

The link volume of 3–manifolds

YO'AV RIECK
YASUSHI YAMASHITA

We view closed orientable 3–manifolds as covers of S^3 branched over hyperbolic links. To a cover $M \xrightarrow{p} S^3$, of degree p and branched over a hyperbolic link $L \subset S^3$, we assign the complexity $p \operatorname{Vol}(S^3 \setminus L)$. We define an invariant of 3–manifolds, called the *link volume* and denoted by $\operatorname{LinkVol}(M)$, that assigns to a 3–manifold M the infimum of the complexities of all possible covers $M \rightarrow S^3$, where the only constraint is that the branch set is a hyperbolic link. Thus the link volume measures how efficiently M can be represented as a cover of S^3 .

We study the basic properties of the link volume and related invariants, in particular observing that for any hyperbolic manifold M , $\operatorname{Vol}(M)$ is less than $\operatorname{LinkVol}(M)$. We prove a structure theorem that is similar to (and uses) the celebrated theorem of Jørgensen and Thurston. This leads us to conjecture that, generically, the link volume of a hyperbolic 3–manifold is much bigger than its volume.

Finally we prove that the link volumes of the manifolds obtained by Dehn filling a manifold with boundary tori are linearly bounded above in terms of the length of the continued fraction expansion of the filling curves.

57M12, 57M50; 57M27

1 Introduction

The study of 3–manifolds as branched covers of S^3 has a long history. In 1920 Alexander [1] gave a very simple argument showing that every closed orientable triangulated 3–manifold is a cover of S^3 branched along the 1–skeleton of a tetrahedron embedded in S^3 . We explain his construction and give basic definitions in Section 2. Clearly, if a 3–manifold M is a finite sheeted branched cover of S^3 , then M is closed and orientable. Moise [14] showed that every closed 3–manifold admits a triangulation, thus we see: a 3–manifold M is closed and orientable if and only if M is a finite sheeted branched cover of S^3 . From this point on, by a *manifold* we mean a connected closed orientable 3–manifold.

Alexander himself noticed one weakness of his theorem: the branch set is not a submanifold. He claimed that this can be easily resolved, but gave no indication of the

proof. In 1986 Feighn [5] substantiated Alexander's claim, modifying the branch set to be a link.

Thurston showed the existence of a *universal link*, that is, a link $L \subset S^3$ so that every 3-manifold is a cover of S^3 branched along L . Hilden, Lozano and Montesinos [7; 8] drastically simplified Thurston's example showing, in particular, that the figure eight knot is universal. Cao and Meyerhoff [3] showed that the figure eight knot is the hyperbolic link of the smallest volume. In this paper, we consider hyperbolic links and consider their volume as a measure of complexity, hence we see that every 3-manifold is a cover of S^3 branched along the simplest possible link.

Our goal is to define and study an invariant that asks: how efficient is the presentation of a 3-manifold as a branched cover of S^3 ? We do this as follows: let M be a p -fold cover of S^3 , branched along the hyperbolic link L . We denote this by $M \xrightarrow{p} (S^3, L)$ (read: M is a p -fold cover of S^3 branched along L). The complexity of $M \xrightarrow{p} (S^3, L)$ is defined to be the degree of the cover times the volume of L , that is:

$$p \operatorname{Vol}(S^3 \setminus L).$$

The *link volume* of M , denoted by $\operatorname{LinkVol}(M)$, is the infimum of the complexities of all covers $M \xrightarrow{p} (S^3, L)$, subject to the constraint that L is a hyperbolic link, that is:

$$\operatorname{LinkVol}(M) = \inf \{ p \operatorname{Vol}(S^3 \setminus L) \mid M \xrightarrow{p} (S^3, L); L \text{ hyperbolic} \}.$$

Given a hyperbolic manifold M we consider its volume, $\operatorname{Vol}(M)$, as its complexity. This is consistent with our attitude towards hyperbolic links, and is considered very natural by many 3-manifold topologists. Why is that? What is it that the volume actually measures? Combining results of Gromov, Jørgensen and Thurston (for a detailed exposition, see Kobayashi and the first author's [9]) we learn the following. Let $t_C(M)$ denote the minimal number of tetrahedra required to triangulate a link exterior in M , that is, the least number of tetrahedra required to triangulate $M \setminus N(L)$, where the minimum is taken over all possible links $L \subset M$ (possibly, $L = \emptyset$) and all possible triangulations of $M \setminus N(L)$. Then there exist constants $a, b > 0$ which are independent of M so that

$$(1) \quad a \operatorname{Vol}(M) \leq t_C(M) \leq b \operatorname{Vol}(M).$$

We consider invariants up to linear equivalence, and so we see that Vol and t_C are equivalent. This gives a natural, topological interpretation of the volume. In this paper we begin the study of the link volume, with the ultimate goal of obtaining a topological understanding of it.

The basic facts about the link volume are presented in Section 4. The most important ones are the following easy observations:

- (1) The link volume is attained, that is, for any manifold M there is a cover $M \xrightarrow{p} (S^3, L)$ so that $\text{LinkVol}(M) = p \text{Vol}(S^3 \setminus L)$.
- (2) For every hyperbolic 3-manifold M we have:

$$\text{LinkVol}(M) > \text{Vol}(M).$$

The second point begs the question: is the link volume of hyperbolic manifolds equivalent to the hyperbolic volume? As we shall see below, the results of this paper lead us to believe that this is not the case (Conjectures 1.2 and 1.3).

The right hand side of the inequality (1) implies that, for fixed V , any hyperbolic manifold of volume less than V can be obtained from a manifold X by Dehn filling, where X is constructed using at most bV tetrahedra. Since there are only finitely many such X , this implies the celebrated result of Jørgensen and Thurston: for any $V > 0$, there exists a finite collection of compact “parent manifolds” $\{X_1, \dots, X_n\}$, so that ∂X_i consists of tori and any hyperbolic manifold of volume at most V is obtained by Dehn filling X_i , for some i . Our first result is:

Theorem 1.1 (Structure Theorem) *There exists a universal constant $\Lambda > 0$ so that for every $V > 0$, there is a finite collection $\{\phi_i: X_i \rightarrow E_i\}_{i=1}^{n_V}$, where X_i and E_i are complete finite volume hyperbolic manifolds and ϕ_i is an unbranched cover, and for any cover $M \xrightarrow{p} (S^3, L)$ with $p \text{Vol}(S^3 \setminus L) < V$ the following conditions hold:*

- (1) *For some i , M is obtained from X_i by Dehn filling, S^3 is obtained from E_i by Dehn filling and the following diagram commutes (where the vertical arrows represent the covering projections and the horizontal arrows represent Dehn fillings):*

$$\begin{array}{ccc} X_i & \longrightarrow & M \\ \downarrow \phi_i & & \downarrow \phi \\ E_i & \longrightarrow & (S^3, L) \end{array}$$

- (2) *E_i can be triangulated using at most $\Lambda V / p$ tetrahedra (hence X_i can be triangulated using at most ΛV tetrahedra so that ϕ_i is simplicial).*

For $V > 0$, let \mathcal{M}_V denote the set of manifolds of link volume less than V . Since the link volume is always attained, applying Theorem 1.1 to covers realizing the link volumes of manifolds in \mathcal{M}_V , we obtain a finite family of “parent manifolds” X_1, \dots, X_n

that give rise to every manifold in \mathcal{M}_V via Dehn filling, much like Jørgensen and Thurston. The extra structure given by the projection $\phi_i: X_i \rightarrow E_i$ implies that the fillings that give rise to manifolds of low link volume are very special:

Fix V , and let X_i be as in the statement of Theorem 1.1. Then for any hyperbolic manifold M that is obtained by filling X_i we have $\text{Vol}(M) < \text{Vol}(X_i)$. On the other hand, it is by no means clear that $\text{LinkVol}(M) < V$, for it is not easy to complete the diagram in Theorem 1.1:

- (1) X_i must cover a manifold E_i .
- (2) The covering projection and the filled slopes must be compatible (see Section 2.3 for definition).
- (3) The slopes filled on E_i must give S^3 , a very unusual situation since E_i is hyperbolic.

These lead us to believe that the link volume, as a function, is much bigger than the volume. Specifically we conjecture:

Conjecture 1.2 *Let X be a compact connected orientable manifold whose boundary consists of a single torus and suppose that the interior of X admits a complete finite volume hyperbolic metric. For a slope α on ∂X , let $X(\alpha)$ denote the closed manifold obtained by filling X along α .*

Then for any $V > 0$, there exists a finite set of slopes \mathcal{F} on ∂X , so that if

$$\text{LinkVol}(X(\alpha)) < V,$$

then α intersects some slope in \mathcal{F} at most $V/2$ times.

As is well known, the volume of the figure eight knot complement is about $2.029\dots$, twice v_3 , the volume of a regular ideal tetrahedron. By considering manifolds that are obtained by Dehn filling the figure eight knot exterior we see that Conjecture 1.2 implies:

Conjecture 1.3 *For every $V > 0$ there exists a manifold M so that*

$$\text{Vol}(M) < 2v_3 = 2.029\dots \quad \text{and} \quad \text{LinkVol}(M) > V.$$

To describe our second result, we first define the knot volume and a few other variations of the link volume; for the definition of simple cover, see Section 2.2.

