic) if and only if they are candidate links.

4.2. Collapsing

At the beginning of Section 4, we defined the *crossing rectangles*: a large bracelet $B_{d,n}$ has d crossing rectangles $R \simeq [0, 1] \times [0, 1]$, such that two opposite sides $\{0, 1\} \times [0, 1]$ of R define the preferred slopes in two consecutive Conway spheres, and the two other sides $[0, 1] \times \{0, 1\}$ belong to the tangle (union of arcs) contained in $B_{d,n}$.

In a candidate link K containing at least one large bracelet, we collapse each crossing rectangle $R \simeq [0, 1] \times [0, 1]$ as above to a segment $\{*\} \times [0, 1]$.

PROPOSITION 4.4. *The space obtained after collapsing the crossing rectangles to segments is still homeomorphic to the manifold $\mathbb{S}^3 \setminus K$.*

Proof. As in Proposition 4.3, it is enough to check that the closed crossing rectangles, before collapsing, are disjointly embedded in the union of ideal tetrahedra and uncollapsed large bracelets. First, consider a gluing between two large bracelets: since the two preferred slopes on the gluing Conway sphere are distinct (by the minimum-distance table, Table 1, of Proposition 3.6), no points of the crossing rectangles adjacent to this Conway sphere get identified (in the product region corresponding to the Conway sphere, all tetrahedron edges are disjoint). Then consider a gluing between a large bracelet (with preferred slope s) and a trivial tangle (with preferred slope s'): in Remark 4.2, we observed that none of the edges that undergo identifications have slope s , because s, s' are not Farey neighbors (being at distance 2 or more in the Farey graph). Therefore, no points of any crossing rectangles are identified. Since the (closed) crossing rectangles are disjointly embedded in $\mathbb{S}^3 \setminus K$, we can collapse each of them to a segment without changing the space $\mathbb{S}^3 \setminus K$ up to homeomorphism. \square

4.3. *Blocks associated to large bracelets*

In this section, we construct the blocks associated to large bracelets. We begin by considering a large *non-augmented* bracelet B_d (where $d \geq 3$) and set out to construct an ideal polyhedron version of the space \mathcal{C} , now defined as $M_d \setminus K_d$ with crossing rectangles collapsed to (ideal) segments. We will construct \mathcal{C} as a polyhedral solid torus (or block) \mathcal{P} with some edge identifications.

Consider a (closed) solid torus $\hat{\mathcal{P}}$ with a preferred, core-parallel, simple closed curve $\hat{\gamma}$ on $\partial\hat{\mathcal{P}}$. We endow both $\hat{\mathcal{P}}$ and $\hat{\gamma}$ with orientations that will remain fixed throughout the paper. Draw d disjoint oriented curves $\hat{\gamma}_1, \dots, \hat{\gamma}_d$ isotopic to $\hat{\gamma}$ in $\partial\hat{\mathcal{P}}$, so that $\partial\hat{\mathcal{P}}$ is the union of d annuli \hat{U}_i , each bounded by two curves, $\hat{\gamma}_i$ and $\hat{\gamma}_{i+1}$ (indices are taken modulo d). For matters of orientation, we assume that the orientation entering \hat{U}_i (from a point of $\hat{\gamma}_i$), followed by the orientation of $\hat{\gamma}_i$, is equal to the orientation on $\partial\hat{\mathcal{P}}$ induced by the orientation of $\hat{\mathcal{P}}$ (Figure 4.3, left). Next, remove one point in each $\hat{\gamma}_i$. We still denote by $\hat{\mathcal{P}}$ the solid torus with these d points removed, and by \hat{U}_i the annuli with two boundary points removed. Each curve $\hat{\gamma}_i$ is now replaced by an ideal (oriented) arc, also denoted by $\hat{\gamma}_i$, joining a puncture to itself.

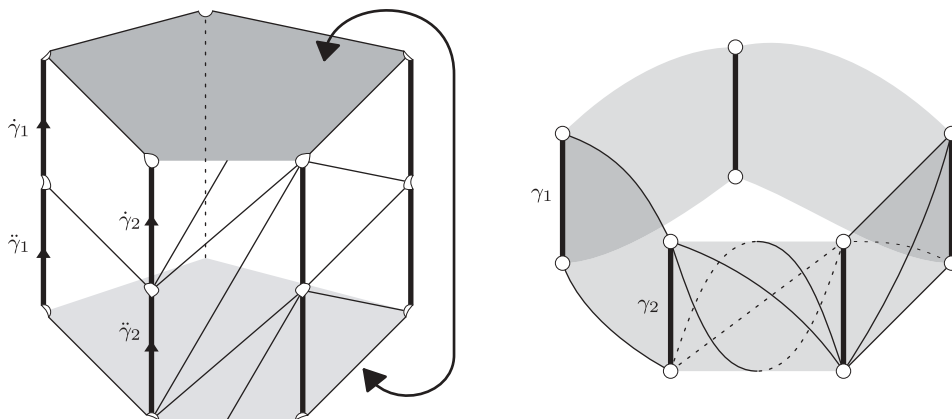


FIGURE 4.3. *Block associated to a large (non-augmented) bracelet of degree 5. The solid torus \mathcal{P} is shown before and after the identifications $\hat{\gamma}_i \sim \check{\gamma}_i$. In the right panel, \mathcal{P} is the space outside the grey ‘ring’.*

To construct an ideal cellulation of $\partial\hat{\mathcal{P}}$, we must decompose each annulus \hat{U}_i into contractible ideal polygons. There are two options for doing so.

(1) Choose an ideal arc across \hat{U}_i , connecting the puncture of $\hat{\gamma}_i$ to the puncture of $\hat{\gamma}_{i+1}$ (for each i , there is a \mathbb{Z} -worth of possible choices for such an arc). Then \hat{U}_i is an ideal square cell with a pair of opposite sides identified.

(2) Choose two disjoint, non-isotopic ideal arcs across \hat{U}_i , decomposing \hat{U}_i into two ideal triangles. There is again a \mathbb{Z} -worth of possible choices.

Finally, consider the two-fold cyclic cover \mathcal{P} of $\hat{\mathcal{P}}$. Each $\hat{\gamma}_i$ lifts to two arcs $\dot{\gamma}_i$ and $\check{\gamma}_i$ in $\partial\mathcal{P}$ with the same pair of (distinct) punctures as end points. We now identify $\dot{\gamma}_i$ with $\check{\gamma}_i$, by an orientation-reversing homeomorphism. The resulting arc is called γ_i , and the quotient of \mathcal{P} under this identification is homeomorphic to \mathcal{C} , as defined at the beginning of this section. (The γ_i are the collapsed crossing rectangles, and the removed tangle K_d is ‘in the ideal vertices’.) Each arc γ_i is a crossing segment, not endowed with any orientation. See Figure 4.3.

Under the identification, each band U_i (the lift of \hat{U}_i in \mathcal{P}) becomes a 4-punctured sphere included in the boundary of \mathcal{C} . Any decomposition of \hat{U}_i , as in the dichotomy above, defines a decomposition of U_i into two ideal squares or four ideal triangles. The two cases of the dichotomy correspond, as in Subsection 4.1, to the neighboring bracelet either (1) being glued directly to \mathcal{C} , using a very ‘simple’ homeomorphism, or (2) needing an interface of ideal tetrahedra or being in fact a 1-bracelet (trivial tangle).

Recall that in each case of the dichotomy, the involved choices in \mathbb{Z} are forced by the gluing homeomorphisms (see Subsection 4.1), and essentially reflect the number of half-twists in the band of the bracelet (M_d, K_d) that defined \mathcal{C} (and the block \mathcal{P}).

4.3.1. *Augmented polyhedral d -bracelets* ($d \geq 1$). Topologically, the construction above is valid for all $d \geq 1$, not just $d \geq 3$, and an augmented d -block (associated to an augmented d -bracelet) is obtained by drilling out the core of the solid torus \mathcal{P} . If we denote by \mathcal{P}_{aug} the result of the drill-out, then any decomposition of $\partial\mathcal{P}$ into ideal triangles and squares (as in the dichotomy above) induces, by coning off to the core of \mathcal{P}_{aug} , a decomposition of \mathcal{P}_{aug} into contractible ideal polyhedra (tetrahedra and square-based pyramids): the coning-off is induced by the product structure $\mathbb{T}^2 \times (0, 1)$ of the interior of \mathcal{P}_{aug} . For the purpose of finding angle structures, it will be convenient to regard \mathcal{P}_{aug} as such a union of ideal polyhedra, rather than an elementary block *per se*.

