

FILLING KNOTS AND ESSENTIAL GEODESIC SYSTOLE

CHRISTOPHER J LEININGER AND YANDI WU

ABSTRACT. We prove that every closed, orientable 3-manifold contains hyperbolic knots with arbitrarily large essential geodesic systole. That is, for any 3-manifold M , and $R > 0$ there is a hyperbolic knot $K \subset M$ such that every closed geodesic in the knot complement which is non-null-homotopic in M has length at least R . The construction also produces filling knots, in the sense of Freedman-Krushkal [FK23]. This answers two questions of Biringer.

1. INTRODUCTION

The *systole* of a Riemannian manifold N , $\text{sys}(N)$, is the length of the shortest closed geodesic in N . In hyperbolic geometry, there are interesting relationships between the systole and the topology of the manifold. For example, if N is a closed *hyperbolic* 3-manifold, White found that $\text{sys}(M)$ is bounded by a function of the *rank*; i.e. the minimal number of generators for $\pi_1(M)$ [Whi02]; see also Kapovich-Weidmann [KW03].

Another interesting connection (relevant for the current paper) was proved by Adams and Reid in [AR00] for link complements. They showed that if M is *any* closed 3-manifold which does not admit a hyperbolic metric, then for *any* link $L \subset M$ such that $M \setminus L$ is hyperbolic, one has $\text{sys}(M \setminus L) \leq 7.35534\dots$. The hypothesis that M admits no hyperbolic metric is necessary. On the other hand, for a closed hyperbolic 3-manifold, Lakeland and the first author found that for *any* link $L \subset M$ with $M \setminus L$ hyperbolic, $\text{sys}(M \setminus L)$ is bounded above by a function of the volume of M [LL14]. In either case, for a fixed M , there is a uniform upper bound on $\text{sys}(M \setminus L)$ that does not depend on the link L itself.

We note one similarity between the proofs of [AR00] and [LL14]: the short geodesics which illustrate the bound on $\text{sys}(M \setminus L)$ are null-homotopic in the ambient manifold, M . This inspires the following definition.

Definition 1 (Essential geodesic systole). The *essential geodesic systole* of a link $L \subset M$ is the length of the shortest closed geodesic in $M \setminus L$ which is not null-homotopic in M ,

$$\text{sys}_e^*(K) = \inf\{\ell(\gamma) \mid \gamma \text{ a closed geodesic in } M \setminus L \text{ with } i \circ \gamma \text{ non-null-homotopic in } M\}.$$

In an earlier paper, the authors of this paper proved the following:

Theorem 1.1 (Theorem 1.2, [LW25]). *Let M be a closed, orientable 3-manifold with $\pi_1(M) \neq \{*\}$. Then for every $R > 0$, there exists some hyperbolic link $L \subset M$ with*

$$\text{sys}_e^*(L) > R.$$

CJL was supported by NSF grant DMS-2305286.

Remark 1.2. Definition 1 differs from the definition of *essential systole*, denoted $\text{sys}_e(L)$, given in [LW25]; there, the infimum was taken over all closed *curves* rather than closed geodesics, although Theorem 1.1 holds for both definitions (see [LW25] Remark 1.3). We use this new definition since it allows the knot itself to be non-null-homotopic in M , which will be more convenient for the current paper. Thus, when we say *essential systole* in the present paper, we implicitly mean essential geodesic systole.

The purpose of this paper is to strengthen the theorem above, as follows, answering a question of Biringer for essential geodesic systole (see [LW25, Question 4.1]).

Theorem 1.3. *Let M be an orientable 3-manifold. Let $R > 0$. Then there exists some hyperbolic knot $K \subset M$ such that*

$$\text{sys}_e^*(K) > R.$$

In [LW25], we used Theorem 1.1 to prove the existence of *filling links* in any closed, orientable 3-manifold with nontrivial fundamental group. A filling link, originally defined by Freedman and Krushkal in [FK23], is a link that “cuts through all the topology” of its ambient 3-manifold by killing the relations in *any* group presentation of $\pi_1(M)$ with respect to any minimal generating set, *viewed inside* $M \setminus L$.

More precisely, for a closed 3-manifold, M , a *1-spine* (also called a *carrier graph*) for M is a graph Γ with $\text{rank}(\pi_1(\Gamma)) = \text{rank}(\pi_1(M))$ together with a map $f: \Gamma \rightarrow M$ such that $f_*: \pi_1(\Gamma) \rightarrow \pi_1(M)$ is surjective. Given a link $L \subset M$, let $i: M \setminus L \rightarrow M$ be the inclusion. We say that $f: \Gamma \rightarrow M \setminus L$ is an *L -relative 1-spine* if $i \circ f: \Gamma \rightarrow M$ is a 1-spine. Freedman and Krushkal declare that a link $L \subset M$ is a *filling link* if for every L -relative 1-spine $f: \Gamma \rightarrow M \setminus L$, we have

$$f_*: \pi_1(\Gamma) \rightarrow \pi_1(M \setminus L)$$

is injective. They then asked in [FK23] whether filling links exist in every closed 3-manifold, and in [LW25], we proved that the answer is ‘yes’ for all closed, orientable 3-manifolds.

After [LW25], Biringer asked the authors whether filling *knots* exist (see [LW25, Question 4.1]), and we again prove that the answer is ‘yes’.

Theorem 1.4. *Any closed, orientable 3-manifold M contains a filling knot $K \subset M$.*

We recall that the filling links with large essential systole from [LW25] are constructed using sequences of triangulations of the ambient manifold M , and the number of link components grows with the number of tetrahedra in the associated triangulation. Thus, the construction of knots with large essential geodesic systole is significantly harder and requires new ideas. Moreover, (as noted above) in [LW25], we prove that the links are filling as a direct consequence of having large essential systole, applying the main result of [KW03]. The distinction between large essential systole and large essential *geodesic* systole (see Theorem 1.2), means that the existence of filling knots in the current paper also requires new ideas.

Remark 1.5. As in [LW25] (also as suggested by Biringer), for any $m \in \mathbb{N}$, we prove the existence of a *full rank- m filling* knot, which is a stronger version of Theorem 1.4. When $m = \pi_1(M)$, a full rank- m filling knot is also filling. However,

for ease of exposition, we defer this technicality to Section 3.5; see Theorem 3.9. Similarly, there is a more precise version of Theorem 1.3; see Theorem 3.8.

We now give a brief outline of the paper. In Section 2, we include some important tools used in the rest of the paper. In Section 3, given a manifold M , we construct a family of knots used to prove both main theorems. First, we build “augmented” versions of the links seen in [LW25] by replacing certain totally geodesic subsurfaces of the original link complements with hyperbolic tangle complements, which we call *connectors*. These newly glued-in connectors allow us “join” together link components by applying half-twists to their totally geodesic subsurfaces; c.f. [Ada86, Lac04, Pur10]. We then apply large Dehn fillings and surgeries to reduce the number of link components to just 1.

Finally, in Sections 4 and 5, we prove Theorems 1.3 and 1.4 respectively. The proofs use similar approaches, but slightly different tools. The existence of knots with large essential geodesic systole relies on Thurston’s Hyperbolic Dehn Surgery Theorem and other tools from hyperbolic geometry. While it is probably true that these hyperbolic knots are also filling, it is more natural and convenient to endow our knot complements with negatively-curved metrics. Thus, we instead prove the existence of filling links using locally $\text{CAT}(-1/2)$ metrics coming from negatively-curved Dehn fillings derived from the Gromov–Thurston 2π Theorem.

Acknowledgments. The authors are very grateful to Alan Reid for many useful conversations about the project, and especially for suggesting the proof of Theorem 3.3 using the `verify_hyperbolicity` function in SnapPy. They also thank Ian Biringer for initially suggesting the problem, and for helpful discussions during the second author’s visit to Boston College.

AI disclosure. No AI was used for any part of this work, other than that imposed by *Google* in its search engine.

2. PRELIMINARIES

We recall some basic tools used in our constructions in the rest of this paper. We let $\mathbb{H}^3$ denote hyperbolic 3-space with metric $d_{\mathbb{H}^3}$. In the upper-half space model, the orientation preserving isometry group is $\text{PSL}_2(\mathbb{C})$ acting by extension of Möbius transformation. A (n orientable) *hyperbolic 3-manifold* is the quotient of $\mathbb{H}^3$ by a discrete, torsion-free subgroup of $\text{PSL}_2(\mathbb{C})$. Given a closed 3-manifold M , we say that a link $L \subset M$ is *hyperbolic* if $M \setminus L$ admits a complete hyperbolic structure of finite volume. See [Thu82, Thu78] for more on hyperbolic geometry and hyperbolic manifolds.

Recall that an n -manifold X with a geodesic metric is *locally* $\text{CAT}(-1)$ if every point has a neighborhood in which geodesic triangles are at least as thin as those in the hyperbolic plane, $\mathbb{H}^2$; if the metric is complete, then by the Cartan–Hadamard Theorem [BH99, Theorem II.4.1], its universal cover $\tilde{X}$ with its pull-back metric is $\text{CAT}(-1)$ (*all* triangles satisfy the above comparison). For $\kappa < 0$, the (locally) $\text{CAT}(\kappa)$ condition is similarly defined, by comparing with triangles in $\mathbb{H}^2$ with the metric scaled by $\sqrt{|\kappa|}$. In this paper, we will make abundant use of the following consequence of Thurston’s Hyperbolization Theorem [Thu82, Thu86, Mor84, McM92], and which allows us to endow the locally $\text{CAT}(-1)$

link complements constructed in this paper with hyperbolic metrics. These locally CAT(-1) and hyperbolic link complements will have the added property that their universal covers are quasi-isometric. A proof, appealing to Hyperbolization, is sketched in [LW25].

Theorem 2.1 ([LW25] Theorem 2.2). *Suppose a 3-manifold N admits a complete, finite-volume, locally CAT(-1), piecewise hyperbolic metric which is hyperbolic in the cusps. Then N admits a complete, finite-volume hyperbolic metric. Moreover, this metric can be chosen so that the identity on the universal covers is a quasi-isometry.*

2.1. Thick-thin decomposition. The ϵ -thin (resp. thick) part of a manifold M , denoted $M_{(0,\epsilon)}$ (resp. $M_{[\epsilon,\infty)}$), is the set of points $x \in M$ such that the injectivity radius of M at x is less than (resp. greater than or equal to) ϵ . One important consequence of the Margulis Lemma is that there is a value $\epsilon_0 > 0$, known as the (3-dimensional) *Margulis constant*, such that for any $\epsilon \leq \epsilon_0$ and 3-manifold M , $M_{(0,\epsilon)}$ is either a *Margulis tube*, the tubular neighborhood of a short curve, or a cusp neighborhood; see e.g. [Thu78, Chapter 5].

Another consequence of the Margulis Lemma is the following, which is an important ingredient in our arguments. Given a discrete, torsion-free subgroup $\Gamma < \mathrm{PSL}(2, \mathbb{C})$, the ϵ -thin part of $\mathbb{H}^3$ is the preimage of the ϵ -thin part of M/Γ . For $\epsilon > \epsilon_0$, the ϵ -thin part is a disjoint union of horoballs and neighborhoods of geodesics. Note that the ϵ -thin part is precisely the set of points which are moved a distance less than 2ϵ by some nontrivial element of Γ .

Lemma 2.2. *Given $K > 0$ and $0 < \epsilon < \epsilon_0$, there exists $\eta_0 > 0$ with the following property. Suppose $\Gamma < \mathrm{PSL}(2, \mathbb{C})$ is a discrete group and $\tilde{z}_1, \dots, \tilde{z}_\eta \in \mathbb{H}^3$ are points in distinct ϵ -thin parts for Γ . If there is some $\tilde{z}_0 \in \mathbb{H}^3$ so that $d_{\mathbb{H}^3}(\tilde{z}_0, \tilde{z}_i) \leq K$ for all $i \geq 1$, then $\eta < \eta_0$.*

Proof. Since $0 < \epsilon < \epsilon_0$, then for $\delta = \frac{\epsilon_0 - \epsilon}{2}$, the ball B_i of radius δ about $\tilde{z}_i$ is contained in the ϵ_0 -thin part for Γ . Moreover, $\{B_i\}_i$ is a collection of pairwise disjoint balls of radius δ contained in the ball of radius $K + \delta$. Thus we can take η_0 to be an integer larger than the ratio of the volume of a ball of radius $K + \delta$ to that of a ball of radius δ . $\square$

Another key ingredient in this paper is the fact that geodesics intersecting thick parts of 3-manifolds that enter Margulis tubes or cusps are long. This fact is made precise by the following lemma.

Lemma 2.3 ([BM82, Mey87]). *There is a continuous positive function*

$$R : [0, \epsilon_0] \rightarrow \mathbb{R}_{\geq 0} \text{ such that } \lim_{\epsilon \rightarrow 0} R(\epsilon) = \infty$$

with the following property. Given $\epsilon \leq \epsilon_0$, let $T_\epsilon \subset T_{\epsilon_0}$ be connected components of the ϵ - and ϵ_0 -thin parts, respectively, of a hyperbolic manifold M . Then

$$d_M(\partial T_\epsilon, \partial T_{\epsilon_0}) \geq R(\epsilon).$$

2.2. Dehn filling/surgery. A crucial ingredient in our construction of knots rather than links is Thurston's hyperbolic Dehn filling theorem; see [Thu78, Chapter 5] and [BP92, Chapter E]. This allows cusps of a finite volume hyperbolic 3-manifold to be "filled in" by solid tori while preserving hyperbolicity and approximating some geometric features of the original manifold.