Definitions 1.4 (1) The *knot volume* of a 3-manifold M is obtained by considering only hyperbolic knots in the definition of the link volume, that is,

$$\text{KnotVol}(M) = \inf \{ p \text{Vol}(S^3 \setminus K) \mid M \xrightarrow{p} (S^3, K); K \text{ is a hyperbolic knot} \}.$$

(2) The *simple knot volume* of a 3-manifold M is obtained by considering only simple covers in the definition of the knot volume, that is,

$$\text{KnotVol}_s(M) = \inf \left\{ p \text{Vol}(S^3 \setminus K) \mid M \xrightarrow{p} (S^3, K); \begin{array}{l} K \text{ is a hyperbolic knot} \\ \text{and the cover is simple} \end{array} \right\}.$$

(3) For an integer $d \geq 3$, the *simple d -knot volume* is obtained by restricting to p -fold covers for $p \leq d$ in the definition of the simple knot volume, that is,

$$\text{KnotVol}_{s,d}(M) = \inf \left\{ p \text{Vol}(S^3 \setminus K) \mid M \xrightarrow{p} (S^3, K); \begin{array}{l} K \text{ is a hyperbolic knot,} \\ \text{the cover is simple} \\ \text{and } p \leq d \end{array} \right\}.$$

Similarly, one can play with various restrictions on the covers considered. However, one must ensure that the definition makes sense. For example, the *regular link volume* can be defined using only regular covers. This makes no sense, as not every manifold is the regular cover of S^3 . It follows from Hilden [6] and Montesinos [15] that every 3-manifold is a simple 3-fold cover of S^3 branched over a hyperbolic knot; hence the definitions above make sense. Our next result is an upper bound, and holds for any of the variations listed in Definitions 1.4. Since these definitions are obtained by adding restrictions to the covers considered, it is clear that $\text{KnotVol}_{s,3}(M)$ is greater than or equal to any of the others, including the link volume. We therefore phrase Theorem 1.6 below for that invariant. But first we need:

Definition 1.5 Let T be a torus and μ, λ generators for $H_1(T; \mathbb{Z})$. By identifying μ with $1/0$ and λ with $0/1$, we get an identification of the slopes of T with $\mathbb{Q} \cup \{1/0\}$, where by a *slope* we mean a nontrivial element of $H_1(T; \mathbb{Z})$, defined up to sign, that can be represented by a connected simple closed curve on T . Then the *depth* of a slope α , denoted by $\text{depth}(\alpha)$, is the length of the shortest continued fraction expansion representing p/q . For a collection of tori T_1, \dots, T_n with bases chosen for $H_1(T_i; \mathbb{Z})$ for each i we define

$$\text{depth}(p_1/q_1, \dots, p_n/q_n) = \sum_{i=1}^n \text{depth}(p_i/q_i).$$

We are now ready to state:

Theorem 1.6 *Let X be a connected compact orientable 3-manifold, ∂X consisting of n tori T_1, \dots, T_n and fix μ_i, λ_i , generators for $H_1(T_i; \mathbb{Z})$ for each i .*

Then there exist a universal constant B and a constant A that depends on X and the choice of bases for $H_1(T_i; \mathbb{Z})$, so that for any p_i/q_i ($i = 1, \dots, n$),

$$\text{KnotVol}_{s,3}(X(p_1/q_1, \dots, p_n/q_n)) < A + B \text{ depth}(p_1/q_1, \dots, p_n/q_n),$$

where $X(p_1/q_1, \dots, p_n/q_n)$ denotes the manifold obtained by filling X along the slopes p_i/q_i .

Remark 1.7 Both in Theorem 1.6 and in Corollary 1.8 below we can change the depths of p_i/q_i arbitrarily by changing the bases used for $H_1(T_i; \mathbb{Z})$; this will result in a change to the constant A . If Conjecture 1.2 holds, then the values needed for A are unbounded.

As noted above, $\text{KnotVol}_{s,3}(M)$ is greater than or equals to all the invariants defined in Definition 1.5 and the link volume. Hence Theorem 1.6, which gives an upper bound, holds for all these invariants, and in particular:

Corollary 1.8 *With the hypotheses of Theorem 1.6, there exist a universal constant B and a constant A that depends on X and the choice of bases for $H_1(T_i; \mathbb{Z})$, so that for any slopes p_i/q_i ($i = 1, \dots, n$),*

$$\text{LinkVol}(X(p_1/q_1, \dots, p_n/q_n)) < A + B \text{ depth}(p_1/q_1, \dots, p_n/q_n).$$

Organization This paper is organized as follows. In Section 2 we go over necessary background material. In Section 3 we explain some possible variations on the link volume. Notably, we define the *surgery volume* (definition due to Kimihiko Motegi) and an invariant denoted by $pB(M)$ (definition due to Ryan Blair). We show that, *conjecturally in contrast to the link volume*, the surgery volume of hyperbolic manifolds is bounded in terms of their volumes. We also show that $pB(M)$ is linearly equivalent to $g(M)$, the Heegaard genus of M . In Section 4 we explain basic facts about the link volume and list some open questions. In Section 5 we prove Theorem 1.1. In Section 6 we prove Theorem 1.6.

Acknowledgement We thank Ryan Blair, Mike Hilden, Tsuyoshi Kobayashi, Kimihiko Motegi, Hitoshi Murakami and Jair Remigio-Juárez for helpful conversations. We are very grateful to the anonymous referee for a careful reading of the paper and numerous helpful suggestions.

2 Background

By a *manifold* we mean a connected closed orientable 3-manifold. In some cases, we consider a connected compact orientable 3-manifolds; then we explicitly say *compact manifold*. By a *hyperbolic manifold* X we mean a complete, finite volume Riemannian 3-manifold locally isometric to \mathbb{H}^3 . It is well known that any orientable hyperbolic manifold X is the interior of a compact manifold \bar{X} and $\bar{X} \setminus X = \partial\bar{X}$ consists of tori. To simplify notation, we do not refer to \bar{X} explicitly and call $\partial\bar{X}$ the *boundary of X* . We assume familiarity with the basic concepts of 3-manifold theory and hyperbolic manifolds and in particular the Margulis constant. By *volume* we mean the hyperbolic volume. The volume of a hyperbolic manifold M is denoted by $\text{Vol}(M)$.

We follow standard notation. In particular, by *Dehn filling* (or simply *filling*) we mean attaching a solid torus to a torus boundary component.

2.1 Branched covering

We begin by recalling Alexander's Theorem [1]; because its proof is very short and elegant we include a sketch below.

Theorem 2.1 (Alexander) *Let \mathcal{T} be a triangulation of S^n obtained by doubling an n -simplex. Let M be a closed orientable triangulated n -manifold. Then M is a cover of S^n branched along $\mathcal{T}^{(n-2)}$, the $(n-2)$ -skeleton of \mathcal{T} .*

Sketch of proof Let M be as above. Given \mathcal{T}_M , a triangulation of M , let \mathcal{T}'_M denote its barycentric subdivision. Each vertex v of \mathcal{T}'_M is the center of a k -face of \mathcal{T}_M , for some k . Label v with the label k . By construction, there are exactly $n+1$ labels, $0, \dots, n$, and no two adjacent vertices have the same label.

Note that the 1-skeleton of \mathcal{T} is K_{n+1} , the complete graph on $n+1$ vertices. Label these vertices with the labels $0, \dots, n$ so that every label appears exactly once.

We define a function from $\mathcal{T}'_M^{(n-1)}$ (the $(n-1)$ -skeleton of \mathcal{T}'_M) to S^n by sending each k -face simplicially to the unique k -face of S^n with the same labeling (for $k < n$); it is easy to see that this function is well defined. However, each n -cell of M can be sent to either of the two n simplices of \mathcal{T}'_M . We pick the simplex so that the map is orientation preserving.

It is left to the reader to verify that this is indeed a cover branched over the $(n-2)$ -skeleton of the triangulation of S^n . \square

Lemma 2.2 For any compact triangulated n -manifold M , $B \subset M^{(n-2)}$ a subcomplex and $d > 0$, there are only finitely many p -fold covers of M branched along B for $p \leq d$.

Proof It is well known that a p -fold cover of M branched along B is determined by a presentation of $\pi_1(M \setminus B)$ into S_p , the symmetric group on p elements (see [18]). The lemma follows from the fact that $\pi_1(M \setminus B)$ is finitely generated and S_p is finite. \square

2.2 Simple covers and the Montesinos move

Definition 2.3 Let $f: M \rightarrow N$ be a cover of finite degree p branched along $B \subset N$. Note that every point of $N \setminus B$ has exactly p preimages, and every point of B has at most p preimages. $f: M \rightarrow N$ is called *simple* if every point of B has exactly $p - 1$ preimages.

Let $M \rightarrow (S^3, L)$ be a 3-fold simple cover branched along the link L . We view L diagrammatically, as projected into $S^2 \subset S^3$ in the usual way. Since the cover is simple, each generator in the Wirtinger presentation of $S^3 \setminus L$ corresponds to a permutation in the symmetric group on 3 elements (that is, $(1\ 2)(3)$ or $(1\ 3)(2)$ or $(2\ 3)(1)$). We consider these as three colors, and color each strand of L accordingly. By assumption, M is connected; hence not all generators correspond to the same permutation. Finally, the relators of the Wirtinger presentation guarantee that at each crossing either all three color appear or only one color does. Thus we obtain a 3-coloring of the strands of L .