An essential feature of our polyhedral realizations of d -blocks and augmented d -blocks is that crossing segments always arise as edges, thus defining a preferred slope on each Conway sphere of the block. If the block comes from an augmented bracelet, this slope is also the one defined by the ‘extra’ link component inside the augmented bracelet.

5. Angle structures for the link complement

In this section, we find dihedral angles (satisfying Definition 2.2 and the hypotheses of Theorem 1.1) for the blocks and ideal tetrahedra constructed in Section 4. The key property will be an explicit description (Subsection 5.2) of the space of angle structures associated to a sequence of tetrahedra forming a product region (see Definition 4.1) or a 1-bracelet (trivial tangle). This description (in large part borrowed from [9, Section 5 and Appendix]) is sufficiently tractable that we can say exactly when the tetrahedra admit dihedral angles that match a given system of angles for the solid tori (Propositions 5.10 and 5.11). As a result, we can show that all candidate links (see Definition 3.7) admit angle structures. We will treat the easier case of non-Montesinos links in Subsection 5.3 and the trickier case of Montesinos links in Subsection 5.4. Montesinos links are tricky because they include the third family of exceptions to Theorem 1.5.

The strategy is as follows: in Subsection 5.1, we choose some dihedral angles to parametrize the deformation space of solid tori. We show (Propositions 5.3, 5.6, 5.7) that these dihedral angles define valid angle structures whenever they are, in some appropriate sense, small enough. On the other hand, the same parameters need to be large enough (Proposition 5.11) for the trivial tangles to admit angle structures. The conflict that arises can be managed for candidate links, but causes the exceptional Montesinos links (ruled out in Definition 3.7) to have no angle structures.

5.1. Angle structures for a non-augmented large block

We consider a solid torus \mathcal{P} of degree d whose boundary is subdivided into ideal triangles and quadrilaterals, as in Subsection 4.3. In this section, we study the space of angle structures for \mathcal{P} . We restrict attention to those angle structures which are invariant under the natural fixed-point-free involution of \mathcal{P} (recall that \mathcal{P} was defined as a 2-fold covering in Subsection 4.3).

Recall the preferred direction of $\partial\mathcal{P}$, defined by the crossing arcs of the corresponding bracelet. As parameters, we will use the (exterior) dihedral angles at all those edges of $\partial\mathcal{P}$ which are *not* along the preferred direction. The angles at the edges along the preferred direction can then be recovered from the requirement that the angles around any vertex of the block add to 2π .

For simplicity, assume that $\partial\mathcal{P}$ is decomposed into ideal triangles only. Recall from Subsection 4.3 that the preferred direction came with an orientation, which we call ‘upwards’. Thus, each of the d bands U_1, \dots, U_d in $\partial\mathcal{P}$ is traversed by two *ascending* and two *descending* edges (for example, in Figure 4.3 (left), though all four edges across U_2 seem to go upwards to the right, we agree to call only the steeper pair ascending, and the other pair descending). By the normalization of markings that precedes Definition 3.7, the slopes of the descending edges in the corresponding Conway spheres are 0, while the slopes of the ascending edges are 1, and the slopes associated to the edges along the preferred direction are ∞ .

For each $1 \leq i \leq d$, we denote by a formal variable $a_i \in [0, \pi)$ the (exterior) dihedral angle at the ascending pair of edges, and $b_i \in [0, \pi)$ the angle at the descending pair. (We can recover the case where U_i is subdivided into two squares by allowing $a_i = 0$ or $b_i = 0$). If the index i is read modulo d , the exterior dihedral angle at an edge (in the preferred direction) of $U_i \cap U_{i+1}$ must be

$$c_{i,i+1} = \pi - \frac{1}{2}(a_i + a_{i+1} + b_i + b_{i+1}), \quad (5.1)$$

so that the angles around each ideal vertex add to 2π . To force all angles to be non-negative, we require that

$$0 \leq a_i < \pi \quad \text{and} \quad 0 \leq b_i < \pi \quad \text{and} \quad 0 < a_i + b_i \leq \pi. \quad (5.2)$$

The particular choice of strong and weak inequalities here implies the following lemma.

LEMMA 5.1. *Any triangular face of \mathcal{P} has at most one edge with dihedral angle 0.*

Proof. The exterior dihedral angles of a triangular face are $a_i, b_i, c_{i,i+1}$ or $a_i, b_i, c_{i-1,i}$. If $a_i = 0$ one easily checks that $b_i, c_{i,i+1}, c_{i-1,i}$ are all positive (because $b_i = a_i + b_i > 0$ and $a_i + b_i = b_i < \pi$). If $b_i = 0$ the argument is the same. If $c_{i,i+1} = 0$ then $a_i, a_{i+1}, b_i, b_{i+1}$ are all positive (because their sum is 2π). \square

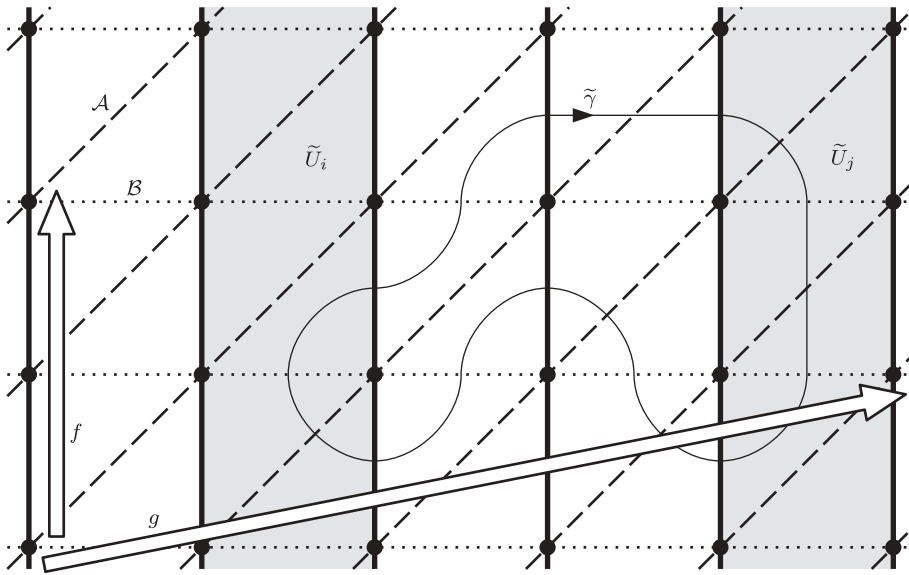


FIGURE 5.1. The cover $\mathbb{R}^2 \setminus \mathbb{Z}^2$ of $\partial\mathcal{P}$, with a closed curve $\tilde{\gamma}$.

Thus a block \mathcal{P} with only triangular faces and *non-negative* dihedral angles a_i, b_i satisfying (5.2) uniquely defines a block with *positive* dihedral angles and, possibly, some quadrilateral faces. By default, we will usually consider that a block \mathcal{P} has only triangular faces and look for angle systems satisfying (5.2).

From Definition 2.2, recall that an *angle structure* on \mathcal{P} requires every normal simple closed curve γ in $\partial\mathcal{P}$ that bounds a disk in \mathcal{P} to have total bending number at least 2π (the total bending number is the sum of the exterior dihedral angles at the edges encountered by γ , counted with multiplicity). Such a curve γ can be defined as a non-backtracking closed path in the dual graph.

Identify $\partial\mathcal{P}$ with $(\mathbb{R}^2 \setminus \mathbb{Z}^2)/\langle f, g \rangle$ where $f(x, y) = (x, y + 2)$ and $g(x, y) = (x + d, y + k)$, where $d \geq 3$ is the degree of \mathcal{P} and $k \in \mathbb{Z}$ is an integer such that any g -invariant straight line in $\mathbb{R}^2 \setminus \mathbb{Z}^2$ projects in $\partial\mathcal{P}$ to the boundary of a compression disk of \mathcal{P} . (In other words, k is, up to a constant, the number of half-twists in the band of the bracelet associated to \mathcal{P} .) The bands $U_i \subset \partial\mathcal{P}$ lift to a subdivision of $\mathbb{R}^2 \setminus \mathbb{Z}^2$ into ‘vertical’ bands $(\tilde{U}_i)_{i \in \mathbb{Z}}$, where $\tilde{U}_i = (i, i + 1) \times \mathbb{R}$ (see Figure 5.1). Consider a normal simple closed curve γ in $\partial\mathcal{P}$ bounding a disk in \mathcal{P} , and lift γ to a curve $\tilde{\gamma}$ in $\mathbb{R}^2 \setminus \mathbb{Z}^2$.