Suppose M is a compact 3-manifold and $\{T_1, \dots, T_k\} \subset \partial M$ its boundary components, each of which is homeomorphic to a torus. Let $\alpha_j \subset T_j$ be an essential simple closed curve in T_j . Let U_j be a solid torus, with meridian β_j . Consider a homeomorphism

$$\phi_j : (\partial U_j, \beta_j) \rightarrow (T_j, \alpha_j)$$

which sends each β_j homeomorphically to each α_j . The resulting manifold obtained by “filling in” each T_j via the map ϕ_j (that is, gluing U_j to T_j) is the *Dehn filling of M along α_j* . Up to homeomorphism, the result of Dehn filling depends only on the isotopy class of α_j in T_j . Such an isotopy class is called a *slope* on T_j , and we often use the same notation for the simple closed curve and its isotopy class. We can parameterize the isotopy classes α_j by considering framings of each T_j by a choice of longitude μ_j and meridian ν_j (i.e. a basis for $\pi_1(T_j)$). We can encode α_j by the tuple $(p_j, q_j) \in \mathbb{R}^2$ (up to sign) in $\pi_1(T_j)$, if (after orienting α_j), $\alpha_j = p_j\mu_j + q_j\nu_j$. Setting $u_j = (p_j, q_j)$, we write both $M(u_1, \dots, u_k) = M(\alpha_1, \dots, \alpha_k)$ as the manifold obtained from Dehn filling along each α_j . We also allow α_j or u_j to be “ ∞ ”, by which we mean that we perform no Dehn filling on the boundary component.

If M is the interior of a compact manifold with torus boundary, we similarly define Dehn filling by deleting product neighborhoods of the ends to produce a compact 3-manifold whose boundary components are all homeomorphic to tori; these tori, viewed in the original non-compact manifold M , will be called *peripheral tori*. If a filling slope is ∞ , then we glue the product neighborhood of the end back in. If $L \subset M$ is a (tame) link in compact 3-manifold M , then a *Dehn surgery* on L is the process of drilling out L (i.e. removing L), and performing Dehn filling on the resulting open manifold (or equivalently, removing tubular neighborhoods of L and Dehn filling torus boundary components of the resulting manifold).

Let H be the space of all finite-volume 3-manifolds equipped with the *geometric topology*. That is, a sequence of 3-manifolds $\{N_i\} \subset H$ converges to $N \in H$ if and only if for some sequence of positive numbers $\epsilon_i \rightarrow 0$, there is a sequence $\{\phi_i\}$, where each ϕ_i is an $(1 + \epsilon_i)$ -biLipschitz diffeomorphism between the ϵ_i -thick part of N_i and the ϵ_i -thick part of N . We are now ready to state the following theorem due to Thurston; see [Thu78, Theorem 5.8.2], [BP92, Sections E.5-E.6].

Theorem 2.4 (Hyperbolic Dehn Surgery Theorem [Thu78]). *Suppose N is a finite volume, cusped hyperbolic 3-manifold. Then $N(u_1, \dots, u_k)$ is hyperbolic, provided each u_i avoids a finite set of (non- ∞) slopes on the i^{th} cusp. Furthermore, if $u_{i,n} \rightarrow \infty$ in $\mathbb{R}^2 \cup \{\infty\}$ for each i , then $\{N(u_{1,n}, \dots, u_{k,n})\}_n$ converges to N .*

2.3. A theorem of Kapovich–Weidmann. Lastly, we introduce a result which will be used to show that the knots we construct are filling knots. Let G be a free group of rank n . Two generating sets for G are *Nielsen equivalent* if they differ by a sequence of *Nielsen transformations*, which are analogs of elementary row operations on $GL_n(\mathbb{Z})$ (see [Nie24] for details), and which generate the automorphism group of G . In other words, Nielsen transformations map ordered an (ordered) basis of G to another basis of G . Like in [LW25], in Section 5, we will make use of the following version of a theorem due to Kapovich–Weidmann, which roughly states that a group that acts by isometries on a δ -hyperbolic space ($\delta > 0$) is either free or contains a nontrivial element with short translation length.

Theorem 2.5 ([KW03]). *For any integer $n > 0$, there exists a constant $C(n)$ with the following property. Suppose a group $G = \langle g_1, \dots, g_n \rangle$ acts by isometries on a*

δ -hyperbolic geodesic metric space (X, d) , where $\delta > 0$. Then one of the following holds:

- (1) The group G is free with basis $(g_1, g_2, \dots, g_n)$, or
- (2) The tuple $(g_1, g_2, \dots, g_n)$ is Nielsen-equivalent to a tuple $(g'_1, g'_2, \dots, g'_n)$ such that for some $x \in X$, $d_X(g'_1.x, x) < \delta C(n)$.

3. CONSTRUCTIONS OF LINKS

We fix a closed, orientable 3-manifold M for the remainder of this paper. We first recall the construction of links $\{L\}$ in M with large essential systole, built from special triangulations of M , and remind the reader of a few relevant facts about them. We then define a new tangle, and describe features of a nice hyperbolic metric on it. Next, we explain how to modify L to obtain a new link L' in M which will serve as the primary ingredient, via Dehn surgery, for the construction of knots that prove the main theorems. We end with the construction of the knots and state more precise versions of the theorems from the introduction.

3.1. Triangulations and combinatorial systole. To describe the construction of the knots in a closed 3-manifold M needed to prove both Theorem 1.3 and Theorem 1.4, we recall two definitions from [LW25].

Given a triangulation τ of M , we define the *combinatorial systole* of τ to be the length of the shortest non-null-homotopic loop in M , measured with respect to the graph metric on the 1-skeleton of τ . There are various ways to make this precise, and the one that is most useful for our purposes is to consider the subgraph $\hat{\tau}$ of the 1-skeleton of the barycentric subdivision of τ consisting of all 1-simplices that meet the barycenters of the tetrahedra of τ . This is a bipartite graph, with vertices partitioned into barycenters of tetrahedra and all other barycenters. A pair of tetrahedra-vertices are distance two apart if and only if the tetrahedra intersect. The inclusion induces a surjective homomorphism $\pi_1(\hat{\tau}) \rightarrow \pi_1(M)$, and we define

$$\text{combsys}(\tau) = \min\{\ell_{\hat{\tau}}(\gamma) \mid \gamma \text{ a loop in } \hat{\tau}, \gamma \neq * \text{ in } M\}.$$

We will utilize a special type of triangulation τ of our ambient manifold M which we call a *Cooper-Thurston triangulation*. These were defined by Cooper and Thurston in [CT88], and have the property that the link of every vertex is one of 5 particular 2-complexes (each of which is flag), and each edge has 4, 6, 8, or 10 tetrahedra glued around it. In [LW25, Proposition 3.2], we proved the following.

Proposition 3.1. *For any closed, orientable 3-manifold M with $\pi_1 M \neq \{*\}$ and $r > 0$, there exists a Cooper-Thurston triangulation τ with $\text{combsys}(\tau) > r$.*

3.2. Link construction and metrics. Given a Cooper-Thurston triangulation τ of M , in [LW25] we constructed an associated hyperbolic link $L = L(\tau) \subset M$ whose complement $M_L = M \setminus L$ admits an explicit, complete, locally CAT(-1), singular hyperbolic metric, with singular locus contained in the 1-skeleton of τ . In particular, by Theorem 2.1, the link is itself hyperbolic, and we write M_L^h for the hyperbolic metric coming from that theorem. The essential systole of L is bounded below by a function of the combinatorial systole of τ (see [LW25]). Before we describe the new constructions of links and knots, we first recall the construction from [LW25].

First, we fix a “model” tetrahedron T_0 , and consider the tangle $J_0 \subset T_0$ shown on right of Figure 1. Here $J_0 \subset T_0$ consists of six properly embedded arcs, one

connecting every pair of faces, and four circles, one contained in each face. In particular, we can index the components of J_0 by the set of faces (for circle components) and edges (for arc components, index an arc by the edge shared by the faces it meets). In [LW25] we showed that $T_0 \setminus J_0$, viewed as a manifold with corners along the edges, admits a hyperbolic structure for which the boundary is totally geodesic so that the dihedral angles along each edge are equal to $\frac{\pi}{2}$. The tangle and hyperbolic structure are constructed to be invariant by the full symmetry group $\mathfrak{Tet} \cong S_4$ of T_0 . Specifically, we take a fundamental domain Δ for the action of $\mathfrak{Tet}$ on T_0 , which is one tetrahedron of the first barycentric subdivision of T_0 . By inspection, Δ meets the tangle J_0 in exactly two arcs as indicated on the left in Figure 1. We then explicitly constructed a partially ideal hyperbolic polyhedron which is homeomorphic to $\Delta \setminus J_0$ so that the dihedral angles are as indicated in the figure. The hyperbolic structure on $T_0 \setminus J_0$ is such that the restriction to $\Delta \setminus J_0$ (or any image by $\mathfrak{Tet}$) is isometric to this partially ideal hyperbolic polyhedron. The link $L \subset M$ associated to the triangulation τ is defined by the property that for every tetrahedron T , there is a homeomorphism of pairs $(T, T \cap L) \cong (T_0, J_0)$. The locally CAT(-1) metric on M_L has $T \setminus L$ isometric to the fixed hyperbolic metric on $T_0 \setminus J_0$, for every tetrahedron T of τ .




FIGURE 1. **Left:** Dihedral angles for hyperbolic polyhedron homeomorphic to $\Delta \setminus (J_0 \cap \Delta)$. **Right:** The tangle $J_0 \subset T_0$.

Suppose $L \subset M$ is a link associated to a Cooper-Thurston triangulation, τ , as above. There are three families of surfaces in M , pictured in Figure 2, described as follows. Let $\mathcal{E}_\tau$, $\mathcal{F}_\tau$, and $\mathcal{T}_\tau$ denote the set of edges, faces, and tetrahedra of τ .

- (1) For every $f \in \mathcal{F}_\tau$ of τ , the *face surface/disk* associated with f , denoted Σ_f , is the disk bounded by the component of L embedded in f .
- (2) For every $e \in \mathcal{E}_\tau$, there is a component of L that encircles e , built from a subset of the arcs of L intersected with the tetrahedra containing e . The *edge surface/disk* associated with e , denoted Σ_e , is the disk bounded by this link component.
- (3) For every $T \in \mathcal{T}_\tau$, let $f_1, \dots, f_4$ be the faces of T . The *tetrahedron surface* associated with T , denoted Σ_T , is obtained by “pushing” the interior of the closure of the complement of the face surfaces in the boundary, $\partial T \setminus \bigcup_{i=1}^4 \Sigma_{f_i}$, into the interior of T .

For each edge, face, or tetrahedron x , we write Σ_x° for the intersection of Σ_x with M_L . Each Σ_x° is homeomorphic to a sphere minus some number of points (or equivalently closed disks) and are pictured in Figure 2.




FIGURE 2. **Left:** A face f and its associated surface Σ_f . **Middle:** An edge e , and its associated surface Σ_e . **Right:** The four-holed sphere Σ_T associated with a tetrahedron T .

Remark 3.2. In either M_L (the locally CAT(-1) manifold) or M_L^h (the complete hyperbolic manifold) all punctures and (removed) boundary components of face, edge, and tetrahedron surfaces are cusps. Thus, although it is convenient to view the edge and face surfaces as punctured disks from a topological perspective (and the tetrahedron surface as having four “missing” boundary components), from a geometric point of view, they are more naturally all punctured spheres. In this paper, we will move freely between these two perspectives depending on which is more descriptive.

Because every tetrahedron T of τ has $T \setminus L$ isometric to $T_0 \setminus J_0$ in the locally CAT(-1) metric, as explained in [LW25], these surfaces are all totally geodesic. We recall the reasoning and point out some additional features that we will need.

First, the face surfaces and edge surfaces are constructed by gluing together polygons in the faces of the ideal polyhedron $\Delta \setminus J_0$ along sides that have dihedral angle $\frac{\pi}{2}$ in Δ , and so one can see explicitly that they are totally geodesic. The number of faces of $\Delta \setminus J_0$ needed to construct an edge surface in M_L is precisely twice the number of tetrahedra that the edge meets, and because L comes from a Cooper-Thurston triangulation, these surfaces are built from 8, 12, 16, or 20 polygons. Additionally, in $T_0 \setminus J_0$, the isometric $\mathfrak{Tet}$ action implies every face surface is isometric to a single model face surface we denote Σ_F° . This is a hyperbolic surface, and the stabilizer of a face surface via the isometric action of $\mathfrak{Tet}$ shows that it admits a three-fold rotational symmetry, with quotient a $(3, \infty, \infty)$ triangle orbifold (topologically a sphere with two punctures and one orbifold point of order 3). In particular, every face surface in M_L admits such a symmetry. The group $\mathfrak{Tet}$ also leaves the tetrahedral surface invariant, and quotienting this surface by the orientation preserving subgroup yields a triangle orbifold. Theorem 2.1 of [AS05] implies that the tetrahedron surface is also totally geodesic (to apply the theorem, we can treat the surfaces in the boundary as reflectors). See [LW25] for details.

