

UNIFORM MODELS AND SHORT CURVES FOR RANDOM 3-MANIFOLDS

PETER FELLER, ALESSANDRO SISTO, AND GABRIELE VIAGGI

ABSTRACT. We provide two constructions of hyperbolic metrics on 3-manifolds with Heegaard splittings that satisfy certain topological conditions, which both apply to random Heegaard splitting with asymptotic probability 1.

These constructions provide a lot of control on the resulting metric, allowing us to prove various results about the coarse growth rate of geometric invariants, such as diameter and injectivity radius, and about arithmeticity and commensurability in families of random 3-manifolds. For example, we show that the diameter of a random Heegaard splitting grows coarsely linearly in the length of the associated random walk.

The constructions only use tools from the deformation theory of Kleinian groups, that is, we do not rely on the solution of the Geometrization Conjecture by Perelman. In particular, we give a proof of Maher's result that random 3-manifolds are hyperbolic that bypasses Geometrization.

1. INTRODUCTION

Every closed orientable 3-manifold M can be presented as a *Heegaard splitting*. This means that M is diffeomorphic to a 3-manifold M_f obtained by gluing together two handlebodies (taking the second one with opposite orientation) of the same genus H_g along an orientation preserving diffeomorphism $f \in \text{Diff}^+(\Sigma)$ of their boundary $\Sigma := \partial H_g$

$$M_f = H_g \cup_{f: \partial H_g \rightarrow \partial H_g} H_g.$$

For a fixed genus g , however, not all 3-manifolds can arise. In this article, we restrict our attention to the family of those 3-manifolds that can be described as Heegaard splittings of a fixed genus $g \geq 2$.

The problem of finding hyperbolic metrics on “most” 3-manifolds with a splitting of a fixed genus $g \geq 2$ was originally raised by Thurston (as Problem 24 in [Thu82]) and made more precise by Dunfield and Thurston (see Conjecture 2.11 of [DT06]) by introducing the notion of random Heegaard splittings.

This notion is based on the observation that the diffeomorphism type of M_f only depends on the isotopy class of the gluing map f , so it is well-defined for elements in the *mapping class group*

$$[f] \in \text{Mod}(\Sigma) := \text{Diff}^+(\Sigma)/\text{Diff}_0(\Sigma).$$

Therefore, Heegaard splittings of genus $g \geq 2$ are naturally parametrized by mapping classes $[f] \in \text{Mod}(\Sigma)$.

2020 *Mathematics Subject Classification*. 57K30, 57K32, 30F60, 20P05.

Key words and phrases. Heegaard splittings, random 3-manifolds, hyperbolization, short curves, model metrics.

A family $(M_n)_{n \in \mathbb{N}}$ of random Heegaard splittings of genus $g \geq 2$, or *random 3-manifolds*, is one of the form $M_n = M_{f_n}$ where $(f_n)_{n \in \mathbb{N}}$ is a random walk on the mapping class group $\text{Mod}(\Sigma)$ driven by some initial probability measure μ whose finite support generates $\text{Mod}(\Sigma)$. If $(f_n)_{n \in \mathbb{N}}$ is such a random walk, we will denote by \mathbb{P}_n the distribution of the n -th step f_n and by \mathbb{P} the distribution of the path $(f_n)_{n \in \mathbb{N}}$.

Exploiting work of Hempel [Hem01] and the solution of the Geometrization Conjecture by Perelman, Maher showed in [Mah10b] that a random Heegaard splitting of genus $g \geq 2$ admits a hyperbolic metric, thus answering Dunfield and Thurston's conjecture.

The main goal of this article is to provide a constructive and effective approach to the hyperbolization of random 3-manifolds. Before describing our main contributions (Theorems 5 and 6), we note that this approach yields a proof of Maher's result that does not rely on Perelman's work, which we can informally state as follows.

Theorem 1. *There is a Ricci flow free hyperbolization for random 3-manifolds.*

By Ricci flow free hyperbolization we mean that we construct the hyperbolic metric only using tools from the deformation theory of Kleinian groups. More specifically, the main tools we use are the model manifold technology by Minsky [Min10] and Brock, Canary and Minsky [BCM12], as well as the effective version of Thurston's Hyperbolic Dehn Surgery by Hodgson and Kerckhoff [HK05].

We develop two different approaches to Theorem 1, both bringing new and more refined information than the mere existence of a hyperbolic metric.

One approach provides a so-called model metric that captures, up to uniform bilipschitz distortion, the geometry of 3-manifolds admitting suitable Heegaard splitting and allows the computation of its geometric invariants. The hyperbolic metric is constructed explicitly by gluing elementary building blocks. The other approach provides a purely topological criterion for a Heegaard splitting to admit a hyperbolic metric for which a specific simple closed curve on the Heegaard surface is short. Both criteria apply to a large class of Heegaard splittings, not just random splittings.

We describe these two approaches and our main results (Theorem 5 and Theorem 6) in the next two subsections. But first we list some applications to random 3-manifolds, which will be discussed in more detail later.

The first two results describe the behaviour of geometric invariants, namely diameter and injectivity radius.

Theorem 2. *There exists $c > 0$ such that*

$$\mathbb{P}_n[M_f \text{ is hyperbolic and } \text{diam}(M_f) \in [n/c, cn]] \xrightarrow{n \rightarrow \infty} 1.$$

Theorem 3. *There exists $c > 0$ such that*

$$\mathbb{P}_n[M_f \text{ is hyperbolic and } \text{inj}(M_f) \leq c/\log(n)^2] \xrightarrow{n \rightarrow \infty} 1.$$

Our methods would also allow us to obtain a result similar to Theorem 2 for the volume, but this (and more) is known already in that case, as we discuss below.

As described in [ST19], $1/\log(n)^2$ is exactly the coarse decay rate for the length of the shortest curve in random mapping tori. Our methods, however, only give an upper bound in the case of random Heegaard splittings.

The third application has a more algebraic flavor, showing that random 3-manifolds are not arithmetic and belong to multiple commensurability classes.

Theorem 4. *With asymptotic probability 1 the following holds*

- (1) M_f is not arithmetic.
- (2) M_f is not in a fixed commensurability class \mathcal{R} .

Uniform bilipschitz models for random 3-manifolds. The notion of model manifold that we use is similar to the ones considered by Brock, Minsky, Namazi and Souto in [Nam05, NS09, BMNS16]. A Riemannian metric (M_f, ρ) is a ε -model metric for $\varepsilon < 1/2$ if there is a decomposition into five pieces $M_f = H_1 \cup \Omega_1 \cup Q \cup \Omega_2 \cup H_2$ satisfying the three requirements

- (1) H_1 and H_2 are homeomorphic to genus g handlebodies while Q, Ω_1 and Ω_2 are homeomorphic to $\Sigma \times [0, 1]$.
- (2) ρ has negative curvature $\text{sec} \in (-1 - \varepsilon, -1 + \varepsilon)$, but outside the region $\Omega = \Omega_1 \cup \Omega_2$ the metric is purely hyperbolic, i.e. $\text{sec} = -1$.
- (3) The piece Q is almost isometrically embeddable in a complete hyperbolic 3-manifold diffeomorphic to $\Sigma \times \mathbb{R}$.

The importance of the last requirement is due to the fact that we understand explicitly hyperbolic 3-manifolds diffeomorphic to $\Sigma \times \mathbb{R}$ thanks to the work of Minsky [Min10] and Brock, Canary and Minsky [BCM12].

The following is a more precise version of Theorem 1.

Theorem 5. *For every $0 < \varepsilon < 1/2$ and $K > 1$ we have*

$$\mathbb{P}_n[M_f \text{ has a hyperbolic metric } K\text{-bilipschitz to a } \varepsilon\text{-model metric}] \xrightarrow{n \rightarrow \infty} 1.$$

We remark that ε -model metrics on random Heegaard splittings, similar to the ones that we build here, are constructed in [HV22]. There, the existence of a underlying hyperbolic metric is guaranteed by Maher's result, and it is unclear whether the ε -model metrics are uniformly bilipschitz to it.

As an aside, we also mention that, using an unpublished result by Tian [Tia], the mere fact that a metric ρ is a ε -model metric and that the regions Ω_1, Ω_2 where it is not hyperbolic have uniformly bounded diameter (as follows from [HV22]), implies, if $\varepsilon > 0$ is sufficiently small, that ρ is uniformly close up to third derivatives to a hyperbolic metric. However, we do not rely on Tian's result. Instead, in order to provide a uniform bilipschitz control, we exploit ergodic properties of the random walk and drilling and filling theorems by Hodgson and Kerckhoff [HK05] and Brock and Bromberg [BB04].

Our methods follow closely [BMNS16] and [BD15] where uniform ε -model metrics are constructed for special classes of 3-manifolds.

The idea is the following. We can obtain a hyperbolic metric on M_f by Hodgson and Kerckhoff's effective Hyperbolic Dehn Surgery [HK05] from a complete finite volume hyperbolic metric on a *drilled* manifold \mathbb{M} which has the following form. Let $\Sigma \times [1, 4]$ be a tubular neighborhood of $\Sigma \subset M_f$. We consider 3-manifolds

diffeomorphic to

$$\mathbb{M} = M_f - (P_1 \times \{1\} \cup P_2 \times \{2\} \cup P_3 \times \{3\} \cup P_4 \times \{4\})$$

where P_j is a pants decomposition of the surface $\Sigma \times \{j\}$. A finite volume hyperbolic metric on such a manifold can be constructed explicitly by gluing together the convex cores of two maximally cusped handlebodies H_1, H_2 and three maximally cusped I-bundles Ω_1, Q, Ω_2 .

$$\mathbb{M} = H_1 \cup \Omega_1 \cup Q \cup \Omega_2 \cup H_2.$$

Most of our work consists of finding suitable pants decompositions for which the Dehn surgery slopes needed to pass from \mathbb{M} to M_f satisfy the assumptions of the effective Hyperbolic Dehn Surgery Theorem [HK05]. In order to find them we crucially need two major tools. On one hand an explicit control on the geometry of hyperbolic handlebodies similar to the one obtained in [HV22], and on the other hand ergodic properties of the random walks proved by Baik, Gekhtman, and Hamenstädt [BGH20].

Short curves via knots on Heegaard surfaces. We now discuss short closed geodesics in hyperbolic Heegaard splittings and random 3-manifolds. We identify purely topological conditions on a simple closed curve on the Heegaard surface $\gamma \subset \Sigma$ that ensure that M_f has a hyperbolic metric and γ is a very short geodesic in it.

Let \mathcal{D} and $f\mathcal{D}$ be the *disk sets* of the Heegaard surface $\Sigma \subseteq M_f$, i.e. \mathcal{D} and $f\mathcal{D}$ are the subsets of the curve graph \mathcal{C} of the Heegaard surface Σ given by the essential simple closed curves $\delta \subset \Sigma$ that compress in the first and second handlebody of the Heegaard surface, respectively. For a subsurface $W \subset \Sigma$, we denote by $d_W(\mathcal{D}, f\mathcal{D})$ the distance in the curve graph of W of the subsurface projections of \mathcal{D} and $f\mathcal{D}$.

Theorem 6. *Let $\Sigma := \partial H_g$ for a fixed $g \geq 2$. There exists a constant $C_\Sigma > 0$ such that the following holds. Let $\gamma \subset \Sigma$ be a non-separating simple closed curve with complement $W := \Sigma - \gamma$. Let $f \in \text{Mod}(\Sigma)$ be a mapping class. Suppose that*

- (a) *Both (H_g, γ) and $(H_g, f(\gamma))$ are pared acylindrical handlebodies.*
- (b) *We have a large subsurface projection $d_W(\mathcal{D}, f\mathcal{D}) \geq C_\Sigma$.*

Then M_f has a hyperbolic metric. Moreover, the length of γ in M_f is bounded by

$$\ell_{M_f}(\gamma) \leq C_\Sigma / d_W(\mathcal{D}, f\mathcal{D}).$$

We recall that, informally speaking, the pair (H_g, γ) is called a *pared acylindrical handlebody* if $\Sigma - \gamma$ is incompressible and there are no non-trivial essential cylinders in (H_g, γ) and $(H_g, \Sigma - \gamma)$. These objects arise naturally in the study of cusped hyperbolic metrics on H_g (see Thurston [Thu86a]). Many pairs (H_g, γ) have this property. For example, if $\gamma \subset \Sigma$ satisfies $d_{\mathcal{C}}(\gamma, \mathcal{D}) \geq 3$, then (H_g, γ) is pared acylindrical. As the disk set \mathcal{D} is a small quasi-convex subset of \mathcal{C} by Masur and Minsky [MM04], non-separating curves $\gamma \in \mathcal{C}$ that are far from \mathcal{D} are abundant.

Theorem 6 builds upon the groundbreaking work of Minsky [Min10] on hyperbolic metrics on $\Sigma \times \mathbb{R}$. In that setting, the collection of simple closed curves on Σ that are isotopic to very short closed geodesics on a hyperbolic metrics on $\Sigma \times \mathbb{R}$ can be read off the list of subsurface coefficients associated to the end invariants of such a metric. To some extent, in a complicated Heegaard splitting M_f , the role of the end invariants can be, as a first approximation, replaced by the disk

sets \mathcal{D} and $f\mathcal{D}$ of the splitting. For sufficiently complicated hyperbolic Heegaard splittings we have the following conjectural description of length of curves. A curve $\gamma \subset \Sigma$ is isotopic to a short geodesic if and only if it lies on the boundary $\gamma \subset \partial W$ of a proper essential subsurface $W \subset \Sigma$ where the disk sets \mathcal{D} and $f\mathcal{D}$ have large subsurface projection $d_W(\mathcal{D}, f\mathcal{D})$. We understand Theorem 6 as a first step in the direction of making precise this conjectural description.

The idea for the proof of Theorem 6 is the following. We associate to $\gamma \subset \Sigma$ the 3-manifold $M_f - \gamma$. Notice that it decomposes as

$$M_f - \gamma = (H_g - \gamma) \cup (H_g - f(\gamma)).$$

If both (H_g, γ) and $(H_g, f(\gamma))$ are pared acylindrical handlebodies, then the JSJ theory tells us that the complement $M_f - \gamma$ is irreducible and atoroidal. Moreover it is also Haken since we can choose $\Sigma - \gamma \subset M_f - \gamma$ as a Haken surface. Hence, by Thurston's Hyperbolization Theorem, the 3-manifold $M_f - \gamma$ admits a complete finite volume hyperbolic metric.

As in the proof of Theorem 5, we find a hyperbolic metric on M_f for which γ is a short curve provided that the filling slope in the complete finite volume hyperbolic metric on $M_f - \gamma$ satisfies the assumptions of the effective Hyperbolic Dehn Surgery Theorem [HK05]. To this extent we argue that the size of a standard torus horosection of the cusp of $M_f - \gamma$ is comparable with the subsurface projection $d_W(\mathcal{D}, f\mathcal{D})$. In order to check this we use tools from the model manifold technology by Minsky [Min10] and Brock, Canary and Minsky [BCM12].

Ergodic properties of random walks on the mapping class group imply that the condition of large subsurface projection of \mathcal{D} and $f_n\mathcal{D}$ on the complement W_n of some non-separating curves $\gamma_n \subset \Sigma$ holds with asymptotic probability one, that is, with $\mathbb{P}_n \rightarrow 1$. Thus, Theorem 6 applies to random 3-manifolds and gives a proof of Theorem 1 that does not rely on Perelman's work.

We briefly comment on the assumption (a) of Theorem 6. Its use is twofold: It implies that the complement $M_f - \gamma$ is hyperbolizable and it also implies that the inclusions $\Sigma - \gamma, \Sigma - f(\gamma) \subset H_g$ are doubly incompressible (as defined by Thurston [Thu86a]). The latter plays a central role in the proof via Thurston's Uniform Injectivity [Thu86a]. As mentioned above, since \mathcal{D} and $f\mathcal{D}$ are a small quasi-convex subsets of the curve graph, there are plenty of non-separating curves that satisfy $d_{\mathcal{C}}(\gamma, \mathcal{D} \cup f\mathcal{D}) \geq 3$ and, hence, condition (a).

In a direction opposite to Theorem 6, one can ask whether every very short curve on a complicated hyperbolic Heegaard splitting arises from a Hyperbolic Dehn Surgery on a complete hyperbolic manifold of the form $M_f - \gamma$ with $\gamma \subset \Sigma$ a simple closed curve. This is the case for strongly irreducible hyperbolic Heegaard splittings M_f as proved by Souto [Sou08] and Breslin [Bre11]. They show that there is a constant $\varepsilon_\Sigma > 0$ such that, every closed geodesic of length at most ε_Σ in M_f is isotopic to a simple closed curve γ on the Heegaard surface $\Sigma \subset M_f$.

On the applications. Now that we discussed our constructions, we can explain how we obtain the applications that we listed earlier, i.e. Theorems 2, 3, and 4.

We exploit the geometric control given by the ε -model metric to compute the coarse growth or decay rate of the geometric invariants along the family $(M_{f_n})_{n \in \mathbb{N}}$ (see also [Riv14, Section 11.4]). Our general strategy is the following. We use the model manifold technology [Min10], [BCM12] and compute the geometric invariants

for the middle piece Q of the ε -model metric. Then, we argue that the invariants of Q are uniformly comparable with those of M_f . In this way, Theorem 5 allows for a uniform approach to several results.

For example, combined with a result of Brock [Bro03], Theorem 5 allows the computation of the coarse growth rate of the volume, which is well-known to be linear as explained in [Mah10b] (see also [HV22]). In fact, there is even a law of large numbers for the volume [Via21]. Combined with results of Baik, Gekhtman and Hamenstädt [BGH20] Theorem 5 shows that the smallest positive eigenvalue of the Laplacian behaves like $1/n^2$ as computed in [HV22]. We do not carry out those computations because they are already well established.

To control the diameter of random 3-manifolds (Theorem 2), the ingredients of the proof are again Theorem 5 and a result by White [Whi01].

Instead, to control the injectivity radius (Theorem 3), we use work of the second author and Taylor [ST19].

Finally, the proof of Theorem 4 about arithmeticity and commensurability classes combines a study of geometric limits of random 3-manifolds, see Proposition 8.3, with arguments by Biringer and Souto [BS11].

Overview. The paper is organized as follows. Section 2 introduces some ingredients and tools that we need in our constructions. The rest of the article is divided into three parts as follows. Part 1: Short curves on Heegaard splittings (Sections 3 and 4). Part 2: The construction of the model metric (Sections 5 and 6). Part 3: The application to random 3-manifolds (Section 7) and the computation of the coarse growth rate of geometric invariants (Section 8). Next, we briefly describe the content of each section.

In Section 3 we describe the topological part of the proof of Theorem 6, in particular, we check that $M_f - \gamma$ is hyperbolizable. The geometric part is developed, instead, in Section 4 where we use the model manifold technology to check that the cusp of $M_f - \gamma$ has a large horosection.

In Section 5 we outline the construction of the ε -model metric. In Section 6 we build many examples to which the model metric construction applies.

In Section 7 we prove Theorems 1 and 5 by showing that the examples of Section 6 and the ones described by Theorem 6 are generic from the point of view of a random walk. Lastly, in Section 8 we prove Theorems 2, 3 and 4.

Acknowledgements. Part of this work is contained in the PhD thesis of GV. He wishes to thank Ursula Hamenstädt for her help and support. He also acknowledges the financial support of the Max Planck Institute for Mathematics of Bonn and partial funding by the DFG 427903332.

PF and AS want to follow up on GV's acknowledgement above. A previous arxiv version [Via19] of this article contained GV's path to Theorem 1, a result PF and AS had independently announced in [FMST18]. PF and AS want to acknowledge that the GV had independently pursued and worked out a complete proof of Theorem 1 via Theorem 5 as part of his impressive PhD-work. They are most thankful for GV's insight and flexibility in merging the two works into a coherent work, which exceeds the sum of its parts and which, in particular, includes further applications we would not have found separately.

We thank an anonymous referee for feedback that led to substantial reworking and improvement of Section 6.

PF and AS gratefully acknowledge support by the Swiss National Science Foundation (grants #181199 and #182186, respectively).

2. PRELIMINARIES

In this section we review some of the main ingredients used in one or both of the two constructions that we are going to describe.

In both constructions, geometrically finite hyperbolic structures on handlebodies H_g and I-bundles $\Sigma \times [0, 1]$ play an important role: They appear as building blocks (in Sections 5 and 6) or coverings (in Sections 3 and 4) of our drilled Heegaard splittings M_f . We introduce in this section some terminology and facts from the deformation theory of hyperbolic metrics on 3-manifolds (as discussed, for example, in Chapter 7 of [CM04]).

The curve graph \mathcal{C} and the disk set $\mathcal{D} \subset \mathcal{C}$ are also objects that appear in both constructions. By fundamental work of Minsky [Min10] and Brock, Canary and Minsky [BCM12], the geometry and combinatorics of these graphs captures the internal geometry of hyperbolic manifolds diffeomorphic to $\Sigma \times \mathbb{R}$. The combinatorics of \mathcal{C} also carries a lot of information on the topology of Heegaard splittings M_f : For example, by work of Hempel [Hem01], if the distance $d_{\mathcal{C}}(\mathcal{D}, f\mathcal{D})$ is at least 3, then the splitting M_f satisfies the topological assumptions of the Geometrization Conjecture. In this section we will briefly describe both aspects of these graphs.

The last common ingredient that we present is the Hodgson and Kerckhoff [HK05] deformation theory of hyperbolic 3-manifolds M with cone singularities along geodesic links $\Gamma \subset M$. The Drilling Theorem of Brock and Bromberg [BB04] allows to quantitatively keep under control the change in the hyperbolic metric along the deformation.