Montesinos proved that if we replace a positive crossing where *all three colors appear* by 2 negative crossings the cover is not changed. This is called the *Montesinos move*. The reason is simple: the neighborhood of a 3-colored crossing is a ball, and its cover is a ball as well. (This is false if only one color appears at the crossing!) More generally, when all three colors appear we can replace n half twists with $n + 3k$ half twists ($n, k \in \mathbb{Z}$). The case $n = 0$ is allowed, but then we must require that the two strands in question have distinct colors. We denote such a move by $n \mapsto n + 3k$ Montesinos move. In Figure 1 we show a few views of the Montesinos move.

Finally, we record the following fact for future reference. It is easy to see that a p -fold cover $f: M \rightarrow S^3$ branched along $B \subset S^3$ is connected if and only if the image of $\pi_1(S^3 \setminus B)$ in S_p acts transitively on the set of p letters. For simple 3-fold covers this means:

Lemma 2.4 Let M be a 3-manifold and $f: M \rightarrow S^3$ a simple 3-fold cover branched along the link $L \subset S^3$. Then M is connected if and only if at least two colors appear in the 3-coloring of L .




Figure 1: Montesinos move

2.3 Slopes on tori and coverings

Recall that a *slope* on a torus T is a nontrivial class in $H_1(T; \mathbb{Z})$, defined up to sign, that can be represented by connected simple closed curve. For this subsection we fix the following: Let X and E be complete hyperbolic manifolds of finite volume and $\phi: X \rightarrow E$ an unbranched cover. Let T be a boundary component of X ; note that ϕ induces an unbranched cover $T \rightarrow \phi(T)$.

Let α be a slope on T realized by a connected simple closed curve $\gamma \subset T$. Then $\phi(\gamma)$ is a (not necessarily simple) connected essential closed curve on $\phi(T)$. Since $\phi(T)$ is a torus, there is a simple closed curve $\bar{\beta}$ on $\phi(T)$ so that $\phi(\gamma)$ is homotopic to $\bar{\beta}^m$, for some $m \neq 0$. Let β be the slope defined by $\bar{\beta}$. Define the function ϕ_\downarrow from the slopes on T to the slopes on $\phi(T)$ by setting $\phi_\downarrow(\alpha) = \beta$.

Conversely, let α be a slope on $\phi(T)$ realized by a connected simple closed curve $\gamma \subset \phi(T)$. Then $\phi^{-1}(\gamma)$ is a (not necessarily connected) essential simple closed curve. Each component of $\phi^{-1}(\gamma)$ defines a slope on T , and since these curves are disjoint, they all define the same slope, say β . Define the function ϕ_\uparrow from the slopes on $\phi(T)$ to the slopes on T by setting $\phi_\uparrow(\alpha) = \beta$. It is easy to see that ϕ_\downarrow is the inverse of ϕ_\uparrow . We say that α and $\phi_\downarrow(\alpha)$ are *corresponding slopes*.

Suppose that we fill T and $\phi(T)$. If the slopes filled are not corresponding, then the curve filled on T maps to a curve of $\phi(T)$ that is not null homotopic in the attached solid torus. Thus the map ϕ cannot be extended into that solid torus.

Conversely, suppose that corresponding slopes are filled. We parametrize the attached solid tori as $S^1 \times D^2$ and extend ϕ into the solid tori by coning along each disk $\{p\} \times D^2$ ($p \in S^1$). It is easy to see that if we do this for every component of $\phi^{-1}(\phi(T))$ then the extended map is a cover, branched (if at all) along cores of attached solid tori. (The local degree at the core of a solid torus is the number denoted by m in the construction of ϕ_\downarrow above.)

In conclusion, ϕ induces a correspondence between slopes of T and slopes of $\phi(T)$, and ϕ can be extended to the attached solid tori to give a branched cover after filling if and only if corresponding slopes are filled.

Next, let $T_1, T_2 \subset \partial X$ be tori that project to the same component of ∂E . Then two bijections ϕ_{\downarrow} from the slopes of T_1 and T_2 to the slopes of $\phi(T_1) = \phi(T_2)$ induce a bijection between the slopes of T_1 and the slopes of T_2 ; again we call slopes that are interchanged by this bijection *corresponding*. Filling T_1 and T_2 along corresponding slopes is called *consistent*, *inconsistent* otherwise. Note that after filling X there is a filling of E so that the cover $X \rightarrow E$ extends to a branched cover if and only if the filling of X is consistent on every pair of components of ∂X that project to the same component of ∂E .

2.4 Hyperbolic alternating links

In this subsection we follow Lickorish [12, Chapter 4]. We begin with the following standard definitions:

Definitions 2.5 Let L be a link and D a diagram for L . The projection sphere is denoted by S^2 . Then D is called *alternating* if, for each component K of L , when traversing the projection of K the crossing occur as ...over, under, over, under, ... L is called an *alternating link* if it admits an alternating diagram. A link diagram D is called *strongly prime* if any simple closed curve that intersects it transversely in two simple points (that is, two points that are not crossings) bounds a disk that D intersects in a single arc with no crossings. A link L is called *split* if its exterior admits an essential sphere, that is, if there is an embedded sphere $S \subset S^3 \setminus L$ so that each of the balls obtained by cutting S^3 open along S contains at least one component of L . A link diagram $D \subset S^2$ is called a *split diagram* if there is a circle γ embedded in S^2 , disjoint from D , so that each disk obtained by cutting S^2 open along γ contains at least one component of D . Note that a split diagram is necessarily a diagram for a split link, but the converse does not hold. A link is called *simple* if its exterior does not admit an essential surface of nonnegative Euler characteristic. A link L is called *hyperbolic* if $S^3 \setminus L$ admits a complete, finite volume, hyperbolic metric.

Menasco [13] (see also Lickorish [12]) proved:

Theorem 2.6 *Let D be an alternating link diagram for a link L . If D is strongly prime and is not split, then either L is simple, or it is a torus link.*

Thurston proved:

Theorem 2.7 *Any simple link is hyperbolic.*

Since the only alternating torus links are $2, n$ -torus links, combining these results gives us a way, in certain situations, of identifying hyperbolic links by looking at their diagrams:

Corollary 2.8 *If a link L has a nonsplit, strongly prime, alternating diagram which is not the alternating diagram of a torus $2, n$ -link, then L is hyperbolic.*

2.5 Twist number and hyperbolic volume

For the definition of twist number see, for example, [11]. We briefly recall it here. Let D be a link diagram. Let \sim be the equivalent relation on the crossings of D generated by $c \sim c'$ if c and c' lie on the boundary of a bigon of D . This equivalence relation can be visualized as follows: if c_1, \dots, c_n form an equivalence class of crossings, then after reordering if necessary, there is a chain of $n - 1$ bigons in D with c_i and c_{i+1} on the boundary the i -th bigon.

The *twist number* of a link L , denoted by $t(L)$, is the smallest number of equivalence classes in any diagram for L . Thus, for example, the obvious diagram of twist knots show they have twist number at most 2.

Lackenby [11] gave upper and lower bounds on the hyperbolic volume of link exteriors in terms of their twist number. We emphasize that the lower bound holds for alternating links (or, more precisely, for alternating diagrams), while the upper bound holds *for any diagram of any link*. It is the upper bound that we will need in this work, hence we need not assume the diagram alternates. We will need:

Theorem 2.9 (Lackenby [11]) *There exists a constant c so that for any hyperbolic link L ,*

$$\text{Vol}(S^3 \setminus L) < ct(L).$$

3 Variations

In this section we discuss two variations of the link volume. The first variation is obtained by replacing the volume by another knot invariant (note that one can use any invariant with values in $\mathbb{R}_{\geq 0}$). This variation was suggested by Ryan Blair. Let $L \subset S^3$ be a link and let $b(L)$ denote its bridge index. We consider the complexity of $M \xrightarrow{p} (S^3, L)$ to be $pb(L)$. Define $\text{pB}(M)$ to be the infimum of $pb(L)$, taken over all possible covers.

It is easy to see that the preimage of a bridge surface S for L is a Heegaard surface for M , say Σ . Since S is a $2b$ punctured sphere, $\chi(S \setminus L) = 2 - 2b$. Its preimage has Euler characteristic $p(2 - 2b)$. We obtain Σ by adding some number of points, say $n \geq 0$. Then $\chi(\Sigma) = p(2 - 2b) + n$. Thus we get:

$$\begin{aligned} 2g(\Sigma) - 2 &= -\chi(\Sigma) \\ &= p(2b - 2) - n \\ &= 2pb - (2p + n) \\ &\leq 2pb - 2. \end{aligned}$$

Since $pB(M)$ is positive integer valued, the infimum is attained. By considering a cover that realizes $pB(M)$, we obtain a surface Σ so that $g(\Sigma) \leq pB(M)$. Thus $g(M) \leq g(\Sigma) \leq pB(M)$.

The converse is highly nontrivial. Given an arbitrary manifold M , Hilden [6] constructed a 3-fold cover $M \xrightarrow{3} (S^3, L)$. The construction uses an arbitrary Heegaard surface $\Sigma \subset M$. One feature of Hilden's construction is that $b(L) \leq g(\Sigma) + 2$. Since Σ was an arbitrary Heegaard surface, we may assume that $g(\Sigma) = g(M)$. Thus we see that $pB(M) \leq 3g(M) + 6$. Combining the inequalities we got we obtain:

$$g(M) \leq pB(M) \leq 3g(M) + 6.$$

Thus we see that the Heegaard genus and pB are equivalent.