3.3. An auxiliary tangle. Next we describe an auxiliary tangle in the closed 3-ball, $J_{\text{aux}} \subset \bar{\mathbb{B}}^3$, as well as its geometry. Given a Cooper-Thurston triangulation τ and associated link $L \subset M$, we will use copies of this tangle to “join together” the components of L that intersect the face surfaces to produce a new link $L' \subset M$.

The tangle J_{aux} is shown in Figure 3. It consists of three properly embedded arcs, and four simple closed curves. One of the simple closed curves is on the boundary of $\bar{\mathbb{B}}^3$, and is illustrated as the equator in the figure. The other three, which we refer to as the *internal components* of J_{aux} , encircle the arcs as shown. The picture of J_{aux} shown on the right more clearly illustrates a 3-fold rotational symmetry. The complement $\bar{\mathbb{B}}^3 \setminus J_{\text{aux}}$ has two boundary components which are each thrice punctured disks (or four-times punctured spheres; c.f. Remark 3.2).




FIGURE 3. Two views of the tangle $J_{\text{aux}} \subset \bar{\mathbb{B}}^3$, with the view “from above” on the right, where the arc components are visible only as “dots”.

We first make precise some terminology that will be used throughout the rest of this section. Suppose $\Sigma \subset N$ is a totally geodesic surface in a complete, singular hyperbolic, locally CAT(−1) manifold N .

Definition 2. To *cut a manifold N open along a surface Σ* means to remove Σ , and then take the path-metric completion of the result. We write $N//\Sigma$ to denote N cut along Σ .

The manifold $N//\Sigma$ has two new, totally geodesic boundary components, which may be isometrically glued back together in a way that is canonically isometric to the original manifold N . A similar definition makes sense for cutting open any manifold along a codimension 1 submanifold.

Theorem 3.3. *The manifold $\bar{\mathbb{B}}^3 \setminus J_{\text{aux}}$ has a hyperbolic structure for which both boundary components and the three twice-punctured disks bounded by the interior components are totally geodesic. Furthermore, the 3-fold symmetry group acts by isometry, and each boundary component is isometric to Σ_F° , the model face surface in $T_0 \setminus J_0$, in a way that respects the three-fold symmetry in each.*

The thrice punctured disks have empty boundary as surfaces in $\bar{\mathbb{B}}^3 \setminus J_{\text{aux}}$, so from a geometric point of view, they are more naturally four-times punctured spheres, but this nomenclature of “punctured disks” is helpful for visualizing the surfaces (c.f. Remark 3.2).

Proof. We consider the auxiliary link $L_{\text{aux}} \subset S^3$, shown in Figure 4. The link L_{aux} can be seen to be hyperbolic by using the `verify_hyperbolicity` function in SnapPy [HIK⁺16]. Let D_0 be the disk in $S^3 \setminus L_{\text{aux}}$ bounded by the blue component, as indicated, and write $D_0^\circ = D_0 \cap S^3 \setminus L_{\text{aux}}$.




FIGURE 4. Two views of the link L_{aux} . Note that $(L_{\text{aux}} \setminus \partial D_0) \subset S^3$ is a 6 component chain link, $C_6 \subset S^3$.

The peripheral loops of D_0° are homotopic to distinct cusps of $S^3 \setminus L_{\text{aux}}$, and it follows that all non-null homotopic simple closed curves in D_0° are non-null homotopic in $S^3 \setminus L_{\text{aux}}$; hence, D_0° is essential. Cutting open along D_0° yields $\mathbb{B}^3 \setminus J_{\text{aux}}$, and hence hyperbolicity of $\mathbb{B}^3 \setminus J_{\text{aux}}$ follows.

To see that the boundary components of $\mathbb{B}^3 \setminus J_{\text{aux}}$ are totally geodesic with respect to the given hyperbolic structure, it suffices to show that D_0° is isotopic to a totally geodesic surface. For this, we note that the 3-fold symmetry of $\mathbb{B}^3 \setminus J_{\text{aux}}$ extends to a 3-fold symmetry of $S^3 \setminus L_{\text{aux}}$, which can be realized as an isometry by Mostow's Rigidity Theorem. Moreover, D_0° is invariant, and the quotient is a $(3, \infty, \infty)$ -triangle 2-orbifold Q in the quotient $\mathcal{O}$ of $S^3 \setminus L_{\text{aux}}$. Since D_0° is essential in $S^3 \setminus L_{\text{aux}}$, Q is essential in $\mathcal{O}$. By Theorem 2.1 in [AS05], $\mathcal{O}$ is isotopic to a totally geodesic 2-orbifold Q' . This isotopy can be lifted to an isotopy of D_0° to a totally geodesic surface $D_0^{\circ'}$. Furthermore, because the hyperbolic structure on a triangle orbifold is unique by Corollary 13.3.7 of [Thu78], it follows that $D_0^{\circ'}$ is isometric to Σ_F° , as one can lift the isometry between Q' to the quotient of Σ_F° by its 3-fold symmetry group.

Each of the three internal components of J_{aux} bound disks, D_{12} , D_{23} , and D_{31} in $\mathbb{B}^3$ with boundary components in J_{aux} . The intersection of D_{ij} with $\mathbb{B}^3 \setminus J_{\text{aux}}$ is a thrice punctured sphere, D_{ij}° in the hyperbolic structure in $\mathbb{B}^3 \setminus J_{\text{aux}}$, which are also thrice-punctured spheres in $S^3 \setminus L_{\text{aux}}$ (c.f. Remark 3.2). By [Ada85, Theorem 2.1], after an isotopy we can assume these surfaces are totally geodesic. Since totally geodesic surfaces have the minimal number of components of transverse intersections in their isotopy classes, and since the D_{ij}° can be realized to be pairwise disjoint after an isotopy, and disjoint from D_0° , the same is true of their totally geodesic representatives. In particular, each D_{ij}° can be assumed to be totally geodesic in $\mathbb{B}^3 \setminus J_{\text{aux}}$. $\square$

3.4. New Links. We now return to the construction of a new link associated with a Cooper-Thurston triangulation τ of M . We start with the associated link

$L = L(\tau) \subset M$ described in §3.2. For each face f of τ , we consider two copies of the disk Σ_f in M with the same boundary, and slightly push their interiors into the two tetrahedra sharing the face f . The union of each such pair bounds a ball $\bar{\mathbb{B}}^f$, and $(\bar{\mathbb{B}}^f, \bar{\mathbb{B}}^f \cap L)$ is homeomorphic to the tangle $(\bar{\mathbb{B}}^3, J'_{\text{aux}})$, where J'_{aux} is the union of components of J_{aux} *not including the three internal components*. The new link $L'' = L''(\tau)$ is obtained from L by adding three components for each face f , so that $(\bar{\mathbb{B}}^f, \bar{\mathbb{B}}^f \cap L'') \cong (\bar{\mathbb{B}}^3, J_{\text{aux}})$. We note that L'' can be viewed as a kind of *augmentation* of L as in Adams [Ada86], Lackenby [Lac04], Purcell [Pur10], and many others, and we will utilize this perspective.

Another perspective is that $M_{L''} = M \setminus L''$ is homeomorphic to the 3-manifold obtained by first cutting M_L along the union of the face surfaces (see Definition 2), producing a 3-manifold with two totally geodesic boundary components for each face $f \in \mathcal{F}_\tau$, each isometric to Σ_f° , then gluing in a copy of $\bar{\mathbb{B}}^3 \setminus J_{\text{aux}}$ by isometry along the faces. In particular, the link complement $M_{L''}$ contains *two* face surfaces for every face f , which we denote $\Sigma_{f\pm}^\circ$, in addition to the tetrahedra surfaces Σ_T° for each tetrahedron $T \in \mathcal{T}_\tau$. For every face f of τ , there are also three twice-punctured disks D_{f12}° , D_{f23}° , and D_{f31}° , coming from the punctured disks D_{12}° , D_{23}° , D_{31}° in $\bar{\mathbb{B}}^3 \setminus J_{\text{aux}}$. The following is immediate from the discussion above.

Theorem 3.4. *Suppose $L'' = L''(\tau)$ is the link associated to a Cooper-Thurston triangulation, τ , as above. Then $M_{L''}$ admits a locally CAT(−1) metric so that for all tetrahedra $T \in \mathcal{T}_\tau$, faces $f \in \mathcal{F}_\tau$, and $ij = 12, 23, 31$, the surfaces Σ_T° , $\Sigma_{f\pm}^\circ$, and D_{fij}° are totally geodesic. Furthermore, these totally geodesic surfaces are pairwise disjoint in $M_{L''}$. $\square$*

We now describe the construction of a third link $L' = L'(\tau) \subset M$ associated with τ , from which the knots used to prove the main theorem will be obtained by Dehn surgery. For this construction, for any face $f \in \mathcal{F}_\tau$ and any ij , we let $D_{fij} \subset M$ denote the closure of the punctured disk D_{fij}° ; we assume D_{fij} is a disk whose boundary is a component of L'' , and whose interior meets L'' in two points. We can modify L'' locally in a neighborhood of D_{fij} , by a *half twist* along D_{fij} , which has the effect of replacing the local picture of L'' near D_{fij} on the left of Figure 5 with the one on the right of that same figure, and keeping the link unchanged elsewhere.




FIGURE 5. Twisting in a disk.

Since the set of disks $\{D_{fij}\}_{ij,f}$ are pairwise disjoint, we can independently perform half-twists on any collection of these disks to produce a new link which

depends (up to isotopy) only on the choices of disks in which we perform a half-twist, and not the order in which we performed the twists. Let $\mathcal{D}$ denote some set of such disks and $L' \subset M$ the link obtained by half-twisting in those disks. Observe that we can obtain $M_{L'}$ from $M_{L''}$ by cutting $M_{L''}$ open along the set of punctured disks D_{fij}° whose closures are in $\mathcal{D}$, and then regluing by precomposing “the identity” gluing with the half twist in the disk. We consider all D_{fij}° , $\Sigma_{f\pm}^\circ$, and Σ_T° as surfaces in $M_{L'}$ (and their closures as surfaces in M). The following lemma specifies a choice of $\mathcal{D}$ from which we will construct a specific link L' .

Lemma 3.5. *There is a choice of disks $\mathcal{D}$ as above, so that the resulting link $L' \subset M$ consists of the following union,*

$$L' = K' \sqcup \bigsqcup_{f \in \mathcal{F}_\tau} \left(\partial \Sigma_{f+} \sqcup \left(\bigsqcup_{ij} \partial D_{fij} \right) \right),$$

where K' is a single component.

The link complement $M_{L'}$ admits a complete, singular hyperbolic, locally CAT(−1) metric isometric to $M_{L''}$ outside the union of all the punctured disks D_{fij}° (in both manifolds). In particular, the totally geodesic surfaces Σ_T° for $T \in \mathcal{T}_\tau$ and $\Sigma_{f\pm}^\circ$ for $f \in \mathcal{F}_\tau$ in $M_{L''}$ define totally geodesic surfaces of the same name in $M_{L'}$.

The last disjoint union is over the indices $ij \in \{12, 23, 31\}$. We note that $\partial \Sigma_{f+} = \partial \Sigma_{f-}$. Assuming the conclusion of the lemma, we observe that for every disk D of the form $\Sigma_{f\pm}$ or D_{fij} , the intersection of L' with the interior of D is precisely $K' \cap D$. This observation is the basis of the proof of the lemma.

Proof. Given a choice of disks $\mathcal{D}$ and associated link L' , we say that a disk D_{fij} is *good* for L' if the two intersection points with the interior of D_{fij} are from the same component of L' . For example, observe that L'' is the link associated to $\mathcal{D} = \emptyset$, and that by inspection, no disk D_{fij} is good for L'' .

We first prove the following.

Claim 3.6. *There is a collection $\mathcal{D}$ such that every disk D_{fij} is good for the associated link L' .*

Proof of claim. First, enumerate all the disks, over all faces $f \in \tau$ and all $ij \in \{12, 23, 31\}$, writing the list as $D_1, \dots, D_m$. We will define the collection $\mathcal{D}$ as a non-decreasing union

$$\mathcal{D}_1 \subseteq \mathcal{D}_2 \subseteq \mathcal{D}_3 \subseteq \dots \subseteq \mathcal{D} = \bigcup_{k=1}^m \mathcal{D}_k.$$

The sets $\mathcal{D}_k$ will be defined recursively, having the property that

$$\mathcal{D}_k \subset \{D_1, \dots, D_k\},$$

and so that $\mathcal{D}_k$ is obtained from $\mathcal{D}_{k+1}$ by either adding $\overline{D}_k$ to $\mathcal{D}_{k-1}$ or not. For each k , we let L'_k be the link associated to $\mathcal{D}_k$. The properties of the sequence $\mathcal{D}_k$ described above implies that L'_k is obtained from L'_{k-1} by either twisting in D_k or not, depending on whether it is included in $\mathcal{D}_k$ or not.