2.1. Hyperbolic 3-manifolds. Let M be a connected oriented 3-manifold (without boundary). A *hyperbolic metric* on M is a complete Riemannian metric of constant sectional curvature -1 . A *hyperbolic structure* on M is a pair (M, ρ) , where ρ is a complete hyperbolic metric. Every hyperbolic structure is isometric to a quotient \mathbb{H}^3/Γ of the hyperbolic 3-space \mathbb{H}^3 by a discrete torsion free subgroup $\Gamma < \text{Isom}^+(\mathbb{H}^3)$. We refer to a connected oriented 3-manifold that admits a hyperbolic structure as a *hyperbolic manifold*. Hyperbolic structures of finite volume are unique, up to homotopy, by the Mostow–Prasad rigidity theorem. Hyperbolic manifolds with totally geodesic boundary can also be defined and will play a role in Part 2.

The group Γ acts on the boundary $\mathbb{CP}^1 = \partial\mathbb{H}^3$ by Möbius transformations. The action $\Gamma \curvearrowright \mathbb{CP}^1$ preserves the *limit set* $\Lambda \subset \mathbb{CP}^1$, which is the set of accumulation points of an orbit $\Gamma o \subset \mathbb{H}^3 \cup \partial\mathbb{H}^3$, and its complement $\Omega := \mathbb{CP}^1 - \Lambda$, the *domain of discontinuity*. The action $\Gamma \curvearrowright \Omega$ is free and properly discontinuous. The quotient $(\mathbb{H}^3 \cup \Omega)/\Gamma$ is a 3-manifold with boundary Ω/Γ and the boundary has a natural conformal structure. If Γ is finitely generated and non-abelian, then, by Ahlfors Finiteness Theorem [Ahl64], the quotient Ω/Γ is a surface of finite type. In particular, Ω/Γ admits a *Poincaré metric* which is the unique complete finite area hyperbolic metric in the same conformal class.

2.2. Geometrically finite structures. We describe some flexible classes of geometrically finite hyperbolic structures on a pair of fixed topological spaces, namely handlebodies H_g and I-bundles $\Sigma \times [a, b]$. When dealing with such hyperbolic structures with cusps, the notion of pared (acylindrical) manifold naturally occurs (see Thurston [Thu86a], [Thu86b]).

DEFINITION (Pared Acylindrical). Let N be either H_g or $\Sigma \times [a, b]$. Set $A := S^1 \times [0, 1]$. Let $P \subset \partial N$ be a multicurve with regular neighborhood $U \subset \partial N$. We say that (N, P) is a *pared 3-manifold* provided that the following holds.

- (1) The fundamental group of each component of P injects in $\pi_1(N)$.
- (2) Every π_1 -injective homotopy $(A, \partial A) \rightarrow (N, U)$ is homotopic as a map of pairs to a map that sends A to U .

We say that (N, P) is also *acylindrical* if furthermore

- (3) Every π_1 -injective map $(A, \partial A) \rightarrow (N, \partial N - U)$ is homotopic as a map of pairs into $\partial N - U$.

REMARK 2.1. If (N, P) is pared acylindrical then the components of $\partial N - P$ are π_1 -injective in N , see [Thu86a, Section 7] and Appendix B.

We refer to Chapter 5 of [CM04] for the topological features of pared (acylindrical) manifolds. We can now give the following definition:

DEFINITION (Geometrically Finite). Let N be either H_g or $\Sigma \times [a, b]$. Let $P \subset \partial N$ be a (possibly empty) multicurve such that (N, P) is pared. A hyperbolic metric ρ on $M = \text{int}(N)$ is *geometrically finite* with parabolic locus $P \subset \partial N$ if there is an orientation preserving diffeomorphism $f : N - P \rightarrow (\mathbb{H}^3 \cup \Omega)/\Gamma$ that is isometric on $M = \text{int}(N)$. Such an identification induces a conformal structure on the boundary via $f : \partial N - P \rightarrow \Omega/\Gamma$.

If we have different identifications $f : N - P \rightarrow \mathbb{H}^3 \cup \Omega/\Gamma$ and $f' : N - P \rightarrow \mathbb{H}^3 \cup \Omega'/\Gamma'$, the composition $g = f'f^{-1} : \mathbb{H}^3 \cup \Omega/\Gamma \rightarrow \mathbb{H}^3 \cup \Omega'/\Gamma'$ restricts to an isometry $g : \mathbb{H}^3 \cup \Omega/\Gamma \rightarrow \mathbb{H}^3 \cup \Omega'/\Gamma'$ and, hence, the restriction to the boundary $g : \Omega/\Gamma \rightarrow \Omega'/\Gamma'$ is conformal. Thus, every geometrically finite hyperbolic metric induces a well-defined Riemann surface structure on $\partial N - P$, called the *conformal boundary*, which we can think as a point in the Teichmüller space of $\partial N - P$.

Two geometrically finite hyperbolic metrics ρ, ρ' on $M = \text{int}(N)$ are said to be *equivalent* if there is a diffeomorphism isotopic to the identity $\phi : N - P \rightarrow N - P$ such that $\rho = \phi^*\rho'$. The restriction of the diffeomorphism to the boundary $\phi : \partial N - P \rightarrow \partial N - P$ is isotopic to the identity and conformal with respect to the Riemann surface structures induced by ρ and ρ' . Therefore, equivalent geometrically finite hyperbolic metrics have equivalent conformal boundaries.

DEFINITION (Convex Core). Every hyperbolic metric on $M = \text{int}(N)$ admits a *convex core* $\mathcal{CC}(N) \subset M$, which is the smallest convex closed subset whose inclusion in N is a homotopy equivalence.

2.3. Maximally cusped structures. Maximally cusped structures play a crucial role in Part 2. Let N be either H_g or $\Sigma \times [a, b]$.

DEFINITION (Maximally Cusped). A hyperbolic structure on $M = \text{int}(N)$ is *maximally cusped* or *maximally cusped at P* if it is geometrically finite and has parabolic locus P which is a pants decomposition of ∂N .

Work of Maskit [Mas83], Keen, Maskit and Series [KMS93], and Ohshika [Ohs98] gives the following.

THEOREM 2.2 (see Theorem 7.2.9 in [CM04]). *The following holds.*

- For every pants decomposition P of $\Sigma = \partial H_g$ such that (H_g, P) is a pared acylindrical 3-manifold there exists a unique (up to isotopy) hyperbolic metric on $\text{int}(H_g)$ which is maximally cusped at P . We denote such a metric by $H(P)$.
- For every pants decompositions P, R of Σ such that $(\Sigma \times [a, b], P \times \{a\} \sqcup R \times \{b\})$ is a pared acylindrical 3-manifold there exists a unique (up to isotopy) hyperbolic metric on $\Sigma \times [a, b]$ which is maximally cusped at $P \times \{a\} \sqcup R \times \{b\}$. We denote such a metric by $Q(P, R)$.

Regarding convex cores, work of Keen, Maskit and Series [KMS93] provides the following description.

THEOREM 2.3. *The following holds.*

- Let P be a pants decomposition of $\Sigma = \partial H_g$ such that (H_g, P) is pared acylindrical. Let $H(P)$ be the corresponding maximally cusped structure. The convex core $\mathcal{CC}(H(P)) \subset H_g$ is isotopic to $H_g - P$ and has totally geodesic boundary.
- Let P, R be pants decompositions of Σ such that $(\Sigma \times [a, b], P \times \{a\} \sqcup R \times \{b\})$ is pared acylindrical. Let $Q(P, R)$ be the corresponding maximally cusped structure. The convex core $\mathcal{CC}(Q(P, R)) \subset \Sigma \times [a, b]$ is isotopic to $\Sigma \times [a, b] - (P \times \{a\} \cup R \times \{b\})$ and has totally geodesic boundary.

2.4. Curve graph and disk set. We now introduce the curve graph and the disk sets and relate them to the topology of Heegaard splittings and handlebodies. We also describe some of their geometric properties.

DEFINITION (Curve Graph and Disk Set). The *curve graph* $\mathcal{C} := \mathcal{C}(\Sigma)$ of a compact orientable surface Σ is the graph whose vertices are the isotopy classes of essential non-peripheral simple closed curves. Two vertices are joined by an edge of length 1 if the corresponding curves can be realized disjointly on Σ . We endow the graph \mathcal{C} with the intrinsic path metric $d_{\mathcal{C}}$. In fact, we will only consider the vertex set of \mathcal{C} endowed with its induced metric, which we also denote by \mathcal{C} by abuse of notation.

If $\Sigma = \partial H_g$ is the boundary of a handlebody H_g , we have an associated *disk set* $\mathcal{D} \subset \mathcal{C}$ consisting of those essential simple closed curves of Σ that bound a properly embedded disk in H_g .

If $f : \Sigma \rightarrow \Sigma$ is a gluing map defining the Heegaard splitting $M_f = H_g \cup_f H_g$, then the *Hempel distance* of f is the quantity

$$d_{\mathcal{C}}(\mathcal{D}, f\mathcal{D}).$$

Splittings with large Hempel distance have strong topological properties:

THEOREM 2.4 (Hempel [Hem01]). *Let f be a gluing map. If the Hempel distance of f is at least 3, then the Heegaard splitting M_f is irreducible, does not contain any essential torus or Klein bottle, and is not Seifert fibered.*

Combined with the solution of the Geometrization Conjecture by Perelman, Theorem 2.4 implies that if $d_{\mathcal{C}}(\mathcal{D}, f\mathcal{D}) \geq 3$, then the splitting M_f is hyperbolic.

In [Mah10b], Maher proved that the Hempel distance of a random mapping class grows coarsely linearly and, hence, that random Heegaard splittings are hyperbolic.

In our random 3-manifolds setup, we will also exploit the curve graph to control the topology of our building blocks. For example, we will use it to check that the blocks satisfy the assumptions of Theorems 2.2 and 4.1: From a curve graph point of view, we have the following useful criterion.

LEMMA 2.5. *Let $\gamma, \gamma' \subset \Sigma$ be essential multicurves. We have*

- (i) *If $d_{\mathcal{C}}(\gamma, \mathcal{D}) \geq 2$ then (H_g, γ) is pared.*
- (ii) *If $d_{\mathcal{C}}(\gamma, \mathcal{D}) \geq 3$ then (H_g, γ) is pared acylindrical.*
- (iii) *If $d_{\mathcal{C}}(\gamma, \gamma') \geq 1$ then $(\Sigma \times [0, 1], \gamma \times \{0\} \sqcup \gamma' \times \{1\})$ is pared.*
- (iv) *If $d_{\mathcal{C}}(\gamma, \gamma') \geq 3$ then $(\Sigma \times [0, 1], \gamma \times \{0\} \sqcup \gamma' \times \{1\})$ is pared acylindrical.*

The proof can be found in Appendix A.

From a geometric point of view, Masur and Minsky proved that the curve graph \mathcal{C} is a Gromov hyperbolic space (see [MM99]) and the disk set $\mathcal{D} \subset \mathcal{C}$ is a uniformly quasi-convex subspace (see [MM04]).

As discovered by Minsky [Min00, Min01, Min10], and Brock, Canary and Minsky [BCM12] in groundbreaking work that led to the solution of the Ending Lamination Conjecture, the geometry of the curve graph is related to the internal geometry of hyperbolic metrics on $\Sigma \times [0, 1]$: The relation is established via *end invariants* and the devices of *subsurface projections* and *hierarchies of tight geodesics* introduced in [MM00]. We recall the definition of subsurface projection in the case that is most relevant for us.

DEFINITION (Non-annular Subsurface Projection). Let $W \subset \Sigma$ be a proper compact connected non-annular subsurface of Σ that is not a three-holed sphere. Let $\alpha \in \mathcal{C}$ be any curve. The *subsurface projection* of α to W is the (possibly empty) subset $\pi_W(\alpha)$ of the curve graph $\mathcal{C}(W)$ of all possible essential surgeries of $\alpha \cap W$ (see Section 2 of [MM00] for more details).

If both α and β intersect W essentially, we define

$$d_W(\alpha, \beta) = \text{diam}_{\mathcal{C}(W)}(\pi_W(\alpha) \cup \pi_W(\beta)).$$

2.5. Filling and drilling. One of the most important tools needed in our constructions is the Universal Dehn Surgery Theorem by Hodgson and Kerckhoff [HK05]. We briefly recall the basic setup and statements.

Let M be a compact oriented 3-manifold with boundary. For a component of the boundary T that is a torus, a choice μ of a homotopy class of a non-null homotopic simple closed curve in T is called a *slope*. Gluing a solid torus $S^1 \times D^2$ to M via a gluing homeomorphism $\partial(S^1 \times D^2) = S^1 \times S^1 \rightarrow T$ that sends the homotopy class of $\{1\} \times S^1$ to the slope μ is called *Dehn filling of M along the slope μ* and yields an compact oriented 3-manifold that contains M as a submanifold. We call the slope along which we perform a Dehn filling the *filling slope*.

Let $\eta_3 > 0$ be a Margulis constant which we fix once and for all.

Let \mathbb{M} be a complete finite volume hyperbolic 3-manifold with cusps. Each cusp γ of \mathbb{M} has a standard η_3 -Margulis neighborhood $\mathbb{T}_{\eta_3}(\gamma) \subset \mathbb{M}$ isometric to a quotient \mathcal{O}/\mathbb{Z}^2 where $\mathcal{O} \subset \mathbb{H}^3$ is a horoball and $\mathbb{Z}^2 < \text{Isom}^+(\mathbb{H}^3)$ consists of

parabolic isometries stabilizing \mathcal{O} . The boundary $T_\gamma := \partial\mathbb{T}_{\eta_3}(\gamma) \subset \mathbb{M}$ is a *standard torus horosection* of γ . The intrinsic metric is of T_γ is flat.

DEFINITION (Normalized Length). Let $\mu \subset T_\gamma$ be a simple closed geodesic. The *normalized length* of μ , is the quantity defined by

$$\text{nl}(\mu) := \ell_{T_\gamma}(\mu) / \sqrt{\text{Area}(T_\gamma)}.$$

By Dehn filling \mathbb{M} along a slope μ of T_γ , we mean Dehn filling the underlying manifold of $\mathbb{M} \setminus \text{int}(\mathbb{T}_{\eta_3}(\gamma))$ along a slope μ of the boundary component T_γ of M . Hodgson and Kerckhoff proved the following:

THEOREM 2.6 (Hodgson-Kerckhoff [HK05]). *Let $\eta_3 > 0$ be a Margulis constant. Let \mathbb{M} be a complete finite volume hyperbolic 3-manifold with cusps. For each cusp γ of \mathbb{M} , let μ be a simple closed geodesic on the standard torus horosection $T_\gamma = \partial\mathbb{T}_{\eta_3}(\gamma)$ of γ . Suppose that the normalized length of each μ is at least 10.6273. Then there exists a hyperbolic metric on the result of Dehn filling M along the slopes defined by μ such that core curves of the added solid tori are geodesics.*

If we have a geodesic link $\Gamma \subset M$ in a complete finite volume hyperbolic 3-manifold M and the complement $M - \Gamma$ admits a complete finite volume hyperbolic metric as well, then, by the techniques from [HK05, HK08], there is a relation between the length of the link $\Gamma \subset M$ and the normalized length of the standard meridians of Γ in $M - \Gamma$ made fully explicit in [FPS19].

THEOREM 2.7 ([FPS19, Corollary 6.13]). *Let $\Gamma \subset M$ be a geodesic link in a complete finite volume hyperbolic 3-manifold, and denote by ℓ the total length of Γ . Then $M - \Gamma$ has a complete finite volume hyperbolic metric, and we denote by L the total normalized length L of the meridians of Γ . Moreover, suppose that either $L \geq 7.823$, or that each component of Γ has length at most 0.0996 and $\ell \leq 0.1396$. Then*

$$\frac{2\pi}{L^2 + 16.17} < \ell < \frac{2\pi}{L^2 - 28.78}.$$

Brock and Bromberg studied the change in the geometry of a family M_t of hyperbolic cone manifold structures on a topological model M in a very general setup and proved that away from the singular locus Γ one can get uniform bilipschitz control only depending on the length of Γ . This is the content of the following Drilling Theorem.

THEOREM 2.8 ([BB04, Theorem 6.2]). *For every $K \in (1, 2)$ there exists $\eta_{\text{drill}} \in (0, \eta_3)$ such that the following holds. Let M be a complete finite volume hyperbolic 3-manifold. Let $\Gamma \subset M$ be a geodesic link and let M' be a complete hyperbolic structure on $M - \Gamma$. Suppose that $\ell_M(\Gamma) < \eta_{\text{drill}}$. Then, there exists a K -bilipschitz diffeomorphism of pairs*

$$\left(M' - \bigsqcup_{\alpha \subset \Gamma} \mathbb{T}_{\eta_3}(\alpha), \bigsqcup_{\alpha \subset \Gamma} T_\alpha \right) \rightarrow \left(M - \bigsqcup_{\alpha \subset \Gamma} \mathbb{T}_{\eta_3}(\alpha), \bigsqcup_{\alpha \subset \Gamma} T_\alpha \right)$$

where $\mathbb{T}_{\eta_3}(\alpha)$ denotes a standard η_3 -Margulis neighborhood for α and $T_\alpha = \partial\mathbb{T}_{\eta_3}(\alpha)$ its boundary.

Short curves on Heegaard splittings

3. OUTLINE AND THE TOPOLOGY OF HEEGARD SPLITTINGS

In this and the next section we discuss short curves on Heegaard splittings. The goal is to prove Theorem 6. The family of examples that arise from Theorem 6 is shown to be generic from the point of view of the random walk in Section 7.

3.1. Outline. We now outline the strategy to produce a hyperbolic metric on M_f for which $\gamma \subset \Sigma$ is a short geodesic.

We start by associating to γ the 3-manifold

$$M_f - \gamma = (H_g - \gamma) \cup_f (H_g - f(\gamma)).$$

The Heegaard splitting M_f is obtained from $M_f - \gamma$ by Dehn filling along a filling slope which is completely determined by the topology.

According to Thurston's Hyperbolization Theorem, the manifold $M_f - \gamma$ admits a complete hyperbolic metric provided that it is irreducible, atoroidal and Haken; see, for example, Theorem 1.42 in [Kap09]. We take [Kap09] as our general reference for basic 3-manifold topology. Notice, however, that, since it is the interior of a compact manifold with non-empty boundary, by basic 3-manifold topology, if $M_f - \gamma$ is irreducible, then it will also be automatically Haken (see Corollary 1.24 in [Kap09]).

Both irreducibility and atoroidality will follow from our assumption that both (H_g, γ) and $(H_g, f(\gamma))$ are so-called pared acylindrical handlebodies. This is condition (a) of Theorem 6 and is borrowed from the theory of Jaco, Shalen and Johannson (see Chapter 5 of [CM04]).

Once we have a hyperbolic metric on $M_f - \gamma$ we use it to find a hyperbolic metric on M_f . The tool for such an operation is Hodgson and Kerckhoff effective version of the Hyperbolic Dehn Surgery Theorem (see Theorem 2.6). We need to certify that the canonical filling slope has large normalized length on a standard torus horosection $T := \partial\mathbb{T}_{\eta_3}(\gamma)$ of the cusp of $M_f - \gamma$.

Actually, we will check something stronger, namely that, if condition (b) of Theorem 6 holds, then the flat torus T itself will be long and skinny so that *every* filling slope different from the one corresponding to $(\Sigma - \gamma) \cap T$ will have large normalized length.

The central point of the proof is to show that T is long and skinny provided that the subsurface projection $d_{\Sigma - \gamma}(\mathcal{D}, f\mathcal{D})$ is large. The main ingredient for the argument is the model manifold technology by Minsky [Min10] and Brock, Canary and Minsky [BCM12].

At this point, before going on, we need a further consequence of the condition (a): If (H_g, γ) is pared acylindrical, then the inclusion $\Sigma - \gamma \subset H_g$ is *doubly incompressible*. Double incompressibility allows us to use Thurston's Uniform Injectivity for Pleated Surfaces [Thu86a] and its consequences. In particular, we show that there are two simple closed curves $\alpha \subset \Sigma - \gamma$ and $\beta \subset \Sigma - f(\gamma)$ which are represented by geodesics in $M_f - \gamma$ with moderate length and which are combinatorially close to the disk set projections as follows: $d_{\Sigma - \gamma}(\mathcal{D}, \alpha) \leq 2$ and $d_{\Sigma - \gamma}(f\mathcal{D}, \beta) \leq 2$. Notice that, since $d_{\Sigma - \gamma}(\mathcal{D}, f\mathcal{D})$ is large, also $d_{\Sigma - \gamma}(\alpha, \beta)$ will be large. In order to produce the moderate length curves α and β it turns out to be convenient to work with the coverings of $M_f - \gamma$ corresponding to $\pi_1(H_g - \gamma)$ and $\pi_1(H_g - f(\gamma))$ instead of $M_f - \gamma$ itself.

As a last step, incompressibility of $\Sigma - \gamma \subset M_f - \gamma$ and the fact that the curves $\alpha, \beta \subset \Sigma - \gamma$ are represented by moderate length curves in $M_f - \gamma$ and are combinatorially far apart in the curve graph $\mathcal{C}(\Sigma - \gamma)$ implies, via the model manifold technology, that $T = \partial\mathbb{T}_{\eta_3}(\gamma)$ is long and skinny.

This concludes the outline. The remainder of this section is dedicated to the topological part of the proof. We establish the topological properties that imply that $M_f - \gamma$ is hyperbolic and check that $\Sigma - \gamma$ is doubly incompressible. The geometric part of the argument will be discussed in the next section.

3.2. Hyperbolizable drilled Heegaard splittings. Gluing together two pared acylindrical handlebodies produces an irreducible and atoroidal 3-manifold.

PROPOSITION 3.1. *If (H_g, γ) and $(H_g, f(\gamma))$ are both pared acylindrical then $M_f - \gamma$ is irreducible, atoroidal and Haken.*

In particular, by Thurston's Hyperbolization Theorem for Haken manifolds, the drilled splitting $M_f - \gamma$ admits a hyperbolic structure, which is also unique by Mostow-Prasad rigidity.

We split the proof of Proposition 3.1 into small steps, Lemmas 3.2, 3.3 and 3.4. First we observe that $\Sigma - \gamma \subset M_f - \gamma$ is incompressible.