Another variation, suggested by Kimihiko Motegi, is the *surgery volume*. Given a manifold M , it is well known that M is obtained by Dehn surgery on a link in S^3 , say L . By Myers [16], every compact 3-manifold admits a simple knot. Applying this to $S^3 \setminus N(L)$ it is easy to obtain a knot K so that $L' = L \cup K$ is a hyperbolic link. Since M is obtained from S^3 via surgery along L' (with the original surgery coefficients on L and the trivial slope on K), we conclude that M is obtained from S^3 via surgery along a hyperbolic link. The surgery volume of M is then

$$\text{SurgVol}(M) = \inf\{\text{Vol}(S^3 \setminus L) \mid M \text{ is obtained by surgery on } L, L \text{ is hyperbolic}\}.$$

Results of Gromov, Jørgensen and Thurston show that if a hyperbolic manifold N_1 is obtained by filling a hyperbolic manifold N_2 , then $\text{Vol}(N_1) < \text{Vol}(N_2)$ (see [19, Theorem 6.5.6]). Applying this in our setting (with $S^3 \setminus L'$ as N_2 and M as N_1) we see that for any hyperbolic manifold M , $\text{Vol}(M) < \text{SurgVol}(M)$.

We note that there exists a function $f: (0, \infty) \rightarrow (0, \infty)$ so that any hyperbolic manifold M is obtained by surgery on a hyperbolic link $L \subset S^3$ with $\text{Vol}(S^3 \setminus L) \leq f(\text{Vol}(M))$. To see this, fix V and let X_1, \dots, X_n be the set of parent manifolds of all hyperbolic

manifolds of volume at most V . For each X_i there is a link L_i in S^3 , so that X_i is obtained by surgery on some of the components of L_i and drilling the rest. Therefore, any hyperbolic manifold M with volume at most V is obtained on surgery on some L_i ($i = 1, \dots, n$). Set

$$f(V) = \max_{1 \leq i \leq n} \{\text{Vol}(S^3 \setminus L_i)\}.$$

We get:

$$\text{Vol}(M) \leq \text{SurgVol}(M) \leq f(\text{Vol}(M)).$$

The surgery volume and the hyperbolic volume are equivalent if there is a *linear* function f bounding SurgVol ; we do not know if this is the case.

4 Basic facts and open questions

Basic facts about the link volume are:

The link volume is attained: That is, for every M there exists a cover $M \xrightarrow{p} (S^3, L)$ so that $\text{LinkVol}(M) = p \text{Vol}(S^3 \setminus L)$. Recall that the link volume was defined as an infimum. To see that there is a cover realizing it, we need to show that the infimum is attained. Fix a manifold M , and let $M \xrightarrow{p_n} (S^3, L_n)$ be a sequence of covers that approximates $\text{LinkVol}(M)$. By Jørgensen and Thurston [19] there exists $\epsilon > 0$ so that for every n , $\text{Vol}(S^3 \setminus L_n) > \epsilon$. Hence for large enough n , $p_n \leq \text{LinkVol}(M)/\epsilon$; we see that there are only finitely many values for p_n . By subsequencing if necessary we may assume that p_n is constant, say p . For any collection of covers $M \xrightarrow{p} (S^3, L_n)$ of fixed degree p , the infimum of $\{p \text{Vol}(S^3 \setminus L_n)\}$ is attained since the set of hyperbolic volumes is well-ordered. It follows that the link volume is realized by some cover in $\{M \xrightarrow{p_n} (S^3, L_n)\}$.

The link volume is the volume of a link exterior: That is, for any M , there exists $\tilde{L} \subset M$ so that $\text{LinkVol}(M) = \text{Vol}(M \setminus \tilde{L})$. This follows easily from the previous point. Let $M \xrightarrow{p} (S^3, L)$ be a cover realizing the link volume. Let \tilde{L} be the preimage of L . Then the cover $M \rightarrow S^3$ induces a cover $M \setminus \tilde{L} \rightarrow S^3 \setminus L$. Since the cover $M \setminus \tilde{L} \rightarrow S^3 \setminus L$ is not branched, we can lift the hyperbolic structure on $S^3 \setminus L$ to $M \setminus \tilde{L}$. We obtain a complete finite volume hyperbolic structure on $M \setminus \tilde{L}$ of volume $p \text{Vol}(S^3 \setminus L) = \text{LinkVol}(M)$.

The link volume is bigger than the volume: If the manifold M is hyperbolic then $\text{LinkVol}(M) > \text{Vol}(M)$: this follows immediately from the previous point and the fact the volume always goes down under Dehn filling [19].

The spectrum of link volumes is well ordered: It follows from the second point above that the spectrum of link volumes is a subset of the spectrum of hyperbolic volumes. Since the spectrum of hyperbolic volumes is well ordered, so are all of its subsets.

The spectrum of link volumes is “small”: The reader can easily make sense of the claim that the spectrum of link volumes is a very small subset of the spectrum of hyperbolic volume. In fact, since the link volume is attained, every link volume is an integral product of a volume of a hyperbolic link in S^3 . However, the spectrum of link volumes is not too small: there are infinitely many manifolds M with $\text{LinkVol}(M) < 7.22 \dots$. Moreover, in [17] Jair Remigio-Juarez and the first named author showed that there are infinitely many manifold of *the same* link volume, just under $7.22 \dots$. This is in sharp contrast to the hyperbolic volume function, which is finite-to-one.

For the remainder of this paper we will often use these facts without reference.

Basic questions about the link volume include:

- (1) Calculate $\text{LinkVol}(M)$. It is not clear whether or not there exists an algorithm to calculate the link volume of a given manifold M . This would involve some questions about the set of links in S^3 that give rise to M and appears to be quite hard.
- (2) The following question was proposed by Hitoshi Murakami: If $N \xrightarrow{q} M$ is an unbranched cover then $\text{LinkVol}(N) \leq q \text{LinkVol}(M)$. How good is this bound? Even for $q = 2$, the answer is not clear.
- (3) Since the link volume is attained, for every manifold M there is a positive integer d which is the smallest integer so that there exists a cover $M \xrightarrow{d} (S^3, L)$ realizing $\text{LinkVol}(M)$. What is d and how does it reflect the topology of M ? Can d be arbitrarily large? Does every positive integer d occur in this way, for some M ?
- (4) Characterize the set

$$\{\tilde{L} \subset M \mid \exists M \rightarrow S^3, \text{ branched over } L, \text{ and } \tilde{L} \text{ is the preimage of } L\}.$$

The link volume is, of course, the minimal volume of the exteriors of the hyperbolic links in this set, and in this paper we concentrate on it. It is easy to see that there is no upper bound to the volumes of exteriors of hyperbolic links in this set.

- (5) Do there exist hyperbolic manifolds M_1, M_2 with $\text{Vol}(M_1) = \text{Vol}(M_2)$ and $\text{LinkVol}(M_1) \neq \text{LinkVol}(M_2)$?
- (6) Similarly, do there exist hyperbolic manifolds M_1, M_2 with

$$\text{LinkVol}(M_1) = \text{LinkVol}(M_2) \quad \text{and} \quad \text{Vol}(M_1) \neq \text{Vol}(M_2)?$$

We note that the examples of manifolds with the same link volume mentioned above are all Siefert fibered spaces.

5 Proof of Theorem 1.1

Fix $V > 0$. Fix a Margulis constant $\mu > 0$ for \mathbb{H}^3 and $d > 0$. (We remark that the constant Λ that we obtain in this proof depends on these choices.)

Let M be a manifold of $\text{LinkVol}(M) < V$. Let $M \xrightarrow{p} (S^3, L)$ be a cover for which $p \text{Vol}(S^3 \setminus L) < V$. Denote the d neighborhood of the μ -thick part of $S^3 \setminus L$ by E_L . By construction, E_L is obtained from $S^3 \setminus L$ by drilling out certain geodesics; by Kojima [10, Proposition 4], E_L is hyperbolic.

Let X_ϕ denote the preimage of E_L in M . Then the cover $\phi: M \rightarrow S^3$ induces an unbranched cover $\phi: X_\phi \rightarrow E_L$. By lifting the hyperbolic structure from E_L to X_ϕ , we see that X_ϕ is a finite volume hyperbolic manifold.

By construction, the following diagram commutes (where vertical arrows represent the covering projections and horizontal arrows represent inclusions induced by Dehn fillings):

$$\begin{array}{ccc}
 X_\phi & \longrightarrow & M \\
 \downarrow \phi & & \downarrow \phi \\
 E_L & \longrightarrow & (S^3, L)
 \end{array}$$

By Jørgensen and Thurston (see, for example, [9]), there exists a constant Λ (depending on μ and d), so that for any complete, finite volume hyperbolic manifold N , the d -neighborhood of the μ -thick part of N can be triangulated using no more than $\Lambda \text{Vol}(N)$ tetrahedra. Applying this to $N = S^3 \setminus L$, since the d -neighborhood of the μ -thick part of N is E_L , we see that E_L can be triangulated using at most $\Lambda \text{Vol}(S^3 \setminus L) < \Lambda V/p$ tetrahedra.

Since there are only finite many manifolds that can be triangulated using at most $\Lambda V/p$ tetrahedra, there are only finitely many possibilities for E_L .