Either $\tilde{\gamma}(1) = \tilde{\gamma}(0)$ (that is, $\tilde{\gamma}$ is a closed curve), or $\tilde{\gamma}(1) = g^{\pm 1}(\tilde{\gamma}(0))$ (the exponent cannot be larger than 1 in absolute value, because γ is simple).

LEMMA 5.2. *If $\tilde{\gamma}(1) = \tilde{\gamma}(0)$, the total bending number N_γ of γ is more than 2π , unless γ is just a loop around a puncture of $\partial\mathcal{P}$.*

Proof. If $\tilde{\gamma}$ visits exactly the vertical bands $\tilde{U}_i, \tilde{U}_{i+1}, \dots, \tilde{U}_{j-1}, \tilde{U}_j$, then $\tilde{\gamma}$ crosses both ascending and descending edges in U_i and in U_j (because $\tilde{\gamma}$ never crosses the same edge twice

consecutively). Counting edges met by $\tilde{\gamma}$, we thus get

$$\begin{aligned} N_\gamma &\geq (a_i + b_i) + (a_j + b_j) + 2 \sum_{k=i+1}^{j-1} \min\{a_k, b_k\} + 2 \sum_{k=i}^{j-1} c_{k,k+1} \\ &= 2\pi + 2 \sum_{k=i+1}^{j-1} (\pi - \max\{a_k, b_k\}), \end{aligned}$$

using (5.1). This quantity is larger than 2π , unless $j = i + 1$. If $j = i + 1$, we can assume that $\tilde{\gamma}$ visits $r \geq 1$ times \tilde{U}_i and r times \tilde{U}_{i+1} , and hence

$$N_\gamma \geq r(a_i + b_i) + r(a_{i+1} + b_{i+1}) + 2rc_{i,i+1} = 2r\pi :$$

that quantity is larger than 2π unless $r = 1$. If $r = 1$, then $\tilde{\gamma}$ is (homotopic to) the boundary of a regular neighborhood of the union of s consecutive vertical edges (along the preferred direction): in that case however,

$$N_\gamma = (s + 1)(a_i + b_i) + (s + 1)(a_{i+1} + b_{i+1}) + 2c_{i,i+1} = 2\pi + s(a_i + b_i + a_{i+1} + b_{i+1})$$

is larger than 2π , unless $s = 0$ and γ is just a loop around a puncture of $\partial\mathcal{P}$. □

There remains the case where $\tilde{\gamma}(1) = g^{\pm 1}(\tilde{\gamma}(0))$, that is, γ bounds a compression disk of the solid torus \mathcal{P} . Then, the condition $N_\gamma > 2\pi$ (in terms of the a_i, b_i) is in general non-vacuous, and the conjunction of all these conditions (for all normal curves γ) has no nice closed-form expression in terms of the a_i, b_i . However, the following three propositions give various sufficient conditions for $N_\gamma > 2\pi$ to hold, independently of γ .

PROPOSITION 5.3. *If the block \mathcal{P} is of degree $d \geq 3$ and if*

$$(d - 2)\pi > \sum_{i=1}^d \max\{a_i, b_i\} \quad (\text{girth condition}), \tag{5.3}$$

then $N_\gamma > 2\pi$ for all γ bounding a compression disk, that is, the angles $a_i, b_i, c_{i,i+1}$ define an angle structure for \mathcal{P} .

Proof. If γ bounds a compression disk in \mathcal{P} , then γ crosses all the bands U_i . Therefore,

$$N_\gamma \geq \sum_{i=1}^d \min\{a_i, b_i\} + \sum_{i=1}^d c_{i,i+1} = d\pi - \sum_{i=1}^d \max\{a_i, b_i\},$$

using (5.1). The conclusion follows. □

Consider a normal curve $\gamma \subset \partial\mathcal{P}$ bounding a compression disk in \mathcal{P} . Let \mathcal{A} and \mathcal{B} be the union of all the ascending arcs and descending arcs, respectively, across the annuli $U_i \subset \partial\mathcal{P}$, each oriented from U_{i-1} to U_{i+1} . The torus $\mathbb{T} := \mathbb{R}^2 / \langle f, g \rangle$ naturally contains $\partial\mathcal{P}$ as a subset, and the closure $\overline{\mathcal{A}}$ (respectively $\overline{\mathcal{B}}$ of \mathcal{B}) in \mathbb{T} defines a union of oriented, parallel simple closed curves in \mathbb{T} (see Figure 5.1).

DEFINITION 5.4. We denote by $n_{\mathcal{A}}$ and $n_{\mathcal{B}}$ the absolute values of the (algebraic) intersection numbers of γ with $\overline{\mathcal{A}}$ and $\overline{\mathcal{B}}$, respectively, in \mathbb{T} . This definition clearly does not depend on the choice of compression-disk-bounding curve γ (all such γ are in the same class in $H_1(\mathbb{T}, \mathbb{Z})$). Note that $n_{\mathcal{A}} + n_{\mathcal{B}} \geq d$.

REMARK 5.5. In the special case of Montesinos links, we defined an integer n (up to sign), called the number of half-twists in the main band, in the normalization that precedes Definition 3.7. Since each descending arc has normalized slope 0 in the corresponding Conway sphere, the definition of n implies that $|n| = n_{\mathcal{B}}$ for Montesinos links.

PROPOSITION 5.6. *If $n_{\mathcal{A}} \geq 3$ and $(\alpha_i)_{1 \leq i \leq d}, (\beta_i)_{1 \leq i \leq d}$ are positive numbers such that $\alpha_i \geq \beta_i$, then setting*

$$(a_i, b_i) = (\pi - \varepsilon\alpha_i, \varepsilon\beta_i)$$

defines an angle structure for the block \mathcal{P} for all sufficiently small ε .

Proof. First, the a_i, b_i clearly satisfy condition (5.2) above. Any normal curve γ bounding a compression disk in \mathcal{P} meets at least $n_{\mathcal{A}} \geq 3$ ascending edges, whose pleating angles are all close to π : thus $N_\gamma > 2\pi$ for some small enough ε (independent of γ). A similar proposition holds when $n_{\mathcal{B}} \geq 3$. □

PROPOSITION 5.7. *Suppose the block \mathcal{P} has degree $d = 3$ and $(n_{\mathcal{A}}, n_{\mathcal{B}}) = (2, 1)$. Choose positive numbers $(\alpha_i)_{1 \leq i \leq d}, (\beta_i)_{1 \leq i \leq d}$ such that $\alpha_i \geq \beta_i$. Setting $(a_i, b_i) = (\pi - \varepsilon\alpha_i, \varepsilon\beta_i)$ defines an angle structure on \mathcal{P} (for small ε) whenever*

$$\alpha_{i+1} + \alpha_{i-1} - \beta_i > \alpha_i > \beta_{i+1} + \beta_{i-1}$$

for each $i \in \{1, 2, 3\}$ (taking indices modulo 3).

Proof. Again, condition (5.2) is clearly satisfied. If ε is small enough then, as in Proposition 5.6, it is enough to check $N_\gamma > 2\pi$ for a compression-disk-bounding curve γ that crosses \mathcal{A} exactly twice. Note that such a γ crosses at least one edge (of \mathcal{A} or of \mathcal{B}) within each vertical band U_1, U_2, U_3 . Two cases may arise : if γ crosses \mathcal{A} twice within the same vertical band U_i , then

$$N_\gamma \geq 2a_i + b_{i+1} + b_{i-1} + \sum_{i=1}^3 c_{i,i+1} = 2\pi + (\alpha_{i+1} + \alpha_{i-1} - \beta_i) - \alpha_i.$$

If γ crosses \mathcal{A} twice in distinct bands, say U_{i+1} and U_{i-1} , then

$$N_\gamma \geq a_{i+1} + a_{i-1} + b_i + \sum_{i=1}^3 c_{i,i+1} = 2\pi + \alpha_i - (\beta_{i+1} + \beta_{i-1}).$$

The conclusion follows. A similar proposition holds when $(n_{\mathcal{A}}, n_{\mathcal{B}}) = (1, 2)$. □

Note that the weaker condition $\alpha_{i+1} + \alpha_{i-1} > \alpha_i > \beta_{i+1} + \beta_{i-1}$ would in fact suffice in Proposition 5.7 (by analysing compression-disk-bounding curves γ more carefully). We will not need this (sharp) improvement.