Now we define $\mathcal{D}_1 = \{D_1\}$, and since D_1 is not good for L'' , it is good for $\mathcal{D}_1$: the half-twist has the effect of combining the two components of L'' meeting the interior of D_1 into a single component of L'_1 (see Figure 5).

Next, having defined $\mathcal{D}_{k-1}$ for some $k > 1$, we define

$$\mathcal{D}_k = \begin{cases} \mathcal{D}_{k-1} & \text{if } D_k \text{ is good for } L'_{k-1} \\ \mathcal{D}_{k-1} \cup \{D_k\} & \text{otherwise} \end{cases}$$

We observe that, by induction, for all $1 \leq j \leq k$, D_j is good for L_k , since a half twist in a disk that isn't good makes it good (and doesn't destroy goodness for any other disks). The set $\mathcal{D}$ thus satisfies the property in the claim.

The second half of the statement of Lemma 3.5 follows immediately from Corollary 4.4 of [Ada85]. Indeed, we obtain $M_{L'}$ from $M_{L''}$ by cutting along, applying half-twists to, and regluing totally geodesic thrice-punctured disks. Since each half-twist is an isometry for the hyperbolic metric on the punctured disk, the resulting complete, singular hyperbolic, locally CAT(-1) manifold is isometric to $M_{L''}$ in the complement of those punctured disks along which we performed the half twists. $\square$

For each face f , there are three arcs of L' that connect the interior of Σ_{f-} to the interior of Σ_{f+} , and each one of these arcs intersects two of the disks D_{fij} . Since each disk is good, it follows that all three of these arcs are contained in the same component of L' . From this we see that for any tetrahedron $T \in \mathcal{T}_\tau$, if $f_1, \dots, f_4$ are the faces, then all arcs of L' that meet the interiors of the disks $\Sigma_{f_i\pm}$ ($i = 1, 2, 3, 4$) belong to the same component of L' . Connectivity of M implies that there is a single component K' of L' with the property that for any disk D_{fij} or $\Sigma_{f\pm}$, the intersection of L' with its interior is precisely the intersection of K' with its interior. Thus, L' is the union of K' and the boundaries of these disks. $\square$

Definition 3. Given a Cooper-Thurston triangulation τ of a closed orientable 3-manifold, M , we write $L' = L'(\tau)$ for the link coming from some choice of $\mathcal{D}$ as in Lemma 3.5, and refer to this as a *fully twisted link* associated to τ . We will refer to the component $K' \subset L'$ specified in the lemma as the *focal component* and all other components as *disk components*.

Given a fully twisted link L' associated to τ , it will be useful to have some language to describe how L' sits inside M , and its relation to L and the triangulation τ . Recall that for each face f , $\bar{\mathbb{B}}^f$ is the ball in M bounded by the union of faces $\Sigma_{f+} \cup \Sigma_{f-}$. The collection of balls $\{\bar{\mathbb{B}}^f\}_f$ over all faces is a set of pairwise disjoint balls.

For each face f , the pair $(\bar{\mathbb{B}}^f, \bar{\mathbb{B}}^f \cap L')$ is homeomorphic to $(\bar{\mathbb{B}}^3, J_{\text{aux}}^f)$, where J_{aux}^f is obtained from J_{aux} after possibly applying some half-twists to the disks D_{ij} (namely those specified by the choice of $\mathcal{D}$).

Definition 4. We write

$$\mathbb{B}^f = \bar{\mathbb{B}}^f \setminus (\bar{\mathbb{B}}^f \cap L'),$$

and call this manifold the *connector for f* .

It is useful to think about each tetrahedron as having been cut along the interiors of Σ_f , over its four faces f , and pushed off of the interiors of the associated balls $\bar{\mathbb{B}}^f$. In particular, M is reconstructed by gluing together the disjoint union of the tetrahedra from τ and the balls $\{\bar{\mathbb{B}}^f\}$. Thus, we think of removing the union of the interiors of all $\bar{\mathbb{B}}^f$ as cutting M open along the interiors of the face disks.

Consequently, the way L' meets the complement of the interiors of $\bar{\mathbb{B}}^f$ in M is precisely the way L meets M cut along the interiors of the face disks. More

precisely, let us write $\check{M}$ to denote the complement of the unions of the interiors of $\mathbb{B}^f$ in M ,

$$\check{M} = M \setminus \bigcup_f \text{int}(\mathbb{B}^f),$$

and N to denote M cut along the interiors of the face surfaces,

$$N = M // \bigcup_f \text{int}(\Sigma_f)$$

(as in Definition 2), and then

$$(1) \quad (\check{M}, \check{M} \cap L') \cong (N, N \cap L).$$

In particular, we can view $M_{L'}$ as obtained from M_L by cutting along each Σ_f° , and gluing in the f -connector, $\mathbb{B}^f$,

$$(2) \quad M_{L'} \cong (M_L // \bigcup_f \Sigma_f^\circ) \cup \bigcup_f \mathbb{B}^f \cong N_L \cup \bigcup_f \mathbb{B}^f,$$

where $N_L = N \setminus (N \cap L)$.

3.5. Constructing knots and the theorems revisited. Given a Cooper–Thurston triangulation τ in a closed 3-manifold M , and an integer n , we define a knot $K(n, \tau) \subset M$ as follows. Let L' be the fully twisted link associated to τ and K' its focal component. The disk components of L' (that is, the components in $L' \setminus K'$) bound disks with pairwise disjoint interiors

$$\bigcup_{f, ij} \Sigma_{f+} \cup \Sigma_{f-} \cup D_{fij}.$$

We assume that the peripheral tori of $M_{L'}$ are framed with basis for which the longitude is defined by these disks (and the meridian is the one coming from the solid tori $\mathcal{N}(L')$).

Now consider the link $L' \setminus K' \subset M$ consisting of *only the disk components of L'* . For any integer n , we apply $(1, n)$ -Dehn surgery on *every* component of this link. Since the disks bounded by the disk components have pairwise disjoint interiors, the resulting manifold is homeomorphic to M . The image of the focal component K' by this homeomorphism is a knot which we denote $K(n) = K(n, \tau)$. In fact, we can view $K(n)$ as obtained from K' by adding n full twists at each disk (i.e. $2n$ half twists as in Figure 5); see [Ada86, Lac04, Pur10] for this kind of construction. The homeomorphism sends the disk components to the core curves of the filling solid tori, which results in a link that contains $K(n)$ as one of components and whose components all bound disks.

Remark 3.7. We can view $M_{K(n)} = M \setminus K(n)$ as obtained by applying a $(1, n)$ -Dehn filling on each cusp of $M_{L'}$ coming from a disk component of L' , and leaving the cusp corresponding to the focal component unfilled; that is

$$M_{K(n)} \cong M_{L'}((1, n), \dots, (1, n), \infty).$$

We can now state a more precise version of the theorems from the introduction. The first of these immediately implies Theorem 1.3.

Theorem 3.8. *Suppose M is any closed, orientable 3-manifold with $\pi_1 M \neq \{*\}$ and $R > 0$. Then there exists a Cooper–Thurston triangulation τ so that for all n sufficiently large, the associated knot $K(n, \tau)$ has*

$$\text{sys}_e^*(K(n, \tau)) > R.$$

A map $f: \Gamma \rightarrow M$ has *full rank* if $\text{rank}(f_*(\pi_1(\Gamma))) = \text{rank}(\pi_1(M))$. A knot $K \subset M$ is *full rank- m filling* if for every $f: \Gamma \rightarrow M_K$ such that $\text{rank}(\pi_1(\Gamma)) \leq m$ and $i \circ f: \Gamma \rightarrow M$ is full rank, f_* is injective. The next theorem implies Theorem 1.4 by setting $m = \text{rank}(\pi_1(M))$.

Theorem 3.9. *Suppose M is any closed, orientable 3-manifold and $m > 0$. Then there exists a Cooper–Thurston triangulation τ so that for all n sufficiently large, the associated knot $K(n, \tau)$ is full rank- m filling.*

4. HYPERBOLIC KNOTS WITH LARGE ESSENTIAL GEODESIC SYSTOLE

From now on, the paper diverges in two different directions. In this section, we prove Theorem 3.8 using *hyperbolic* structures on $K(n, \tau)$. In the section that follows, we prove Theorem 3.9 using a locally $\text{CAT}(-\frac{1}{2})$ structure rather than a hyperbolic structure.

The crux of this section is the proof that the link $L' = L'(\tau)$ constructed from τ in Lemma 3.5 has large essential geodesic systole, provided the combinatorial systole is sufficiently large.

Proposition 4.1. *For $R_0 > 0$, there exists $r > 0$ so that if τ is a Cooper–Thurston triangulation of M with $\text{combsys}(\tau) > r$, then the hyperbolic link $L' = L'(\tau) \subset M$ has*

$$\text{sys}_e^*(L') > R_0.$$

We postpone the proof of this proposition, and first use it to prove Theorem 3.8.

Proof of Theorem 3.8. Fix $R_0 > 0$. By Proposition 4.1 and Proposition 3.3 of [LW25], one may take a Cooper–Thurston triangulation of M and its associated link L' with the property that $\text{sys}_e^*(L') > 2R_0$. Given $n \in \mathbb{N}$, recall that

$$M_{K(n)} \cong M_{L'}((1, n), \dots, (1, n), \infty)$$

denotes the manifold obtained by applying the $(1, n)$ -Dehn fillings described in Section 3.5. By Theorem 2.4, since $1 + n^2 \rightarrow \infty$ as $n \rightarrow \infty$, $M_{K(n)}$ is hyperbolic, and converges with respect to the geometric topology to the hyperbolic manifold $M_{L'}$. More precisely, there exists a sequence $\epsilon_n \rightarrow 0$ with the property that there is a $(1 + \epsilon_n)$ -biLipschitz diffeomorphism between the ϵ_n -thick parts of $M_{K(n)}$ and $M_{L'}$. Additionally, observe that for some $\epsilon' < \epsilon_0$ (the Margulis constant) and n sufficiently large, $(M_{K(n)})_{(0, \epsilon')}$ consists of a cusp neighborhood corresponding to the knot $K(n)$ and the union of solid tori which are the filling solid tori (up to isotopy).

For every $n \in \mathbb{N}$, let γ_n denote a geodesic whose length attains the essential geodesic systole of $K(n)$. We claim that there exists some $N \in \mathbb{N}$ such that

$$\text{sys}_e^*(M_{K(N)}) = \ell_{M_{K(N)}}(\gamma_N) > R.$$

We split our proof into two cases.

First, suppose that for all $\epsilon > 0$, there exists $n > 0$ such that γ_n is not contained in $(M_{K(n)})_{[\epsilon, \infty)}$. For each j with $\frac{1}{j} < \epsilon'$, choose some n_j such that γ_{n_j} is not

contained in $(M_{K(n_j)})_{[1/j, \infty)}$. We first claim that for each j , γ_{n_j} cannot be contained in $(M_{K(n_j)})_{(0, \epsilon_0)}$. Indeed, if $\gamma_{n_j} \subset (M_{K(n_j)})_{(0, \epsilon_0)}$, then γ_{n_j} is contained in either a cusp neighborhood corresponding to $K(n_j)$ or a Margulis tube, which would necessarily be a filling solid torus (because it would contain an $\frac{1}{j}$ -thin part, and $\frac{1}{j} < \epsilon'$). If γ_{n_j} is contained in a cusp neighborhood, then it cannot be geodesic, a contradiction. On the other hand, if γ_{n_j} is included in a filling solid torus V , then since the inclusion of V into $M_{K(n_j)}$ is nullhomotopic, γ_{n_j} is null-homotopic in M , another contradiction. Thus it must be the case that for all j , γ_{n_j} intersects $(M_{K(n_j)})_{(0, 1/j)}$, but is not completely contained in $(M_{K(n_j)})_{(0, \epsilon_0)}$. By Lemma 2.3, $\ell_{M_{K(n_j)}}(\gamma_{n_j}) \geq R(\frac{1}{j})$, and $R(\frac{1}{j}) \rightarrow \infty$ as $j \rightarrow \infty$. From this one can find some $N \in \mathbb{N}$ with the property that $\ell_{M_{K(N)}}(\gamma_N) > R_0$, as desired.

Second, suppose that there exists some $0 < \epsilon < \epsilon'$ with the property that for all $n > 0$, $\gamma_n \subset (M_{K(n)})_{[\epsilon, \infty)}$. Choose some $N \in \mathbb{N}$ sufficiently large so that $(M_{K(N)})_{[\epsilon, \infty)} \subset (M_{K(N)})_{[\epsilon_N, \infty)}$ and $\epsilon_N \leq 1$. Note that since γ_N is nontrivial in M , $\ell_{M_{L'}}(\gamma_N) \geq \text{sys}_e^*(L')$. Since there is an $(1 + \epsilon_N)$ -biLipschitz map between $(M_{L'})_{[\epsilon_N, \infty)}$ and $(M_{K(N)})_{[\epsilon_N, \infty)}$, we have

$$\ell_{M_{K(N)}}(\gamma_N) \geq \frac{1}{(1 + \epsilon_N)} \ell_{M_{L'}}(\gamma_N) > \frac{2R_0}{1 + \epsilon_N} \geq R_0.$$

□

4.1. Preliminaries for the proof of Proposition 4.1. We start by recalling some notation and describing the setup for the rest of the section. For any Cooper-Thurston triangulation τ , we denote its associated links $L \subset M$ and $L' \subset M$, with complements M_L and $M_{L'}$, respectively. These implicitly denote the manifolds with their locally CAT(−1) metrics from [LW25, Lemma 3.8] and Lemma 3.5, respectively. Recall that both L and L' are hyperbolic by Theorem 2.1, and we write M_L^h and $M_{L'}^h$ for the hyperbolic metrics on M_L and $M_{L'}$, respectively.