Lifting the triangulation from E_L to X_ϕ , we see that X_ϕ can be triangulated with at most $\Lambda \text{LinkVol}(M) < \Lambda V$ tetrahedra so that $\phi: X_\phi \rightarrow E_L$ is simplicial with respect to this triangulation. This shows that there are only finitely many possibilities for X_ϕ and ϕ . We denote them by $\{\phi_i: X_i \rightarrow E_i\}_{i=1}^{nV}$.

This completes the proof of Theorem 1.1. □

6 The link volume and Dehn filling

In this section we prove Theorem 1.6. The proof is constructive and requires two elements, the first is Hilden’s construction of simple 3–fold covers of S^3 , and the second is the results of Thurston and Menasco that show that an alternating link that “looks like” a hyperbolic link is in fact hyperbolic. For the latter, see Section 2.4. We now explain the former.

In [6] Hilden showed that any 3–manifold is the simple 3–fold cover of S^3 . The crux of his proof is the construction, for any g , of a 3–fold branched cover $p: V_g \rightarrow B$, where V_g is the genus g handlebody and B is the 3–ball. He then proves that any map $f: \partial V_g \rightarrow \partial V_g$ can be isotoped so as to commute with p . Thus f induces a map $\bar{f}: \partial B \rightarrow \partial B$ so that the following diagram commutes (here the vertical arrows denote Hilden’s covering projection):

$$\begin{array}{ccc}
 V_g & \xrightarrow{f} & V_g \\
 \downarrow & & \downarrow \\
 B & \xrightarrow{\bar{f}} & B
 \end{array}$$

Starting with a closed, orientable, connected 3–manifold M , Hilden uses a Heegaard splitting of $M = V_g \cup_f V_g$; the construction above gives a map to $B \cup_{\bar{f}} B \cong S^3$. This is, in a nutshell, Hilden’s construction of M as a cover of S^3 .

Our goal is using a similar construction to get a map from X . Since X has boundary it cannot branch cover S^3 , and we must modify Hilden’s construction. To that end, we first describe the cover $p: V_g \rightarrow B$ in detail. Let S_{3g+2} be the $3g + 2$ times punctured S^2 , viewed as a $3g$ –times punctured annulus. Then $S_{3g+2} \times [-1, 1]$ admits a symmetry of order two (rotation by π about the y –axis) given by $(x, y, t) \mapsto (-x, y, -t)$, where S_{3g+2} is embedded symmetrically in the xy –plane as shown in Figure 2.

$S_{3g+2} \times [-1, 1]$ also admits a symmetry of order 3 by rotating S_{3g+2} about the origin of the xy –plane and fixing the $[-1, 1]$ factor. These two symmetries generate an action of the dihedral group of order 6 on $S_{3g+2} \times [-1, 1]$. It is easy to see that the quotient is a ball. On the other hand, the quotient of $S_{3g+2} \times [-1, 1]$ by the order two symmetry is V_g . This induces the map $f: V_g \rightarrow B$; note that this is a cover, branched along a trivial tangle with $g + 2$ arcs (thus the branch set of the map $M \rightarrow S^3$ described above is a $g + 2$ bridge link, and the braiding is determined by \bar{f}). This is Hilden’s construction, see Figure 3, where the branch set of $V_g \rightarrow B$ is indicated by dashed lines (in B).




Figure 2: S_{3g+2} embedded in the xy -plane




Figure 3: Hilden's covers

A Heegaard splitting for the manifold with boundary X is a decomposition of X into two compression bodies; we assume the reader is familiar with the basic definitions (see, for example, [4]). We use the notation $V_{g,n}$ for a compression body with $\partial_+ V_{g,n}$ a genus g surface and $\partial_- V_{g,n}$ a collection of n tori (so $0 \leq n \leq g$). Since ∂X consists of n tori, any Heegaard splitting of X consists of two compression bodies of the form V_{g,n_1} and V_{g,n_2} , for some g, n_1, n_2 with $n_1 + n_2 = n$. We use the notation V_{g,n_i}^* for the manifold obtained by removing the interior of n_i disjoint balls from the interior of V_{g,n_i} . We use the notation X^* for the manifold obtained by removing the interior of n disjoint balls from the interior of X . Finally, we use the notation $B_{n_i}^*$ for the manifold obtained by removing the interior of n_i disjoint balls from the interior of B .

Since compression bodies do not admit simple 3-fold branched covers of the type we need, we work with V_{g,n_i}^* ; see Figure 4.




Figure 4: Hilden’s covers modified

Figure 4 is very similar to Figure 3, but has a few “decoration” added in blue. The circles added to $S_{3g+2} \times [-1, 1]$ are embedded in $S_{3g+2} \times \{0\}$. There are exactly $3n_i$ such circles. Clearly, they are invariant under the dihedral group action, and their images in V_g and B are shown. By removing an appropriate neighborhood of these circles and their images, we get a simple 3–fold cover from V_{g,n_i}^* to $B_{n_i}^*$.

Applying Hilden’s Theorem to the gluing map $f: \partial_+ V_{g,n_1} \rightarrow \partial_+ V_{g,n_2}$, we obtained a map $\bar{f}: \partial B_{n_1} \rightarrow \partial B_{n_2}$. Clearly, downstairs we see the manifold obtained by removing $n_1 + n_2 = n$ open balls from S^3 ; we denote it by $S^{3,*}$.

Note that the branch set is a tangle (that is, a 1–manifold properly embedded in $S^{3,*}$) that intersects every sphere boundary component in exactly 4 points; we denote this branch set by T . Moreover, the preimage of each component of $\partial S^{3,*}$ consists of exactly two components: a torus that double covers it, and a sphere that projects to it homeomorphically. The map from the torus in ∂X^* to the sphere in $S^{3,*}$ is the quotient under the well-known hyperelliptic involution.

Hilden’s construction, as adopted to our scenario, is the key to everything we do below. We sum up its main properties here:

Proposition 6.1 *Let X be a compact connected orientable manifold with ∂X consisting of n tori. Let X^* be the manifold obtained by removing n open balls from the interior of X . Let $S^{3,*}$ be the manifold obtained by removing n open balls from S^3 . Then there exists a simple 3–fold cover $p: X^* \rightarrow S^{3,*}$. The branch set is a compact 1–manifold, denoted by T , that intersects every boundary component of $S^{3,*}$ in exactly four points.*

The preimage of each component S of $\partial S^{3,}$ consists of one torus component of ∂X that double covers S via a hyperelliptic involution, and one sphere component of $\partial X^* \setminus \partial X$ that maps to S homeomorphically.*

Recall that in Theorem 1.6, X came equipped with a choice of meridian and longitude on each boundary component. $S^{3,*}$ is naturally a subset of S^3 . We isotope $S^{3,*}$ in S^3 so that, after projecting it into the plane, the following conditions hold:

- (1) The balls removed from S^3 are denoted by \bar{B}_i ($i = 1, \dots, n$). The projection of each \bar{B}_i is a round disk; these disks are denoted by B_i ; see Figure 5.




Figure 5: T in a neighborhood of B_i

- (2) T intersects each B_i in exactly four points. Each of these point is an endpoint of a strand of T . The four point are the intersection of the lines of slopes ± 1 through the center of the disk with its boundary, and are labeled (in cyclic order) NE, SE, SW and NW.
- (3) We twist the boundary components of $S^{3,*}$ so that, in addition, the meridian and longitude of the corresponding boundary component of ∂X map to a horizontal and vertical circles, respectively; these curves (slightly rounded) are labeled μ and λ in Figure 5.

Let $T_i \subset \partial X$ be the torus that projects to $\partial \bar{B}_i$. Recall that by *filling* T_i we mean attaching a solid torus V_i to T_i . The choice of meridian and longitude on T_i induces a choice of meridian and longitude on $\partial \bar{B}_i$, viewed as a four times punctured sphere. It is well known that after filling T_i with slope p_i/q_i , V_i double covers a ball attached to B_i and the branch set in this ball forms a p_i/q_i rational tangle. We denote this rational tangle by R_i . In the rest of this paragraph we briefly explain this; for a detailed explanation of rational tangles and their double covers see, for example, [18]. Consider the foliation of V_i by concentric tori with one singular leaf (the core circle). We construct the following explicit model of the hyperelliptic involution: let T_i be the image of \mathbb{R}^2 under the action of \mathbb{Z}^2 given by $(x, y) \mapsto (x + n, y + m)$. Then the

hyperelliptic involution is induced by rotation by π about $(0, 0)$. The four fixed points on T_i are the images of $(0, 0)$, $(\frac{1}{2}, 0)$ (rotate and translate by $(x + 1, y)$), $(0, \frac{1}{2})$ (rotate and translate by $(x, y + 1)$) and $(\frac{1}{2}, \frac{1}{2})$ (rotate and translate by $(x + 1, y + 1)$). Given any slope p_i/q_i (with p_i and q_i relatively prime), it is clear that the foliation of \mathbb{R}^2 by straight lines of slope p_i/q_i is invariant under the rotation by π about $(0, 0)$. The line through $(0, 0)$ goes through $(p_i/2, q_i/2)$, which is the image of one of the other three fixed points as not both p_i and q_i are even. Similarly for the lines through $(\frac{1}{2}, 0)$, $(0, \frac{1}{2})$, $(\frac{1}{2}, \frac{1}{2})$; these lines project to two circles on T_i with exactly two fixed points on each circle. Thus the foliation of \mathbb{R}^2 by lines of slope p_i/q_i induces a foliation of ∂V_i by circles representing the slope p_i/q_i that is invariant under the hyperelliptic involution, in which exactly two leaves that have two fixed points each. We extend this foliation into V_i minus the core by considering the foliation of each concentric torus by circles of slope p_i/q_i . The quotient of each of the concentric tori under the hyperelliptic involution is a sphere with four branch points; this sphere is foliated by circles that represent slope p_i/q_i (as curves on the four times punctured sphere), with two fibers are arcs. The image of the core of the solid torus is an arc. Thus the quotient of V_i is foliated by spheres with one singular leaf that is an arc, and hence it is a ball. By construction the branch set is a two strand tangle; it is not hard to see that this is a rational tangle of slope p_i/q_i .