5.1.1. *Angle structures on augmented blocks.* Finally, we note that an *augmented* block with prescribed non-negative dihedral angles a_i, b_i (where (5.2) holds) can always be realized as a union of tetrahedra with positive angles: the space obtained by coning off the band U_i to the extra component of the tangle in the block (as in Paragraph 4.3.1) is a union of four isometric tetrahedra of interior dihedral angles

$$\frac{\pi - a_i}{2}, \frac{\pi - b_i}{2}, \frac{a_i + b_i}{2}, \tag{5.4}$$

all positive by (5.2). The exterior dihedral angles of the augmented block are recovered as

$$\pi - \left(\frac{\pi - a_i}{2} + \frac{\pi - a_i}{2} \right) = a_i,$$

similarly

$$\pi - \left(\frac{\pi - b_i}{2} + \frac{\pi - b_i}{2} \right) = b_i \quad \text{and} \quad \pi - \left(\frac{a_i + b_i}{2} + \frac{a_{i+1} + b_{i+1}}{2} \right) = c_{i,i+1}.$$

REMARK 5.8. The augmentation component of each augmented d -bracelet bounds d disjoint, homotopically distinct, twice-punctured disks (also known as thrice-punctured spheres), which must be totally geodesic for the hyperbolic metric if one exists.

In fact, suppose that all large blocks are augmented (the candidate link K is called *totally augmented*; a special case of this is the case where there are no large blocks, and K is a 2-bridge link). For totally augmented links, the triangulation constructed above falls into the class studied in [8, Chapter 2]. There, the second author shows that the triangulation not only admits positive angle structures, but that one of these structures (the one with largest volume) actually realizes the hyperbolic metric, and is a refinement of the *geometrically canonical decomposition* in the sense of Epstein and Penner [6]. In other words, a certain (explicit) coarsening of the triangulation is combinatorially dual to the Ford-Voronoi domain of the manifold with respect to horoball neighborhoods of the cusps which are chosen to be pairwise tangent at each thrice-punctured sphere.

5.2. Angle structures for product regions and trivial tangles

In this section, we investigate the space of angle structures for the ideal tetrahedra constructed in Subsection 4.1, that is, not included in augmented blocks. Such tetrahedra live either at the interface of large blocks \mathcal{P} , \mathcal{P}' , or near trivial tangles. While the space of angle structures for a tetrahedron is easy to describe (a triple of positive angles summing to π), the difficulty is to deal with many tetrahedra simultaneously.

We begin by focusing on two large blocks \mathcal{P} , \mathcal{P}' separated by a product region (these come from large bracelets, in the sense of Definition 3.4). Recall from Subsection 4.1 the pleated 4-punctured spheres $(S_j)_{0 \leq j \leq m}$ between \mathcal{P} and \mathcal{P}' : we can endow S_j with a transverse ‘upward’ orientation, from \mathcal{P} to \mathcal{P}' . Suppose that we have solved the problem of finding an angle structure, that is, that the tetrahedra and solid tori are assigned dihedral angles that add up to 2π around each edge. Then we can define the *pleating angle* of the surface S_j at any edge $e \subset S_j$: namely, if the sum of all dihedral angles at e of tetrahedra and/or solid tori above and below S_j are $\pi + \alpha$ and $\pi - \alpha$, respectively, then we say that S_j has pleating angle α at e .

It will turn out that pleating angles of the S_j are very convenient parameters for the space of angle structures: thus, when no angle structure has been defined yet, we will typically look for angle structures realizing a given set of pleating angles of the S_j , and express the dihedral angles of the blocks \mathcal{P} , \mathcal{P}' and ideal tetrahedra in terms of these pleating angles.

We arbitrarily require that *the pleating angles of S_j at the three edges adjacent to any puncture of S_j add up to 0* (note that this property would hold in a true hyperbolic metric, where the holonomy of the loop around any puncture is a parabolic element of $\text{Isom}^+ \mathbb{H}^3$). This property easily implies that the pleating angles of S_j at opposite edges are equal. Restricting to such a subspace of solutions might (in principle) hamper our goal of finding angle structures; however, it is technically very convenient, for reasons we are about to outline now.

Consider the 4-punctured sphere $S^{(i)}$ defined by the vertical band $U_i \subset \partial \mathcal{P}$ (by identifying the edges in ∂U_i to create the crossing arcs, as in Subsection 4.3). The pleating angles of $S^{(i)}$,

in the above convention, are a_i, b_i and $-a_i - b_i$, the latter being the angle at the crossing arcs. Similarly, the 4-punctured sphere $S^{(i+1)}$ defined by the band U_{i+1} (that is, corresponding to the next Conway sphere) has pleating angles $a_{i+1}, b_{i+1}, -a_{i+1} - b_{i+1}$. Let e be the crossing edge $S^{(i)} \cap S^{(i+1)}$: recall that e is obtained by identification of two edges of the solid torus \mathcal{P} , both carrying an interior dihedral angle of $\pi - c_{i,i+1}$. If the interior dihedral angles at e above $S^{(i)}$ and $S^{(i+1)}$ for the transverse orientation add up to $\pi - (a_i + b_i)$ and $\pi - (a_{i+1} + b_{i+1})$, respectively, then the sum of all dihedral angles at e will be

$$[\pi - (a_i + b_i)] + [\pi - (a_{i+1} + b_{i+1})] + 2[\pi - c_{i,i+1}] = 2\pi .$$

Therefore the linear gluing equation at e will automatically be satisfied.

Recall the Farey vertices $s \neq s'$ from Subsection 4.1 associated to the crossing arcs (or preferred slopes) of \mathcal{P} and \mathcal{P}' . If s, s' are Farey neighbors, then \mathcal{P} and \mathcal{P}' are glued directly to one another along a 4-punctured sphere S : the edge pairs of slopes s, s' subdivide S into two ideal squares, and the bands in $\partial\mathcal{P}, \partial\mathcal{P}'$ are traversed by edges exactly as in the first member, (1) vs (2), of the dichotomy of Subsection 4.3 (up to a degree 2 covering).

PROPOSITION 5.9. *At the two parallel edges traversing the band of $\partial\mathcal{P}$, we put an (exterior) dihedral angle $\varepsilon > 0$. We put the same angle ε at the edges traversing the band of \mathcal{P}' . Then, the pleating angles of S at the edge pairs of slope s, s' are $-\varepsilon, \varepsilon$ respectively.*

Proof. This proposition is obvious. □

Moreover, observe that we can artificially select a pair of diagonals in the two squares making up S and define the third pleating angle (along these diagonals) to be 0: then (5.2) is satisfied because the exterior dihedral angles a_i, b_i of \mathcal{P} and \mathcal{P}' at the gluing Conway sphere are 0 and ε , though not necessarily in that order.

We now consider the case where s, s' are not Farey neighbors. The bands of $\partial\mathcal{P}, \partial\mathcal{P}'$ are now subdivided into four triangles each (as in the second member of the dichotomy of Subsection 4.3), defining a pair of ascending and a pair of descending edges in each of the two bands. Fix an arbitrary marking of the Conway sphere along which \mathcal{P} is glued to \mathcal{P}' . Denote by $A \in \mathbb{P}^1\mathbb{Q}$ and $B \in \mathbb{P}^1\mathbb{Q}$ respectively, the slopes of the ascending and descending edge pairs in the band of $\partial\mathcal{P}$, and denote similarly by $A', B' \in \mathbb{P}^1\mathbb{Q}$ the slopes of the edges in $\partial\mathcal{P}'$. We make no assumption on the order of A, A', B, B' in $\mathbb{P}^1\mathbb{Q}$; that is, we favor no convention as to which pair is ascending and which is descending. Denote by a, b the exterior dihedral angles of the block \mathcal{P} at the ascending and descending edges, respectively, and define a', b' in a similar way (relative to \mathcal{P}').

PROPOSITION 5.10. *For any small $\varepsilon > 0$, if $a = a' = b = b' = \varepsilon$, then the tetrahedra between \mathcal{P} and \mathcal{P}' admit positive dihedral angles satisfying the linear gluing equations (at all interior edges).*

Proof. Recall the Farey triangles T_0, \dots, T_m separating s from s' (here, $m \geq 1$). By definition (see Subsection 4.1), we have $T_0 = sAB$ and $T_m = s'A'B'$. Recall also the pleated surface S_i associated to T_i : under our convention (transverse orientation for S_i), the pleating angles of S_0 at the edge pairs of slopes A, B, s are $a, b, -a - b$, respectively. Similarly, the pleating angles of S_m at the edge pairs of slopes A', B', s' are $-a', -b', a' + b'$. We write these numbers in the corresponding corners of T_0 and T_m (Figure 5.2).