A first observation that will simplify the discussion is the following consequence of Lemma 2.3. Recall that ϵ_0 is the Margulis constant, and let $R(\epsilon)$ denote the function of ϵ from Lemma 2.3.

Lemma 4.2. *Given $R_0 > 0$, suppose $0 < \epsilon < \epsilon_0$ is such that $R(\epsilon) > R_0$. If γ^h is a closed hyperbolic geodesic in $M_{L'}^h$ that meets the ϵ -thin part, $(M_{L'}^h)_{(0, \epsilon)}$, then either γ^h is contained in $(M_{L'}^h)_{(0, \epsilon)}$, or else has length greater than R_0 .*

Proof. Suppose γ^h meets $(M_{L'}^h)_{(0, \epsilon)}$, but is not contained in it. Since $(M_{L'}^h)_{(0, \epsilon_0)} \supset (M_{L'}^h)_{(0, \epsilon)}$, and because there are no closed geodesics contained in cusps, γ^h must meet the boundary tori of $(M_{L'}^h)_{(0, \epsilon_0)}$ and $(M_{L'}^h)_{(0, \epsilon)}$. Furthermore, since γ^h travels between them, it must have length at least $R(\epsilon) > R_0$. □

Given the setup of the lemma, we note that a geodesic in the ϵ -thin part is necessarily a multiple of the core curve of a Margulis tube. Thus, if a γ^h is a geodesic in $M_{L'}^h$, then one of the following must hold.

- (1) γ^h has length greater than R_0 ,
- (2) γ^h has length less than ϵ , or
- (3) γ^h is contained in the ϵ -thick part of $M_{L'}^h$.

This is useful to keep in mind for a geodesic, γ^h , realizing the essential systole, as we prepare for the proof of Proposition 4.1. In the first case, we are done, and we

will be able to fairly easily rule out the second case when $\text{combsys}(L')$ is sufficiently large. Thus, the majority of the discussion that follows develops ideas to deal with the third case.

4.2. Crossing surfaces. Recall that for a Cooper–Thurston triangulation τ of M , we write N to denote the manifold obtained by cutting M open along the interiors of the face surfaces. Then M is obtained from N by gluing in the union of the balls $\mathbb{B}^f$ to the exposed sphere boundary components; in particular, M is the union of tetrahedra in $\mathcal{T}_\tau$ and the balls $\mathbb{B}^f$ for $f \in \mathcal{F}_\tau$. We consider a cell structure for which (the interiors of) the tetrahedra and the balls are 3-cells, and call this a τ -cell structure. This is imprecise as we have not described the 2-skeleton, but this will be irrelevant (and one can easily be constructed by appropriately subdividing the 2-skeleton of τ).

The next lemma follows easily from [LW25, Lemma 3.2].

Lemma 4.3. *Given any closed loop γ in M which is non-nullhomotopic, γ intersects at least $D = \lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$ tetrahedra of a τ -cell structure. Furthermore, there exists a collection of such tetrahedra $T_1, \dots, T_D$ with the property that for all $1 \leq i, j \leq D$, $d_\tau(T_i, T_j) \geq 8|j - i|$.*

Proof. There is a map $M \rightarrow M$ obtained by collapsing the balls onto the face disks, mapping each tetrahedron homeomorphically onto the corresponding tetrahedron in τ . This sends γ to a loop γ' which is non-null-homotopic. By [LW25, Lemma 3.1], γ' meets tetrahedra $T_1, \dots, T_D$ with the required properties. Hence γ meets the corresponding tetrahedra in the τ -cell structure. $\square$

Next, recall that by Lemma 3.5, there are face and tetrahedron surfaces in $M_{L'}$, denoted

$$\mathcal{S}_\tau^\circ(\mathcal{F}, \mathcal{T}) = \{\Sigma_{f^\pm}^\circ\}_{f \in \mathcal{F}_\tau} \cup \{\Sigma_T^\circ\}_{T \in \mathcal{T}_\tau}.$$

The edge surfaces from M_L are not defined in $M_{L'}$, but we describe a modification. Recall that N_L is M_L cut open along the union of its face surfaces, and from (2), $M_{L'}$ is obtained from N_L by gluing in the connectors $\mathbb{B}^f$. For any edge surface Σ_e° in M_L , we can cut Σ_e° open and view it inside of N_L , and hence inside $M_{L'}$. More precisely, we write

$$\check{\Sigma}_e^\circ \subset N_L \subset M_{L'}$$

to denote the surface obtained by cutting $\Sigma_e^\circ \subset M_L$ along the union of the face surfaces. This surface is isometric to an ideal hyperbolic n -gon, possibly with a cone singularity in the center; see Figure 6. Because of the constraints on the number of tetrahedra containing any edge e in a Cooper–Thurston triangulation, n is equal to 8, 12, 16, or 20 (see [LW25, Lemma 3.4]), and the cone angle at the center is $\frac{n}{4}\pi$. We call $\check{\Sigma}_e^\circ$ the *cut edge surface*, and we write $\mathcal{S}_\tau^\circ(\mathcal{E}) = \{\check{\Sigma}_e^\circ\}_e$ to denote the set of all cut edge surfaces, over all edges e of τ .

We write $\mathcal{S}_\tau^\circ = \mathcal{S}_\tau^\circ(\mathcal{F}, \mathcal{T}) \cup \mathcal{S}_\tau^\circ(\mathcal{E})$. Each of these surfaces is totally geodesic in the locally CAT(−1) metric in $M_{L'}$. Given a tetrahedron T , consider the family of surfaces $\mathcal{S}_\tau^\circ(T) \subset \mathcal{S}_\tau^\circ$ consisting of the following:

- (1) $\Sigma_{f^\pm}^\circ$, where f is a face of T ,
- (2) $\Sigma_{T'}^\circ$, where T' is a tetrahedron and $T' \cap T \neq \emptyset$, and
- (3) $\check{\Sigma}_e^\circ$, where e is an edge of T .




FIGURE 6. A cut edge surface $\check{\Sigma}_e^\circ$ may be viewed as an ideal polygon with a uniformly bounded number of sides.

Lemma 4.4. *If γ is a closed geodesic in $M_{L'}$ which is non-null-homotopic in M , and has non-empty intersection with some tetrahedron $T \in \mathcal{T}_\tau$, then γ has non-empty transverse intersection with some surface in $\mathcal{S}_\tau^\circ(T)$.*

Remark 4.5. Since we have written $M_{L'}$, rather than $M_{L'}^h$, we note that γ is thus assumed to be geodesic in the locally $\text{CAT}(-1)$ metric.

Proof. The closures of the surfaces in $\mathcal{S}_\tau^\circ$ cut M up into simply connected complementary regions in M , each of which is one of the following three types: (i) one of the balls $\bar{\mathbb{B}}^f$, (ii) a region contained in a single tetrahedron, or (iii) a neighborhood of a vertex (thus contained in a collection of tetrahedra that have non-empty pairwise intersection). Since γ is assumed non-null-homotopic in M , it cannot be contained in any one of these simply connected regions of M . Likewise, because the (totally geodesic) surfaces in $\mathcal{S}_\tau^\circ$ are null-homotopic in M , γ cannot be contained in any of these surfaces, and hence intersects them transversely. Assuming γ intersects $T \in \mathcal{T}_\tau$, it follows that it must intersect one of the simply connected regions of type (ii) contained in T , or (iii) which is a neighborhood of a vertex of T . As γ must escape this region, it has to intersect one of the surfaces in $\mathcal{S}_\tau^\circ$, and by definition, such a surface is in $\mathcal{S}_\tau^\circ(T)$. $\square$

4.3. Pleating surfaces in $\mathcal{S}_\tau^\circ$. When we pass from the locally $\text{CAT}(-1)$ metric to the hyperbolic metric, the surfaces in $\mathcal{S}_\tau^\circ$ are no longer necessarily totally geodesic. Given a geodesic γ in $M_{L'}$ that intersects a surface in $\mathcal{S}_\tau^\circ(\mathcal{F}, \mathcal{T})$, we will show that the homotopic hyperbolic geodesic, γ^h in $M_{L'}^h$, will intersect a pleated representative (Lemma 4.6). This is not necessarily true for the edge surfaces, but we still deduce some useful information (Lemma 4.7).

We start by describing the kinds of pleated surfaces we will consider for each $\Sigma^\circ \in \mathcal{S}_\tau^\circ$. First, each cut edge surface, $\check{\Sigma}_e^\circ$, is an ideal hyperbolic polygon in the hyperbolic metric which we denote Σ_e^h , and we represent it by a pleated surface $\check{\Sigma}_e^h \rightarrow M_{L'}^h$ pleated along an ideal triangulation of $\check{\Sigma}_e^h$. Given a face surface $\Sigma_{f\pm}^\circ$, there are three sides of cut edge surfaces which meet $\Sigma_{f\pm}^\circ$, and these are pairwise disjoint in $\Sigma_{f\pm}^\circ$. We extend this union of geodesics to an ideal triangulation, and take the pleated representative we denote $\Sigma_{f\pm}^h \rightarrow M_{L'}^h$, pleated along this ideal triangulation. We realize each tetrahedron surface by any pleated surface, denoted $\Sigma_T^h \rightarrow M_{L'}^h$.

Lemma 4.6. *If γ is a closed geodesic in $M_{L'}$ that transversely intersects a face or tetrahedron surface $\Sigma^\circ \in \mathcal{S}_\tau(\mathcal{F}, \mathcal{T})$, then the hyperbolic geodesic γ^h homotopic to γ has non-empty intersection with the pleated representative $\Sigma^h \rightarrow M_{L'}^h$.*

Proof. By Theorem 2.1, the identity lifts to a quasi-isometry $\widetilde{M}_{L'} \rightarrow \widetilde{M}_{L'}^h \cong \mathbb{H}^3$ between the universal covers, which extends to a homeomorphism of the Gromov boundaries. Let $\widetilde{\Sigma}^\circ$ be a component of the preimage of Σ° , which is a totally geodesic hyperbolic plane in the CAT(-1) space $\widetilde{M}_{L'}$. The boundary (quasi-)circle at infinity separates the boundary at infinity, $\partial\widetilde{M}_{L'}$, into two components, and there is an elevation $\widetilde{\gamma}$ of γ with one endpoint in each component.

Now with respect to the hyperbolic metric, there is a pleated surface $\widetilde{\Sigma}^h \rightarrow \widetilde{M}_{L'}^h$ covering $\Sigma^h \rightarrow M_{L'}^h$, which is a bounded Hausdorff distance from the image of $\widetilde{\Sigma}^\circ$. The (quasi-)circle at infinity is the image of the circle which is the boundary at infinity of $\widetilde{\Sigma}^\circ$. There is also an elevation $\widetilde{\gamma}^h$ of γ^h which is a bounded Hausdorff distance from the image of $\widetilde{\gamma}$, and hence has the same endpoints at infinity. In particular, these are in complementary components of the boundary at infinity of the pleated surface, and hence $\widetilde{\gamma}^h$ meets that surface. Consequently, γ^h has non-empty intersection with $\Sigma^h \rightarrow M_{L'}^h$. $\square$

For each edge $e \in \mathcal{E}_\tau$, each side of the cut edge surface $\check{\Sigma}_e^\circ$ is contained in a unique face surface $\Sigma_{f\pm}^\circ$. Observe that the surface Σ_{f+}° is such a surface if and only if Σ_{f-}° is. We write $\mathcal{F}(e) \subset \mathcal{F}_\tau$ to denote the set of faces f that arise in this way (i.e. so that $\Sigma_{f\pm}^\circ$ contains a side of $\check{\Sigma}_e^\circ$), and note that $\mathcal{F}(e)$ is precisely the set of faces that contain e .