Definition 6.2 We assume the rational tangles R_i have been isotoped to be alternating (it is well known that this can be achieved). Two rational tangles R_i and $R_{i'}$ are considered *equivalent* if the following two conditions hold:

- (1) The first crossing of the strand of R_i coming in from the NE is the same as the first crossing of the strand of $R_{i'}$ coming in from the NE, that is, either in both tangles the first crossing is an over crossing, or in both tangles it is an undercrossing. Since R_i and $R_{i'}$ are both alternating, this implies that for each of the remaining three corners the analogous condition is satisfied.
- (2) The strands of R_i and $R_{i'}$ that start at NE end at the same point (SE, SW, or NW).

Note that the crossing information is ill-defined for the two tangles $1/0$ and $0/1$, as they have no crossings. We arbitrarily choose an equivalent class for each of these tangle so that the second condition is fulfilled. We obtain 6^n possible equivalence classes (recall that $n = |\partial X|$).

Given slopes on T_1, \dots, T_n , we get rational tangles R_1, \dots, R_n , as described above. In each \bar{B}_i we place a rational tangle, denoted by \widehat{R}_i , so that $\widehat{R}_i \in \{\pm 1, \pm 2, \pm \frac{1}{2}\}$,




Figure 6: \widehat{R}_i , the representatives for equivalence classes

representing the same equivalence class as R_i . We assume that their projections into B_i are as in Figure 6.

We thus obtain a link, denoted by \widehat{T} , and a diagram for \widehat{T} , denoted by \widehat{D} . Since \widehat{T} and \widehat{D} only depend on the equivalence classes of the slopes, when considering all possible slopes, we obtain finitely many links and diagrams (specifically, 6^n).

In order to obtain hyperbolic branch set, we will, eventually, apply Mensaco [13] as explained in Section 2.4. To that end we will need to make the branch set alternating. As we shall see below, we do this using a $1 \rightarrow -2$ and $-2 \rightarrow 1$ Montesinos moves; these moves can be used to make the link alternating in a way that is very similar to crossing changes. Below we will show that we can apply Montesinos moves to T , however, we may not apply these moves to the rational tangles inside B_i . This causes the following trouble: let $\alpha \subset T$ be an interval connecting two punctures, say $\partial \bar{B}_i$ and $\partial \bar{B}_{i'}$ (possibly, $i = i'$). Assume that the last crossing of R_i before α is an overcrossing, and that the number of crossings along α is even. Then if we make T alternate, the last crossing along α will be an overcrossing. This means that the first crossing of $R_{i'}$ after α must be an undercrossing. This may or may not be the case, and we have no control over it.

In order to encode this, we consider the following graph Γ : Γ has n vertices, and they correspond to B_1, \dots, B_n . The edges of Γ correspond to intervals of T that connect B_i to $B_{i'}$ (again, i and i' may not be distinct). Inspired by the discussion above, we assign signs to the edges of Γ as follows (in essence, good edges get a $+$ and bad edges get a $-$):

- (1) Let $I \subset T$ be an interval connecting B_i to $B_{i'}$ (possibly $i = i'$) so that the last crossing before I and the first crossing after I are the same (that is, both are overcrossings or both are undercrossings) and the number of crossings along I is odd. Then the corresponding edge get the sign $+$.
- (2) Let $I \subset T$ be an interval connecting B_i to $B_{i'}$ (possibly $i = i'$) so that the last crossing before I and the first crossing after I are the opposite (that is, one is an overcrossing and the other an undercrossing) and the number of crossings along I is even. Then the corresponding edge get the sign $+$.
- (3) All other edges get the sign $-$.

If Γ is connected, we pick a spanning tree $\widehat{\Gamma}$ for Γ . That is, $\widehat{\Gamma}$ is a tree obtained from Γ by removing edges, so that every vertex of Γ is adjacent to some edge of $\widehat{\Gamma}$. In general, we take $\widehat{\Gamma}$ to be a *maximal forest* in Γ . A forest is a collection of trees, that is, a (possibly disconnected) graph without cycles. A *maximal forest* in Γ is a graph obtained from Γ by removing a minimal (with respect to inclusion) set of edges so that a forest is obtained; equivalently, it is the union of spanning trees for the connected components of Γ . Clearly a maximal forest $\widehat{\Gamma}$ has the following two properties: first, $\widehat{\Gamma}$ contains no cycles. Second, any edge from Γ that we add to $\widehat{\Gamma}$ closes a cycle.

Lemma 6.3 *There is a sign assignment to the vertices of $\widehat{\Gamma}$ so that an edge of $\widehat{\Gamma}$ has a plus sign if and only if the vertices it connects have the same sign.*

Proof of Lemma 6.3 We induct on the number of edges in $\widehat{\Gamma}$. If there are no edges there is nothing to prove.

Assume there are edges. In that case at least one component of $\widehat{\Gamma}$ is a tree with more than one vertex. Such a tree must have a leaf, say v . Remove v and e , the unique edge of $\widehat{\Gamma}$ connected to v . By induction, there is a sign assignment for the remaining vertices fulfilling the conditions of the lemma. We now add v and e . Clearly, we can give v a sign so that the condition of the lemma holds for e . The lemma follows. \square




Figure 7: Modifying \widehat{D} near B_i (negative sign)

We now isotope \widehat{T} and accordingly modify \widehat{D} as shown in Figure 7 at each puncture that corresponds to a vertex with a minus sign. Since this changes the number of crossings on some of strands of Γ , we recalculate the signs on the corresponding edges. Note that the isotopy above adds one or three crossing to every strand of T that corresponds to an edge of $\widehat{\Gamma}$ with sign $-$, and zero, two, four or six crossings to every strand of T that corresponds to an edge with sign $+$. We easily conclude that every edge of $\widehat{\Gamma}$ has sign $+$. Moreover:

Lemma 6.4 Every edge of Γ has sign $+$.

Proof The proof is very similar to the proof that every link projection can be made into an alternating link projection via crossing change and is left for the reader, with the following hint: Suppose there exists an edge, say e , whose sign is $-$. Since we used a maximal forest, there is a cycle in Γ (say e_1, \dots, e_k) so that $e_1 = e$ and e_i belongs to the maximal forest for $i > 1$; in particular, exactly one edge of the cycle has sign $-$. Use this cycle to produce a (not necessarily simple) closed curve in S^2 that intersects the link \widehat{T} transversely an odd number of times. This is absurd, in light of the Jordan curve theorem. \square

Next we prove that X^* can be obtained as a 3-fold cover of $S^{3,*}$ with a particularly nice branch set. We begin with T , \widehat{T} and \widehat{D} described above; their properties are summed up in condition (1) of Lemma 6.5 below. Parts of this argument are similar to Blair [2]. Recall Definitions 2.5 for standard terms in knot theory.

Lemma 6.5 There exists a link \widehat{T} in S^3 , with projection into S^2 denoted by \widehat{D} , so that X^* is a simple 3-fold cover of $S^3 \setminus \bigcup_{i=1}^n \text{int}(\overline{B}_i)$ branched along the tangle $T = \widehat{T} \cap (S^3 \setminus \bigcup_i \text{int}(B_i))$ and the following conditions hold:

- (1) $\widehat{T} \cap \overline{B}_i = \widehat{R}_i$ (recall that \widehat{R}_i projects into B_i as shown in Figure 6), the projection of $\text{int}(T \setminus \widehat{T})$ is disjoint from $\bigcup_{i=1}^n B_i$, and the meridian and longitude of $T_i \subset \partial X$ project to horizontal and vertical circles about B_i (respectively, recall Figure 5).
- (2) \widehat{D} is not a split diagram.
- (3) Every simple closed curve in $S^2 \setminus \bigcup B_i$ that intersects \widehat{D} transversely in two simple points bounds a disk that intersects \widehat{D} in a single arc with no crossing.
- (4) Let $\alpha \subset S^2$ be an arc with one endpoint on $B_{i'}$ and the other on $B_{i''}$ (for $i', i'' = 1, \dots, n$, possibly $i' = i''$) so that $\text{int}(\alpha) \cap (\bigcup_i B_i) = \emptyset$. Then one of the following conditions holds:
 - (a) $i' = i''$, and α cobounds a disk with ∂B_i with no crossings.
 - (b) $|\text{int}(\alpha) \cap \widehat{D}| > 2$.
- (5) In the three-coloring of $\widehat{D} \cap (S^2 \setminus \bigcup_{i=1}^n B_i)$ induced by the cover $X^* \rightarrow S^{3,*}$, every crossing is three colored.
- (6) \widehat{D} is alternating.
- (7) \widehat{T} is a knot.
- (8) \widehat{D} has at least two twist regions, and one of them is not adjacent to B_i for any i .