LEMMA 3.2. *The surface $\Sigma - \gamma \subset M_f - \gamma$ is π_1 -injective. The fundamental group $\pi_1(M_f - \gamma)$ decomposes as $\pi_1(H_g - \gamma) *_{\pi_1(\Sigma - \gamma)} \pi_1(H_g - f(\gamma))$.*

Proof. The first sentence is a restatement of Remark 2.1, and the rest follows from Seifert-van Kampen's Theorem. \square

We now use transversality arguments together with acylindricity of the handlebodies to establish irreducibility and atoroidality of $M_f - \gamma$.

LEMMA 3.3. *$M_f - \gamma$ is irreducible.*

Proof. Assume towards a contradiction that $M_f - \gamma$ is reducible, and let S be an embedded 2-sphere that does not bound a ball. We may and do assume that S intersects $\Sigma - \gamma$ transversally and such that any component of $(\Sigma - \gamma) \cap S$ that is innermost in S is not homotopically trivial in $\Sigma - \gamma$. Indeed, otherwise this innermost component bounds a disc in $\Sigma - \gamma$ and a disc in S such that their union is contained in one of the two handlebodies. However, this sphere bounds a ball B in that handlebody by irreducibility of handlebodies, and thus B can be used to guide an isotopy of S that removes that component of $(\Sigma - \gamma) \cap S$. Note that $\Sigma \cap S$ is nonempty since otherwise S would be contained in one handlebody and

bound a ball, again by irreducibility of handlebodies. Hence, any component α of $\Sigma \cap S$ that is innermost in S compresses in one of the two handlebodies via the disc D in S with $D \cap \Sigma = \alpha$. This contradicts π_1 -injectivity of $\Sigma - \gamma$ into both handlebodies. \square

LEMMA 3.4. $M_f - \gamma$ is atoroidal.

Proof. We show that $M_f - \gamma$ is topologically atoroidal (i.e. every incompressible torus is boundary parallel), which implies atoroidal for Haken manifolds (which $M_f - \gamma$ is); see Section 1.2 in [Kap09]. Let T be an incompressible torus in $M_f - \gamma$. We show that T is boundary parallel. Arrange that T is in general position with respect to $\Sigma - \gamma$, and take $\{\alpha_i\}_{i=1, \dots, n}$ to be the simple closed curves that are the components of the intersection $(\Sigma - \gamma) \cap T$. By an innermost argument, if some α_i was null-homotopic in T , then some α_j would bound a disk in one of the handlebodies H bounded by Σ . However, this implies that α_j is trivial in $\Sigma - \gamma$ since $\Sigma - \gamma$ maps π_1 -injectively to both handlebodies. Then, the union of the disks in S_K and T with boundary α_j bounds a ball in H by irreducibility of H , and we can use this ball to reduce the number of components of $(\Sigma - \gamma) \cap T$. In view of this argument, we can assume that each α_i is an essential curve on T . Moreover, notice that each α_i is then automatically an essential curve on $\Sigma - \gamma$, whence on Σ , too, for otherwise some α_j would bound a compressing disk. We reindex the α_i to make sure that consecutive ones (modulo n) bound an annulus in T .

Consider now some α_i , and let H be the handlebody containing the annulus $A \subseteq T$ bounded by α_i and α_{i+1} (notice that if there is one α_i , then in fact there are at least two because $\Sigma - \gamma$ separates $M_f - \gamma$, so that $(\Sigma - \gamma) \cap T$ needs to have at least two connected components).

Claim: A is boundary parallel, that is, together with an annulus in Σ , A cobounds an interval bundle.

Proof of the Claim. Using that (H, γ) is pared acylindrical, we note that α_i and α_{i+1} are two boundaries of an embedded annulus A' in Σ that is homotopic rel boundary to A in H . Since $\partial H \setminus (\partial A)$ is disconnected, also $H \setminus A$ is disconnected. Let N be the 3-manifold given as the closure of the component with boundary $A \cup A'$. To establish that A is boundary parallel, we show that N is a solid torus.

Let $I \subset A$ be a properly embedded interval connecting the two boundary components of A . And let $I' \subset A'$ be an arc that is homotopic rel boundary to I in H . An innermost circle argument gives that $I' \cup I$ is also null-homotopic in N , hence by Dehn's Lemma there exists a disc $D \subset N$ with boundary $I' \cup I$. Let S be the sphere obtained as the union of the disc given by boundary compression of A along D and the disc in Σ with the same boundary. This sphere S bounds a ball in H , and thus N , by irreducibility of handlebodies. Thus N compresses to a ball; hence, N is a solid torus as desired. \square

Let A' be the annulus in Σ that union A forms a torus that bounds a solid torus in H . If A' does not contain γ , then we can reduce the number of components of $(\Sigma - \gamma) \cap T$ by isotoping A in $M_f - \gamma$ to the other handlebody. Otherwise, A' is (up to isotopy in Σ) $N(\gamma)$, and we can apply an isotopy in M_f to move A inside a regular neighborhood N of γ in M_f . Notice that there is at least one α_j for otherwise T would be contained in one of the handlebodies; however, handlebodies do not contain incompressible tori since incompressible tori are π_1 -injective and the

fundamental groups of handlebodies do not contain \mathbb{Z}^2 subgroups. In particular, T is a union of annuli as above, and hence, after applying finitely many isotopies, we reduce to the case that every annulus of T is entirely contained in N (we can assume that the isotopies we found above move N inside itself). Hence, T can be thought of as an incompressible surface in $N - \gamma$. There is a classification of incompressible surfaces in S^1 -bundles; see [Wal67, Satz 2.8], which in our case implies that T is boundary parallel, as required. \square

Hence $M_f - \gamma$ is irreducible and atoroidal. As mentioned before, irreducibility is already enough to ensure that it is also Haken. Incidentally, we observe that $\Sigma - \gamma \subset M_f - \gamma$ is a Haken surface.

Combining with Lemma 2.5, we get the following.

COROLLARY 3.5. *Let $\gamma \subset \Sigma$ be a non-separating simple closed curve. If $d_{\mathcal{C}}(\gamma, \mathcal{D} \cup f\mathcal{D}) \geq 3$, then $M_f - \gamma$ is hyperbolizable.*

3.3. Double incompressibility. The second crucial topological property of a pared acylindrical handlebody (H_g, γ) that we need is the fact that the inclusion of the boundary $\Sigma - \gamma \subset H_g$ is doubly incompressible.

Double incompressibility was defined by Thurston in [Thu86a] for maps from hyperbolic surfaces to hyperbolic 3-manifolds. In the setting we are interested in, we can rephrase Thurston's definition purely topologically as in the definition below.

DEFINITION (Topologically Doubly Incompressible). Let $\gamma \subset \Sigma$ be an essential simple closed curve with a tubular neighborhood $N(\gamma)$ in Σ and $U(\gamma)$ in H_g . The inclusion $\Sigma - \gamma \subset H_g$ is *topologically doubly incompressible* if it satisfies the following.

- (a) The inclusion $\Sigma - \gamma \subset H_g$ is π_1 -injective.
- (b) Essential relative homotopy classes of maps $(I, \partial I) \rightarrow (\Sigma - N(\gamma), \partial N(\gamma))$ are mapped injectively to relative homotopy classes of maps $(I, \partial I) \rightarrow (H_g - U(\gamma), \partial U(\gamma))$.
- (c) There are no essential cylinders in $\Sigma - N(\gamma)$: This means that every essential map $f : (A, \partial A) \rightarrow (H_g, \Sigma - N(\gamma))$ is either homotopic into $U(\gamma)$ or the restriction of f to ∂A extends to a map into $\Sigma - N(\gamma)$.
- (d) Each maximal abelian subgroup of $\pi_1(\Sigma - \gamma)$ is mapped to a maximal abelian subgroup of $\pi_1(H_g)$.

Since maximal abelian subgroups of a free group are infinite cyclic groups, the last condition is equivalent to

- (e) each maximal cyclic subgroup of $\pi_1(\Sigma - \gamma)$ is mapped to a maximal cyclic subgroup of $\pi_1(H_g)$.

REMARK 3.6. A direct comparison of the definition above and the definition of double incompressibility in [Thu86a] yields the following technical fact that we need later on (to apply facts about doubly incompressible pleated surfaces). Let γ be a simple closed curve on the boundary Σ of the handlebody H_g . Let N be a hyperbolic 3-manifold containing H_g and which deformation retracts onto H_g in such a way that the cusp retracts on a neighborhood of γ . Then the inclusion $\Sigma - \gamma \rightarrow N$ is doubly incompressible.

As Thurston essentially observes in Section 7 of [Thu86a], we have the following:

PROPOSITION 3.7. *If (H_g, γ) is pared acylindrical, then the inclusion $\Sigma - \gamma \subset H_g$ is topologically doubly incompressible.*

Proposition 3.7 is well-known to experts and follows from JSJ theory. However, it might not be easy to extract from the literature. For this reason, and for the sake of being more self-contained, we include a proof in Appendix B.

4. LONG SKINNY CUSP HOROSECTION

In this section we show that a standard torus horosection T of the cusp of $M_f - \gamma$ is long and skinny, so that every filling slope different from the one coming from $(\Sigma - \gamma) \cap T$ will have large normalized length, provided that $d_W(\mathcal{D}, f\mathcal{D})$ is sufficiently large. This is the geometric part of the proof of Theorem 6 and rests on the model manifold technology of Minsky [Min10] and Brock, Canary and Minsky [BCM12].

Here, the standard torus horosection T is $\partial\mathbb{T}_{\eta_3}(\gamma)$, where $\eta_3 > 0$ is a fixed Margulis constant and $\mathbb{T}_{\eta_3} \subset M_f - \gamma$ is the cusp of $M_f - \gamma$ that forms a connected component of the η_3 -thin part of $M_f - \gamma$.

The proof is divided into two steps. The first one consists of finding simple closed curves $\alpha, \beta \subset W := \Sigma - \gamma$ that are represented by closed geodesics in $M_f - \gamma$ with moderate length and such that $d_W(\alpha, \mathcal{D}), d_W(\beta, f\mathcal{D}) \leq 2$. This is the content of Proposition 4.2 and Corollary 4.3. As a second step, once we have such curves α and β , we argue that $d_W(\alpha, \beta)$ gives a coarse lower bound for the length of any slope on T that does not come from the Heegaard surface. We prove this in Proposition 4.7.

4.1. Handlebody covering. In order to find the curves α and β , we work with *handlebody coverings*, which we now describe.

Consider first the pared handlebody (H_g, γ) . The fundamental group of $H_g - \gamma \subset M_f - \gamma$ injects into $\pi_1(M_f - \gamma)$ (see Lemma 3.2) and determines a covering N of $M_f - \gamma$ to which $H_g - \gamma$ lifts homeomorphically. By slight abuse of notation we will not distinguish between $H_g - \gamma$ and its homeomorphic lift to N . Recall from the outline, that we want to find a simple closed curve $\alpha \subset \Sigma - \gamma$ that has moderate length representative in N and satisfies $d_W(\alpha, \mathcal{D}) \leq 2$.

As both $H_g - \gamma$ and N are aspherical, the inclusion $H_g - \gamma \subset N$ is a homotopy equivalence. Since the pair $(N, H_g - \gamma)$ has also the homotopy extension property (e.g. since it can be given the structure of a CW-pair), it follows that the manifold N deformation retracts to $H_g - \gamma$ (this is a general fact; see, for example, [Hat02, Corollary 0.20]). Therefore, according to Proposition 3.7 together with Remark 3.6 to relate the topological and hyperbolic versions of double incompressibility, the inclusion $W = \Sigma - \gamma \subset N$ is doubly incompressible in the sense of [Thu86a].

Even though we will not need it, for the sake of clarity, we provide a more complete description of the covering. The hyperbolic structure on N is isometric to a geometrically finite structure on $\text{int}(H_g)$ with a rank one cusp at γ and the submanifold $H_g - \gamma \subset N$ is a so-called relative Scott core for N [Sco73, McC86, KS89]. The fact that N is homeomorphic to the interior of H_g follows from Bonahon's Tame-ness Theorem [Bon86]. Geometric finiteness is, instead, a consequence of Canary and Thurston's Covering Theorem [Can96].

We now come back to double incompressibility. Crucially, it allows to use Thurston's Uniform Injectivity for Pleated Surfaces [Thu86a].

THEOREM 4.1 (Uniform Injectivity, [Thu86a, Theorem 5.7]). *Fix $\eta > 0$, a Margulis constant. For every $\varepsilon > 0$ there exists $\delta > 0$ such that for any type preserving doubly incompressible pleated surface $g : W \rightarrow N$ with pleating locus λ and induced metric σ , if $x, y \in \lambda$ lie in the η -thick part of (W, σ) , then*

$$d_{\mathbf{P}(N)}(\mathbf{p}_g(x), \mathbf{p}_g(y)) \leq \delta \implies d_\sigma(x, y) \leq \varepsilon.$$

Here $\mathbf{p}_g : \lambda \rightarrow \mathbf{P}(N)$ denotes the map induced by g from the lamination λ to the projective unit tangent bundle of N .

We use Theorem 4.1 to prove the following

PROPOSITION 4.2. *There exists $L > 0$ such that the following holds. For every $\delta \in \mathcal{D}$ there exists a pleated surface $g : (W, \sigma) \rightarrow N$ in the proper homotopy class of the inclusion $W = \Sigma - \gamma \subset N$ that realizes $\lambda := \pi_W(\delta)$ as a sublamination of its pleating locus and such that one of the arcs δ_0 of $\lambda \cap W_0$ satisfies $\ell_\sigma(\delta_0) \leq L$. Here W_0 is the η_0 -non cuspidal part of (W, σ) where $\eta_0 > 0$ is a universal constant.*

Finally, the moderate length segments provided by Proposition 4.2 can be promoted to moderate length curves α and β in a simple way

COROLLARY 4.3. *There exists $L > 0$, only depending on Σ such that there are simple closed curves $\alpha, \beta \subset W$ with $d_W(\alpha, \pi_W(\mathcal{D})) \leq 2$ and satisfying $\ell_{M_f - \gamma}(\alpha^*), \ell_{M_f - \gamma}(\beta^*) \leq L$, where α^*, β^* are the geodesic representatives of α, β in $M_f - \gamma$.*

Proof. We only provide the argument for α since the one for β is completely analogous. Choose $\delta \in \mathcal{D}$ arbitrarily. Let $g : (W, \sigma) \rightarrow N$ and δ_0 be the pleated surfaces and the moderate length arc in $\lambda \cap W_0$ provided by Proposition 4.2. The length of ∂W_0 is bounded by $2\eta_0$. One of the boundary components of a regular neighborhood of $\delta_0 \cup \partial W_0$ is an essential simple closed curve α of length at most $2L + 2\eta_0$ with the property $d_W(\alpha, \pi_W(\delta)) \leq 2$. \square

The proof of Proposition 4.2 is where we fully exploit the assumption that W is the complement of a non-separating simple closed curve. The main ingredients of the proof are Thurston's Uniform Injectivity and a technical lemma about quasi geodesic concatenations in \mathbb{H}^3 . We also use the following general property of pleated surfaces observed by Thurston in [Thu86a].

LEMMA 4.4 ([Min00, Lemma 3.1]). *For every Margulis constant $\eta_0 > 0$ there exists η_3 , only depending on η_0 and the topological type of W , such that, if $g : (W, \sigma) \rightarrow N$ is a π_1 -injective pleated surface, then only the η_0 -thin part of (W, σ) can enter the η_3 -thin part of N .*

4.2. Quasi geodesic concatenations. Beside the use of Uniform Injectivity, our proof of Proposition 4.2 is elementary and rests on the following fact about piecewise broken geodesics in \mathbb{H}^3 , which we state without proof. In the statement, the breaking angle of a break point $I(t)$ of a piecewise geodesic $\gamma : I \rightarrow \mathbb{H}^3$ is the angle formed by the positively oriented left and right tangent at $I(t)$.

LEMMA 4.5. *There exists $L > 0$ and $A = A(L) > 0$ such that the following holds. Let $\gamma : I \rightarrow \mathbb{H}^3$ be a broken piecewise geodesic $\gamma = \gamma_1 * \dots * \gamma_m$ with breaking angles $\angle \gamma_{i-1} \gamma_i$ contained in $(0, \pi/2)$ and geodesic segments of length at least L . Then γ is an A -quasi geodesic.*

In particular, if the length of the geodesic segments is large enough then γ cannot be a loop. Lemma 4.5 follows from the fact that, in a hyperbolic space, local quasi geodesics, such as γ (as is not difficult to check), are global quasi geodesics.

Using Lemma 4.5, we find the following.

LEMMA 4.6. *For each $\varepsilon > 0$ there exists $L_0 > 0$ such that the following holds. Let $X = \mathbb{H}^3 - \sqcup_{1 \leq j \leq n} \mathcal{O}_j$ be the complement in \mathbb{H}^3 of a family of pairwise disjoint open horoballs. Let $\gamma = \gamma_1 * \cdots * \gamma_{2n}$ be a concatenation of paths in X such that*

- γ_{2j+1} is a geodesic of length at least ε on the horosphere $\mathcal{H}_j = \partial\mathcal{O}_j$.
- γ_{2j} is the orthogonal segment connecting \mathcal{H}_{j-1} and \mathcal{H}_j . We require that γ_{2j} has length at least L_0 .

Then γ is not a loop.

Proof. We use the following fact that can be easily checked in the upper half space model of \mathbb{H}^3 . If x, y lie on the same horosphere \mathcal{H} , then we have

$$\sinh(d_{\mathbb{H}^3}(x, y)) = d_{\mathcal{H}}(x, y).$$

Denote by $\mathcal{H}_j = \partial\mathcal{O}_j$ the boundary horosphere of the horoball \mathcal{O}_j . Observe that, by assumption, the flat geodesic $\gamma_{2j} \subset \mathcal{H}_j$ has length $\ell(\gamma_{2j}) \geq \varepsilon$.

If we expand \mathcal{H}_j radially from its center at infinity, the intrinsic geometry of the horosphere expands exponentially with exponent equal to the increase in the radius. Hence, if we inflate all the horoballs \mathcal{O}_j by r , then each γ_{2j+1} is shortened to an arc γ_{2j+1}^r of length

$$\ell(\gamma_{2j+1}^r) = \ell(\gamma_{2j+1}) - 2r \geq L - 2r,$$

while the length of the inflated γ_{2j} , denoted by γ_{2j}^r , becomes

$$\ell(\gamma_{2j}^r) = e^r \ell(\gamma_{2j}).$$

We now straighten all the γ_{2j}^r relative to the endpoints and obtain geodesic arcs α_{2j}^r of length $\ell(\alpha_{2j}^r) = \sinh^{-1}(e^r \ell(\gamma_{2j})) \geq \sinh^{-1}(e^r \varepsilon)$.

Let us now consider the angles between the segments γ_{2j-1} and α_{2j} . Observe that, by assumption, γ_{2j-1}^r is orthogonal to \mathcal{H}_j^r , the inflated horosphere. Therefore, as, by convexity of horoballs, α_{2j} is contained in \mathcal{O}_j^r , the angle $\angle \gamma_{2j-1}^r \alpha_{2j}^r$ between the two geodesics γ_{2j-1} and α_{2j} is in $(0, \pi/2)$.

In conclusion, if the length of each γ_{2j+1} is much larger than $L + 2r$ and $\sinh^{-1}(e^r \varepsilon) \geq L$, where L is the constant of Lemma 4.5, then the broken piecewise geodesic $\gamma = \gamma_1^r * \alpha_2^r * \cdots * \gamma_{2n-1}^r * \alpha_{2n}^r$ is a A -quasi geodesic. If it is also sufficiently long compared to A , which can be again achieved by assuming that each γ_{2j+1} is long enough, it cannot be a loop. \square

4.3. Moderate length surgeries. We now prove Proposition 4.2.

Proof of Proposition 4.2. Pick $\delta \in \mathcal{D}$ arbitrarily. Since (H_g, γ) is pared acylindrical, δ intersects γ essentially. Consider $\lambda = \delta \cap W$, it is a multi-arc in $W = \Sigma - \gamma$.

Now, after collapsing parallel components to a single one, λ can be realized as a sub-lamination of the pleating locus of some pleated surface $g : (\Sigma - \gamma, \sigma) \rightarrow N$ in the proper (relative to cusps) homotopy class of the inclusion $\Sigma - \gamma \subset N$ (see for example Theorems I.5.3.6 and I.5.3.9 in [CEG06]).

When regarding δ as a simple closed curve on Σ , we can think of it as a concatenation $\delta = \beta_1 * \cdots * \beta_{2n}$ of arcs β_{2j} in W and arcs β_{2j+1} crossing a regular annular neighborhood of γ . Here we are using the fact that W is the complement of a non-separating simple closed curve.

Using the proper homotopy between the inclusion $W = \Sigma - \gamma \subset N$ and g , we simultaneously straighten all the β_{2j} 's to subarcs α_{2j} of the geodesic leaves of $g(\lambda)$ that start and end in the standard horosection $T = \partial\mathbb{T}_{\eta_3}(\gamma)$. Then, again, using the proper homotopy between $\Sigma - \gamma \subset N$ and the deformation retraction of N to $N - \mathbb{T}_{\eta_3}(\gamma)$, we replace the arcs β_{2j+1} with arcs α_{2j+1} on the horosection T that are geodesic with respect to the intrinsic flat metric and join the endpoints of the previously obtained α_{2j} and $\alpha_{2(j+1)}$.

We have that $\delta \subset \Sigma$ is homotopic in N to a concatenation $\alpha_1 * \cdots * \alpha_{2n}$ of closed arcs α_j , where each α_{2j+1} is a geodesic in the boundary of the standard horosection T , while each α_{2j} is a proper geodesic arc in $N - \mathbb{T}_{\eta_3}(\gamma)$ contained in the pleating locus $g(\lambda)$.

Notice that, by Lemma 4.4, if η_3 is sufficiently small, then, when we look at α_{2j} in the intrinsic metric of the pleated surface, it will join two points of the η_0 -cuspidal part for some universal $\eta_0 > 0$.