Lemma 4.7. *Suppose γ is a closed geodesic in $M_{L'}$ which is non-nullhomotopic in M and has non-empty intersection with a cut edge surface $\check{\Sigma}_e^\circ$. Then for some $f \in \mathcal{F}(e)$, the hyperbolic geodesic γ^h homotopic to γ passes within distance at most 3 from one of the pleated surfaces $\Sigma_{f+}^h \rightarrow M_{L'}^h$ or $\Sigma_{f-}^h \rightarrow M_{L'}^h$, for some $f \in \mathcal{F}(e)$.*

Proof. We consider an elevation of γ to a bi-infinite geodesic $\widetilde{\gamma}$ in the universal cover $\widetilde{M}_{L'}$. Since $\check{\Sigma}_e^\circ$ is an ideal n -gon, which is simply connected, it lifts to the universal cover $\check{\Sigma}_e^\circ \cong \widetilde{\Sigma}_e^\circ \subset \widetilde{M}_{L'}$. For each $f \in \mathcal{F}(e)$, there is a component of the preimage of Σ_{f+}° and Σ_{f-}° that meets $\widetilde{\Sigma}_e^\circ$ along the lifted side. These are n totally geodesic hyperbolic planes denoted $\{\widetilde{\Sigma}_{f\pm}^\circ\}_{f \in \mathcal{F}(e)}$ in $\widetilde{M}_{L'}$. The boundaries at infinity of these hyperbolic planes form a set of (quasi-)circles that chain together in a cycle, with the n -gon $\widetilde{\Sigma}_e^\circ$ connecting the consecutive “tangent” points in the chain. The n circles bound n disks with pairwise disjoint interiors, $\Omega_1, \dots, \Omega_n$. The union of the circles thus divides $\partial\widetilde{M}_{L'}$ into $n+2$ regions, consisting of the n open disks above, and two remaining components $\Omega_{n+1}, \Omega_{n+2}$, “inside and outside” the chain of circles. We write $\Omega = \Omega_1 \sqcup \dots \sqcup \Omega_{n+2}$. See Figure 7. Since γ transversely intersects $\check{\Sigma}_e^\circ$, the endpoints at infinity of $\widetilde{\gamma}$ are in distinct components of Ω .

Again by Theorem 2.1, the identity lifts to a quasi-isometry $\widetilde{M}_{L'} \rightarrow \widetilde{M}_{L'}^h \cong \mathbb{H}^3$. There is an elevation $\widetilde{\gamma}^h$ of γ^h to $\widetilde{M}_{L'}^h$, which is a bounded Hausdorff distance to the image of $\widetilde{\gamma}$. Similarly, each of the surfaces in $\widetilde{M}_{L'}$ above is sent within a bounded Hausdorff distance of a pleated surface covering $\check{\Sigma}_e^h \rightarrow M_{L'}^h$ or $\Sigma_{f\pm}^h \rightarrow M_{L'}^h$. The latter pleated surfaces are pleated hyperbolic planes with (quasi-)circles at infinity in $\partial\widetilde{M}_{L'}^h$ that are images of the circles in $\partial\widetilde{M}_{L'}$. The complement of the union of these circles is the image of Ω , and hence has $n+2$ components. Furthermore, since the endpoints of $\widetilde{\gamma}^h$ are the same as those of the image of $\widetilde{\gamma}$, these endpoints are in different components.




FIGURE 7. Illustration of the chain of circles in $\widetilde{M}_{L'}$ and quasi-circles in $\widetilde{M}_{L'}^h$.

Next, observe that pleated surface covering $\check{\Sigma}_e^h$ is a pleated ideal hyperbolic n -gon with ideal endpoints at the intersection points of the circles. The sides of the n -gon are geodesics in the pleated surfaces that cover the face surfaces. In particular, the union of these $n + 1$ pleated surfaces separate the $n + 2$ regions at infinity in $\widetilde{M}_{L'}^h$, and thus $\tilde{\gamma}^h$ must intersect at least one of them. If it intersects one of the surfaces covering a face surface, then γ^h intersects that face surface, and so comes within distance 0 of it, proving the lemma in this case. Otherwise, γ^h intersects the pleated cut edge surface. Since this is a pleated ideal n -gon with $n \leq 20$, its area is at most 18π . Therefore, the minimal distance ρ from a point of intersection of γ^h with $\check{\Sigma}_e^h$ to a side of that polygon satisfies

$$\text{Area}(B_h(\rho)) = 4\pi \sinh^2(\rho/2) \leq 18\pi,$$

where $B_h(\rho)$ is a hyperbolic ball of radius ρ . Consequently, $\rho \leq 2 \sinh^{-1}(3/\sqrt{2}) < 3$. Therefore, γ^h comes within distance 3 of some $\Sigma_{f\pm}^h$ for $f \in \mathcal{F}(e)$, as required. $\square$

Proof of Proposition 4.1. Recall that we have fixed $R_0 > 0$, and want to show that $\text{sys}_e^*(L') \geq R_0$, provided $\text{combsys}(\tau)$ is sufficiently large, where $L' = L'(\tau)$. How large we must take $\text{combsys}(\tau)$ will depend on other constants, which we now choose. Without loss of generality, we assume $R_0 > 3$.

Let $\delta_0 > 0$ denote any upper bound on the diameter of the ϵ_0 -thick part of a hyperbolic 4-punctured sphere (which exists by an area argument). Next, let $\epsilon < \epsilon_0$ be any positive constant so that $R(\epsilon) > R_0 + \epsilon_0 + \delta_0$. Similar to Lemma 2.3, the results of [BM82, Mey87] imply that there is some positive $\epsilon' < \epsilon$, so a point in any hyperbolic 3-manifold which is within distance 3 of the ϵ -thick part is in the ϵ' -thick part. We also let $\delta > 0$ denote any upper bound on the diameter of the ϵ' -thick part of a 4-punctured sphere.

Suppose $K = R_0 + 3 + \delta$ and let $\eta_0 \geq 2$ be a natural number as in Lemma 2.2 for this value of K and ϵ . Suppose τ is a Cooper-Thurston triangulation with

$$\left\lfloor \frac{\text{combsys}(\tau)}{16} \right\rfloor = \eta \geq \eta_0,$$

and let $L' = L'(\tau)$ be the associated link. We suppose $\text{sys}_e^*(L') \leq R_0$, and derive a contradiction. This will prove the proposition.

Let γ^h be a hyperbolic geodesic in $M_{L'}^h$ that realizes $\text{sys}_e^*(L')$, and γ in $M_{L'}$ the homotopic geodesic in the locally CAT(-1) metric. Let $T_1, \dots, T_\eta$ be the tetrahedra from Lemma 4.3 which γ intersects. By Lemma 4.4, γ has non-empty transverse intersection with some surfaces in $\Sigma_\tau^\circ(T_i)$, for all $i = 1, \dots, \eta$. By Lemmas 4.6 and 4.7, there exist pleated surfaces $\Sigma_{x_1}^h, \dots, \Sigma_{x_\eta}^h \rightarrow M_{L'}^h$ which comes within distance 3 of a point of γ^h , where one of the following holds:

- (1) x_i is a tetrahedron T'_i where $T_i \cap T'_i \neq \emptyset$ or
- (2) $x_i = f_i \pm$, and f_i is a face of some T'_i , where T'_i and T_i share an edge e_i .

In particular, in both cases, we have $d_\tau(T_i, T'_i) \leq 2$. Since $T_1, \dots, T_\eta$ are as in Lemma 4.3, the triangle inequality implies $d_\tau(T'_i, T'_j) \geq 4$, for all $i \neq j$. By considering the totally geodesic representatives of the given surfaces in $M_{L'}$, we see that the only loops in any one of these surfaces that can be freely homotoped into another are the peripheral loops in the cusps of face surfaces, for cusps corresponding to the punctures coming from the focal component $K' \subset L'$; see Definition 3. We call these the *focal cusps* of the pleated surfaces, for cusps corresponding to the punctures coming from the focal component $K' \subset L'$; see Definition 3. We call these the *focal cusps* of the pleated surfaces. We let $z_i \in \Sigma_{x_i}^h$ and $y_i \in \gamma^h$ be points with $d(y_i, z_i) \leq 3$, for all $i = 1, 2, \dots, \eta$.

We divide the proof into two cases, depending on whether or not γ^h meets the ϵ -thin part of $M_{L'}^h$. First, suppose γ^h meets the ϵ -thin part. Since we assumed that $\ell(\gamma^h) \leq R_0$, Lemma 4.2 implies γ^h is contained in the ϵ -thin part, and since γ^h realizes the essential systole, it is the core geodesic of a ϵ -thin Margulis tube, $\mathbb{T}_\epsilon$. Since $d(y_i, z_i) \leq 3 < R(\epsilon)$, Lemma 2.3 implies that $z_i \in \Sigma_{x_i}^h$ is in the component $\mathbb{T}_{\epsilon_0}$ of the ϵ_0 -thin part of $M_{L'}^h$ with $\mathbb{T}_\epsilon \subset \mathbb{T}_{\epsilon_0}$. Within distance δ_0 of z_i is a point in some ϵ_0 -thin part, hence some loop $\alpha_i \subset \Sigma_{x_i}^h$, of length at most $2\epsilon_0$. Every point on α_i is within distance $3 + \delta_0 + \epsilon_0 < R(\epsilon)$ of y_i . Consequently, α_i is contained in $\mathbb{T}_{\epsilon_0}$, and is thus homotopic to a power γ^h (the core of $\mathbb{T}_{\epsilon_0}$), for all $i = 1, \dots, \eta$. This is a contradiction since the only nontrivial loops in distinct pleated surfaces in the set $\Sigma_{x_1}^h, \dots, \Sigma_{x_\eta}^h$ that can be homotopic to each other are loops in focal cusps. Therefore, we can assume that γ^h is completely contained in the ϵ -thick part of $M_{L'}^h$.

Since z_i is distance at most 3 from y_i , it follows from our choice of constants that z_i is in the ϵ' -thick part of $M_{L'}^h$, and hence the ϵ' -thick part of $\Sigma_{x_i}^h$. Consequently, z_i has distance at most δ from some point w_i in a component of the ϵ' -thin part of $\Sigma_{x_i}^h$, which we assume (as we may) is not a focal cusp. In particular, w_i is contained in a component, $\mathbb{T}_\epsilon^i$, of the ϵ -thin part of $M_{L'}^h$. Since these are not focal cusp thin parts, $\mathbb{T}_\epsilon^1, \dots, \mathbb{T}_\epsilon^\eta$ are all distinct thin parts. The distance y_i to w_i is at most $3 + \delta$, and so picking some point $y_0 \in \gamma^h$, we see that $d(y_0, w_i) \leq R_0 + 3 + \delta = K$. Choosing paths from y_0 to w_i whose length is at most K , we can lift these to paths from a point $\tilde{y}_0 \in \mathbb{H}^3$ to points $\tilde{w}_1, \dots, \tilde{w}_\eta \in \mathbb{H}^3$. The points $\tilde{w}_1, \dots, \tilde{w}_\eta$ all lie within distance K of $\tilde{y}_0$ and are in distinct ϵ -thin parts for $\Gamma = \pi_1(M_{L'}^h)$. By Lemma 2.2, $\eta < \eta_0$, which is a contradiction. Since both cases are impossible, it must be that γ^h has length greater than R_0 , as required. $\square$

5. FILLING KNOTS

In this final section, we prove Theorem 3.9 about the knots $K(n, \tau)$ (see §3.5) which in particular proves the existence of filling knots (Theorem 1.4). Like in

[LW25], we apply Theorem 2.5, though we cannot directly appeal to Theorem 3.8, since the theorem guarantees large essential *geodesic* systole, rather than simply large essential systole. In particular, the action of $\pi_1(M \setminus K(n, \tau))$ on hyperbolic 3-space may be parabolic on elements that are non-trivial in $\pi_1(M)$. The action of $\pi_1(M \setminus K(n, \tau))$ on a hyperbolic space for which we apply Theorem 2.5 will be built from the locally CAT(-1) metric on $M \setminus L'(\tau)$ and a construction of Gromov and Thurston (see [BH96]). We begin by briefly recalling this construction.

5.1. Negatively curved Dehn filling. An important ingredient in our construction is the *Gromov-Thurston 2π -Theorem*, or more precisely, its proof, which we now briefly recall; see [BH96]. Given a horoball cusp $U \cong T^2 \times [0, \infty)$, we equip the boundary ∂U with its natural Euclidean metric. If α in $\partial U \cong T^2$ is a slope, we define its *length* to be the length of its geodesic representative. The *transverse distance* of α is the distance between the two boundary components of the Euclidean annulus $\partial U \setminus \alpha$.

If α is a slope on ∂U of length greater than 2π , then there is a negatively curved metric on a solid torus $V \cong S^1 \times D^2$ which is isometric to U in a neighborhood of the boundary, with the property that the slope of the meridian $\{*\} \times \partial D^2 \subset \partial V$ is mapped by this isometry to α in ∂U . Moreover, the metric is *maximally symmetric*: the “obvious” isometric action of $O(2) \times O(2)$ on $\partial V \cong S^1 \times \partial D^2$ extends to an isometric action on V . The fixed point sets of the first-factor-reflections define a foliation by totally geodesic meridian disks of V . The construction is quite flexible, and a minor modification of the proof of [Lei06, Lemma 4] implies the following.

Lemma 5.1. *There exists $r > 2\pi$ with the following property. If U is a standard horoball cusp and α is a slope on ∂U with length $\ell > r$, then there is a maximally symmetric Riemannian metric on a solid torus V with sectional curvature in $[-2, -1/2]$, with the following properties:*

- (1) *There is an isometry $h: \mathcal{N}_1(\partial U) \rightarrow \mathcal{N}_1(\partial V)$ of the 1-neighborhoods of the boundaries such that $h(\alpha)$ is a meridian in ∂V , and*
- (2) *there is a constant ξ , depending only on ℓ , so that the length of the core geodesic of V is ξ times the transverse distance of α .* $\square$

We will use Lemma 5.1 to carry out a *Gromov-Thurston negatively curved Dehn surgery* on a manifold N with a locally CAT($-\frac{1}{2}$), piecewise Riemannian (or hyperbolic) metric containing a hyperbolic cusp $U \subset N$. To do so, suppose we have a slope $\alpha \subset \partial U$ of length $\ell > r$, where r is the constant from Lemma 5.1. Take a solid torus V given by the lemma. Now we remove $U \setminus \mathcal{N}_1(\partial U)$ from N , and glue in V to the result via the isometry $h: \mathcal{N}_1(\partial U) \rightarrow \mathcal{N}_1(\partial V)$. The resulting piecewise Riemannian metric on the Dehn surgered manifold $N(\alpha)$ is again locally CAT($-\frac{1}{2}$) and agrees outside of V with the original piecewise Riemannian metric on N outside of U ; that is, $N(\alpha) \setminus V$ is isometric to $N \setminus U$ (by construction).