Remark 6.6 To obtain conditions (1)–(5) we modify \widehat{T} via isotopy; except for the move shown in Figure 9, the projection of the support of this isotopy is disjoint from $\bigcup_{i=1}^n B_i$. Note that in the move shown in Figure 9 each edge gets an even number of crossings added. Hence the signs of the edges of Γ do not change, and Lemma 6.4 still holds after we obtain conditions (1)–(5). (We use this lemma to obtain condition (6), and never need it again after that.)

Proof (1) This already holds. We note that none of the moves applied in the proof of this lemma changes this. We will not refer to condition (1) explicitly.

(2) \widehat{D} is diagrammatically split if and only if it is disconnected. Suppose \widehat{D} is disconnected, and let K_j and $K_{j'}$ be components of \widehat{T} that project to distinct components of \widehat{D} . Let $\alpha \subset S^2 \setminus \bigcup_{i=1}^n B_i$ be an embedded arc with one endpoint on K_j and the other on $K_{j'}$ (note that $K_j, K_{j'} \not\subset \bigcup_{i=1}^n B_i$, hence α exists; α may intersect \widehat{D} in its interior). We perform an isotopy along α , as shown in Figure 8.




Figure 8: Isotopy along α

After that K_j crosses $K_{j'}$ outside $\bigcup_{i=1}^n B_i$; clearly, this reduces the number of components of \widehat{D} . Repeating this process if necessary, condition (2) is obtained.

(3) For each B_i , let $N(B_i)$ be a collar neighborhood of B_i so that $\widehat{D} \cap N(B_i)$ consists of the tangle in B_i and four short segments as in the left hand side of Figure 9. We assume further that for $i \neq j$, $N(B_i) \cap N(B_j) = \emptyset$. Inside each $N(B_i)$ perform the isotopy shown in Figure 9.

Next we count the number of simple closed curves in $S^2 \setminus \bigcup_i B_i$ that intersect \widehat{D} in two points and do not bound a disk Δ with $\widehat{D} \cap \Delta$ is a single arc with no crossings. These curves are counted up to *diagrammatic isotopy*, that is, an isotopy via curves that are transverse to \widehat{D} at all time and in particular are disjoint from the crossings.

Let C_1, \dots, C_k be the closures of the components of $S^2 \setminus (\widehat{D} \cup (\bigcup_i B_i))$. Let $\gamma, \gamma' \subset S^2 \setminus \bigcup_i B_i$ be two simple closed curves that intersects \widehat{D} transversely in two simple points. Then \widehat{D} cuts γ into 2 arcs, say one in the region C_j and one in $C_{j'}$. Note that if $j = j'$, then C_j is adjacent to itself, and in particular there is a simple closed curve in S^2 that intersects \widehat{D} transversely in one point, which is absurd. Condition (2)




Figure 9: Isolating B_i

(connectivity of \widehat{D}) is equivalent to all regions being disks, and hence implies that γ and γ' are diagrammatically isotopic if and only if both curves traverse the same regions C_j and $C_{j'}$, and $\gamma \cap \partial C_j$ is contained in the same segments of $C_j \cap C_{j'}$ as $\gamma' \cap \partial C_{j'}$. (See Figure 10; here a *segment* means an interval $I \subset S^2 \setminus \bigcup_i \text{int}(B_i)$, so that $I \subset C_j \cap C_{j'}$, ∂I are crossings or lie on ∂B_i for some i , and I contains no crossings in its interior.) For any pair of regions C_j and $C_{j'}$, let $n_{j,j'}$ be the number of segments in $C_j \cap C_{j'}$ (for example, in Figure 10, $n_{j,j'} = 4$). Then we see that the number of simple closed curves that intersect \widehat{D} in two simple points, traverse C_j and $C_{j'}$, and do not bound a disk containing a single arcs of \widehat{D} (counted up to diagrammatic isotopy) is $\binom{n_{j,j'}}{2}$, where $\binom{0}{2}$ and $\binom{1}{2}$ are naturally understood to be 0. Hence the total number of such curves (counted up to diagrammatic isotopy) is

$$(2) \quad \sum_{1 \leq j < j' \leq k} \binom{n_{j,j'}}{2}.$$




Figure 10: Segments

Now assume that condition (3) does not hold; then there exist regions C_j and $C_{j'}$ with $n_{j,j'} \geq 2$. Let I be an interval of $C_j \cap C_{j'}$. Since we isolated B_i (for all i) as shown in Figure 9, the endpoints of I cannot lie on ∂B_i and must therefore both be crossings. The move shown in Figure 11 reduces $n_{j,j'}$ by one. This move introduces several new




Figure 11: Isotopy to reduce $C_{i,i'}$

regions, and those are shaded in Figure 11. Inspecting Figure 11, we see that for any pair of regions $C_j, C_{j'}$ that existed prior to the move, $n_{j,j'}$ does not increase, and for any pair of regions $C_j, C_{j'}$ with at least one new region, $n_{j,j'}$ is 1 or 0. Hence the sum in Equation (2) is reduced, and repeated application of this move yields a diagram \widehat{D} for a link \widehat{T} for which condition (3) holds; by construction, condition (2) still holds.

(4) This holds thanks to the isotopy performed in the previous step and shown in Figure 9.

(5) Since \widehat{D} is the branch set of the simple 3-fold cover $X^* \rightarrow S^{3,*}$ it inherits a 3-coloring as explained in Section 2.1, where the colors are transpositions in S_3 . Since X^* is connected, at least two colors appear in the coloring of T (recall Lemma 2.4; that lemma was stated for covers of S^3 but it is easy to see that it holds for covers of $S^{3,*}$ as well).

Assume there exists a one colored crossing of \widehat{D} outside $\bigcup_{i=1}^n B_i$, say c , and let p be a point on a strand of \widehat{D} that is of a different color than c and so that $p \notin \bigcup_{i=1}^n B_i$. Let α be an arc connecting p and c so that $\alpha \cap (\bigcup_{i=1}^n B_i) = \emptyset$. If $\text{int}(\alpha)$ intersects a strand of \widehat{D} whose color is different than the color of c , we cut α short at that intersection. Thus we may assume that any point of $\text{int}(\alpha) \cap \widehat{D}$ has the same color as c . We apply the move shown in Figure 12 (often used by Hilden, Montesinos and others).

This move reduces the number of one colored crossings outside $\bigcup_i B_i$, and hence repeating this move gives condition (5).

We now verify that conditions (2)–(4) still hold. Inspecting Figure 12, we see that condition (2), which is equivalent to connectivity of \widehat{D} , clearly holds. A simple closed curves that intersects \widehat{D} twice after this moves, intersects it at most twice before the move. By considering these curves and Figure 12 we conclude that condition (3) holds as well (in checking this, note that $\text{int} \alpha \cap \widehat{D}$ maybe empty; to rule out one case, you need to use the coloring: a red arc cannot be connected to a blue arc without a crossing). For each i , the preimage of $\partial \bar{B}_i$ is disconnected; hence the four segments of \widehat{D} on




Figure 12: Making the crossings 3-colored

the left side of Figure 9 are all the same color. Since \widehat{D} is connected and has more than one color, it must have a three-colored crossing, which cannot be contained in $N(B_i)$ for any i . We can take the point p in the construction above to be a point near the three colored crossing, and in particular, we may assume that $p \notin N(B_i)$ for any i . Therefore this move effects $\widehat{D} \cap N(B_i)$ by adding arcs that traverse $N(B_i)$ without intersecting B_i itself, but not changing any of the existing diagram in the right hand side of Figure 9. Therefore condition (4) holds.

(6) Recall that the tangles \widehat{R}_i are alternating ($i = 1, \dots, n$). It is well known that any link projection can be made into an alternating projection by reversing some of its crossings. We mark the crossings of \widehat{D} by \pm , marking a crossing $+$ if we do not need to reverse it and $-$ otherwise. By reversing all the signs if necessary, we may assume that the signs in B_1 are $+$. Since the signs of all the edges of Γ are $+$ (Lemma 6.4 and Remark 6.6), the signs in every B_i are all $+$. Thus all the crossings that are marked $-$ are outside $\bigcup_{i=1}^n B_i$, and hence three colored. We change each of this crossing using the Montesinos move $+1 \mapsto -2$ or $-1 \mapsto +2$, as in the top row of Figure 1, noting that this does not change the double cover. It is clear that now \widehat{D} is an alternating diagram fulfilling conditions (1)–(6).

(7) Assume \widehat{T} is a link. If there is a crossing outside $\bigcup_i B_i$ that corresponds to two distinct components of \widehat{T} , we perform a $+1 \mapsto +4$ or $-1 \mapsto -4$ Montesinos move; this reduces the number of components of T . Assume there is no such crossing, and let α be an arc connecting strands (say s_1 and s_2) that correspond to two distinct components of \widehat{T} . Since no B_i contains a closed component, we may assume $\alpha \cap (\bigcup_i B_i) = \emptyset$; furthermore, by truncating α if necessary, we may assume that $\text{int}(\alpha) \cap \widehat{D} = \emptyset$. By condition (4) at least one endpoint of s_2 is a crossing outside $\bigcup_i B_i$, say c . If s_1 and s_2 have the same color, we replace α with an arc that connects s_1 with a strand adjacent to s_2 at c . By condition (5) c is three-colored, and by assumption, both its strands

correspond to the same component of \widehat{T} . Thus we obtain an arc that connects distinct components and has endpoints of different colors. Finally, we assume without loss of generality that the crossing information at s_1 is as shown in Figure 13. Since \widehat{D} is connected and alternating, considering the face of the projection sphere S^2 containing α , we conclude that the crossing information on s_2 is as shown in that figure. We change \widehat{D} using a $0 \mapsto \pm 3$ Montasinos move (as shown in the bottom of Figure 1), obtaining a diagram fulfilling conditions (1)–(6) that corresponds to a link with fewer components; see Figure 13.