Now, by Theorem 4.1, which applies to g because it is doubly incompressible, there is a uniform $\varepsilon > 0$ such that the length of each α_{2j+1} , is at least ε . Let L_0 be as in Lemma 4.6, for the given ε . Notice that this constant only depends on Σ . Being homotopically trivial in N , the loop $\alpha_1 * \cdots * \alpha_{2n}$ lifts to a closed loop in \mathbb{H}^3 . Hence, by Lemma 4.6, some α_{2i} must have length less than L_0 . \square

4.4. Size of the standard horosection. Consider $T = \partial\mathbb{T}_{\eta_3}$, the boundary of the standard horosection. We show that T is long and skinny, meaning that the length of any slope $\mu \subset T$ different from the one coming from $(\Sigma - \gamma) \cap T$ is very long, provided that $d_W(\alpha, \beta)$ is sufficiently large.

PROPOSITION 4.7. *There exists $c = c(\Sigma) > 0$ such that, for α and β any pair of curves as in Corollary 4.3, we have*

$$d_W(\alpha, \beta) \leq c\ell_T(\mu) + c,$$

for every slope μ in T that is not homotopic to a component of $(\Sigma - \gamma) \cap T$.

Proof. Notice that there is a bound D' , depending only on Σ , on the length of a component of $(\Sigma - \gamma) \cap T$. In fact, the boundary of the cusp of a complete hyperbolic surface of finite area can be bounded in terms of the topological type of the surface only, as the area of the cusp is an increasing function of the length of its boundary, while the total area of the surface only depends on the topological type (e.g. by the Gauss-Bonnet theorem for non-compact surfaces).

Observe now that the intrinsic diameter of T satisfies

$$\text{diam}(T) \leq D := (\inf_{\mu} \ell_T(\mu) + D')/2,$$

where the inf runs over all μ as in the statement. Hence, any two points on T can be joined by a flat geodesic of length at most D .

Consider the covering $p : Q \rightarrow M_f - \gamma$ corresponding to $\pi_1(W = \Sigma - \gamma) < \pi_1(M_f - \gamma)$. By Bonahon's Tameness [Bon86] Q is geometrically and topologically tame. Moreover, since $\Sigma - \gamma$ is not a virtual fiber (for example, because it separates

$M_f - \gamma$), the covering Q is a geometrically finite hyperbolic metric on $W \times \mathbb{R}$ (without accidental parabolics) by Thurston-Canary Covering Theorem [Can96]. Denote by $X \sqcup Y = \partial\mathcal{C}\mathcal{C}(Q)$ the boundary of the convex core.

By work of Minsky [Min10] (see also Theorem 2.1.3 of Bowditch [Bow11]), there exists a uniform quasi-geodesic $l := \alpha_0, \alpha_1, \dots, \alpha_n$ in $\mathcal{C}(W)$, that is

$$\frac{1}{C}|i - j| - C \leq d_W(\alpha_i, \alpha_j) \leq C|i - j| + C$$

for some uniform constant $C > 0$, with the following properties

- (i) Every α_j has a geodesic representative α_j^* in Q of moderate length, that is $\ell(\alpha_j^*) \leq L$ for some uniform $L > 0$.
- (ii) The initial and terminal curves have moderate length on the X and Y boundary components, that is $\ell_X(\alpha_0), \ell_Y(\alpha_n) \leq L$, with L as before.
- (iii) Every curve $\beta \in \mathcal{C}(W)$ such that $\ell_Q(\beta) \leq L$ lies uniformly close to the quasi-geodesic l , that is $d_W(\beta, l) \leq R$ for some uniform $R > 0$.

We notice that

$$n \geq c_0 d_W(\alpha, \beta) - c_0$$

for some uniform $c_0 > 0$. In fact, on one hand, by property (iii), we have $d_W(\alpha, l), d_W(\beta, l) \leq R$. On the other hand, the uniform quasi-geodesic property of $l := \alpha_0, \dots, \alpha_n$ and hyperbolicity of the curve graph $\mathcal{C}(W)$ gives $d_W(\alpha_0, \alpha_n) \geq d_W(\alpha, \beta) - R_1$ for some uniform $R_1 > 0$. Combined together the two properties give the estimate above.

The moderate length geodesics α_j^* are well spaced in Q . This is a consequence of the model manifold technology, which we use in the form of the following result of Bowditch [Bow11] and Brock and Bromberg [BB11].

THEOREM 4.8 (see Theorem 2.1.4 of [Bow11] and Theorem 7.16 of [BB11]). *For every $L > 0$ there exists $A > 1$ such that the following holds. Let Q a hyperbolic structure on $W \times \mathbb{R}$ for which the boundary ∂W is parabolic. Let $Q_0 = Q - \mathbb{T}_{\eta_3}(\partial W)$ be the complement of the standard cusp neighborhoods. Suppose that $\alpha, \beta \in \mathcal{C}(W)$ are simple closed curves represented in Q by closed geodesics of length at most L . Then*

$$\frac{1}{A} d_{\mathcal{C}(W)}(\alpha, \beta) - A \leq \rho_{Q_0}(\mathbb{T}_{\eta_3}(\alpha), \mathbb{T}_{\eta_3}(\beta)) \leq A d_{\mathcal{C}(W)}(\alpha, \beta) + A$$

Where ρ_{Q_0} denotes the η_3 -electric distance in Q_0 .

The η_3 -electric distance $\rho_{Q_0}(x, y)$ between two points $x, y \in Q_0$ is defined to be the infimum of the electric lengths of all paths joining them in Q_0 . The *electric length* of a path δ in Q_0 is the length of the portion of δ that lies in the η_3 -thick part of Q_0 . Observe that, by definition, the electric distance satisfies $\rho_{Q_0} \leq d_{Q_0}$.

The electric distance is also defined on hyperbolic surfaces (W, σ) , and it is a fact (bounded diameter lemma; see Lemma 1.10 of [Bon86]) that for any fixed Margulis constant $\eta > 0$, if (W, σ) has finite area, then the η -electric diameter of (W, σ) is uniformly bounded only in terms of η and $\chi(W)$. This fact applies in our setting: The η_3 -electric diameter of any pleated surface is uniformly bounded.

It follows from Theorem 4.8 and the fact that the sequence $\alpha_0, \dots, \alpha_n$ is a uniform quasi-geodesic that

$$\frac{1}{B}|i - j| - B \leq \rho_{Q_0}(\alpha_i^*, \alpha_j^*) \leq B|i - j| + B.$$

for some uniform $B > 0$. Also notice that X and Y are uniformly close to α_0^* and α_n^* respectively, meaning that

$$\rho_{Q_0}(\alpha_0^*, X), \rho_{Q_0}(\alpha_n^*, Y) \leq \log(2L/\eta_3).$$

This follows from the following standard fact, which we state without proof.

LEMMA 4.9 (see e.g. [BS17, Lemma 4.2]). *Let $\alpha \subset Q$ be a closed curve homotopic to a geodesic α^* . Then*

$$d_Q(\alpha, \mathbb{T}_{\eta_3}(\alpha^*)) \leq \log(2\ell(\alpha)/\eta_3).$$

For simplicity, from now on, we assume $n = 2m$ since it does not affect the argument and simplifies the notation. Consider a pleated surface $G : (W, \sigma) \rightarrow Q$ realizing the middle curve $\alpha_{m=n/2}$ in Q . Under the covering projection G descends to the pleated surface $g := pG : (W, \sigma) \rightarrow M_f - \gamma$ realizing α_m in $M_f - \gamma$.

The manifold Q has two cusps that cover $\mathbb{T}_{\eta_3}(\gamma)$; we fix one of those and denote it by \mathbb{T}_Q^1 , with boundary $T_Q^1 = \partial\mathbb{T}_Q^1$. Then, we choose points $x \in T_Q^1 \cap X$ and $w \in W$ such that $G(w) \in T_Q^1$. Observe that $\text{inj}_x(X), \text{inj}_w(W, \sigma) \geq \eta_3$. As a consequence, since the electric diameter of pleated surfaces is uniformly bounded, we have

$$\rho_{Q_0}(G(w), \alpha_m^*), \rho_{Q_0}(x, \alpha_0^*) \leq K$$

for some uniform $K > 0$. Now connect $p(x)$ to $g(w)$ via a shortest flat geodesic ξ between them on T . Denote by δ the lift of ξ to Q with basepoint $G(w)$. We have $\ell(\delta) \leq D = (\ell_T(\mu) + D')/2$.

In order to conclude the proof of Proposition 4.7 it remains to show that

Claim: We have $\ell(\delta) \geq cn - c$ for some uniform constant $c > 0$.

We divide the proof of this claim into two cases.

Case I. The endpoint z of δ different from $G(w)$ coincides with x .

In this case, we have

$$\ell(\delta) \geq \rho_{Q_0}(x, G(w)) \geq \rho_{Q_0}(\alpha_0^*, \alpha_m^*) - 2K \geq \frac{m}{B} - B - 2K,$$

as desired.

Case II. The endpoint z of δ different from $G(w)$ differs from x .

In this case, let τ be a non peripheral loop on X based at x that has moderate length, say $\ell_X(\tau) \leq L_1$ for some uniform $L_1 > 0$ only depending on W .

We now observe that $p(\tau)$ does not lift to a loop based at z in Q . In fact, we claim that $p(\tau)$ admits a unique lift which is a loop based at a point on the chosen cusp of Q and such a lift is τ , which is based at $x \neq z$. In general, lifts of $p(\tau)$ based at a point on $p^{-1}(p(x)) \cap T_Q^1$ correspond to elements $\kappa \in \pi_1(T, p(x))$ such that $\kappa p(\tau) \kappa^{-1} \in p_*\pi_1(X, x)$.

LEMMA 4.10. *If $c \in \pi_1(\Sigma - \gamma)$ is not peripheral and $\kappa \in \pi_1(M_f - \gamma) - \pi_1(\Sigma - \gamma)$, then*

$$\kappa c \kappa^{-1} \notin \pi_1(\Sigma - \gamma).$$

Proof. Recall that $\pi_1(M_f - \gamma)$ is a free product with amalgamation

$$\pi_1(M_f - \gamma) = \pi_1(H_g - \gamma) *_{\pi_1(\Sigma - \gamma)} \pi_1(H_g - f(\gamma)).$$

Write κ as a reduced word $\kappa = a_1 b_1 \dots a_n b_n$ with $a_j \in \pi_1(H_g - \gamma) - \pi_1(\Sigma - f(\gamma))$ for $j > 1$ and $b_j \in \pi_1(H_g - \gamma) - \pi_1(\Sigma - \gamma)$ for $j < n$. Since $\kappa \notin \pi_1(\Sigma - \gamma)$, either $b_n \notin \pi_1(\Sigma - \gamma)$, or we can take b_n to be the identity and $a_n \notin \pi_1(\Sigma - \gamma)$.

The two cases can be dealt with in the same way, so we only consider the first case, that is, we assume that $b_n \notin \pi_1(\Sigma - \gamma)$. We have

$$\kappa \kappa^{-1} = a_1 b_1 \dots a_n b_n c b_n^{-1} a_n^{-1} \dots b_1^{-1} a_1^{-1}.$$

We claim that $b_n c b_n^{-1} \in \pi_1(H_g - \gamma)$ is not in $\pi_1(\Sigma - \gamma)$ provided that c is not a peripheral element: This follows from the fact that (H_g, γ) is pared acylindrical. In fact, suppose that $b_n c b_n^{-1} \in \pi_1(\Sigma - \gamma)$ and consider the homotopy between c and $b_n c b_n^{-1}$ which takes place in $H_g - \gamma$. We have the following possibilities: If the homotopy is deformable into the cusp, then c would be peripheral, which is ruled out by our initial assumption. As (H_g, γ) is pared acylindrical, if the homotopy is not deformable to the cusp, then it is deformable to the boundary $\Sigma - \gamma$. In this case b_n would be contained in $\pi_1(\Sigma - \gamma)$, which is again a contradiction. Therefore, if c is not peripheral, the word $\kappa \kappa^{-1} = a_1 b_1 \dots a_n (b_n c b_n^{-1}) a_n^{-1} \dots b_1^{-1} a_1^{-1}$ is still reduced, and it contains a term not in $\pi_1(\Sigma - \gamma)$. Hence it cannot represent an element in $\pi_1(\Sigma - \gamma)$. \square

We now return to the main argument for Proposition 4.7. By Lemma 4.10, the loop $g(\tau)$ lifts to an arc η with basepoint z on the preferred cusp and another endpoint u on a different component of $p^{-1}(\mathbb{T}_{\eta_3}(\gamma))$. We now observe that, if η_3 has been chosen sufficiently short in the beginning, no component of $p^{-1}(\mathbb{T}_{\eta_3}(\gamma))$ different from the cusps of Q intersects the convex core $\mathcal{CC}(Q)$.

LEMMA 4.11. *If η_3 is sufficiently small, only depending on the topological type of W , then*

$$p^{-1}\mathbb{T}_{\eta_3}(\gamma) \cap \mathcal{CC}(Q) = \text{cusp}(Q).$$

Proof. Let $p^{-1}\mathbb{T}_{\eta_3}(\gamma) = \bigsqcup_{j \in I} \mathcal{O}_j$ be the full preimage of the cusp under the covering projection $p : Q \rightarrow M_f - \gamma$. Suppose that a component $X \subset \partial\mathcal{CC}(Q)$ of the boundary of the convex core intersects one of the components \mathcal{O}_j of the lift of the Margulis tube. Note that $p : X \rightarrow M_f - \gamma$ is a type preserving pleated surface in the homotopy class of the inclusion $\Sigma - \gamma \subset M_f - \gamma$ and that $p(\mathcal{O}_j) = \mathbb{T}_{\eta_3}(\gamma)$. By Lemma 4.4, the pleated surface $p(X)$ can only intersect $\mathbb{T}_{\eta_3}(\gamma)$ in its η_0 -cuspidal part, for some uniform η_0 , if η_3 has been chosen sufficiently small in the beginning. This means that \mathcal{O}_j intersects the cuspidal part on X and, hence, \mathcal{O}_j is one of the cusps of Q . \square

We are now able to conclude: By Lemmas 4.10 and 4.11, the arc $\delta * \eta$ has an endpoint $G(w) \in T_Q^1 \cap \mathcal{CC}(Q)$ and another one $u \in p^{-1}(\mathbb{T}_{\eta_3}(\gamma)) - \mathcal{CC}(Q)$ outside the convex core. Therefore, it must intersect $\partial\mathcal{CC}(Q) = X \sqcup Y$. Say it intersects X . In particular $\rho_{Q_0}(G(w), X) \leq d_{Q_0}(G(w), X) \leq \ell(\delta * \eta)$, which, combined with the previously established inequalities gives us

$$\begin{aligned} \ell(\delta) + L_1 &\geq \ell(\delta * \eta) \\ &\geq \rho_{Q_0}(G(w), X) \\ &\geq \rho_{Q_0}(\alpha_m^*, \alpha_0^*) - \rho_{Q_0}(X, \alpha_0^*) - 2K \\ &\geq m/B - B - 2K - \log(2L/\eta_3). \end{aligned}$$

This concludes the proof of the claim. \square

4.5. The proof of Theorem 6. Combining Proposition 4.7 and Theorem 2.6, we complete the proof of Theorem 6 as follows.

Proof of Theorem 6. We endow $M_f - \gamma$ with a complete finite volume hyperbolic structure, which exists by the assumption of (H_g, γ) and $(H_g, f(\gamma))$ being pared acylindrical.

Let μ be the flat geodesic on $T = \partial\mathbb{T}_{\eta_3}(\gamma)$ that represents the filling slope needed to pass from $M_f - \gamma$ to M_f (also known as the meridian of γ). By Theorem 2.6 and Theorem 2.7, if $\text{nl}(\mu) \geq 11$, then there exists a hyperbolic metric on M_f for which γ is a geodesic of length $\ell_{M_f}(\gamma) \leq a/\text{nl}(\mu)^2$ for some universal constant $a > 0$.

By Proposition 4.7, we have

$$\ell(\mu) \geq cd_W(\alpha, \beta) - c,$$

which is larger than $\frac{1}{2}cd_W(\mathcal{D}, f\mathcal{D})$ provided that $d_W(\mathcal{D}, f\mathcal{D})$ is sufficiently large. Notice that $\text{Area}(T) \leq \eta_3\ell(\mu)$ whenever $\eta_3 < \ell(\mu)$ (which we have given that $d_W(\mathcal{D}, f\mathcal{D})$ is sufficiently large); hence,

$$\text{nl}(\mu) = \ell(\mu)/\sqrt{\text{Area}(T)} \geq \sqrt{\ell(\mu)/\eta_3}.$$

Thus $\text{nl}(\mu) \geq 11$ if the subsurface projection of the disk sets to W is sufficiently large. This shows that M_f is hyperbolic.

In order to conclude, it remains to bound the length of γ in M_f . This follows again from Theorem 2.6:

$$\ell(\gamma) \leq \frac{a}{\text{nl}(\mu)^2} \leq \frac{a}{\ell(\mu)/\eta_3} \leq \frac{2a\eta_3/c}{d_W(\mathcal{D}, f\mathcal{D})}. \quad \square$$

PART 2

Model manifolds

5. A GLUING SCHEME

Here we outline a construction for the ε -model metric which follows closely ideas of Brock and Dunfield [BD15] and Brock, Minsky, Namazi, and Souto [BMNS16]. At the end of the discussion we formulate a criterion of applicability. In the entire section we fix a gluing map $f \in \text{Diff}^+(\Sigma)$.

5.1. Assembling simple pieces. We identify a tubular neighborhood of the Heegaard surface $\Sigma \subset M_f = H_g \cup_f H_g$ with $\Sigma \times [1, 4]$ such that Σ is identified with $\Sigma \times \{2.5\}$.

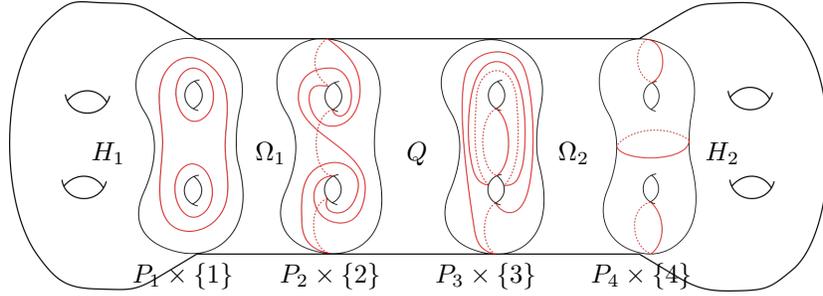


FIGURE 1. Gluing.

Given pants decompositions P_1 , P_2 , P_3 , and P_4 in Σ , we consider

$$\mathbb{M} := M_f - (P_1 \times \{1\} \sqcup P_2 \times \{2\} \sqcup P_3 \times \{3\} \sqcup P_4 \times \{4\}).$$

We isolate the five pieces (see Figure 1)

$$\begin{aligned} H_1 &:= H_g - (\Sigma \times (1, 2.5] \sqcup P_1 \times \{1\}), \\ \Omega_1 &:= \Sigma \times [1, 2] - (P_1 \times \{1\} \sqcup P_2 \times \{2\}), \\ Q &:= \Sigma \times [2, 3] - (P_2 \times \{2\} \sqcup P_3 \times \{3\}), \\ \Omega_2 &:= \Sigma \times [3, 4] - (P_3 \times \{3\} \sqcup P_4 \times \{4\}), \text{ and} \\ H_2 &:= H_g - (\Sigma \times [2.5, 4] \sqcup P_4 \times \{4\}). \end{aligned}$$

LEMMA 5.1. *Provided that*

$$(A) \quad \begin{aligned} & (H_g - \Sigma \times (1, 2.5], P_1 \times \{1\}), \\ & (\Sigma \times [1, 2], P_1 \times \{1\} \sqcup P_2 \times \{2\}), \\ & (\Sigma \times [2, 3], P_2 \times \{2\} \sqcup P_3 \times \{3\}), \\ & (\Sigma \times [3, 4], (P_3 \times \{3\} \sqcup P_4 \times \{4\})), \text{ and} \\ & (H_g - \Sigma \times [2.5, 4], P_4 \times \{4\}) \end{aligned}$$

are all pared acylindrical manifolds, there exists a hyperbolic metric on \mathbb{M} such that the restriction of the metric to each of the pieces is isometric to the convex core of the unique maximally cusped hyperbolic structure on the piece.

Proof. As all the pieces are all pared acylindrical manifolds, we can endow each of them with a complete hyperbolic metric with totally geodesic boundary and rank one cusps at the pants decompositions. Indeed, we can take the convex cores of the (up to isotopy unique) maximally cusped hyperbolic structure on the pieces, which exist by Theorem 2.3.

Since there is only one hyperbolic metric up to isotopy on a triply punctured sphere (such as the connected components of the complement of the pared loci in the boundaries of the pieces), we can arrange the metrics on the pieces such that they agree on the intersections and match together to give a complete finite volume hyperbolic metric on \mathbb{M} . \square

In order to pass from \mathbb{M} to the closed 3-manifold M_f , we have to perform Dehn fillings on each cusp. This is the second step of the construction. The filling slopes are completely determined by the identification of \mathbb{M} with the drilled M_f : They are the meridians γ of small tubular neighborhoods of the curves in $\alpha \times \{j\} \subset P_j \times \{j\}$ inside $\Sigma \times [1, 4]$. Theorem 2.6 gives us sufficient conditions to guarantee that M_f has a hyperbolic metric. Theorem 2.8 gives sufficient conditions to ensure that \mathbb{M} is K -bilipschitz to M_f away from its cusps. We make this precise in the next subsection.

5.2. A filling criterion. To be able to apply Theorems 2.6, 2.8, and 2.7, we have to certify that the filling slopes have large normalized length. We now discuss a criterion to check this condition.