We will also need an orbifold version of the lemma and this Dehn surgery construction. For this, we consider a cusp U as before and a positive integer $q > 1$. Now suppose α is a slope on ∂U with length ℓ such that $q \cdot \ell > r$. Consider a q -fold cyclic cover of $\tilde{U} \rightarrow U$ corresponding to a homomorphism $\pi_1(U) \rightarrow \mathbb{Z}/q\mathbb{Z}$ for which α is sent to a generator. Note that α^q lifts to a slope $\tilde{\alpha}$ in $\partial \tilde{U}$ with length $q \cdot \ell > r$. Let $\tilde{V}$ be the solid torus from the lemma applied to $\tilde{U}$ and $\tilde{\alpha}$. Also let $\tilde{h}: \mathcal{N}_1(\partial \tilde{U}) \rightarrow \mathcal{N}_1(\partial \tilde{V})$ be the isometry.

The action of $\mathbb{Z}/q\mathbb{Z}$ on $\mathcal{U}$ with quotient $\tilde{U} \rightarrow U$ determines an action on $\tilde{V}$ via conjugation by h (and extending): indeed, $\tilde{\alpha}$ is invariant by the action, and $h(\tilde{\alpha})$ bounds a meridian disk in $\tilde{V}$. Let $\tilde{V} \rightarrow V$ denote the quotient by this action. The quotient V is also a solid torus *orbifold* with Riemannian metric with curvature in $[-2, -\frac{1}{2}]$ and orbifold locus the geodesic core curve for V , with local group $\mathbb{Z}/q\mathbb{Z}$ (see e.g. [Thu78, Chapter 13]). The map $\tilde{h}$ descends to an isometry $h: \mathcal{N}_1(\partial U) \rightarrow \mathcal{N}_1(\partial V)$. See Figure 8.




FIGURE 8. Metric near the orbifold locus. Here, $q = 3$, so the core curve of V is a singular curve of order 3.

As with the Gromov–Thurston negatively curved Dehn surgery above, given a locally $\text{CAT}(-\frac{1}{2})$, piecewise Riemannian manifold N containing a hyperbolic cusp U , we can perform Gromov–Thurston negatively curved q -orbifold surgery, for a slope α of length ℓ on ∂U , as long as $q \cdot \ell > r$. The result in this case is a locally $\text{CAT}(-\frac{1}{2})$, piecewise Riemannian orbifold, whose underlying space is $N(\alpha)$ and which has orbifold locus a geodesic core of the filling solid torus, with local group $\mathbb{Z}/q\mathbb{Z}$.

5.2. Constructing an appropriate action. Recall from Lemma 3.5 that, given any Cooper–Thurston triangulation τ , we have an associated link $L' = L'(\tau) \subset M$. The complement $M_{L'}$ is equipped with a complete, piecewise hyperbolic, locally $\text{CAT}(-1)$ metric. Furthermore, from (2), $M_{L'}$ is obtained by isometrically gluing copies of the model tangle complement $T_0 \setminus J_0$ (see §3.2) according to τ , with the connector $\mathbb{B}^f$ inserted between tetrahedra glued along the face f , for each $f \in \mathcal{F}_\tau$.

We also recall from Lemma 3.5 that the link L' contains a component $K' = K'(\tau) \subset L'$ that meets every tetrahedron $T \in \mathcal{T}_\tau$ in the union of the 6 arcs, and each connector $\mathbb{B}^f$ in a union of 3 arcs. We choose a fixed (pairwise disjoint) union of cusp neighborhoods $U_0 \subset T_0 \setminus J_0$ of the ends corresponding to these arc components, and let $U_T \subset T \setminus L'$ be the image under the isometry $T \setminus L' \cong T_0 \setminus J_0$. For each $f \in \mathcal{F}_\tau$, we similarly choose pairwise disjoint cusp neighborhoods of the ends of each connector $U^f \subset \mathbb{B}^f$, for the ends corresponding to the three arcs in $\mathbb{B}^f$. There are only 8 isometry types of connectors (one for each choice of whether the twist in the disk D_{fij} is performed or not in the construction from Lemma 3.5). We choose the cusps $U^f \subset \mathbb{B}^f$ to be such that the isometry to one of the “8 model connectors” sends U^f to some fixed cusp of the model. Finally, we make all these choices (in the models, independent of M and τ) so that upon gluing, we obtain a cusp neighborhood of the end of $M_{L'}$ corresponding to K' . We denote this cusp neighborhood $U_{K'} \subset M_{L'}$. For any $R > 0$, we let $U_{K'}^R \subset U_{K'}$ be the smaller cusp

whose distance to $\partial U_{K'}$ is R . Let μ denote the meridian on $\partial U_{K'}$, i.e. the isotopy class that bounds a disk in M . We will view μ as a slope on both $\partial U_{K'}$ and $\partial U_{K'}^R$, for any R . Observe that μ is the isotopy class of any component of intersection of $\partial U_{K'}$ with any face surface $\Sigma_{f\pm}^\circ$. Since $\Sigma_{f\pm}^\circ$ is totally geodesic, its intersection with $\partial U_{K'}$ consists of geodesic representatives of the isotopy class in the flat metric, and its length is the length of μ in $\partial U_{K'}$. In particular, this length, which we denote by ℓ_0 , is independent of M and τ . The length of μ in $\partial U_{K'}^R$ is also independent of M and τ , and in fact is just given by $e^{-R}\ell_0$, since $\partial U_{K'}^R$ is isometric to $\partial U_{K'}$, scaled by e^{-R} .

Lemma 5.2. *There exists a constant $C > 0$ with the following property. For any M and τ , the transverse distance of μ in $\partial U_{K'}^R$ is at least*

$$e^{-R}C|\mathcal{T}_\tau|.$$

Moreover, the length of α on $U_{K'}^R$ is $e^{-R}\ell_0$.

Proof. We already observe the last statement before the statement of the lemma, so we need only prove the first one.

Let ∂U_0 denote the horospherical part of the boundary of U_0 in $T_0 \setminus J_0$ (i.e. it does not include the part of ∂U_0 contained in $\partial T_0 \setminus J_0$). We observe that $\partial U_0 \cap (T_0 \setminus J_0)$ is the disjoint union of 6 isometric Euclidean annuli. Let $c > 0$ denote the distance between the boundary components inside any one of these annuli. Since $U_{K'}$ meets $T \setminus L'$ in a cusp isometric to U_0 , it follows that every tetrahedron contributes $6c$ to the transverse distance of μ in $\partial U_{K'}$. Setting $C = 6c$, and again noting that $\partial U_{K'}^R$ is isometric to $\partial U_{K'}$ scaled by e^{-R} , completes the proof. $\square$

Lemma 5.3. *Given $R \geq 0$, there exists $\varepsilon > 0$ with the following property. Suppose M is any closed 3-manifold and τ any Cooper–Thurston triangulation. Then for any two tetrahedra in $T, T' \in \mathcal{T}_\tau$, with $d_\tau(T, T') \geq 4$, we have*

$$d_{M_{L'} \setminus U_{K'}^R}(x, y) > \varepsilon$$

for all $x \in T \setminus (L' \cup U_{K'}^R)$ and $y \in T' \setminus (L' \cup U_{K'}^R)$.

That is, in $M_{L'} \setminus U_{K'}^R$, the distance between points in disjoint tetrahedra is at least ε .

Proof. Except for the end corresponding to K' , all other ends have cusp neighborhoods either contained in a single connector, or meeting a single connector and the two tetrahedra to which it is glued (note that these tetrahedra are adjacent). Thus, after removing $U_{K'}$, for any tetrahedron $T \in \mathcal{T}_\tau$, there is a $\varepsilon_T > 0$ so that the ε_T -neighborhood of $T \setminus (L' \cup U_{K'}^R)$ in $M_{L'} \setminus U_{K'}^R$ only meets adjacent tetrahedra and adjacent connectors. Since the link of any vertex of a Cooper–Thurston triangulation is isomorphic to one of five 2-complexes, and since there are only 8 isometry types of connectors, it follows that the isometry type of

$$\left(\bigcup_{d_\tau(T, T') \leq 2} T' \right) \setminus (L' \cup U_{K'}^R)$$

is one of finitely many types, independent of M and τ . In particular, we can choose ε_T depending only on the isometry type of the above submanifold. Doing so, there is a uniform lower bound $\varepsilon' > 0$ to such ε_T , and setting $\varepsilon = \varepsilon'/2$, proves the lemma. $\square$

A non-null-homotopic loop in $M_{L'}$ must pass through $\lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$ tetrahedra that are pairwise distance at least 8 apart, by Lemma 4.3. In particular, these tetrahedra are disjoint, and so such a curve which is contained in $M_{L'} \setminus U_{K'}^R$ must travel distance ε in passing from one tetrahedron to the next. Thus, combined with Lemma 5.3 we immediately find the following.

Corollary 5.4. *If γ is a closed curve in $M_{L'} \setminus U_{K'}^R$, which is non-null-homotopic in M , then its length in the locally CAT(-1) metric is bounded below by $\varepsilon \cdot \lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$.* $\square$

Given M , τ , and $n \in \mathbb{Z}$, we write $M_{K(n,\tau)} = M \setminus K(n,\tau)$ (see §3.5).

Corollary 5.5. *For any Cooper–Thurston triangulation τ of a closed 3-manifold M , and $R \geq 0$, there exists n_0 so that for all n with $|n| > n_0$, $M_{K(n,\tau)}$ admits a CAT($-\frac{1}{2}$) metric with the following properties.*

- (1) *There is a cusp neighborhood U_n^R in $M_{K(n,\tau)}$ and isometry $U_n^R \cong U_{K'}^R \subset M_{L'}$ for which the meridian μ_n is the image of μ under the isometry, and*
- (2) *if γ is a closed curve in $M_{K(n,\tau)} \setminus U_n^R$ which is non-null-homotopic in M , then the length of γ is greater than $\varepsilon \cdot \lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$.*

Note that the case $R = 0$ gives us U_n^0 is isometric to $U_{K'}^0 = U_{K'}$. We write $U_n = U_n^0$.

Proof. Choose pairwise disjoint cusps $W_1, \dots, W_k$ for the ends of $M_{L'}$, different than the one corresponding to K' , which are also disjoint from $U_{K'}$. Inside of each, we choose smaller cusps $W'_i \subset W_i$, for all i , whose distance to ∂W_i is at least $\varepsilon \cdot \lfloor \text{combsys}(\tau)/16 \rfloor$. Choose $n_0 > 0$ so that for all n with $|n| > n_0$ the slope $(1, n)$ on $\partial W'_i$ has length at least r , where r is the constant from Lemma 5.1. Now perform the $(1, n)$ negatively curved Dehn filling on each W'_i (as explained in the discussion following the Lemma 5.1). Suppose V_i is the filling solid torus from the lemma. The result is a locally CAT($-\frac{1}{2}$) metric on $M_{K(n,\tau)}$, which outside the union of V_i is isometric to $M_{L'}$, outside the union of the W'_i . In particular, since all the cusps $U_{K'}^R$ are disjoint from all V_i , the given isometry restricts to an isometry from a cusp neighborhood $U_n^R \subset M_{K(n,\tau)}$ to $U_{K'}^R$. Since $K(n,\tau)$ is the image of K' by the homeomorphism $M \rightarrow M$ as described in §3.5, the meridian μ_n in ∂U_n^R and μ are related by the isometry, proving the first claim.

Now suppose γ in $M_{K(n,\tau)} \setminus U_n^R$ is any curve which is non-null-homotopic in M . We observe that if γ meets any of the solid tori V_i , then since it cannot be contained in any larger solid torus containing V_i (since these are null homotopic in M), it must intersect ∂W_i to get to V_i . This requires distance at least $\varepsilon \cdot \lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$ by our choice of cusps W_i , and so γ must have length at least this number, as required. Otherwise, γ is disjoint from the filling solid tori, hence it is related by an isometry to a closed curve in $M_{L'} \setminus U_{K'}^R$. By Corollary 5.5, the length must be at least $\varepsilon \cdot \lfloor \frac{\text{combsys}(\tau)}{16} \rfloor$, completing the proof of the second part, hence the proof of the corollary. $\square$

We now describe the key proposition needed to prove the theorem.