Figure 13: Making the branch set into a knot

Iterating this process, we obtain a knot.

(8) If this is not already satisfied, we take an arc $\alpha \subset \bigcup_{i=1}^n N(B_i)$ connecting two strands of distinct colors and perform the move described in Figure 14, consisting of two $0 \rightarrow \pm 6$ Montesinos moves. Similar to condition (7) above, since the diagram is alternating and connected, an appropriate choice of sign will result in an alternating diagram. It is straightforward to see that conditions (1)–(7) are satisfied.

This completes the proof of Lemma 6.5 □

We are now ready to complete the proof of Theorem 1.6. Fix X as in the statement of the theorem and pick a slope on each components of ∂X , say p_i/q_i on the torus $T_i \subset \partial X$; note that we are using the meridian-longitudes to express the slopes as rational numbers (possibly, $1/0$). Construct a 3-fold simple cover $X^* \rightarrow S^{3,*}$ as in Lemma 6.5 that corresponds to the appropriate equivalence classes of the slopes (recall Definition 6.2). For convenience we work with \widehat{D} , the diagram of \widehat{T} , that fulfills the conditions of Lemma 6.5.

We now change the diagram \widehat{D} by replacing the rational tangle \widehat{R}_i in B_i (that represents the equivalence class of p_i/q_i) with the rational tangle R_i (that realizes the slope




Figure 14: Adding twist regions

p_i/q_i), $i = 1, \dots, n$. By construction the four strands of \widehat{D} that connect to B_i are single colored, and we color the R_i by the same color. Thus we obtain a diagram of a three colored link denoted by K .

We claim that K has the following properties:

- (1) K is a knot.
- (2) K admits an alternating projection.
- (3) This projection is nonsplit.
- (4) This projection is strongly prime.
- (5) K is not a torus knot.

We prove each claim in order:

(1) Since the tangles \widehat{R}_i and R_i are equivalent they connect the same points on ∂B_i (Definition 6.2). By Lemma 6.5(7), \widehat{T} is a knot. Hence K , which is obtained from \widehat{T} by replacing \widehat{R}_i by R_i , is a knot as well.

(2) By Lemma 6.5(6), \widehat{D} is alternating. By the definition of the equivalence classes of rational tangles, K (which is obtained by replacing \widehat{R}_i by R_i) admits an alternating projection.

(3) Let $\gamma \subset S^2$ be a simple closed curve disjoint from the diagram for K . If γ is diagrammatically isotopic (that is, an isotopy through curves that are transverse to the diagram at all times) to a curve that is disjoint from $\bigcup_i B_i$ then by Lemma 6.5(2) γ bounds a disk disjoint from \widehat{D} ; this disk is also disjoint from the diagram of K . If γ is diagrammatically isotopic into B_i , then γ bounds a disk disjoint from the diagram for

K since rational tangles are prime. Finally, if γ is not isotopic into or out of $\bigcup_i B_i$, we violate condition (4b) of Lemma 6.5. Hence the diagram for K is nonsplit.

(4) This is very similar to (3) and is left to the reader.

(5) By condition (8) the diagram \widehat{D} has more than one twist region, and at least one is not adjacent to B_i for any i . The diagram obtained for K is alternating, and hence no crossing cancellation can occur. Note that any alternating torus knot is a $2, n$ -torus knot, and every alternating diagram for it has only one twist region. We conclude that K is not a torus knot.

By Menasco and Thurston (see Corollary 2.8), conditions (2)–(5) imply that K is hyperbolic.

Next we note that the 3-coloring of K defines a 3-fold cover of S^3 ; by construction, the cover of $S^{3,*}$ is X^* . The cover of each rational tangle is disconnected and consists of a solid torus attached to $T_i \subset \partial X$ with slope p_i/q_i , and a ball attached to a component of $\partial X^* \setminus \partial X$. Thus we obtain $X(p_1/q_1, \dots, p_n/q_n)$ as a simple 3-fold cover of S^3 branched over K .

We now isotope each rational tangle R_i to realize its depth, that is, realizing the twist number of each rational tangle (recall Section 2.5). The twist number of R_i is exactly $\text{depth}(p_i/q_i)$. The tangle T (which is the projection of K outside $\bigcup_i B_i$) has a fixed number of twist regions, say t . Hence the total number of twist regions is at most $t + \sum_{i=1}^n \text{depth}(p_i/q_i) = t + \text{depth}(\alpha)$ (where $\alpha = (\alpha_1, \dots, \alpha_n)$ denotes the multislope on ∂X , as in Section 1). This gives an upper bound for the twist number for K :

$$t(K) \leq t + \text{depth}(\alpha).$$

Lackenby [11] (recall Section 2.5 and in particular Theorem 2.9) showed that there exists a constant c so that:

$$\text{Vol}(S^3 \setminus K) \leq ct(K).$$

Hence we get:

$$\begin{aligned} \text{KnotVol}_{s,3}(X(\alpha_1, \dots, \alpha_n)) &\leq 3 \text{Vol}(S^3 \setminus K) \\ &\leq 3ct(K) \\ &\leq 3ct + 3c(\text{depth}(\alpha)). \end{aligned}$$

By setting $A = 3c$ and $B = 3ct$, we obtain constants fulfilling the requirements of Theorem 1.6 that are valid for any multislope $\alpha' = (\alpha'_1, \dots, \alpha'_n)$, with α'_i in the same equivalence class as α_i . As there are only finitely many (specifically, 6^n) equivalence classes, taking the maximal constants A and B for these classes completes the proof of Theorem 1.6. □

References

- [1] **J W Alexander**, *Note on Riemann spaces*, Bull. Amer. Math. Soc. 26 (1920) 370–372 MR1560318
- [2] **R C Blair**, *Alternating augmentations of links*, J. Knot Theory Ramifications 18 (2009) 67–73 MR2490017
- [3] **C Cao, G R Meyerhoff**, *The orientable cusped hyperbolic 3-manifolds of minimum volume*, Invent. Math. 146 (2001) 451–478 MR1869847
- [4] **A J Casson, C M Gordon**, *Reducing Heegaard splittings*, Topology Appl. 27 (1987) 275–283 MR918537
- [5] **M E Feighn**, *Branched covers according to J W Alexander*, Collect. Math. 37 (1986) 55–60 MR882791
- [6] **H M Hilden**, *Every closed orientable 3-manifold is a 3-fold branched covering space of S^3* , Bull. Amer. Math. Soc. 80 (1974) 1243–1244 MR0350719
- [7] **H M Hilden, M T Lozano, J M Montesinos**, *The Whitehead link, the Borromean rings and the knot 9_{46} are universal*, Collect. Math. 34 (1983) 19–28 MR747855
- [8] **H M Hilden, M T Lozano, J M Montesinos**, *On knots that are universal*, Topology 24 (1985) 499–504 MR816529
- [9] **T Kobayashi, Y Rieck**, *A linear bound on the tetrahedral number of manifolds of bounded volume (after Jørgensen and Thurston)*, from: “Topology and geometry in dimension three”, (W Li, L Bartolini, J Johnson, F Luo, R Myers, JH Rubinstein, editors), Contemp. Math. 560, Amer. Math. Soc. (2011) 27–42 MR2866921
- [10] **S Kojima**, *Isometry transformations of hyperbolic 3-manifolds*, Topology Appl. 29 (1988) 297–307 MR953960
- [11] **M Lackenby**, *The volume of hyperbolic alternating link complements*, Proc. London Math. Soc. 88 (2004) 204–224 MR2018964
- [12] **W B R Lickorish**, *An introduction to knot theory*, Graduate Texts in Mathematics 175, Springer, New York (1997) MR1472978
- [13] **W Menasco**, *Closed incompressible surfaces in alternating knot and link complements*, Topology 23 (1984) 37–44 MR721450
- [14] **E E Moise**, *Affine structures in 3-manifolds. V. The triangulation theorem and Hauptvermutung*, Ann. of Math. 56 (1952) 96–114 MR0048805
- [15] **J M Montesinos**, *A representation of closed orientable 3-manifolds as 3-fold branched coverings of S^3* , Bull. Amer. Math. Soc. 80 (1974) 845–846 MR0358784
- [16] **R Myers**, *Simple knots in compact, orientable 3-manifolds*, Trans. Amer. Math. Soc. 273 (1982) 75–91 MR664030
- [17] **J Remigio-Juárez, Y Rieck**, *The link volumes of some prism manifolds*, Algebr. Geom. Topol. 12 (2012) 1649–1665 MR2966698

- [18] **D Rolfsen**, *Knots and links*, Mathematics Lecture Series 7, Publish or Perish, Houston, TX (1990) MR1277811
- [19] **W P Thurston**, *The geometry and topology of three-manifolds* Available at <http://library.msri.org/books/gt3m/PDF>

*Department of Mathematics, University of Arkansas
Fayetteville, AR 72701, USA*

*Department of Information and Computer Sciences
Nara Women's University, Nara 630-8506, Japan*

yoav@uark.edu, yamashita@ics.nara-wu.ac.jp

Received: 7 May 2012 Revised: 21 September 2012