We define

$$\begin{aligned} \mathbb{N}_1 &:= H_1 \cup \Omega_1 = H_g - P_2 \times \{2\}, \\ \mathbb{N}_2 &:= H_2 \cup \Omega_2 = H_g - P_3 \times \{3\}, \text{ and} \\ \mathbb{Q} &:= \Omega_1 \cup Q \cup \Omega_2 = \Sigma \times [1, 4] - (P_1 \times \{1\} \sqcup P_2 \times \{2\} \sqcup P_3 \times \{3\} \sqcup P_4 \times \{4\}). \end{aligned}$$

The curves in P_1 and P_4 represent rank one cusps on $\partial\mathbb{Q}$ while the curves in P_2 and P_3 represent rank two cusps. Similarly, the curves P_2 and P_3 represent rank one cusps of \mathbb{N}_1 and \mathbb{N}_2 while the curves in P_1 and P_4 are cusps of rank two.

We now try to understand what happens when we Dehn fill the rank two cusps in $\mathbb{Q}, \mathbb{N}_1, \mathbb{N}_2$. Choosing the filling slopes to be the canonical meridians of P_1, P_2, P_3, P_4

in M_f , the Dehn filled manifolds are

$$\begin{aligned}\mathbb{Q}^{\text{fill}} &= \Sigma \times [1, 4] - (P_1 \times \{1\} \sqcup P_4 \times \{4\}), \\ \mathbb{N}_1^{\text{fill}} &= H_g - (\Sigma \times (2, 2.5] \sqcup P_2 \times \{2\}), \text{ and} \\ \mathbb{N}_2^{\text{fill}} &= H_g - (\Sigma \times [2.5, 3) \sqcup P_3 \times \{3\}).\end{aligned}$$

We endow $\mathbb{Q}^{\text{fill}}, \mathbb{N}_1^{\text{fill}}, \mathbb{N}_2^{\text{fill}}$ with a complete hyperbolic metric with totally geodesic boundary and rank one cusps at the curves $P_1 \times \{1\} \sqcup P_4 \times \{4\}, P_2 \times \{2\}, P_3 \times \{3\}$. Again, this is possible provided that

$$(B) \quad \begin{aligned} &(\Sigma \times [1, 4], P_1 \times \{1\} \sqcup P_4 \times \{4\}), \\ &(H_g - \Sigma \times (2, 2.5], P_2 \times \{2\}), \text{ and} \\ &(H_g - \Sigma \times [2.5, 3), P_3 \times \{3\})\end{aligned}$$

are all pared acylindrical.

Recall that our goal is to show that the filling slopes we singled out on the rank two cusps of $\mathbb{Q}, \mathbb{N}_1, \mathbb{N}_2$ have very large normalized length. First, observe that $\mathbb{Q}, \mathbb{N}_1, \mathbb{N}_2$ isometrically embed in their doubles $D\mathbb{Q}, D\mathbb{N}_1, D\mathbb{N}_2$ which are finite volume hyperbolic 3-manifolds. Checking that the filling slopes of the rank two cusps of $\mathbb{Q}, \mathbb{N}_1, \mathbb{N}_2$ have large normalized length is the same as checking that they have large normalized length in $D\mathbb{Q}, D\mathbb{N}_1, D\mathbb{N}_2$.

Also notice that

$$\mathbb{Q}, \mathbb{N}_1, \mathbb{N}_2 \subset \mathbb{Q}^{\text{fill}}, \mathbb{N}_1^{\text{fill}}, \mathbb{N}_2^{\text{fill}} \subset D\mathbb{Q}^{\text{fill}}, D\mathbb{N}_1^{\text{fill}}, D\mathbb{N}_2^{\text{fill}}.$$

Again, the doubles are finite volume hyperbolic 3-manifolds.

The idea is as follows. Suppose that we can find hyperbolic metrics on $\mathbb{Q}^{\text{fill}}, \mathbb{N}_1^{\text{fill}}, \mathbb{N}_2^{\text{fill}}$ such that the curves

$$P_2 \times \{2\} \sqcup P_3 \times \{3\}, P_1 \times \{1\}, P_4 \times \{4\} \subset \mathbb{Q}^{\text{fill}}, \mathbb{N}_1^{\text{fill}}, \mathbb{N}_2^{\text{fill}}$$

are very short geodesics. Then, $D\mathbb{Q}, D\mathbb{N}_1, D\mathbb{N}_2$ are exactly the complete finite volume hyperbolic structures on the complement of those geodesic links $D\mathbb{Q}^{\text{fill}} - P_1 \times \{1\} \sqcup P_4 \times \{4\}, D\mathbb{N}_1^{\text{fill}} - P_2 \times \{2\}, D\mathbb{N}_2^{\text{fill}} - P_3 \times \{3\}$.

By Theorem 2.7, the condition of having large normalized length for the meridians of $P_1 \times \{1\} \sqcup P_4 \times \{4\}, P_2 \times \{2\}, P_3 \times \{3\}$ is equivalent to the condition of being very short for those geodesics.

The previous discussion leads us to make the following definition.

DEFINITION (Filling Criterion). We say that pants decompositions P_1, P_2, P_3, P_4 satisfy the filling criterion with parameter $\eta > 0$, if (A) and (B) are all pared acylindrical and the following hold:

- (1) $P_1 \times \{1\}$ is isotopic to its geodesic realization in $\mathbb{N}_1^{\text{fill}}$ and each component of this geodesic link has length at most η ,
- (2) $P_4 \times \{4\}$ is isotopic to its geodesic realization in $\mathbb{N}_2^{\text{fill}}$ and each component of this geodesic link has length at most η , and
- (3) $P_2 \times \{2\} \sqcup P_3 \times \{3\}$ is isotopic to its geodesic realization in \mathbb{Q}^{fill} and each component of this geodesic link has length at most η .

By the above discussion we have the following. For $L > 0$, if P_1, P_2, P_3, P_4 satisfy the filling criterion with parameter

$$\eta := \min \left\{ \frac{2\pi}{L^2 + 16.17}, \frac{0.1396}{3g - 3} \right\},$$

then the normalized length of the filling slopes of \mathbb{M} corresponding to $P_1 \times \{1\}, P_2 \times \{2\}, P_3 \times \{3\}, P_4 \times \{4\} \subset M_f$ is at least L . Hence, by Theorems 2.6 and 2.7 (combined with Theorem 2.8 for the furthermore-part), we established the following.

PROPOSITION 5.2. *Fix $K \in (1, 2)$. There exists $\eta > 0$ such that the following holds. Suppose that there are four pants decompositions P_1, P_2, P_3, P_4 such that the filling criterion with parameter η is satisfied. Then, M_f admits a hyperbolic metric such that $\Gamma := \bigsqcup_{j=1,2,3,4} P_j \times \{j\}$ is a geodesic link and each component has length at most η . Furthermore, if \mathbb{M} denotes the unique finite volume hyperbolic structure on $M_f - \Gamma$, then we have a K -bilipschitz diffeomorphism of pairs*

$$\left(\mathbb{M} - \bigsqcup_{\alpha \in P_1 \cup P_2 \cup P_3 \cup P_4} \mathbb{T}_{\eta_3}(\alpha) \right) \cong \left(M_f - \bigsqcup_{\alpha \in P_1 \cup P_2 \cup P_3 \cup P_4} \mathbb{T}_{\eta_3}(\alpha) \right).$$

We conclude with a small remark. The model manifold technology of Minsky [Min10] and Brock, Canary and Minsky [BCM12], provides several tools to locate and measure the length of the geodesic representatives of $P_2 \times \{2\}$ and $P_3 \times \{3\}$ in \mathbb{Q}^{fill} . However, the same technology is not available for handlebodies. This is the place where the difficulties arise. We get around the lack of a model for handlebodies by restricting ourselves to the geometry of the collars of $\mathbb{N}_1^{\text{fill}}, \mathbb{N}_2^{\text{fill}}$ where, in some cases, we have a good amount of control (as in [HV22]).

6. A FAMILY OF EXAMPLES

In this section we construct many examples satisfying Proposition 5.2 (see Proposition 6.4). Later, in Section 7, we will show that this family is generic from the point of view of random walks.

6.1. A family of hyperbolic mapping tori. For every pants decomposition $P \subset \Sigma$, we first construct hyperbolic mapping tori

$$T_\psi = \Sigma \times [0, 1] / (x, 0) \sim (\psi(x), 1)$$

with the property that the multicurve $P \times \{0\} \subset \Sigma \times \{0\}$ is a very short geodesic link. In particular, the infinite cyclic covering of T_ψ (see Figure 2) is a hyperbolic structure \hat{T}_ψ on $\Sigma \times \mathbb{R}$ where the multicurves $\psi^n P \times \{n\} \subset \Sigma_n := \Sigma \times \{n\}$ are very short geodesic links for every $n \in \mathbb{Z}$.

DEFINITION (Mapping Class with a Short Pants Decomposition). Let $\eta > 0$ and let P be a pants decomposition of Σ . If $\psi \in \text{Mod}(\Sigma)$ is a mapping class such that its mapping torus

$$T_\psi := (\Sigma \times [0, 1] / (x, 0) \sim (\psi x, 1))$$

has a hyperbolic metric with respect to which $P \times \{0\} \subset T_\psi$ is a geodesic link of length at most η , then we call ψ a *pseudo-Anosov mapping class with a short pants decomposition P of length at most η* .

These objects are abundant. Concretely, we have the following.

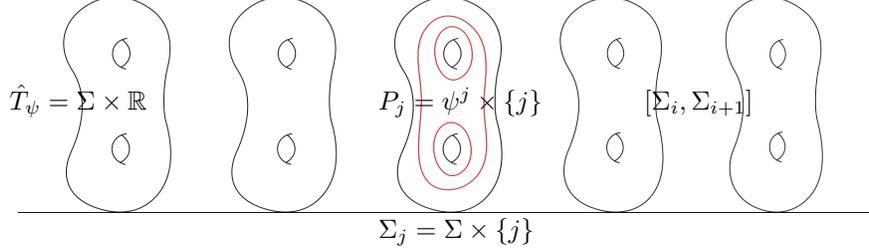


FIGURE 2. Infinite cyclic covering.

LEMMA 6.1. *For every pants decomposition $P \subset \Sigma$ and $\varepsilon > 0$, there exists a pseudo-Anosov mapping class ψ with a short pants decomposition P of length at most ε .*

Proof. Let $\phi \in \text{Mod}(\Sigma)$ be a mapping class such that

$$(\Sigma \times [0, 1], P \times \{0\} \sqcup \phi P \times \{1\})$$

is pared acylindrical. For example, one can choose ϕ to be a large power of any pseudo-Anosov. Consider the convex core Q of the maximally cusped structure $Q(P, \phi P)$. The boundary ∂Q consists of totally geodesic hyperbolic three punctured spheres that are paired according to ϕ . We glue them together isometrically as prescribed by the pairing. The glued manifold is a finite volume hyperbolic 3-manifold diffeomorphic to

$$T_\phi - P \times \{0\} = (\Sigma \times [0, 1]/(x, 0) \sim (\phi x, 1)) - P \times \{0\}.$$

The curves in $P \times \{0\}$ represent rank two cusps. Observe that for each $\alpha \in P$ any choice of corresponding boundary torus $\partial\mathbb{T}(\alpha) \subseteq T_\phi - P \times \{0\}$ has a preferred meridian m_α (given as the boundary of an essential disc in $\mathbb{T}(\alpha) \subseteq T_\phi$) and a preferred longitude l_α (given as a curve in the fiber $\Sigma \times \{0\} \subseteq T_\phi$ that is parallel to α).

If we perform Dehn surgeries with slopes $m_\alpha + kl_\alpha$, the resulting closed manifold will be diffeomorphic to the mapping torus T_ψ , where $\psi = \phi\delta_P^k$ and δ_P is a Dehn multitwist about the pants decomposition P . By Thurston's Hyperbolic Dehn Surgery [BP92, Chapter E.6]) (or by invoking Theorems 2.6 and 2.7), if k is large enough, then the resulting manifold carries a hyperbolic metric for which the core curves of the added solid tori are very short geodesics. Hence, for k large enough, ψ is a pseudo-Anosov mapping class with a short pants decomposition P of length at most ε . \square

6.2. The criterion. We now describe some criteria (Propositions 6.2 and 6.3) to copy and paste the geometry of a hyperbolic mapping torus with a short pants decomposition into the collar geometry of maximally cusped handlebodies H and trivial bundles $\Sigma \times [0, 1]$. These structures will automatically satisfy the filling criterion for a suitable choice of pants decompositions (Proposition 6.4).

In the statement of the following proposition we will think of the repelling lamination of a pseudo-Anosov as an element of the Gromov boundary $\partial\mathcal{C}$ of \mathcal{C} , see [Kla18].

PROPOSITION 6.2. *There exists $\eta > 0$ with the following property. Let ψ be a pseudo-Anosov mapping class with repelling lamination $\lambda \in \partial\mathcal{C}$. There exists $m = m(\psi) > 0$ such that the following holds. Suppose that*

- P is a η -short pants decomposition for ψ ,
- H is a handlebody with boundary $\partial H = \Sigma$ and disk set $\mathcal{D}_H \subset \mathcal{C}$,
- for every $\delta \in \mathcal{D}_H$ we have

$$(\lambda|\delta)_P \geq m$$

where the left hand side is the Gromov product of δ and λ based at P .

Consider the maximally cusped structure $N = H(P)$ on $\text{int}(H)$. Identify a collar of $\Sigma = \partial H$ with $\Sigma \times [0, 2]$ where $\partial H = \Sigma \times \{0\}$. Then the geodesic realization of $\psi^{-1}P \times \{1\}$ in N is isotopic to $\psi^{-1}P \times \{1\}$ and has length at most 2η .

We will prove Proposition 6.2 at the end of the section. The idea is that if m is large enough then the geometry of N near the boundary of its convex core looks like the geometry of the region $[\Sigma_{-2}, \Sigma_0] \subset \hat{T}_\psi$ and there the geodesic realization of $\psi^{-1}P$ lies as a short geodesic link on the level surface Σ_1 (see Figure 3).

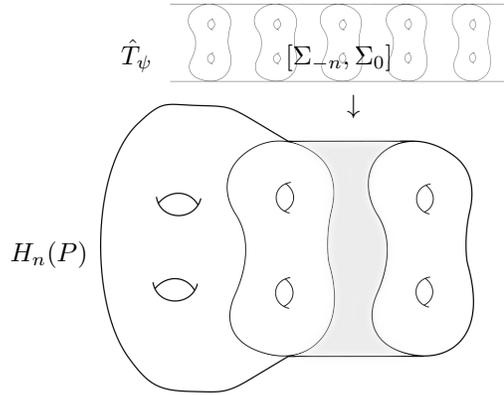


FIGURE 3. Collars handlebodies.

Similarly, we will also use the following variation (see also Figure 4):

PROPOSITION 6.3. *There exists $\eta > 0$ with the following property. Let ψ be a pseudo-Anosov mapping class with repelling lamination $\lambda \in \partial\mathcal{C}$. There exists $m = m(\psi) > 0$ such that the following holds: Suppose that*

- P is a η -short pants decomposition for ψ ,
- R is a pants decomposition such that

$$(\lambda|R)_P \geq m$$

where the left hand side is the Gromov product of R and λ based at P .

Consider the maximally cusped structure $Q = Q(P, R)$ on $\text{int}(\Sigma \times [0, 3])$. Then the geodesic realization of $\psi P \times \{1\}$ in Q is isotopic to $\psi P \times \{1\}$ and has length at most 2η .

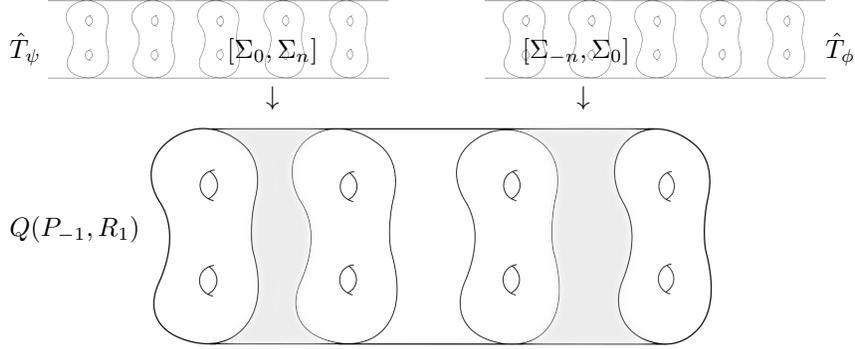


FIGURE 4. Collars I-bundles.

Having in mind the application to the random 3-manifold setup, we combine the two propositions in the following criterion which produces many examples of pants decompositions that satisfy the assumptions of Proposition 5.2:

PROPOSITION 6.4. *There exists $\eta > 0$ with the following property. Let $o \in \mathcal{D}$ be a basepoint. Let ψ, ϕ be pseudo-Anosov mapping classes with short pants decompositions P, R of length at most η . For every $\delta > 0$ there exist $B = B(\delta)$ and $h = h(\delta, \psi, \phi)$ such that the following holds.*

Let f be a gluing map. Set $S_0 := o, S_5 := fo$. Suppose that there are four points $S_1, S_2, S_3, S_4 \in [o, fo] \subset \mathcal{C}$ with the following properties:

- (i) *Along the geodesic segment $[o, fo]$ we have $S_j < S_{j+1}$ and $d_{\mathcal{C}}(S_j, S_{j+1}) \geq h$.*
- (ii) *There exist mapping classes τ and ν such that $[S_1, S_2]$ and $[S_3, S_4]$ δ -fellow travel the translates τl_ψ and νl_ϕ of the uniform quasi-axes $l_\psi, l_\phi : \mathbb{R} \rightarrow \mathcal{C}$ of ψ and ϕ respectively.*
- (iii) *We have $d_{\mathcal{C}}([S_1, S_2], \mathcal{D}) \geq h$ and $d_{\mathcal{C}}([S_3, S_4], f\mathcal{D}) \geq h$.*

Then, there exist four pants decompositions

$$(P_1, P_2, P_3, P_4)$$

such that:

- *They satisfy the filling criterion with parameter η for the gluing map f .*
- *We have $d_{\mathcal{C}}(P_2, [S_1, S_2]), d_{\mathcal{C}}(P_3, [S_3, S_4]) \leq B$,*
- *We have $P_2 = \tau\psi\tau^{-1}P_1$ and $P_4 = \nu\phi\nu^{-1}P_3$.*

In Section 7, we will use Proposition 6.4 to detect whether, for a random gluing map f , M_f can be described as one of the examples we constructed simply by staring at the geometry of the segment $[o, fo]$, where $o \in \mathcal{D}$ is some basepoint that we will fix once and for all.

6.3. The proof of Proposition 6.4. Assuming Propositions 6.2 and 6.3, we give a proof of Proposition 6.4.

Proof of Proposition 6.4. Let $P_n, R_n := \tau\psi^n P, \nu\phi^n R$.

Let $L := \max\{L(\psi), L(\phi)\} > 0$ be the maximum of the translation length of ψ, ϕ on \mathcal{C} (it is positive as ψ, ϕ are pseudo-Anosov). Note that $L(\tau\psi\tau^{-1}) = L(\psi)$ and

$L(\nu\phi\nu^{-1}) = L(\phi)$. Let $h := 2nL$, with n much larger than the $m(\psi)$'s and $m(\phi)$'s as given by Propositions 6.2 and 6.3.

As the geodesic representatives of P_n, R_n are very short in the infinite cyclic coverings $\hat{T}_\psi, \hat{T}_\phi$ of the hyperbolic mapping tori T_ψ, T_ϕ , it follows from work of Minsky [Min10] combined with the structural results of Masur and Minsky [MM99], [MM00], that there exists a uniform constant B , only depending on Σ , and $t, r \in \mathbb{R}$ such that

$$d_{\mathcal{C}}(P, l_\psi(t)), d_{\mathcal{C}}(R, l_\phi(r)) \leq B$$

for some t, r and thus, by equivariance

$$d_{\mathcal{C}}(P_j, \tau l_\psi(t + jL(\psi))), d_{\mathcal{C}}(R_i, \tau l_\phi(r + iL(\phi))) \leq B$$

for every $j, i \in \mathbb{Z}$. Up to reparametrization, we can assume $t = r = 0$ and simplify the notation by introducing

$$\alpha := l_\psi(0), \beta := l_\phi(0)$$

and

$$\alpha_j := \tau l_\psi(jL(\psi)), \beta_i := \nu l_\phi(iL(\phi)).$$

By the above discussion $d_{\mathcal{C}}(P_j, \alpha_j), d_{\mathcal{C}}(R_i, \beta_i) \leq B$ for every $i, j \in \mathbb{Z}$. (Notice that the α_j and β_i are linearly aligned along τl_ψ and νl_ϕ .)

By assumption (ii) of Proposition 6.4 and the choice of $h \geq 2nL(\psi), 2nL(\phi)$, we can find $a, b \in \mathbb{Z}$ such that the following hold:

- (1)
 - $\tau l_\psi[(a - n)L(\psi), (a + n)L(\psi)]$ is in a δ -neighborhood of $[S_1, S_2]$.
 - $\nu l_\phi[(b - n)L(\phi), (b + n)L(\phi)]$ is in a δ -neighborhood of $[S_3, S_4]$.

Again, up to renumbering, we can assume $a = b = 0$.

We now show that

$$(P_{-1}, P_0, R_0, R_1)$$

satisfies the filling criterion with parameter $2\eta > 0$, where η is the minimum of the η s in Propositions 6.2 and 6.3.

We first make sure that the eight building blocks (A) and (B) are pared acylindrical (and hence admit maximally cusped structures). As the maps ψ, ϕ are pseudo-Anosov, the pared acylindrical assumptions are automatically satisfied for (P_{-1}, P_0) and (R_0, R_1) . In order to check the others, we use Lemma 2.5.