Proposition 5.6. *Suppose M is closed orientable 3-manifold. Given $R > 0$, there exists τ and n_0 , such that for all n with $|n| > n_0$, there is an action of $\pi_1(M_{K(n,\tau)})$ on a CAT($-\frac{1}{2}$) space, X . Under this action, every element of $\pi_1(M_{K(n,\tau)})$ which is non-null-homotopic in M acts loxodromically with translation length at least R .*

Proof of Proposition 5.6. Given $R > 0$, let $\varepsilon > 0$ be as in Lemma 5.3. Let C and ℓ_0 be the constants from Lemma 5.2. Fix $q > 0$ large enough so that $q \cdot e^{-R}\ell_0 > r$, where $r > 0$ is the constant from Lemma 5.1. Let ξ be the constant from Lemma 5.1, depending on the length, $q \cdot e^{-R}\ell_0 > r$.

Suppose τ is any Cooper–Thurston triangulation of M with

$$\min \left\{ \xi e^{-R} C |\mathcal{T}_\tau|, \varepsilon \cdot \left\lfloor \frac{\text{combsys}(\tau)}{16} \right\rfloor \right\} > R$$

and suppose $n_0 = n_0(R, \tau)$ from Corollary 5.5. Let n be any integer with $|n| > n_0$ and $U_n^R \subset M_{K(n, \tau)}$ the cusp neighborhood from that corollary. Since U_n^R is isometric to $U_{K'}^R$ (also from that corollary), Lemma 5.2, implies that the transverse distance of μ_n is $e^{-R} C |\mathcal{T}_\tau|$ and the length of μ_n is $e^{-R} \ell_0$. Let V denote the q -orbifold solid torus for which there is an isometry $\mathcal{N}_1(\partial V) \cong \mathcal{N}_1(\partial U_n^R)$ for the slope μ_n as in Lemma 5.1. By the lemma, the length of the core geodesic, K_q , (i.e. the orbifold locus) of V given by

$$\xi e^{-R} C |\mathcal{T}_\tau| > R.$$

Now, we perform the negatively curved q -orbifold Gromov–Thurston surgery along μ_n on the cusp U_n^R of $M_{K(n, \tau)}$ with the solid torus V . For simplicity, we denote the result $M_{K(n, \tau)}(\mu_n)_q$, simply as $N_q \cong M_{K(n, \tau)}(\mu_n)_q$ for the remainder of the proof. Given a loop γ in $M_{K(n, \tau)}$, we can take a sequence of representatives infimizing the lengths, and extract a limit, γ^* in N_q (which exists since N_q is compact). We denote the length of γ^* by $l(\gamma) = l(\gamma^*)$.

Claim 5.7. *If γ is any closed loop in $M_{K(n, \tau)}$ which is non-null-homotopic in M , then $l(\gamma) > R$.*

Proof. Since $U_n^R \subset U_n$, it follows that $(U_n \setminus U_n^R) \cup V$ is also an orbifold solid torus we denote $V' \subset N_q$. If γ^* is contained in V' , then γ^* must be (a multiple of) the core geodesic K_q of V' (any other loop is null-homotopic in M). This is also the core of V , and thus has length greater than R as noted above. If γ^* is not contained in V' , but intersects V , then it intersects both $\partial V'$ and ∂V , but the distance between these tori is R , hence $l(\gamma) > R$ again. Therefore, we can assume that γ^* is disjoint from V . Then γ^* is contained in $N_q \setminus V$ which is isometric to $M_{K(n, \tau)} \setminus U_n^R$. By Corollary 5.5, $l(\gamma^*)$ is greater than $\varepsilon \cdot \left\lfloor \frac{\text{combsys}(\tau)}{16} \right\rfloor \geq R$, completing the proof of the claim. $\square$

Now observe that $M_{K(n, \tau)}$ is homeomorphic to $N_q \setminus K_q$. Now consider the universal cover of $N_q \setminus K_q$, equipped with the incomplete pull-back metric from its inclusion into N_q , and let X denote the metric completion. The action of $\pi_1(M_{K(n, \tau)}) = \pi_1(N_q \setminus K_q)$ on the universal cover extends to X . The quotient is precisely N_q . Since N_q is equipped with a locally $\text{CAT}(-\frac{1}{2})$ orbifold metric, the universal cover of $N_q \setminus K_q$ with the incomplete metric is locally $\text{CAT}(-\frac{1}{2})$. Its metric completion, X , is obtained by adding bi-infinite geodesics, and is locally $\text{CAT}(-\frac{1}{2})$ by [All13, Theorem 1.2]. By the Cartan–Hadamard Theorem, X is $\text{CAT}(-\frac{1}{2})$ (see [BH99, Theorem II.4.1]).

If $\gamma \in \pi_1(M_{K(n, \tau)})$ is non-null-homotopic in M , then γ acts loxodromically: indeed, the axis projects to γ^* . Consequently, the translation length is greater than R by the claim, as required. $\square$

With Proposition 5.6 in hand, Theorem 3.9 now follows easily from Theorem 2.5.

Proof of Theorem 3.9. Let M be any closed 3-manifold and $m > 0$, as in the statement of the theorem. Let $C = \max\{C(k) \mid k \leq m\}$, where $C(k)$ is the constant from Theorem 2.5.

By Proposition 1.2 of [BH99, Section III.H.1], any $\text{CAT}(-\frac{1}{2})$ space is δ -hyperbolic for some uniform $\delta > 0$. Choose any $R > \delta C$. By Theorem 5.6, there exists a Cooper–Thurston triangulation τ and $n \geq 0$, associated with a knot $K := K(n, \tau)$, such that $\pi_1(M_K)$ acts on a $\text{CAT}(-\frac{1}{2})$ space X , and every loop which is non-null-homotopic in M acts with translation length at least R in M_K .

Let $i : M_K \hookrightarrow M$ be the inclusion map. Consider a graph Γ such that $\text{rank}(\pi_1(\Gamma)) = k \leq m$ and a map

$$f : \Gamma \rightarrow M_K$$

with the property that $i \circ f : \Gamma \rightarrow M$ has full rank (see §3.5). Let $B = \langle h_1, \dots, h_k \rangle$ be a basis for $\pi_1(\Gamma)$. Write $g_i := f_*(h_i)$ for all $1 \leq i \leq k$. Observe that $f_*(\pi_1(\Gamma))$ acts on X by isometries, and since $i \circ f$ has full rank, $\langle g_1, \dots, g_k \rangle$ is a generating set for $f_*(\pi_1(\Gamma))$.

Suppose $\langle g_1, \dots, g_k \rangle$ is Nielsen equivalent to some $\langle g'_1, \dots, g'_k \rangle$, and

$$d_{\bar{X}}(g'_1.x, x) \leq \delta C \leq \delta C(k)$$

for some $x \in X$. Then $i_*(g'_1)$ is trivial in $\pi_1(M)$ by our choice of R . By Lemma 2.1 of [LW25], the Nielsen moves on $\langle g_1, \dots, g_k \rangle$ lift to Nielsen moves on $\langle h_1, \dots, h_k \rangle$, creating a new basis $\langle h'_1, \dots, h'_k \rangle$ on $\pi_1(\Gamma)$ with the property that $f_*(h'_i) = g'_i$ for all $1 \leq i \leq k$. But then

$$(i \circ f)_*(h'_1) = i_*(g'_1) = 0.$$

This means $(i \circ f)(\pi_1(\Gamma))$, which has generating set $i_*(g'_2), \dots, i_*(g'_k)$, which has fewer than k elements. This contradicts the assumption that $i \circ f$ was full rank.

By Theorem 2.5, $f_*(\pi_1(\Gamma))$ is free. By the Hopfian property for free groups, f_* is injective. Since f and $k \leq m$ were arbitrary, K is m -rank filling in M . $\square$

REFERENCES

- [Ada85] Colin C. Adams. Thrice-punctured spheres in hyperbolic 3-manifolds. *Trans. Amer. Math. Soc.*, 287(2):645–656, 1985. [10](#), [13](#)
- [Ada86] Colin C. Adams. Augmented alternating link complements are hyperbolic. In *Low-dimensional topology and Kleinian groups (Coventry/Durham, 1984)*, volume 112 of *London Math. Soc. Lecture Note Ser.*, pages 115–130. Cambridge Univ. Press, Cambridge, 1986. [3](#), [11](#), [14](#)
- [All13] Daniel Allcock. Completions, branched covers, Artin groups, and singularity theory. *Duke Math. J.*, 162(14):2645–2689, 2013. [26](#)
- [AR00] Colin C. Adams and Alan W. Reid. Systoles of hyperbolic 3-manifolds. *Math. Proc. Cambridge Philos. Soc.*, 128(1):103–110, 2000. [1](#)
- [AS05] Colin Adams and Eric Schoenfeld. Totally geodesic Seifert surfaces in hyperbolic knot and link complements. I. *Geom. Dedicata*, 116:237–247, 2005. [8](#), [10](#)
- [BH96] Steven A. Bleiler and Craig D. Hodgson. Spherical space forms and Dehn filling. *Topology*, 35(3):809–833, 1996. [22](#)
- [BH99] Martin R. Bridson and André Haefliger. *Metric spaces of non-positive curvature*, volume 319 of *Grundlehren der mathematischen Wissenschaften [Fundamental Principles of Mathematical Sciences]*. Springer-Verlag, Berlin, 1999. [3](#), [26](#), [27](#)
- [BM82] Robert Brooks and J. Peter Matelski. Collars in Kleinian groups. *Duke Math. J.*, 49(1):163–182, 1982. [4](#), [20](#)
- [BP92] Riccardo Benedetti and Carlo Petronio. *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992. [4](#), [5](#)
- [CT88] D. Cooper and W. P. Thurston. Triangulating 3-manifolds using 5 vertex link types. *Topology*, 27(1):23–25, 1988. [6](#)

- [FK23] Michael Freedman and Vyacheslav Krushkal. Filling links and spines in 3-manifolds. *Comm. Anal. Geom.*, 31(10):2307–2333, 2023. With appendix by Christopher J. Leininger and Alan W. Reid. [1](#), [2](#)
- [HIK⁺16] Neil Hoffman, Kazuhiro Ichihara, Masahide Kashiwagi, Hidetoshi Masai, Shin’ichi Oishi, and Akitoshi Takayasu. Verified computations for hyperbolic 3-manifolds. *Exp. Math.*, 25(1):66–78, 2016. [10](#)
- [KW03] Ilya Kapovich and Richard Weidmann. Nielsen methods and groups acting on hyperbolic spaces. *Geom. Dedicata*, 98:95–121, 2003. [1](#), [2](#), [5](#)
- [Lac04] Marc Lackenby. The volume of hyperbolic alternating link complements. *Proc. London Math. Soc. (3)*, 88(1):204–224, 2004. With an appendix by Ian Agol and Dylan Thurston. [3](#), [11](#), [14](#)
- [Lei06] Christopher J. Leininger. Small curvature surfaces in hyperbolic 3-manifolds. *J. Knot Theory Ramifications*, 15(3):379–411, 2006. [22](#)
- [LL14] Grant S. Lakeland and Christopher J. Leininger. Systoles and Dehn surgery for hyperbolic 3-manifolds. *Algebr. Geom. Topol.*, 14(3):1441–1460, 2014. [1](#)
- [LW25] Christopher J. Leininger and Yandi Wu. Filling links and essential systole, 2025. [1](#), [2](#), [3](#), [4](#), [5](#), [6](#), [7](#), [8](#), [15](#), [16](#), [17](#), [22](#), [27](#)
- [McM92] Curt McMullen. Riemann surfaces and the geometrization of 3-manifolds. *Bull. Amer. Math. Soc. (N.S.)*, 27(2):207–216, 1992. [3](#)
- [Mey87] Robert Meyerhoff. A lower bound for the volume of hyperbolic 3-manifolds. *Canad. J. Math.*, 39(5):1038–1056, 1987. [4](#), [20](#)
- [Mor84] John W. Morgan. On Thurston’s uniformization theorem for three-dimensional manifolds. In *The Smith conjecture (New York, 1979)*, volume 112 of *Pure Appl. Math.*, pages 37–125. Academic Press, Orlando, FL, 1984. [3](#)
- [Nie24] Jakob Nielsen. Die Isomorphismengruppe der freien Gruppen. *Math. Ann.*, 91(3-4):169–209, 1924. [5](#)
- [Pur10] Jessica S. Purcell. Hyperbolic geometry of multiply twisted knots. *Comm. Anal. Geom.*, 18(1):101–120, 2010. [3](#), [11](#), [14](#)
- [Thu78] William P. Thurston. The geometry and topology of 3-manifolds, 1978. [3](#), [4](#), [5](#), [10](#), [23](#)
- [Thu82] William P. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982. [3](#)
- [Thu86] William P. Thurston. Hyperbolic structures on 3-manifolds. I. Deformation of acylindrical manifolds. *Ann. of Math. (2)*, 124(2):203–246, 1986. [3](#)
- [Whi02] Matthew E. White. Injectivity radius and fundamental groups of hyperbolic 3-manifolds. *Comm. Anal. Geom.*, 10(2):377–395, 2002. [1](#)

UNIVERSITY OF NOTRE DAME

Email address: cleining@nd.edu

URL: <https://sites.google.com/view/chris-leiningers-webpage/home>

RICE UNIVERSITY

Email address: yw220@rice.edu

URL: <https://yandiwu.github.io/>