For each $0 < i < m$, the oriented line Λ from s to s' enters T_i across some Farey edge $e_i = T_i \cap T_{i-1}$, and exits through another edge e_{i+1} , either to the left or to the right: we say

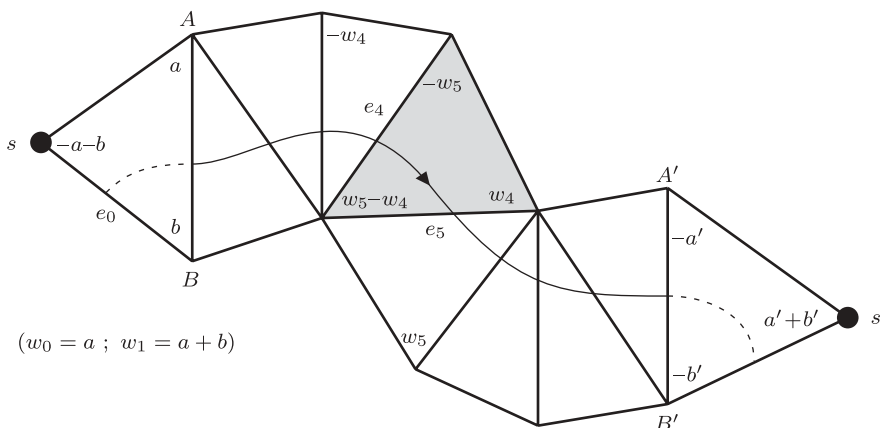


FIGURE 5.2. The pleating angles written in the corners of the Farey triangles T_i associated to the pleated surfaces S_i .

that Λ makes a *Left* or makes a *Right* at T_i , and encode the combinatorics of Λ into a word $\Omega = RLL\dots R$ of length $m - 1$.

No letter (*R* or *L*) is associated *a priori* to the Farey triangles T_0 and T_m . However, we will posit that the path enters T_0 through the Farey edge $e_0 := sB$, and exits T_m through the edge $e_{m+1} := s'B'$, and associate the relevant letter (*R* or *L*) to T_0 and to T_m . Hence, Ω becomes a word of length $m + 1$. This convention is totally artificial, but it will allow us to streamline the notation in our argument.

For $1 \leq i \leq m$, we denote by $\dot{\Delta}_i$ and $\ddot{\Delta}_i$ the two ideal tetrahedra separating the pleated surfaces S_i and S_{i-1} : note that $\{\dot{\Delta}_i, \ddot{\Delta}_i\}$ is naturally associated to the Farey edge $e_i = T_i \cap T_{i-1}$. Our goal is to define dihedral angles for $\dot{\Delta}_i$ and $\ddot{\Delta}_i$ such that the linear gluing conditions around all edges are satisfied. We will in fact restrict to solutions invariant under the Klein group V_4 , that is, such that $\dot{\Delta}_i$ and $\ddot{\Delta}_i$ are isometric[†] for all i (this implies in particular that the angles of $\dot{\Delta}_i$ and $\ddot{\Delta}_i$ at any shared edge are equal). In what follows, Δ_i refers to any one of the ideal tetrahedra $\dot{\Delta}_i, \ddot{\Delta}_i$.

Denote by $\pi - w_i$ the dihedral angle of Δ_i at the pair of opposite edges that is not in $S_i \cap S_{i-1}$ (that is, the pair of edges involved in the diagonal exchange that Δ_i represents). Then S_i has one pleating angle equal to w_i while S_{i-1} has one pleating angle equal to $-w_i$. By translating indices, we find that for all $1 \leq i \leq m - 1$ the pleating angles of S_i must be

$$-w_{i+1}, w_i \text{ and } w_{i+1} - w_i$$

(the value of the third pleating angle is forced upon us by the condition that the pleating angles add up to 0). Further, we can write these three pleating angles in the corners of the Farey triangle T_i associated to S_i (this was partially done in Figure 5.2). In Figure 5.3 (top), denoting by e_i the Farey edge $T_i \cap T_{i-1}$ associated to the tetrahedra $\{\dot{\Delta}_i, \ddot{\Delta}_i\}$, we see that w_i is in the corner of T_i opposite e_i , and $-w_{i+1}$ is in the corner of T_i opposite e_{i+1} .

We repeat the same procedure for all indices $1 \leq i \leq m - 1$. It also extends naturally to $i = 0$ and $i = m$ if we just set $(w_0, w_1) = (a, a + b)$ and $(w_m, w_{m+1}) = (a' + b', a')$: we then recover the pleating angles of S_0 and S_m defined previously.

The bottom part of Figure 5.3 shows the result of the labeling for two consecutive Farey triangles T_{i-1} and T_i , whose corresponding pleated surfaces S_{i-1} and S_i bound the pair of

[†]The graph carried by each 4-punctured sphere S_i is that of the edges of a tetrahedron, and its combinatorial symmetry group is A_4 ; the group $V_4 \subset A_4$ acts on these graphs in the usual way, by pairs of disjoint transpositions of ideal vertices.

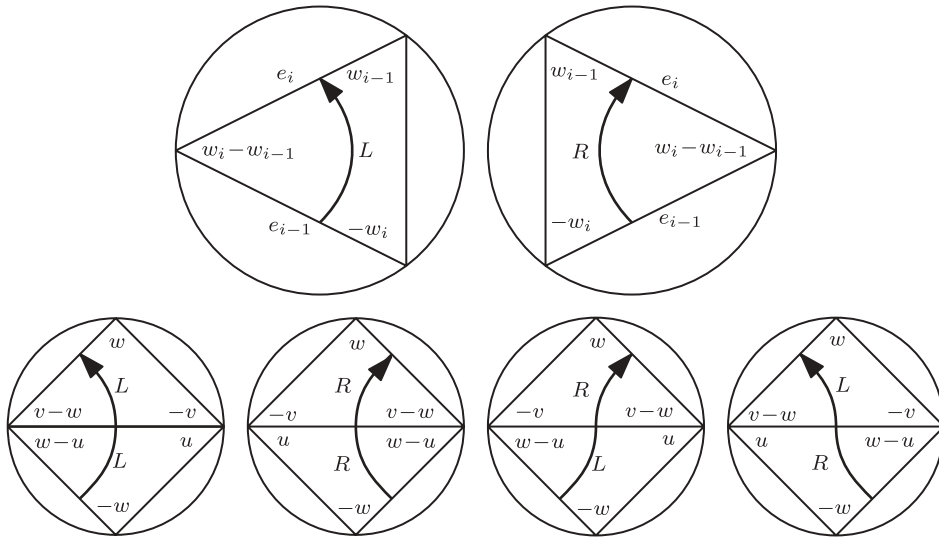


FIGURE 5.3. Top: two consecutive Farey edges e_{i-1}, e_i . Bottom: e_i is the horizontal edge and $(w_{i-1}, w_i, w_{i+1}) = (u, w, v)$.

isometric tetrahedra $\{\check{\Delta}_i, \ddot{\Delta}_i\}$ (where $1 \leq i \leq m$). There are four possible cases, depending on the letters (R or L) of Ω on the Farey triangles T_{i-1} and T_i . In order for Δ_i to have positive angles, assuming that $0 < w_i < \pi$ (for all $1 \leq i \leq m$), it is necessary and sufficient that each pleating angle written just below the horizontal edge e_i be larger than the pleating angle written just above, in Figure 5.3 (bottom): namely, the difference between these two pleating angles (of S_i and S_{i-1}) is twice a dihedral angle of Δ_i .

Suppose that $(w_{i-1}, w_i, w_{i+1}) = (u, w, v)$. Denoting by x_i (y_i) the angle of Δ_i at the edge whose slope is given by the right (left) end of e_i in Figure 5.3 (bottom), we thus find the following formulas for x_i, y_i, z_i (still depending on the neighboring letters of Ω):

Ω	$L \quad L$	$R \quad R$	$L \quad R$	$R \quad L$
x_i	$\frac{1}{2}(u + v)$	$\frac{1}{2}(-u + 2w - v)$	$\frac{1}{2}(u + w - v)$	$\frac{1}{2}(-u + w + v)$
y_i	$\frac{1}{2}(-u + 2w - v)$	$\frac{1}{2}(u + v)$	$\frac{1}{2}(-u + w + v)$	$\frac{1}{2}(u + w - v)$
z_i	$\pi - w$	$\pi - w$	$\pi - w$	$\pi - w$

(5.5)

Define a *hinge index* i as an index such that the Farey triangles adjacent to the Farey edge e_i carry different letters (R and L). From (5.5), we see that Δ_i has positive angles if and only if

- $0 < w_i < \pi$ for all $1 \leq i \leq m$ (range condition);
 - $w_{i+1} + w_{i-1} < 2w_i$ if i is not a hinge index (concavity condition);
 - $|w_{i+1} - w_{i-1}| < w_i$ if i is a hinge index (hinge condition).
- (5.6)