Let us check that $d_{\mathcal{C}}(P_{-1}, \mathcal{D}), d_{\mathcal{C}}(P_0, \mathcal{D}) \geq 3$: We know that P_{-1}, P_0 are uniformly close to α_{-1}, α_0 which, in turn, are δ -close to $[S_1, S_2]$. By the assumption (iii), we have $d_{\mathcal{C}}([S_1, S_2], \mathcal{D}) \geq h$. Thus, provided that h is large enough, the claim holds. The same argument shows that $d_{\mathcal{C}}(R_0, f\mathcal{D}), d_{\mathcal{C}}(R_1, f\mathcal{D}) \geq 3$ provided that h is sufficiently large.

Let us now check that $d_{\mathcal{C}}(P_{-1}, R_1), d_{\mathcal{C}}(P_0, R_0) \geq 3$: Consider (P_0, R_0) first, the argument for the other pair (P_{-1}, R_1) is completely analogous. We know that P_0, R_0 are uniformly close to α_0, β_0 which in turn are δ -close to points on $[S_1, S_2]$ and $[S_3, S_4]$. Thus, up to a uniform additive constant, $d_{\mathcal{C}}(P_0, R_0)$ is larger than $d_{\mathcal{C}}(S_2, S_3) \geq h$.

Hence, the pared acylindrical assumptions are satisfied and the eight pieces admit maximally cusped structures. We now check the length-isotopy properties as given by the filling criterion.

Consider the maximally cusped structures N_1, N_2 , and Q , where N_1, N_2 are hyperbolic structure on the interior of handlebodies with boundary Σ and disk sets $\mathcal{D}, f\mathcal{D}$ with cusps at P_0, R_0 , while Q is a hyperbolic structure on the interior of $\Sigma \times [0, 3]$ with cusps at $P_{-1} \times \{0\} \sqcup R_1 \times \{3\}$. Identify a collar of the boundary Σ of the two handlebodies with $\Sigma \times [0, 2]$ in such a way that the boundary is $\Sigma \times \{0\}$. Then

- (1) The collection $P_{-1} \times \{1\}$ is isotopic to its geodesic realization in N_1 and has length at most η .
- (2) The collection $R_1 \times \{1\}$ is isotopic to its geodesic realization in N_2 and has length at most η .
- (3) The collection $P \times \{1\} \sqcup R \times \{2\}$ is isotopic to its geodesic realization in Q and has length at most η .

We use Proposition 6.2 to check (1) and (2). We use Proposition 6.3 to check (3).

Note that τ^{-1} (resp. ν^{-1}) induces a natural orientation preserving isometry between the maximally cusped hyperbolic structure on the handlebody with boundary Σ and disk set \mathcal{D} (resp. $f\mathcal{D}$) associated with the pants decomposition P_0 (resp. R_0) and the maximally cusped hyperbolic structure on the handlebody with boundary Σ and disk set $\tau^{-1}\mathcal{D}$ (resp. $\nu^{-1}f\mathcal{D}$) associated with the pants decomposition $\tau^{-1}P_0$ (resp. $\tau^{-1}R_0$).

Therefore, according to Proposition 6.2 to prove (1) and (2) we only need that:

- (a) $(\lambda_\psi^- | \tau^{-1}\mathcal{D})_P \geq m(\psi)$.
- (b) $(\lambda_\phi^+ | \nu^{-1}f\mathcal{D})_R \geq m(\phi)$.

Here $\lambda_\psi^-, \lambda_\phi^+$ are the repelling laminations of ψ, ϕ^{-1} thought as points on the Gromov boundary of \mathcal{C} and the left hand sides of the inequalities are Gromov products based at P, R .

Properties (a) and (b) are similar. We just prove the first. By properties of the Gromov product on Gromov hyperbolic spaces (such as \mathcal{C}) we have

$$(\lambda_\psi^- | \tau^{-1}\mathcal{D})_P \geq \min\{(\tau^{-1}P_{-n} | \tau^{-1}\mathcal{D})_P, (\lambda_\psi^- | \tau^{-1}P_{-n})_P\} - K$$

for some uniform constant K . Regarding $(\tau^{-1}P_{-n} | \tau^{-1}\mathcal{D})_P = (P_{-n} | \mathcal{D})_{P_0}$, we claim that, up to some uniform additive constant, it is larger than $d_{\mathcal{C}}(P_{-n}, P_0)$. The latter, in turn, is comparable to $d_{\mathcal{C}}(\alpha_{-n}, \alpha_0) \simeq nL(\psi)$ by the above observations. In order to prove the claim, observe that $[P_{-n}, P_0]$ lies uniformly close to $[S_1, S_2]$ (with P_{-n} closer to S_1 and P_0 closer to S_2), which is far from \mathcal{D} (with S_2 further away than S_1). By hyperbolicity of the curve graph and quasi-convexity of the disk set, this implies that the segment $[P_0, P_{-n}]$ uniformly fellow travels the shortest geodesic connecting P_0 to \mathcal{D} . This implies the claim. Regarding the second term, we have

$$(\lambda_\psi^- | \tau^{-1}P_{-n})_P = (\lambda_\psi^- | \psi^{-n}P)_P.$$

since, by definition, $\tau^{-1}P_{-n} = \psi^{-n}P$. The Gromov product is computed up to a uniform additive error by

$$(\lambda_\psi^- | \psi^{-n}P)_P = \frac{1}{2} [d_{\mathcal{C}}(l_\psi(-kL(\psi)), P) + d_{\mathcal{C}}(P, \psi^{-n}P) - d_{\mathcal{C}}(\psi^{-n}P, l_\psi(-kL(\psi)))]$$

where k is sufficiently large. We recall that $l_\psi(-kL(\psi)) = \psi^{-k}l_\psi(0) = \psi^{-k}\alpha$ and, by our choices, $d_{\mathcal{C}}(\alpha, P) \leq B$ is uniformly bounded. Thus the right hand side is,

up to a uniform additive constant greater or equal to

$$\frac{1}{2} [d_{\mathcal{C}}(\psi^{-k}\alpha, \alpha) + d_{\mathcal{C}}(\alpha, \psi^{-n}\alpha) - d_{\mathcal{C}}(\psi^{-n}\alpha, \psi^{-k}\alpha)]$$

which is comparable to $L(\psi)n$ as $\alpha = l_{\psi}(0)$ and ψ acts on l_{ψ} by translations by $L(\psi)$. Thus we proved that up to a uniform additive constant we have that $(\lambda_{\psi}^{-}|\tau^{-1}\mathcal{D})_P$ is at least $nL(\psi)$.

We now discuss point (3).

Again, observe that τ^{-1} (resp. ν^{-1}) induces a natural orientation preserving isometry between the maximally cusped hyperbolic structure on $\Sigma \times [0, 3]$ associated with the pants decomposition P_1, R_1 and the maximally cusped hyperbolic structure on the same manifold associated with the pants decomposition $\tau^{-1}P_1, \tau^{-1}R_1$ (resp. $\nu^{-1}P_1, \nu^{-1}R_1$). Hence, according to Proposition 6.3, for point (3) in order to know that $P \times \{1\}$ and $R \times \{2\}$ are isotopic to their geodesic realizations in $Q = Q(P_{-1}, R_1)$ we need that:

- (a') $(\lambda_{\psi}^{+}|\tau^{-1}R_1)_P \geq m(\psi)$.
- (b') $(\lambda_{\phi}^{-}|\nu^{-1}P_1)_R \geq m(\phi)$.

Properties (a') and (b') are also similar and follow exactly the same lines as the previous computation done for (a) and (b). In order to avoid repetitions, we omit these computations and focus instead on an additional issue, namely that we also need to know that $P \times \{1\}$ and $R \times \{2\}$ can be *simultaneously* isotoped to their geodesic realizations. This can be done as follows: Let β_P, β_R be the geodesic realizations of $P \times \{1\}, R \times \{2\}$ in $Q(P_{-1}, R_1)$. By Theorem 4.8, the distance between two ε -Margulis neighborhoods $\mathbb{T}_{\varepsilon}(\alpha), \mathbb{T}_{\varepsilon}(\beta)$ is bounded by $d_{\mathcal{C}}(\alpha, \beta)$.

This implies that the ε -Margulis neighborhoods around the curves of β_P have (electric) distance from $\partial_+\mathcal{CC}_0(Q)$ uniformly bounded in terms of

$$D_P := \max\{d_{\mathcal{C}}(\gamma, \gamma') \mid \gamma \subset P_{-1}, \gamma' \subset P\}.$$

Similarly, the (electric) distance of the ε -Margulis tubes around the curves in β_R from $\partial_-\mathcal{CC}_0(Q)$ is bounded in terms of

$$D_R := \max\{d_{\mathcal{C}}(\gamma, \gamma') \mid \gamma \subset R_1, \gamma' \subset R\}.$$

On the other hand the (electric) distance between $\partial_+\mathcal{CC}_0(Q)$ and $\partial_-\mathcal{CC}_0(Q)$ is bounded from below by $d_{\mathcal{C}}(P_{-1}, R_1)$. Recall that $d_{\mathcal{C}}(P_{-1}, R_1) \geq d_{\mathcal{C}}(S_2, S_3) \geq h$ up to some uniform additive constant (as P_{-1}, R_1 are uniformly close to $[S_1, S_2], [S_3, S_4]$ respectively). If h is much larger than $D_P + D_R$, then we can isotope β_P into $\partial_+\mathcal{CC}_0(Q)$ and β_R into $\partial_-\mathcal{CC}_0(Q)$ simultaneously. This comes from the following:

Claim. It is possible to find a boundary parallel surface $S \subset \mathcal{CC}_0(Q)$ that separates β_P from β_R .

Sketch of proof. First, as the distance between $\partial_+\mathcal{CC}_0(Q)$ and $\partial_-\mathcal{CC}_0(Q)$ is very large, we can find a pleated surface that separates β_P from β_R . This can be done as follows: By the Filling Theorem of Canary [Can96] and Thurston [Thu80], one can find a pleated surface S' passing uniformly near a point in $\mathcal{CC}_0(Q)$ whose distance from $\partial\mathcal{CC}_0(Q)$ is equal to $d(\partial_+\mathcal{CC}_0(Q), \partial_-\mathcal{CC}_0(Q))/2$. Such a pleated surface has uniformly bounded intrinsic (electric) diameter, in particular, it is disjoint from the D_P, D_R -(electric) neighborhoods U_+, U_- of $\partial_+\mathcal{CC}_0(Q), \partial_-\mathcal{CC}_0(Q)$ provided that $d(\partial_+\mathcal{CC}_0(Q), \partial_-\mathcal{CC}_0(Q)) \asymp d_{\mathcal{C}}(P_{-1}, R_1) \geq h$ is large enough. As $\mathcal{CC}_0(Q) \simeq \Sigma \times \{0\}$

and the pleated surface S' is homotopic to $\Sigma \times \{0\}$, we have that S' separates such neighborhoods U_+, U_- . Using standard techniques (see Freedman, Hass, and Scott [FHS83]), we can find in a small neighborhood of S' , still disjoint from U_+, U_- , an embedded surface $S \subset \mathcal{CC}_0(Q)$ homotopic to S' . The fact that S is boundary parallel comes from the classification of incompressible surfaces in $\mathcal{CC}_0(Q) \simeq \Sigma \times [0, 1]$. \square

This concludes the proof of the proposition. \square

6.4. The proof of Proposition 6.2. The proof of Proposition 6.3 follows exactly the same strategy as the proof of Proposition 6.2, except that the latter has additional topological issues that need to be addressed, so we only fully spell out the proof of Proposition 6.2.

We adapt a strategy of Namazi and Souto [NS09] and Brock, Minsky, Namazi, and Souto [BMNS16] and argue by contradiction.

We briefly outline the argument. Suppose we have a sequence of handlebodies H_n with boundary $\partial H_n = \Sigma$ and disk sets $\mathcal{D}_n \subset \mathcal{C}$ such that the sequence of Gromov products diverges

$$(2) \quad \inf_{\delta_n \in \mathcal{D}_n} \{(\lambda|\delta_n)_P\} \rightarrow \infty,$$

but the geodesic realization of the pants decomposition $\psi^{-1}P$ in the maximally cusped handlebody $N_n = H_n(P)$ does not satisfy the conclusion of the proposition. (In order to prove Proposition 6.3 one analogously considers a sequence $Q_n = Q(P, R_n)$ with $(\lambda|R_n)_P \rightarrow \infty$.) We show that the sequence of hyperbolic manifolds N_n (or Q_n for Proposition 6.3) converges to a hyperbolic manifold Q diffeomorphic to $\Sigma \times (-\infty, 0)$ on which the curves $\psi^{-j}P \times \{j\}$ are geodesic links of length at most η for $j > 0$ and the curves in $P \times \{0\}$ represent rank one cusps. Convergence is essentially in the sense of the algebraic and geometric convergence. This implies the following. For fixed $T > 0$, if n is large enough, then there is a 2-bilipschitz embedding $\phi_n : \Sigma \times [-T, -1/T] \rightarrow N_n$ such that the restriction of ϕ_n to $\Sigma \times \{-1\}$ is homotopic to the inclusion $f_n : \Sigma = \partial H_n \rightarrow H_n$. As the map ϕ_n is 2-bilipschitz, a small variation of stability of quasi-geodesics (Lemma 6.5) gives that $\phi_n(\psi^{-1}P \times \{1\})$ is isotopic to its geodesic realization in N_n which has length at most 2η . Notice that $\phi_n(\psi^{-1}P \times \{-1\})$ is homotopic to $f_n(\psi^{-1}P)$. A topological argument then shows that $\phi_n(\psi^{-1}P \times \{-1\})$ is isotopic to $\psi^{-1}P \subset \partial H_n$. Thus, we have shown that, for n large, the geodesic representatives of $\psi^{-1}P \subset \partial H_n$ in N_n are isotopic to $\psi^{-1}P$ and have length at most 2η . This contradicts the initial assumptions.

For convenience, the proof, which will take the rest of the section, is divided into small steps. In the proof we use the general fact, which we prove in Appendix C.

LEMMA 6.5. *There exists $\eta < \eta_3/2$ such that the following holds. Let \mathbb{T}_{η_3} be a Margulis tube with core geodesic α of length $l(\alpha) \leq \eta$. Suppose that there exists a 2-bilipschitz embedding of the tube in a hyperbolic 3-manifold $f : \mathbb{T}_{\eta_3} \rightarrow M$. Then $f(\alpha)$ is homotopically non-trivial and it is isotopic to its geodesic representative within $f(\mathbb{T}_{\eta_3})$.*

Proof of Proposition 6.2. We choose the constant η as the minimum of the constant from Lemma 6.5 and the constant η_{dri11} for $K = 2$ from Theorem 2.8. Let ψ be a pseudo-Anosov mapping class with a short pants decomposition P of length at most η . For a fixed $K > 0$, assume towards a contradiction that there exist arbitrarily

large n for which (2) holds and either the geodesic realization of $\psi^{-1}P \times \{1\}$ in N_n is not isotopic to $\psi^{-1}P \times \{1\}$ or it is isotopic, but has length strictly larger than 2η .

We begin by discussing the algebraic convergence of the manifolds N_n .

Let $\rho_n : \pi_1(\partial H) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ be the composition of the holonomy $\pi_1(H_n) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ of N_n with the map $\pi_1(\partial H_n) \rightarrow \pi_1(H_n)$ induced by the inclusion of the boundary $\partial H_n \subset H_n$.

LEMMA 6.6. *For every $\gamma \in \pi_1(\Sigma)$ there exists $n_\gamma > 0$ such that for every $n \geq n_\gamma$ we have $\gamma \notin \ker(\rho_n)$.*

Proof. We can assume that γ is primitive. By the Loop Theorem, if $\gamma \in \ker(\rho_n)$, then, after representing γ as a 4-valent graph on Σ , there exists an essential simple cycle $\alpha_n \subset \gamma$ which is also compressible. Since there are only finitely many simple cycles on γ , we can assume that $\alpha_n = \alpha$ with $\alpha \subset \gamma$ a fixed simple cycle. Therefore, $\alpha \in \ker(\rho_n)$ which means $\alpha \in \mathcal{D}_n$. However, by the assumption on the Gromov products, we have $(\lambda|\alpha)_P \rightarrow \infty$, which is not possible. \square

We now obtain algebraic convergence of the sequence ρ_n in the following sense.

LEMMA 6.7. *Up to subsequences and conjugating, ρ_n algebraically converges to a discrete and faithful representation $\rho : \pi_1(\Sigma) \rightarrow \mathrm{PSL}_2(\mathbb{C})$.*

In different language, Lemma 6.7 states that ρ_n has a subsequence that converges in the character variety.

Proof. The proof is an application of [BMNS16, Theorem 3.1]. We briefly recall the statement in our setup. Let $\rho_n : \pi_1(\Sigma \times [-1, 1]) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ be an *eventually faithful* sequence of representations, that is, a sequence of representations with discrete image satisfying the conclusion of Lemma 6.6. Suppose that

- Each ρ_n maps every curve in $P \times \{1\}$ to a parabolic isometry.
- There exists a sequence of simple closed curves $\zeta_n \subset \Sigma \times \{-1\}$ converging to a filling lamination λ and such that $\ell_{\rho_n}(\zeta_n)$ remains bounded.

Then there exists a subsequence that converges algebraically to a discrete and faithful representation ρ . (In fact [BMNS16, Theorem 3.1] allows for ζ_n to be certain multicurves called markings, but we only apply the theorem for simple closed curves.)

In our situation we can choose $\zeta_n := \delta_n$, where $\delta_n \in \mathcal{D}$ is a disk-bounding curve, so, by construction, $\ell_{\rho_n}(\zeta_n) = \ell_{\rho_n}(\delta_n) = 0$. We further note that for every sequence of disk bounding curves $\delta_n \in \mathcal{D}_n$ we have $\delta_n \rightarrow \lambda \in \partial\mathcal{C}$ as we assumed that $(\lambda|\delta_n)_P \rightarrow \infty$. Hence, we can apply [BMNS16, Theorem 3.1]. \square

We fix once and for all a basepoint $o \in \mathbb{H}^3$.

Let $Q := \mathbb{H}^3/\rho(\pi_1(\Sigma))$. Let $y \in Q$ be the projection of o to Q . By work of Thurston [Thu80, Chapter 8] and Bonahon [Bon86], the manifold Q is diffeomorphic to $\Sigma \times \mathbb{R}$. Let $f : \Sigma \rightarrow Q$ be a marking for Q inducing ρ at the level of the holonomy. We now compute the end invariants of Q in terms of the reference marking $f : \Sigma \rightarrow Q$. For a comprehensive discussion of ends and end invariants we refer to Section 2 of [Min10].

We already know that Q , has rank one cusps at the curves of P since, by construction as a limit of the ρ_n, ρ maps every curve in P to a parabolic isometry.

We now compute the other end invariants. In order to do so, we discuss the geometric convergence of the sequence N_n .

Consider the discrete groups $\Gamma_n := \rho_n(\pi_1(\Sigma))$ and the hyperbolic 3-manifolds $N_n = \mathbb{H}^3/\Gamma_n$. Let x_n be the projection of o to N_n .

As $\inf_n \{\text{inj}_{x_n}(N_n)\} > 0$, we have that, up to subsequences, the sequence of groups Γ_n converges in the Chabauty topology to a discrete group Γ (see Chapter E of [BP92]). This means that

- (1) For every $\gamma \in \Gamma$ there exists a sequence $\gamma_n \in \Gamma_n$ such that $\gamma_n \rightarrow \gamma$.
- (2) If $\gamma_{n_j} \in \Gamma_{n_j}$ converges to γ in $\text{PSL}_2(\mathbb{C})$, then $\gamma \in \Gamma$.

In particular, for our sequence Γ_n we have $\rho(\pi_1(\Sigma)) < \Gamma$ by the second property. Consider $N := \mathbb{H}^3/\Gamma$ and let $x \in N$ be the projection of $o \in \mathbb{H}^3$. As $\rho(\pi_1(\Sigma)) < \Gamma$, we have a covering projection $\pi : Q \rightarrow N$.

It is well-known that convergence of discrete groups $\Gamma_n \rightarrow \Gamma$ in the Chabauty topology implies convergence in the pointed geometric topology of the associated hyperbolic manifolds $(N_n, x_n) \rightarrow (N, x)$ (see Theorem E.1.13 of [BP92]). This means that given R and $L > 1$, for every n large enough there exists a smooth map $\phi_n : B(x, R) \subset N \rightarrow B(x_n, R) \subset N_n$ which is L -bilipschitz and maps $\phi_n(x) = x_n$.

Suppose that $f(\Sigma) \subset B(y, R) \subset Q$ so that $\pi f(\Sigma) \subset B(x, R) \subset N$.

Fact. We can assume that $\phi_n \pi f$ induces ρ_n at the level of the holonomy.

In order to explain this, we briefly recall the construction of the maps ϕ_n as in the proof of Theorem E.1.13 in [BP92]: Consider $B(o, R) \subset \mathbb{H}^3$. For every $\sigma > 0$, define the (finite) set

$$\Gamma(R + \sigma) := \{\gamma \in \Gamma \mid \gamma B(o, R + \sigma) \cap B(o, R + \sigma) \neq \emptyset\} = \{\gamma_1, \dots, \gamma_t\}.$$

Let $\{\alpha_1, \dots, \alpha_m\}$ be a set of generators for $\pi_1(\Sigma)$. By our assumption $\pi f(\Sigma) \subset B(x, R)$, we have $\{\rho(\alpha_1), \dots, \rho(\alpha_m)\} \subset \Gamma(R)$. By convergence in the Chabauty topology, for every $\gamma_j \in \Gamma(R)$ we can find $\gamma_j^n \in \Gamma_n$ such that $\gamma_j^n \rightarrow \gamma_j$ as $n \rightarrow \infty$. Notice that, by algebraic convergence $\rho_n \rightarrow \rho$, for each $\rho(\alpha_j) \in \Gamma(R)$ we can choose $\rho_n(\alpha_j)$ as an approximating element.