Recall that $m \geq 1$. It is clear that there exist sequences (w_0, \dots, w_{m+1}) satisfying the above conditions such that $(w_0, w_1, w_m, w_{m+1}) = (\varepsilon, 2\varepsilon, 2\varepsilon, \varepsilon)$: for example, set all $(w_i)_{1 \leq i \leq m}$ equal to 2ε , then perturb the non-hinge parameters among $\{w_i\}_{1 < i < m}$ to obtain strong (piecewise) concavity. Proposition 5.10 is proved. \square

5.2.1. *Trivial tangles.* As in [9, Appendix], this method of constructing angle structures extends to the case when \mathcal{P} is glued to a trivial tangle (1-bracelet) of slope s' , realized by

tetrahedra. Then, Figure 4.2 on page 23 shows the surface S_{m-1} associated to the next-to-last Farey triangle T_{m-1} . As in Figure 4.2, we assume (up to changing the marking of the 4-punctured Conway sphere) that $s' = \frac{1}{2}$ and $T_{m-1} = 01\infty$. Gluing the faces of S_{m-1} in pairs amounts to requiring that S_{m-1} have pleating angle $-\pi$ at the edges of slope ∞ : in other words, that $w_m = \pi$. Thus if we put $w_m = \pi$ (the Farey edge associated to w_m is $e_m = T_m \cap T_{m-1}$), Table (5.5) still allows us to derive the angles of the tetrahedra $\Delta_1, \dots, \Delta_{m-1}$, and the positivity conditions are still given by (5.6). (In that case, there is no ‘artificial’ letter on T_m and no ‘artificial’ parameter w_{m+1} .) Angle structures are thus given by sequences

$$(w_0, w_1, w_2, \dots, w_{m-1}, w_m) = (a, a + b, w_2, \dots, w_{m-1}, \pi), \tag{5.7}$$

which satisfy (5.6) for $0 < i < m$. To describe for which pairs (a, b) such a sequence exists, we need some notation.

Given two rationals $q = y/x$ and $q' = y'/x'$ in reduced form in $\mathbb{P}^1\mathbb{Q}$, define

$$q \wedge q' := \left\| \begin{array}{cc} y & y' \\ x & x' \end{array} \right\|$$

(absolute value of the determinant). We will use the following key property: if pqr is a Farey triangle and u, p, r, q are cyclically ordered in $\mathbb{P}^1\mathbb{Q}$, then $u \wedge r = (u \wedge p) + (u \wedge q)$. The property is easily checked in the case $(p, q) = (0, \infty)$ (where u, r have opposite signs), and the general case follows because the \wedge -notation is invariant under $\text{PSL}_2(\mathbb{Z})$, which acts transitively on oriented Farey edges pq .

PROPOSITION 5.11. *Suppose that a trivial tangle complement is glued to a large block \mathcal{P} that has non-negative pleating angles a, b at the edge pairs of slope $A, B \in \mathbb{P}^1\mathbb{Q}$, satisfying (5.2). Suppose that $s, s' \in \mathbb{P}^1\mathbb{Q}$ are the preferred slopes of \mathcal{P} and of the trivial tangle, respectively. Then sAB is the Farey triangle T_0 ; the points s, A, s', B are cyclically ordered in $\mathbb{P}^1\mathbb{Q}$, and the tetrahedra $\Delta_1, \dots, \Delta_{m-1}$ (realizing the trivial tangle complement) admit positive angles whenever*

- $s \wedge s' = 2$ and $a + b = \pi$; or
- $s \wedge s' > 2$ and $a(B \wedge s') + b(A \wedge s') > \pi > a + b$.

Proof. We will show a little more: namely, that the condition is sufficient *and necessary* for the existence of a sequence w satisfying both (5.6) and (5.7). The statements about the relative positions of s, A, B, s' are true by construction and were proved in Subsection 4.1. The case $s \wedge s' = 2$ corresponds to $m = 1$, with the 4-punctured sphere $S_{m-1} = S_0 \subset \partial\mathcal{P}$ being glued directly to itself (as in Figure 4.2). Since $m = 1$, a sequence of the form (5.7) exists if and only if $a + b = \pi$. We now assume that $s \wedge s' > 2$, and consider the sequence of Farey triangles T_0, \dots, T_m from s to s' (where $m \geq 2$). The inequality $\pi > a + b$ must clearly be true in (5.7) by (5.6), so we focus on the other inequality (which says that a, b are *not too small*).

For each $0 \leq i \leq m$, define q_i to be the vertex of T_i not belonging to the edge e_i (where $e_0 = sB$ and $e_i = T_i \cap T_{i-1}$ otherwise). In particular, $q_0 = A$. If

$$\alpha_i = A \wedge q_i \quad \text{and} \quad \beta_i = B \wedge q_i,$$

it is easy to check that both (α_i) and (β_i) make the concavity and hinge conditions of (5.6) *critical* in the following sense: for each $0 < i < m$,

- if i is not a hinge index, then $\alpha_{i+1} + \alpha_{i-1} = 2\alpha_i$ and $\beta_{i+1} + \beta_{i-1} = 2\beta_i$;
- if i is a hinge index, then $\alpha_{i+1} = \alpha_i + \alpha_{i-1}$ and $\beta_{i+1} = \beta_i + \beta_{i-1}$.

(In the first case, observe that $\alpha_{i+1} - \alpha_i = A \wedge p = \alpha_i - \alpha_{i-1}$, where p is the common vertex of the Farey edges e_{i-1}, e_i, e_{i+1} . In the second case, observe that $q_{i-1}q_iq_{i+1}$ is a Farey triangle and A, q_i, q_{i+1}, q_{i-1} are cyclically ordered in $\mathbb{P}^1\mathbb{Q}$.) We say that (α_i) and (β_i) satisfy *the closure*

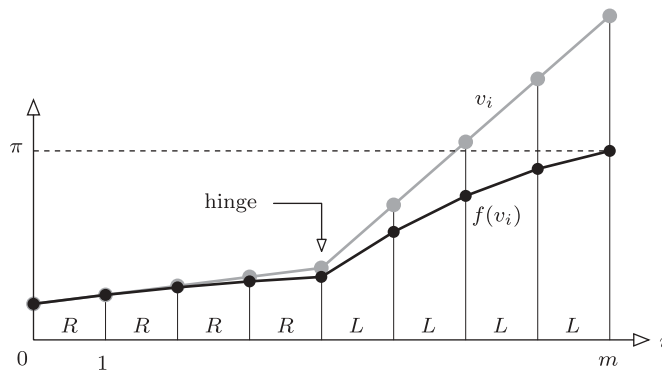


FIGURE 5.4. The sequences $v = (v_i)_{0 \leq i \leq m}$ and $w = f \circ v$.

of the concavity and hinge conditions of (5.6) (the system obtained by turning all the strong inequalities of those conditions into weak ones).

Clearly, any linear combination of the sequences (α_i) and (β_i) also makes the concavity and hinge conditions of (5.6) critical. Define

$$v_i := a\beta_i + b\alpha_i = a(B \wedge q_i) + b(A \wedge q_i),$$

so that $(v_0, v_1) = (a, a + b)$, and (v_i) satisfies the closure of (5.6).

Note that $q_m = s'$, so $v_m = a(B \wedge s') + b(A \wedge s')$ is the left member of the inequality of the Proposition.

Claim: If v' is another sequence that satisfies the closure of (5.6) and $(v'_0, v'_1) = (v_0, v_1)$, then $v'_i \leq v_i$ and $v'_i - v'_{i-1} \leq v_i - v_{i-1}$ for all $1 \leq i \leq m$.