Lemma E.1.16 of [BP92] shows that, up to slightly increasing R and choosing σ appropriately small, for every n large enough we have

$$\Gamma_n(R + \sigma) = \{\gamma \in \Gamma_n \mid \gamma B(o, R + \sigma) \cap B(o, R + \sigma) \neq \emptyset\} = \{\gamma_1^n, \dots, \gamma_t^n\}.$$

In [BP92] it is then explained how to construct a sequence of smooth embeddings $\hat{\phi}_n : B(o, R) \rightarrow B(o, R + \sigma)$ such that

$$\hat{\phi}_n(\gamma_j p) = \gamma_j^n \hat{\phi}_n(p)$$

and $\hat{\phi}_n$ converges to the natural inclusion in the smooth topology.

Such a map $\hat{\phi}_n$ descends to the desired approximating map $\phi_n : B(x, R) \subset N \rightarrow B(x_n, R) \subset N_n$. We are now ready to conclude: Let $\hat{f} : \hat{\Sigma} \rightarrow \mathbb{H}^3$ be the lift of the marking f to the universal covers. We have

$$\hat{\phi}_n \hat{f}(\alpha_i q) = \hat{\phi}_n(\rho(\alpha_j) \hat{f}(q)) = \rho_n(\alpha_j) \hat{\phi}_n(\hat{f}(q))$$

for every generator α_j . This shows that $\phi_n \pi f$ induces ρ_n at the level of the holonomy.

By the fact, we have that $\phi_n \pi f$ is homotopic to $f_n : \Sigma \rightarrow N_n$, where $f_n : \Sigma \rightarrow N_n$ is the marking inducing $\rho_n \circ (\pi_1(\partial H_n) \rightarrow \pi_1(H_n))$ at the level of the holonomy: Both f_n and $\phi_n \pi f$ induce the same map ρ_n at the level of the holonomy. Hence, they induce the same map at the level of fundamental groups. As both source and target are aspherical manifolds, the maps are freely homotopic.

We use this information to prove the following:

LEMMA 6.8. *The lamination λ is not realized in Q .*

Proof. Suppose that λ is realized. Since every sequence $\delta_n \in \mathcal{D}_n$ converges to λ , we can also realize δ_n for every n large enough in a neighborhood of λ . Thus for every n large enough, the curve $f(\delta_n)$ is homotopic to its geodesic representative β_n in Q within $B(y, R)$ for some large R . Let $\phi_n : B(x, R) \subset N \rightarrow N_n$ be the L -bilipschitz approximating map provided by geometric convergence. As L can be chosen arbitrarily close to 1, the curve $\phi_n \pi(\beta_n)$ is not null-homotopic in N_n . However, $\pi(\beta_n)$ is homotopic within $B(x, R)$ to $\pi f(\delta_n)$ and $\phi_n \pi f$ is homotopic to f_n . Hence, $\phi_n \pi f(\delta_n)$ is homotopic to $f_n(\delta_n)$ which is null-homotopic in N_n . This provides a contradiction and finishes the proof. \square

Since each curve in P is a rank one cusp, every complementary components of P is a triply punctured sphere, and λ fills the surface Σ , there is no room for other end invariants. By the solution of the Ending Lamination Conjecture by Minsky [Min10] and Brock, Canary, and Minsky [BCM12], we conclude that Q is the unique hyperbolic structure on $\Sigma \times (-\infty, 0)$ with end invariants (λ, P) . We will also denote such a structure by $Q = Q(\lambda, P)$.

We now show that Q coincides with N .

LEMMA 6.9. *The covering map $\pi : Q \rightarrow N$ is trivial.*

Proof. We prove that every component E of $Q - \mathcal{CC}(Q)$ embeds in N under the covering map π and distinct components E, E' have disjoint images under π . Then we argue that this is enough to conclude that π is trivial.

By the structure of the relative ends of Q , the components of $Q - \mathcal{CC}(Q)$ correspond 1-to-1 to the components W of $\Sigma - P$.

They are all of the following form: Consider a component $W \subset \Sigma - P$. Notice that $\rho_n(\pi_1(W))$ leaves invariant a half space $HS_n \subset \mathbb{H}^3$ which is a lift to \mathbb{H}^3 of the the component of $N_n - \mathcal{CC}(N_n)$ corresponding to W . Moreover, every element $\rho_n(\gamma)$ with $\gamma \in \pi_1(H_n) - \pi_1(W)$ moves HS_n off itself, that is, HS_n is *precisely invariant* under $\rho_n(\pi_1(W))$ in Γ_n .

By assumption, the sequence of fuchsian representations $\rho_n : \pi_1(W) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ of the pair of pants W converges to $\rho : \pi_1(W) \rightarrow \mathrm{PSL}_2(\mathbb{C})$. It is immediate to check that the limit is also a fuchsian representation, that it preserves a half space HS which is the limit of the half spaces HS_n , and that $\rho_n(\pi_1(W))$ converges to $\rho(\pi_1(W))$ in the Chabauty topology.

The quotient $E = HS/\rho(\pi_1(W))$ corresponds to the W -component of $Q - \mathcal{CC}(Q)$.

Claim. The restriction of π to a component E of $Q - \mathcal{CC}(Q)$ is a homeomorphism onto the image.

Proof of claim. Suppose that there exists $\alpha \in \Gamma - \rho(\pi_1(W))$ such that $\alpha HS \cap HS \neq \emptyset$. Let $\alpha_n \in \Gamma_n$ be a sequence that approximates α . Since $HS_n \rightarrow HS$ and $\alpha_n \rightarrow \alpha$ we deduce that $\alpha_n HS_n \cap HS_n \neq \emptyset$ for n large enough. Therefore, $\alpha_n \in \rho_n(\pi_1(W))$ as every element in $\Gamma_n - \rho_n(\pi_1(W))$ moves HS_n off itself. This implies $\alpha \in \rho(\pi_1(W))$ as the peripheral subgroups $\rho_n(\pi_1(W))$ are converging to $\rho(\pi_1(W))$ in the Chabauty topology. As a consequence, the restriction of the projection π to $E = HS/\rho(\pi_1(W))$ is a homeomorphism onto the image. \square

Claim. If E, E' are distinct components of $Q - \mathcal{CC}(Q)$, then $\pi(E) \cap \pi(E') = \emptyset$.

Proof of claim. Suppose that $\pi(E) \cap \pi(E') \neq \emptyset$. Represent the components as $E = HS/\rho(\pi_1(W))$ and $E' = HS'/\rho(\pi_1(W'))$ where $HS, HS' \subset \mathbb{H}^3$ are precisely invariant half-spaces under $\rho(\pi_1(W)), \rho(\pi_1(W'))$ and W, W' are distinct components of $\Sigma - P$.

As $\pi(E) \cap \pi(E') \neq \emptyset$, there exists $\alpha \in \Gamma$ such that $\alpha HS' \cap HS \neq \emptyset$. In this case, we must have $\alpha HS' = HS$: If $\alpha HS' \neq HS$, the circle at infinity bounding $\alpha \partial HS'$, which is contained in the limit set of $\rho(\pi_1(\Sigma))$, would intersect the round disk at infinity bounding HS and such a disk is contained in the domain of discontinuity of $\rho(\pi_1(\Sigma))$.

Therefore, as $\alpha HS' = HS$, we have $\rho(\pi_1(W)) = \alpha \rho(\pi_1(W')) \alpha^{-1}$. Since Γ_n converges in the Chabauty topology to Γ , this implies that $\rho_n(\pi_1(W))$ and $\rho_n(\pi_1(W'))$ are also conjugate in Γ_n for n large enough. However, this can happen only if $W = W'$. \square

We are now ready to prove that the projection $\pi : Q \rightarrow N$ is an isometry: Consider a component E and $\pi(E) \subset N$. If π is not an isometry, then, by the first claim, there is a component of the pre-image of $\pi(E)$ not contained in E . Such component must intersect a component E' of $Q - \mathcal{CC}(Q)$ different from E : For example, because it contains points with arbitrarily large injectivity radius and there is a uniform upper bound on the injectivity radius at points in $\mathcal{CC}(Q)$ (see Canary [Can96]). But distinct components E, E' map to disjoint sets by the second claim. \square

Thus, we have shown that the sequence of (N_n, x_n) converges geometrically to (Q, y) . One end of the manifold $Q(\lambda, P)$ looks very similar to the infinite cyclic covering \hat{T}_ψ of the mapping torus T_ψ . We do not need to make this assertion very precise as, for our purposes, it is sufficient to prove the following:

LEMMA 6.10. *There exists an identification of Q with $\Sigma \times (-\infty, 0)$ that induces $\rho : \pi_1(\Sigma) \rightarrow \mathrm{PSL}_2(\mathbb{C})$ at the level of the holonomy and such that $\psi^{-1}P \times \{-1\}$ for $j < 0$ is a geodesic link of length at most 2η .*

Proof of Lemma 6.10. As ψ is a pseudo-Anosov mapping class with a short pants decomposition P of length at most η (by our choice of η at the beginning of the proof), the mapping torus

$$T_\psi = \Sigma \times [0, 1]/(x, 0) \sim (\psi(x), 1)$$

admits a hyperbolic metric for which the curves $P \times \{0\}$ on the fiber $\Sigma \times \{0\}$ form a geodesic link of length at most η . Let $\hat{T}_\psi = \Sigma \times \mathbb{R}$ be the infinite cyclic covering of T_ψ . Concretely, $T_\psi = \hat{T}_\psi/\Psi$ where

$$\Psi : \hat{T}_\psi \rightarrow \hat{T}_\psi$$

is the isometry $\Psi(x, t) := (\psi(x), t + 1)$ generating the deck group of the covering.

We now show that the convex core of $Q = Q(\lambda, P)$ is geometrically very close to the region $\Sigma \times (-\infty, 0] - P \times \{0\} \subset \hat{T}_\psi$. As the curves in $\psi^{-1}P \times \{-1\} \subset \Sigma \times \{-1\}$ form a geodesic link of length at most 2η , this will imply that the same property holds for Q .

Consider the cyclic coverings

$$T_{\psi^n} = \Sigma \times [0, n]/(x, 0) \sim (\psi^n(x), 1) = \hat{T}_\psi/\Psi^n$$

of T_ψ . Each of the pants decompositions $\psi^j P \times \{j\}$ on the level surfaces $\Sigma \times \{j\}$ is a geodesic link of length at most η .

By our choice of η at the beginning of the proof, there is a (unique) complete finite volume hyperbolic structure on $T_{\psi^n} - P \times \{0\}$, which we denote by T'_{ψ^n} , and by Theorem 2.8 there is a 2-bilipschitz map

$$\Phi_n : T_{\psi^n} - \bigsqcup_{\gamma \subset P \times \{0\}} \mathbb{T}_\varepsilon(\gamma) \rightarrow T'_{\psi^n} - \bigsqcup_{\gamma \subset P \times \{0\}} \mathbb{T}_\varepsilon(\gamma).$$

Let us fix a basepoint $x \in \Sigma \times [0, 1] \subset \hat{T}_\psi$ in the ε -thick part. Let $x_n \in \Sigma \times [n-1, n]$ be the projection of $\Psi^n(x)$ to T_{ψ^n} . Consider the sequence of pointed manifolds (T_{ψ^n}, x_n) and $(T'_{\psi^n}, x'_n := \Phi_n(x_n))$.

By construction, the sequence (T_{ψ^n}, x_n) converges in the pointed geometric topology to the infinite cyclic covering \hat{T}_ψ of T_ψ : For example, we can choose the approximating maps to be the restrictions of the covering map $\hat{T}_\psi \rightarrow T_{\psi^n}$ to larger and larger metric balls around x .

Notice that, as Φ_n is 2-bilipschitz, the injectivity radius of T'_{ψ^n} at $x'_n = \Phi_n(x_n)$ is bounded from below by $\varepsilon/2$. Therefore, the sequence (T'_{ψ^n}, x'_n) also converges in the pointed geometric topology to a hyperbolic manifold (T', y) . Notice that each ε -Margulis neighborhood of a rank 2 cusp $\mathbb{T}_\varepsilon(\gamma) \subset T'_{\psi^n}$ for $\gamma \subset P \times \{0\}$ converges to a rank two cusp $\mathbb{T}_\varepsilon(\gamma) \subset T'$.

Furthermore, by Ascoli-Arzelà, the sequence of 2-bilipschitz maps Φ_n converge to a 2-bilipschitz map

$$\Phi : \hat{T}_\psi - \bigsqcup_{\gamma \subset P \times \{0\}} \mathbb{T}_\varepsilon(\gamma) \rightarrow T' - \bigsqcup_{\gamma \subset P \times \{0\}} \mathbb{T}_\varepsilon(\gamma).$$

The curves in $\psi^{-1}P \times \{-1\} \subset \hat{T}_\psi$ form a geodesic link of length at most 2η contained on some level surface $S \subset \hat{T}_\psi - \bigsqcup_{\gamma \subset P \times \{0\}} \mathbb{T}_\varepsilon(\gamma)$. Applying Lemma 6.5, we deduce that the geodesic representatives of the curves $\Phi(\psi^{-1}P \times \{-1\}) \subset \Phi(S)$ are isotopic to $\Phi(\psi^{-1}P \times \{-1\})$ within their standard ε -Margulis neighborhoods. We lift such isotopy to the $\Phi(S)$ -covering $Q' \rightarrow T'$.

We now prove that Q' is isometric to $Q = Q(\lambda, P)$ with an isometry in the isotopy class of the identity with respect to the natural marking on Q' and Q . This suffices to finish the proof.

The covering Q' is a hyperbolic structure on $\Sigma \times \mathbb{R}$. By construction, the curves in $\Phi(P \times \{0\})$ represent rank one cusps. The restriction of Φ to $\Sigma \times [-1, -\infty)$ lifts to a 2-bilipschitz embedding to Q' . As the end $\Sigma \times [-1, -\infty)$ is simply degenerate with ending lamination $\lambda \subset S$ and Φ is 2-bilipschitz, the image $\Phi(\Sigma \times [-1, -\infty))$ is a simply degenerate end with ending lamination $\Phi(\lambda) \subset \Phi(S)$. By the solution

of the Ending Lamination Conjecture, we conclude that Q' is isometric to $Q = Q(\lambda, P)$. \square

Lastly, we exploit the geometric convergence of N_n to Q combined with some topological argument to show that the geodesic representative of $\psi^{-1}P \times \{1\}$ in N_n is short and isotopic to $\psi^{-1}P \subset \partial H_n$, thus producing the desired contradiction.

By Lemma 6.10, we know that we can identify $Q = \Sigma \times (-\infty, 0)$ so that $\psi^{-1}P \times \{-1\}$ is a geodesic link of length at most 2η .

We consider some large product region $\Sigma \times [-T, -1/T]$ that contains the union of the Margulis neighborhoods of the curves in $\psi^{-1}P \times \{-1\}$. Let $\phi_n : \Sigma \times [-T, -1/T] \rightarrow N_n$ be a 2-bilipschitz approximating map provided by geometric convergence. As in the above discussion, we can assume that the map ϕ_n induces ρ_n at the level of the holonomy.

By Lemma 6.5, $\phi_n(\psi^{-1}P \times \{-1\})$ is isotopic to its geodesic representative. So, in order to conclude the proof it is enough to show the following:

LEMMA 6.11. *$\phi_n(\psi^{-1}P \times \{-1\})$ is isotopic to $\psi^{-1}P \subset \partial H_n$.*

Notice that, as ϕ_n induces ρ_n at the level of the holonomy, we already know that $\phi_n(\psi^{-1}P \times \{-1\})$ is homotopic to $\psi^{-1}P \subset \partial H_n$. What we have to do is to replace being homotopic with being isotopic. The idea of the proof is that $\phi_n(\Sigma \times \{-1\})$ is parallel to ∂H_n in such a way that $\phi_n(\gamma \times \{-1\})$ is parallel to $\gamma \subset \partial H_n$.

The proof is a little bit tedious so we summarize the main steps. Let us point out that the main issue is to keep track of the markings up to isotopy rather than up to homotopy for the convergence $N_n \rightarrow Q$. This is complicated by the fact that the map $\phi_n : \Sigma \times \{-1\} \rightarrow N_n$ is compressible.

Step 1. We show that $\phi_n(\Sigma \times \{-1\})$ is parallel to $\Sigma = \partial H_n$.

This immediately implies that $\phi_n : \Sigma \times \{-1\} \rightarrow N_n$ is isotopic to the composition $f_n h_n$, where $h_n : \Sigma \rightarrow \Sigma$ is a homeomorphism and $f_n : \Sigma \rightarrow H_n$ is the natural inclusion.

Step 2. The homeomorphism h_n has the property that $h_n(\gamma) = \gamma$ for every $\gamma \subset P$.

As a consequence, we have that h_n is a homeomorphism of Σ that leaves invariant the pants decomposition P . Hence, it is a product of Dehn twists around the curves in P .

Step 3. The homeomorphism h_n is isotopic to the identity.

Combining the three steps, we conclude that ϕ_n is isotopic to f_n .

Proof of Lemma 6.11. We now prove the steps outlined above.

Proof of Step 1. The first step comes from geometric convergence and some standard 3-dimensional topology.

Consider the non-cuspidal part of the convex core $\mathcal{CC}_0(N_n) \subset N_n$ obtained by removing from the convex core $\mathcal{CC}(N_n)$ the standard ε -Margulis neighborhoods of the cusps.

We claim that the boundary $\partial \mathcal{CC}_0(N_n)$ lies in a uniform neighborhood of the surface $\phi_n(\Sigma \times \{-1\})$. To see this, notice that the boundary $\partial \mathcal{CC}_0(N_n)$ decomposes as the disjoint union of the ε -non cuspidal parts of the boundary components $\partial \mathcal{CC}(N_n)$

joined by flat annuli on the boundary of the standard ε -Margulis neighborhoods of the cusps.

A component C_W^n of $\partial\mathcal{CC}_0(N_n)$ that comes from $\partial\mathcal{CC}(N_n)$ is homotopic to W for a component $W \subset \Sigma - P$. Moreover, it has the following properties. First, it has uniformly bounded diameter $\text{diam}(C_W^n) \leq D$. Secondly, it contains a (self-intersecting) closed geodesic γ_n^* homotopic to γ_W of uniformly bounded length $\ell(\gamma_n^*) \leq L$. Both properties come from the fact that C_W^n is isometric to the ε -thick part of the unique complete finite area hyperbolic metric on W .

By standard hyperbolic geometry, using the fact that ϕ_n is 2-bilipschitz, the distance between $\phi_n(\Sigma \times \{-1\})$ and C_W^n can be bounded in terms of the length of γ_W on $\Sigma \times \{-1\}$ and L . Thus each component C_W^n lies in a uniform neighborhood of $\phi_n(\Sigma \times \{-1\})$.

In order to conclude, we just have to consider the distance of the annular components A_γ^n of $\partial\mathcal{CC}_0(N_n)$ from $\phi_n(\Sigma \times \{-1\})$. We already know that the distance of the boundary components of such annuli A_γ^n from $\phi_n(\Sigma \times \{-1\})$ is uniformly bounded. In particular, the distance of the boundary components is also bounded (in terms of $\text{diam}(\Sigma \times \{-1\})$). As the intrinsic diameter of A_γ^n can be bounded in terms of the distance of the boundary components, the claim follows.

Therefore, if T is large enough, then $\phi_n(\Sigma \times [-T, -1/T])$ contains $\partial\mathcal{CC}_0(N_n)$.

We now argue that $\partial\mathcal{CC}_0(N_n)$ is incompressible in $\phi_n(\Sigma \times [-T, -1/T])$. By standard 3-manifold topology (see [Wal68, Proposition 3.1 and Corollary 3.2]), this implies that $\partial\mathcal{CC}_0(N_n)$ is parallel to the boundary surfaces of the product. Since, by properties of convex cores, $\partial\mathcal{CC}_0(N_n)$ is also parallel to ∂H_n , this will be sufficient to prove the claim of the first step.

If $\partial\mathcal{CC}_0(N_n)$ was compressible, then one of the following two cases happen. Either, after some compressions, we find a closed incompressible surface in $\phi_n(\Sigma \times [-T, -1/T])$ of genus strictly smaller than the genus of Σ or, we have that the entire convex core $\mathcal{CC}_0(N_n)$ is contained in $\phi_n(\Sigma \times [-T, -1/T])$. The first case is ruled out by the classification of incompressible surfaces in products $\Sigma \times [0, 1]$. The second case is ruled out by the fact that the boundary $\phi_n(\Sigma \times \{-T\})$ is contained in $\mathcal{CC}(N_n)$. Thus, we must have that $\partial\mathcal{CC}_0(N_n)$ is incompressible in $\phi_n(\Sigma \times [-T, -1/T])$. \square

Proof of Step 2. Let us describe the homeomorphism h_n .

Consider the non-cuspidal part of the convex core $\mathcal{CC}_0(N_n)$. By the structure of the ends of hyperbolic manifolds, there is a homeomorphism $\Sigma \times (0, 1) \rightarrow N_n - \mathcal{CC}_0(N_n)$ with the following properties: The inclusion $\Sigma \times \{t\} \rightarrow N_n$ is isotopic to f_n , the standard marking induced by the inclusion $\Sigma = \partial H_n \subset H_n$. Each cusp $\mathbb{T}_\varepsilon(\gamma)$ corresponds to $A_\gamma \times (0, 1)$ where A_γ is an annulus around γ . Let $\text{pr}_1 : \Sigma \times (0, 1) \rightarrow \Sigma$ be the projection of the first factor.

As $\phi_n(\Sigma \times \{-1/T\})$ is contained in $N_n - \mathcal{CC}(N_n)$ and is parallel to $\Sigma \times \{0\}$ the composition $\text{pr}_1 \phi_n : \Sigma \rightarrow \Sigma$ is a degree one homotopy equivalence and $f_n \text{pr}_1 \phi_n$ is homotopic to ϕ_n . Let h_n be an orientation preserving homeomorphism homotopic to $\text{pr}_1 \phi_n : \Sigma \rightarrow \Sigma$.