The claim is true for $i = 1$, and follows in general by induction on i : if i is not a hinge index, we have

$$\begin{aligned} v'_{i+1} &\leq v'_i + (v'_i - v'_{i-1}) &\leq v_i + (v_i - v_{i-1}) &= v_{i+1} ; \\ v'_{i+1} - v'_i &\leq v'_i - v'_{i-1} &\leq v_i - v_{i-1} &= v_{i+1} - v_i \end{aligned}$$

(in each line, the first inequality is true by (5.6), and the second, by induction). Similarly, if i is a hinge index, then

$$\begin{aligned} v'_{i+1} &\leq v'_i + v'_{i-1} &\leq v_i + v_{i-1} &= v_{i+1} ; \\ v'_{i+1} - v'_i &\leq v'_{i-1} &\leq v_{i-1} &= v_{i+1} - v_i . \end{aligned}$$

Thus, if $v_m \leq \pi$, then no sequence $w = v'$ satisfies both (5.6) and (5.7), so there can be no positive dihedral angle assignment of the form (5.5) for the tetrahedra $\Delta_1, \dots, \Delta_{m-1}$. Conversely, if $v_m > \pi$, choose a concave function $f : [0, v_m] \rightarrow \mathbb{R}$ that is equal to the identity on $[0, v_1]$, increasing and strictly concave on $[v_1, v_m]$, with $f(v_m) = \pi$; and define $w_i := f(v_i)$ for all $1 \leq i \leq n$ (see Figure 5.4). Then $(a = w_0, w_1, \dots, w_m = \pi)$ satisfies (5.6) (except of course the range condition at $w_m = \pi$). Indeed, the range condition holds because v is increasing and f is monotonic; the concavity condition holds because the image of an arithmetic sequence under a concave map is concave; and the hinge condition at a hinge index i , which only says that $f(v_{i-1}) + f(v_i) > f(v_{i+1})$, is the sum of the concavity inequalities

$$f(v_{i-1}) > \frac{v_{i-1}}{v_{i+1}} f(v_{i+1}) \quad \text{and} \quad f(v_i) > \frac{v_i}{v_{i+1}} f(v_{i+1}). \quad \square$$

5.2.2. *2-bridge links.* When two trivial tangles are glued together, we obtain a 2-bridge link. The construction preceding Proposition 5.11 can be performed both near s and near s' ,

and we refer to [9] for a much more complete treatment. In the remainder of the paper, we assume that the candidate link K is not a 2-bridge link; that is, it contains at least one large bracelet.

5.3. *Non (strongly) Montesinos links*

Observe that the condition in each of Propositions 5.3, 5.6, 5.7 requires that the angles a_i, b_i be, in a loose sense, small enough, while Proposition 5.11 requires them to be large enough. The conflict that arises causes some arborescent link complements to have no angle structures (essentially, the third family of exceptions in Theorem 1.5).

DEFINITION 5.12. A *strongly Montesinos link* is an arborescent link which, after the reduction of Section 3, consists of one non-augmented d -bracelet ($d \geq 3$) attached to d trivial tangles.

Note that some very simple Montesinos links are not strongly Montesinos: for example, links with tangle slopes $(\pm\frac{1}{2}, \pm\frac{1}{2}, \frac{y}{x})$ were reduced in step (5) of the algorithm of Section 3. Such links consist of an augmented 1-bracelet glued to a non-augmented 1-bracelet, and may or may not be candidate links, depending on whether the preferred slopes of the two bracelets satisfy the minimum-distance table, Table 1, of Proposition 3.6.

PROPOSITION 5.13. *If the candidate link K is not strongly Montesinos, then we can endow all blocks with non-negative dihedral angles satisfying (5.2) such that the girth condition (5.3) holds for all non-augmented blocks, and the condition of Proposition 5.11 holds at each trivial tangle. As a result, the ideal decomposition of the link complement admits angle structures, and the link is hyperbolic.*

Proof. Consider a non-augmented block \mathcal{P} : since the link is not Montesinos, \mathcal{P} is separated by a product region from some other large block (augmented or not). By Propositions 5.9 and 5.10, the dihedral angles a_k, b_k of \mathcal{P} at the boundary of that product region can be taken smaller than or equal to any small $\varepsilon_0 > 0$, and the product region will still admit positive angle structures. (More precisely, we take $a_k = b_k = \varepsilon_0$ (Proposition 5.10) except in the special case where the two blocks are glued directly to one another: then, we use the observation that immediately follows Proposition 5.9 and take a_k, b_k equal to ε_0 and 0, though not necessarily in that order.)

We must now find dihedral angles for \mathcal{P} such that the girth condition (5.3) holds. Consider a trivial tangle attached to \mathcal{P} , along the band U_i . Note that

$$A \wedge s' + B \wedge s' = s \wedge s' \geq 2,$$

where the inequality follows from the table of minimal Farey distances, Table 1, in Proposition 3.6. Therefore, by Proposition 5.11, the tetrahedra in the trivial tangle will admit positive structures as soon as

$$a_i = b_i = \begin{cases} \frac{\pi}{2} & \text{if } s \wedge s' = 2, \\ \varepsilon + \frac{\pi}{s \wedge s'} & \text{if } s \wedge s' > 2, \end{cases} \tag{5.8}$$

for a small $\varepsilon > 0$. Since $\max\{a_k, b_k\} = \varepsilon_0$, under this choice of values, the right member $\sum_{i=1}^d \max\{a_i, b_i\}$ of the girth condition (5.3) is thus at most

$$\varepsilon_0 + (d - 1)\frac{\pi}{2} .$$

If $d \geq 4$, this quantity is already less than the left member $(d - 2)\pi$ of (5.3). If $d = 3$, recall from Section 3 that \mathcal{P} is not attached to two trivial tangles of slope $1/2$ modulo \mathbb{Z} (otherwise, we would have replaced the 3-bracelet associated to \mathcal{P} by an augmented 1-bracelet in step (5) of the algorithm). Thus, the above upper bound can be further improved to

$$\varepsilon_0 + (\pi/2) + (\varepsilon + \pi/3) < \pi = (d - 2)\pi,$$

so the girth condition (5.3) is satisfied.

As for augmented bracelets, there is nothing to check: as soon as the a_i, b_i for an augmented block \mathcal{P} satisfy (5.2) and (near trivial tangles) the condition of Proposition 5.11, \mathcal{P} is realized by a union of tetrahedra with positive angles, by equation (5.4) above.

Finally, Theorem 1.1 implies the existence of a hyperbolic structure. □

5.4. Strongly Montesinos links

Suppose that the candidate link K is strongly Montesinos, and recall the non-negative integers n_A, n_B from Definition 5.4.

PROPOSITION 5.14. *If $n_A \geq 3$ or $n_B \geq 3$, the block decomposition admits angle structures.*

Proof. Assume that $n_A \geq 3$ (it is enough to treat this case). It is straightforward to find a pair (α, β) satisfying the condition of Proposition 5.6 (namely $\alpha \geq \beta$), such that $(a, b) = (\pi - \alpha\varepsilon, \beta\varepsilon)$ satisfies the condition of Proposition 5.11 for small ε . For example, take

$$\beta = \alpha \quad \text{if } s \wedge s' = 2, \quad \text{and} \quad \beta = \frac{2}{3}\alpha \quad \text{if } s \wedge s' \geq 3$$

(recall $(A \wedge s') + (B \wedge s') = s \wedge s'$, so $a(B \wedge s') + b(A \wedge s') \geq \min\{2a + b, a + 2b\} > \pi$). The conclusion now follows from the two quoted propositions. □

The next two propositions deal exactly with the remaining strongly Montesinos links, where

$$\max\{n_A, n_B\} = 2$$

(since $n_A + n_B \geq d \geq 3$, this leaves only the cases $(n_A, n_B) \in \{(1, 2), (2, 2), (2, 1)\}$). In each Proposition, we find a few non-hyperbolic links:

- In Proposition 5.15, it is the link already mentioned in Figure 3.8 and Remark 3.10, which falls into the second class of exceptions of Theorem 1.5.
- In Proposition 5.16, it is exactly the strongly Montesinos links among the third class of exceptions of Theorem 1.5.

Recall (Definition 5.4 and Remark 5.5) that n_B is the number of twists in the ‘main band’ of a strongly Montesinos link K , as defined prior to Definition 3.7: thus it would be a straightforward exercise to translate the current block presentations back into planar link diagrams. Recall as well that we have chosen markings for the Conway spheres of K , in which the large bracelet has preferred slope ∞ and the trivial tangles have preferred slopes in the interval $(0, 1)$. All the exceptions arising in Propositions 5.15 and 5.16 were pre-emptively ruled out by the last condition in the Definition 3.7 of candidate links.

PROPOSITION 5.15. *If $d = 4$ and $n_A = n_B = 2$, the block decomposition admits angle structures, unless all the trivial tangles have slope $1/2$.*

Proof. The quadruple (s, A, B, s') associated to a trivial tangle of slope $y/x \in (0, 1)$ is by definition $(\infty, 1, 0, y/x)$: therefore, the key condition $a(s' \wedge B) + b(s' \wedge A) > \pi$ from

Proposition 5.11 becomes

$$a(y) + b(x - y) > \pi. \tag{5.9}$$

Moreover, the denominator x of the slope y/x of the trivial tangle is the integer

$$\infty \wedge (y/x) = s \wedge s'.$$

If at least one of these denominators is larger than 2, we can set a_i, b_i as in (5.8) above, and immediately obtain the girth condition (5.3) because $\pi/2 + \pi/2 + \pi/2 + (\varepsilon + \pi/3) < 2\pi$. If all denominators are 2, the link is not hyperbolic (Figure 3.8) and not candidate; it belongs to the second family of exceptions of Theorem 1.5. \square

PROPOSITION 5.16. *If $d = 3$ and $(n_A, n_B) = (2, 1)$, then assume that the trivial tangles have slopes $y_1/x_1, y_2/x_2, y_3/x_3 \in (0, 1)$: the block decomposition admits angle structures, unless one has $y_1 = y_2 = y_3 = 1$ and $1/x_1 + 1/x_2 + 1/x_3 \geq 1$.*

Proof. Again, if $\sum 1/x_i < 1$, we can set a_i, b_i as in (5.8) to obtain the girth condition (5.3). Thus, assume that $\sum 1/x_i \geq 1$ and (up to a permutation) $y_3 \geq 2$. This entails in particular that $x_3 \geq 3$.

We will set $(a_i, b_i) = (\pi - \alpha_i, \beta_i)$ for well-chosen *small* positive numbers α_i, β_i . As in (5.9) above, the key condition from Proposition 5.11 is still $a_i y_i + b_i(x_i - y_i) > \pi$. If $y_i > 1$, this condition is vacuous for small α_i, β_i . If $y_i = 1$, it can be written $(x_i - 1)\beta_i > \alpha_i$. Thus the following is the full set of sufficient conditions to be satisfied.

- If $x_i = 2$ then $\alpha_i = \beta_i > 0$ (see Proposition 5.11).
- If $x_i > 2$ then $\alpha_i > \beta_i > 0$ (see Proposition 5.11).
- If $x_i > 2$ and $y_i = 1$ then $(x_i - 1)\beta_i > \alpha_i$ (see Proposition 5.11).

- For all $i \in \{1, 2, 3\}$, we have $\alpha_{i+1} + \alpha_{i-1} > \alpha_i + \beta_i$ and $\alpha_i > \beta_{i+1} + \beta_{i-1}$, taking indices modulo 3 (see Proposition 5.7).

Here the first three conditions ensure the existence of angle structures for the trivial tangles; the last one, for the solid torus.

If $x_1 = x_2 = x_3 = 3$, we take

$$\begin{aligned} (\alpha_1, \alpha_2, \alpha_3) &= (5\varepsilon, 5\varepsilon, 7\varepsilon) \\ (\beta_1, \beta_2, \beta_3) &= (3\varepsilon, 3\varepsilon, \varepsilon) \end{aligned} .$$

While we choose $\alpha_1 < 2\beta_1$ and $\alpha_2 < 2\beta_2$ here, note that α_3 is not subjected to the condition $\alpha_3 < (x_3 - 1)\beta_3$, because $y_3 \geq 2$. The list of conditions above is clearly satisfied.

Finally, if $x_1 = 2$, then $x_2 \geq 3$: otherwise, the 3-bracelet associated to the block \mathcal{P} would have been replaced by an augmented 1-bracelet in step (5) of the algorithm of Section 3 (so the link would not be strongly Montesinos). We thus set

$$\begin{aligned} (\alpha_1, \alpha_2, \alpha_3) &= (7\varepsilon, 9\varepsilon, 13\varepsilon) \\ (\beta_1, \beta_2, \beta_3) &= (7\varepsilon, 5\varepsilon, \varepsilon) \end{aligned} .$$

Again, the list of conditions above is clearly satisfied. \square

Thus all candidate links are hyperbolic. We have proved Theorem 3.8, and hence Theorem 1.5.

Acknowledgements. This project began as the first author's Ph D thesis under the guidance of Steve Kerckhoff, was nourished by advice from Francis Bonahon, and reached its completion while both authors were visiting Osaka University and enjoying the hospitality of Makoto

Sakuma. All three of these mentors deserve our deep gratitude for their help and encouragement. We would also like to thank Frédéric Paulin and the anonymous referee for their careful reading and suggestions.

References

1. E. M. ANDREEV, ‘Convex polyhedra of finite volume in Lobačevskii space’, *Math. USSR-Sb.* 12 (1970) 255–259.
2. MICHEL BOILEAU, BERNHARD LEEB and JOAN PORTI, ‘Geometrization of 3-dimensional orbifolds’, *Ann. of Math.* (2) 162 (2005) 195–290.
3. FRANCIS BONAHON and LAURENT SIEBENMANN, ‘Geometric splittings of knots, and Conway’s algebraic knots’, Draft of a monograph, 1979–1985.
4. DARYL COOPER, CRAIG D. HODGSON and STEVEN P. KERCKHOFF, ‘Three-dimensional orbifolds and cone-manifolds’, *MSJ Memoirs*, With a postface by Sadayoshi Kojima, vol. 5 (Mathematical Society of Japan, Tokyo, 2000).
5. ADRIEN DOUADY, ‘Variétés à bord anguleux et voisinages tubulaires’, (French), *Differential topology*, Semin. H. Cartan 14 (Secrétariat Mathématique, Paris, 1961/1962).
6. D. B. A. EPSTEIN and R. C. PENNER, ‘Euclidean decompositions of noncompact hyperbolic manifolds’, *J. Differential Geom.* 27 (1988) 67–80.
7. DAVID GABAI, ‘Genera of the arborescent links’, *Mem. Amer. Math. Soc.* 59 (1986) i–viii and 1–98.
8. FRANÇOIS GUÉRITAUD, ‘Géométrie hyperbolique effective et triangulations idéales canoniques en dimension trois’, PhD Thesis, University of Paris–XI, Orsay, 2006.
9. FRANÇOIS GUÉRITAUD and DAVID FUTER (appendix), ‘On canonical triangulations of once-punctured torus bundles and two-bridge link complements’, *Geom. Topol.* 10 (2006) 1239–1284.
10. WOLFGANG HAKEN, ‘Theorie der Normalflächen’, *Acta Math.* 105 (1961) 245–375.
11. MARC LACKENBY, ‘Word hyperbolic Dehn surgery’, *Invent. Math.* 140 (2000) 243–282.
12. BRUNO MARTELLI and CARLO PETRONIO, ‘3–Manifolds with complexity at most 9’, *Experimental Math.* 10 (2001) 207–237; arXiv:math.GT/0005104.
13. WILLIAM MENASCO, ‘Closed incompressible surfaces in alternating knot and link complements’, *Topology* 23 (1984) 37–44.
14. WALTER D. NEUMANN and DON ZAGIER, ‘Volumes of hyperbolic three-manifolds’, *Topology* 24 (1985) 307–332.
15. ULRICH OERTEL, ‘Closed incompressible surfaces in complements of star links’, *Pacific J. Math.* 111 (1984) 209–230.
16. YO’AV RIECK and ERIC SEDGWICK, ‘Persistence of Heegaard structures under Dehn filling’, *Topology Appl.* 109 (2001) 41–53.
17. IGOR RIVIN, ‘Euclidean structures on simplicial surfaces and hyperbolic volume’, *Ann. of Math.* (2) 139 (1994) 553–580.
18. IGOR RIVIN, ‘A characterization of ideal polyhedra in hyperbolic 3-space’, *Ann. of Math.* (2) 143 (1996) 51–70.
19. MAKOTO SAKUMA, ‘The geometries of spherical Montesinos links’, *Kobe J. Math.* 7 (1990) 167–190.
20. MAKOTO SAKUMA and JEFFREY WEEKS, ‘Examples of canonical decompositions of hyperbolic link complements’, *Japan. J. Math. (N.S.)* 21 (1995) 393–439.
21. JEAN-MARC SCHLENKER, ‘Hyperbolic manifolds with polyhedral boundary’, Preprint, 2001, arXiv:math.GT/0111136.
22. WILLIAM P. THURSTON, *The geometry and topology of three-manifolds*, Princeton University Mathematics Department Notes (Princeton University, Princeton, NJ, 1980); Available at <http://www.msri.org/gt3m/>.
23. WILLIAM P. THURSTON, ‘Three-dimensional manifolds, Kleinian groups and hyperbolic geometry’, *Bull. Amer. Math. Soc. (N.S.)* 6 (1982) 357–381.
24. JEFF WEEKS, ‘Computation of hyperbolic structures in knot theory’, *Handbook of knot theory* (Elsevier, Amsterdam, 2005) 461–480.
25. YING-QING WU, ‘Dehn surgery on arborescent knots’, *J. Differential Geom.* 43 (1996) 171–197.

David Futer
 Mathematics Department
 Temple University
 Philadelphia, PA 19122
 USA

dfuter@math.temple.edu

François Guéritaud
 Laboratoire de mathématiques
 Université de Paris–Sud
 91405 Orsay cedex
 France

Francois.Gueritaud@normalesup.org