In order to compute the homotopy class of $h_n(\gamma)$ for $\gamma \subset P$ we can proceed as follows: In the product $\Sigma \times [-T, -1/T]$ we can isotope the inclusion $\iota_{-1} : \Sigma \rightarrow \Sigma \times \{-1\}$ to $\iota : \Sigma \rightarrow \Sigma \times [-T, -1/T]$ so that it intersects each of the cusps $\mathbb{T}_{\varepsilon/2}(\gamma)$ in an annulus $\iota(U_\gamma)$ around γ . As ϕ_n is 2-bilipschitz and $\phi_n \iota$ is homotopic to f_n ,

we have that $\phi_n \iota(U_\gamma) \subset \mathbb{T}_\varepsilon(\gamma)$. The homotopy class of $h_n(\gamma)$ is the homotopy class of the curve $\text{pr}_1(\phi_n \iota(\gamma))$ which is a primitive essential curve contained in A_γ . Thus $h_n(\gamma) \simeq \gamma$. \square

Proof of Step 3. Observe that h_n has the following property: Let $\delta_n \in \mathcal{D}_n$ be a disk. We have that $f_n(h_n(\delta_n))$ is homotopic to $\phi_n(\delta_n)$ which, in turn, is homotopic to $f_n(\delta_n)$. Therefore $h_n(\delta_n)$ is again a compressible curve. By the Loop Theorem, we conclude $h_n(\delta_n) \in \mathcal{D}_n$. In other words h_n preserves \mathcal{D}_n . However, by assumption, no component of P is contained in \mathcal{D}_n and this implies that any non-trivial product of Dehn twists around curves in P cannot preserve \mathcal{D}_n . \square

Thus, we proved Lemma 6.11. \square

This concludes the proof of Proposition 6.2. \square

PART 3

Geometric properties of random 3-manifolds

7. RANDOM 3-MANIFOLDS

In this section we prove a precise form of Theorem 5 about the structure of random 3-manifolds. Before we state the theorem, we recall some background and set some notation regarding random walks.

7.1. Random walks. We start by recalling some background material on random walks on the mapping class group. We crucially consider only random walks driven by probability measures μ whose support S is a finite symmetric generating set for the *entire* mapping class group.

DEFINITION (Random Walk). Let $(s_n)_{n \in \mathbb{N}}$ be a sequence of independent random variables with values in S and distribution μ . The *n-th step of the random walk* is the random variable $f_n := s_1 \dots s_n$. We denote by \mathbb{P}_n its distribution. The *random walk driven by μ* is the process $(f_n = s_1 \dots s_n)_{n \in \mathbb{N}} \in \text{Mod}(\Sigma)^\mathbb{N}$. It has a distribution which we denote by \mathbb{P} .

The mapping class group acts on the curve graph $\text{Mod}(\Sigma) \curvearrowright \mathcal{C}$. If we fix a base point $o \in \mathcal{C}$ we can associate to every random walk $(f_n)_{n \in \mathbb{N}}$ an orbit $\{f_n o\}_{n \in \mathbb{N}} \subset \mathcal{C}$.

We choose $o \in \mathcal{C}$ with the following property:

Standing assumption: The base point $o \in \mathcal{C}$ is chosen to lie on the disk set, $o \in \mathcal{D}$.

It is a standard consequence of the subadditive ergodic theorem that there exists a constant $L \geq 0$, called the *drift* of the random walk on the curve graph, such that for \mathbb{P} -almost every sample path $(f_n)_{n \in \mathbb{N}}$ we have

$$\frac{d_{\mathcal{C}}(o, f_n o)}{n} \xrightarrow{n \rightarrow \infty} L.$$

In general, the drift can be 0. However, it has been established by Maher [Mah10a] that, in our case, $L > 0$.

7.2. Statement and discussion. We are now ready to state the precise version of Theorem 5, using the above setup. We say that a sequence of events has asymptotic probability 1 if the probability of the events goes to 1 as n tends to infinity.

Theorem 5. *Fix $K > 1$ and $\varepsilon > 0$. Let $o \in \mathcal{D}$ be a fixed basepoint. Denote by $L > 0$ the drift of the random walk and by τ_n the parametrization of the geodesic segment $[o, f_n o]$ by arc length. With asymptotic probability 1, we have the following for some uniform constant B : There exist pants decompositions P_n and R_n of Σ such that*

- (a) The manifold M_{f_n} admits a hyperbolic metric such that restriction of the metric to a tubular neighborhood of the Heegaard surface is K -bilipschitz to the η_3 -non cuspidal part of the convex core of the maximally cusped structure $Q(P_n, R_n)$.
- (b) $d_C(\tau_n[\varepsilon Ln, 2\varepsilon Ln], P_n), d_C(\tau_n[(1 - 2\varepsilon)Ln, (1 - \varepsilon)Ln], R_n) \leq B$.

In Subsection 7.5 we also describe how to derive Theorem 1 from Theorem 6 and properties of the random walk.

The proof of Theorem 5 does not use 3-dimensional hyperbolic geometry anymore. Rather, via Proposition 6.4, we will only have to work with the dynamics of a random walk the curve graph.

The idea is that thanks to the work done in the previous sections, namely Proposition 5.2 and Proposition 6.4, we only need to check that the geodesic segment $\tau_n = [o, f_n o]$ contains four points $o < S_1 < S_2 < S_3 < S_4 < f_n o$ satisfying the conditions of Proposition 6.4.

The heuristic picture is the following. The endpoints o and $f_n o$ lie on the disk sets \mathcal{D} and $f_n \mathcal{D}$. Hyperbolicity of the curve graph, quasi-convexity of the disk sets together imply that, if \mathcal{D} and $f_n \mathcal{D}$ are sufficiently far away, then the path $[o, f_n o]$ roughly decomposes into three parts: Initially, it fellow travels \mathcal{D} . Then, it follows a shortest geodesic between \mathcal{D} and $f_n \mathcal{D}$. Lastly, it fellow travels $f_n \mathcal{D}$.

Any subsegment of the middle piece automatically satisfies property (iii).

Properties (i) and (ii) follow, instead, from ergodic properties of the random walk, see below for discussion and references. In particular, we will use that for any pseudo-Anosov ϕ , the segment $[o, f_n o]$ often fellow travels a translate of the axis l_ϕ of the pseudo-Anosov. Therefore, we just have to make sure that the two needed long fellow travelings happen on the subsegment lies in the middle piece of $[o, f_n o]$.

We will deduce this combining the aforementioned ergodic properties of random walks with work of Maher [Mah10b] who proved that, with asymptotic probability 1, the distance between \mathcal{D} and $f_n \mathcal{D}$ increases linearly and up to a sublinear error is the distance between the endpoints o and $f_n o$. Hence, the middle piece in the above description takes up almost all of $[o, f_n o]$.

7.3. Ergodic properties of random walks. We can now state the ergodic property of random walks that we need. It is inspired by [BGH20, Proposition 6.9]. In fact, we believe that the following statement can be extracted from its proof, with the exception, perhaps, of the logarithmic size of the fellow traveling. We include a complete proof of the precise form that we need.

THEOREM 7.1. *Let $\phi \in \text{Mod}(\Sigma)$ be a pseudo-Anosov with quasi-axis l_ϕ in the curve graph, and let $0 < a < b < 1$. Denote by $L > 0$ the drift of the random walk. There exists $\varepsilon_0 > 0$ such that with asymptotic probability 1 the following holds. Denote by τ_n the segment $[o, f_n o]$. Then l_ϕ has a subsegment of length $\varepsilon_0 \log(n)$ one of whose translates uniformly fellow-travels a subsegment of $\tau_n[aLn, bLn]$.*

Proof. For g in $\text{Mod}(\Sigma)$, we denote by π^g the closest-point projection to gl_ϕ .

Due to hyperbolicity of the curve graph, such projections have strong contraction properties. In particular, it is well-known that they imply the following: There exists a constant $D > 0$, depending on ϕ , such that if $d_C(\pi^g(x), \pi^g(y)) \geq D$ then the geodesic $[x, y]$ has a subsegment $[x_1, y_1]$ with $d_C(x_1, \pi^g(x)), d_C(y_1, \pi^g(y)) \leq D$. Such a segment $[x_1, y_1]$ δ -fellow travels l_ϕ for some uniform δ .

Therefore, in order to conclude, it suffices to prove the following claim (the theorem follows up to moving a, b an arbitrarily small amount and modifying ε).

Claim: Given a, b, ϕ as in the statement, there exists $\varepsilon > 0$ such that with asymptotic probability 1, there exists g such that $d_{\mathcal{C}}(\pi^g(o), \pi^g(f_n o)) \geq \varepsilon \log(n)$ and $d_{\mathcal{C}}(o, \pi^g(o)), d_{\mathcal{C}}(o, \pi^g(f_n o)) \in [aLn, bLn]$.

The claim is a consequence of the following properties, which can be found in the existing literature as we explain below: There exist $a' < b', \varepsilon \in (0, (b-a)/10)$ and $C > 0$ such that the following hold with asymptotic probability 1

- (i) $d_{\mathcal{C}}(f_j o, \tau_n) \leq C \log(n)$ for every $j \leq n$.
- (ii) $d_{\mathcal{C}}(f_{[a'n]} o, o) \in [(a + \varepsilon)Ln, (a + 2\varepsilon)Ln]$ and $d_{\mathcal{C}}(f_{[b'n]} o, o) \in [(b - 2\varepsilon)Ln, (b - \varepsilon)Ln]$.
- (iii) There are g and $\varepsilon > 0$ such that $d_{\mathcal{C}}(\pi^g(f_{[a'n]} o), \pi^g(f_{[b'n]} o)) \geq \varepsilon \log(n)$.
- (iv) For the same g of point (iii), we have $d_{\mathcal{C}}(\pi^g(f_{[a'n]} o), \pi^g(o)) \leq \varepsilon \log(n)/3$ and $d_{\mathcal{C}}(\pi^g(f_{[b'n]} o), \pi^g(f_n o)) \leq \varepsilon \log(n)/3$.

Assuming (i)-(iv) we prove the claim. Afterwards, we give the references to the literature.

Proof of the claim. In the whole proof, all statements and inequalities are meant to hold with asymptotic probability 1. By properties (iii) and (iv), it follows that $d_{\mathcal{C}}(\pi^g(o), \pi^g(f_n o)) \geq \varepsilon \log(n)/3$, whence the first part of the claim. We now argue that $\pi^g(o)$ and $\pi^g(f_n o)$ have distance within the desired interval from o .

Observe that the geodesic joining $f_{[a'n]} o$ to $f_{[b'n]} o$ fellow travels gl_{ϕ} along the subsegment connecting $x_n := \pi^g(f_{[a'n]} o)$ to $y_n := \pi^g(f_{[a'n]} o)$, because the projections are very far apart by property (iii). In particular, x_n and y_n are uniformly close to points p_n and q_n on $[f_{[a'n]} o, f_{[b'n]} o]$ respectively. By property (iv), the projections x_n and y_n are also logarithmically close to $\pi^g(o)$ and $\pi^g(f_n o)$. Therefore, we have

$$\begin{aligned} d_{\mathcal{C}}(o, \pi^g(o)) &= d_{\mathcal{C}}(o, p_n) + O(\log(n)), \\ d_{\mathcal{C}}(o, \pi^g(f_n o)) &= d_{\mathcal{C}}(o, q_n) + O(\log(n)). \end{aligned}$$

Hence, we can focus on estimating $d_{\mathcal{C}}(o, p_n)$ and $d_{\mathcal{C}}(o, q_n)$. In fact, we will provide an estimate on $d_{\mathcal{C}}(o, p)$ for any point $p \in [f_{[a'n]} o, f_{[b'n]} o]$.

By the triangle inequality, for any point $p \in [f_{[a'n]} o, f_{[b'n]} o]$, we have

$$\begin{aligned} d_{\mathcal{C}}(o, f_{[b'n]} o) - d_{\mathcal{C}}(f_{[a'n]} o, f_{[b'n]} o) &\leq d_{\mathcal{C}}(o, p), \\ d_{\mathcal{C}}(o, p) &\leq d_{\mathcal{C}}(o, f_{[a'n]} o) + d_{\mathcal{C}}(f_{[a'n]} o, f_{[b'n]} o). \end{aligned}$$

We now estimate distances using properties (i) and (ii).

In view of (ii) and the inequalities above, for our purposes it suffices to show that

$$d_{\mathcal{C}}(f_{[a'n]} o, f_{[b'n]} o) \leq d_{\mathcal{C}}(o, f_{[b'n]} o) - d_{\mathcal{C}}(o, f_{[a'n]} o) + O(\log(n)).$$

We obtain this inequality as follows: Let $r_n, s_n \in \tau_n$ be provided by property (i), so that $d_{\mathcal{C}}(r_n, f_{[a'n]} o), d_{\mathcal{C}}(s_n, f_{[b'n]} o) = O(\log(n))$. Using that r_n and s_n lie on a

geodesic originating at o , we have

$$\begin{aligned} d_{\mathcal{C}}(f_{[a'n]}o, f_{[b'n]}o) &\leq d_{\mathcal{C}}(r_n, s_n) + O(\log(n)) \\ &= |d_{\mathcal{C}}(o, s_n) - d_{\mathcal{C}}(o, r_n)| + O(\log(n)) \\ &\leq |d_{\mathcal{C}}(o, f_{[b'n]}o) - d_{\mathcal{C}}(o, f_{[a'n]}o)| + O(\log(n)). \end{aligned}$$

Since $\varepsilon < (b - a)/10$, in view of (ii) we can remove the absolute value, and obtain the required estimate. \square

Now we provide references for the properties (i)-(iv).

Property (i) is a corollary of Theorem 10.7 [MS20] obtained by summing the probabilities that each step of the walk is logarithmically far from τ_n .

Property (ii) follows from positivity of the drift, which implies that for any $\varepsilon > 0$, with probability going to 1 as k tends to infinity we have $d_{\mathcal{C}}(o, f_k o) \in [(L - \varepsilon)k, (L + \varepsilon)k]$.

This easily allows us to choose appropriate a', b' .

For later purposes, we also note that (i) and the aforementioned property imply the following proposition, which is a version of a theorem of Tiozzo [Tio15] and could also be deduced from said theorem.

THEOREM 7.2. *In the setting of the theorem, for any $\varepsilon > 0$ with asymptotic probability 1 we have $d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \leq \varepsilon m$ for all $\varepsilon n \leq m \leq (1 - \varepsilon)n$.*

Properties (iii) and (iv) follow from Theorem 2.3 and Proposition 3.2 of [ST19], where a general framework is provided to show that random walks create logarithmically large projections. We explain how: In the terminology of [ST19] we want to show that

$$(\mathcal{S} := \{gl_\phi\}_{g \in \text{Mod}(\Sigma)}, Y_0 := l_\phi, \{\pi^g\}_{g \in \text{Mod}(\Sigma)}, \natural),$$

where we define the projections on the group $\pi^g : \text{Mod}(\Sigma) \rightarrow gl_\phi$ to be

$$\pi^g(h) := \pi^g(ho),$$

forms a *projection system* (as in [ST19, Definition 2.1]), where \natural is the relation on the translates of having bounded projection to each other, and that the probability measure μ is *admissible* (as in [ST19, Definition 2.2]).

The fact that the 4-tuple is a projection system follows from the contraction property of the projections π^g and well-known arguments. More specifically, referring to the requirements (1)-(5) of [ST19, Definition 2.1], we have that: Properties (1)-(3) are straightforward. Property (4) follows from the contraction property and, e.g., [Sis17, Lemma 2.5]. Property (5) follows instead from the fact that there are finitely many cosets of $\langle g \rangle$ such that if the projection of hl_ϕ on l_ϕ is unbounded, then h belongs to one of these cosets. This follows from, e.g., [Sis17, Corollary 4.4].

The fact that μ is admissible is also not difficult to be checked: Among the requirements perhaps only property (4) is not immediate. This property says, in our context, that the probability that the random walk ends up in one of the cosets of $\langle g \rangle$ for which the projection of hl_ϕ on l_ϕ is unbounded is exponentially small in the length of the walk. This holds because after n steps the random walk can only possibly visit linearly many of the elements of those cosets (they are undistorted), while the probability of ending up at any one of them is exponentially small, just because $\text{Mod}(\Sigma)$ is non-amenable. Now that we explained why [ST19] applies, the

third item follows from the [ST19, Theorem 2.3], while the fourth one follows from [ST19, Proposition 3.2], with $R = 0$. \square

As a different application of the same projection systems framework used above, we have the following statement whose proof is rather similar to the previous one. It will be used in the proof of Theorem 1 via Theorem 6 and in the application to the decay rate of the shortest geodesics for random 3-manifolds.

Denote by $L = \lim d_{\mathcal{C}}(o, f_n o)/n > 0$ the drift of the random walk.

THEOREM 7.3. *Denote by L the curve graph drift of the random walk, and let $0 < a < b < 1$. Then, there exist $\varepsilon_0 > 0$ and $C > 0$ such that with asymptotic probability 1 the following holds. There exists a non-separating simple closed curve $\gamma_n \subset \Sigma$ such that*

- $d_{\gamma_n}(o, f_n o) \geq \varepsilon_0 \log(n)$.
- $d_Y(o, f_n o) \leq C$ for every proper subsurface $Y \subset \Sigma - \gamma_n$.
- $d_{\mathcal{C}}(o, \gamma_n) \in [aLn, bLn]$.

Proof. If we exclude the requirement about the location of the curve γ_n with respect to $[o, f_n o]$, that is, the third requirement in the list, then the statement of the theorem is exactly the content of Proposition 7.1 of [ST19]. Here we want to control simultaneously the presence of a curve γ_n with large annular projection and bounded projections to the subsurfaces disjoint from it together with the position of γ_n on the segment $[o, f_n o]$ in order to make sure that it lies far away from the disk sets \mathcal{D} and $f_n \mathcal{D}$.

We have that the following properties hold with asymptotic probability 1:

- (i)' $d_{\mathcal{C}}(f_j o, [o, f_n o]) \leq C \log(n)$.
- (ii)' $d_{\mathcal{C}}(f_{[a'n]} o, o) \in [(a + \varepsilon)Ln, (a + 2\varepsilon)Ln]$ and $d_{\mathcal{C}}(f_{[b'n]} o, o) \in [(b - 2\varepsilon)Ln, (b - \varepsilon)Ln]$.

Observe that, if there is a large annular subsurface projection between $f_{[a'n]} o$ and $f_{[b'n]} o$ on some γ_n , then the curve γ_n lies on the 1-neighborhood of $[f_{[a'n]} o, f_{[b'n]} o]$ by the Bounded Geodesic Image Theorem [MM99]. Just like in the proof of the claim in Theorem 7.1, properties (i)' and (ii)' ensure then that γ_n is at the appropriate distance from o . That is, it satisfies the last requirement of Theorem 7.3.

Regarding the size of the subsurfaces projections we proceed as follows. We need three “buffer projections”, as in Proposition 7.1 [ST19], whose proof yields the following: There are $\varepsilon_1 > 0$ and $C_1 > 0$ such that the following holds with asymptotic probability 1. There exist non-separating curves $\gamma_1^n, \gamma_2^n, \gamma_3^n$ such that

- (iii)' $d_{\gamma_3^n}(f_{[a'n]} o, f_{[b'n]} o) \geq \varepsilon_1 \log(n)$.
- (iv)' $d_Y(f_{[a'n]} o, f_{[b'n]} o) \leq C_1$ for all subsurfaces $Y \subset \Sigma - \gamma_2^n$.
- (v)' $d_{\gamma_1^n}(f_{[b'n]} o, \gamma_2^n), d_{\gamma_3^n}(f_{[a'n]} o, \gamma_2^n), d_{\gamma_2^n}(f_{[a'n]} o, \gamma_1^n), d_{\gamma_2^n}(f_{[b'n]} o, \gamma_3^n) \leq C_1$.

Similar to property (iv) used in the proof of Theorem 7.1, we have the following replacement which holds with asymptotic probability 1 and follows again from Proposition 3.2 of [ST19]

- (vi)' $d_{\gamma_1^n}(f_{[a'n]} o, o) \leq \varepsilon_1 \log(n)/3$ and $d_{\gamma_3^n}(f_{[b'n]} o, o) \leq \varepsilon \log(n)/3$.

A consequence of these properties is that the annular projections $\pi_{\gamma_2^n}(o)$ and $\pi_{\gamma_2^n}(f_{[a'n]}o)$ coarsely coincide. This is a routine application of the Behrstock Inequality [Beh06], which states that there is a constant B such that for all curves α, β, γ we have $\min\{d_\gamma(\alpha, \beta), d_\alpha(\gamma, \beta)\} \leq B$ (provided that the quantities are well-defined). Here is the argument: By properties (iii)' and (v)', $d_{\gamma_1^n}(f_{[a'n]}o, \gamma_2^n)$ is large. Hence, by property (vi)', $d_{\gamma_1^n}(o, \gamma_2^n)$ is also large. Therefore, by the Behrstock Inequality, $d_{\gamma_2^n}(o, \gamma_1^n)$ is bounded and, by property (v)', the same holds for $d_{\gamma_2^n}(o, f_{[a'n]}o)$ as required.

The same argument also applies to projections $\pi_Y(o)$ and $\pi_Y(f_{[a'n]}o)$ for all sub-surfaces Y in the complement of γ_2^n . In fact, we have the following: Since $d_{\gamma_1^n}(o, \gamma_2^n)$ and $d_{\gamma_1^n}(f_{[a'n]}o, \gamma_2^n)$ are both large and Y is a subsurface of the complement of γ_2^n , also $d_{\gamma_1^n}(o, \partial Y)$ and $d_{\gamma_1^n}(f_{[a'n]}o, \partial Y)$ are large. Hence, by the Behrstock Inequality, $d_Y(o, \gamma_1^n)$ and $d_Y(f_{[a'n]}o, \gamma_1^n)$ are both uniformly bounded.

Changing the roles of δ and $f_{[a'n]}o$ with $f_n o$ and $f_{[b'n]}o$ concludes the proof. \square

7.4. The proof of Theorem 5. Consider the geodesic segment $[o, f_n o]$. Fix $\delta > 0$ large enough. We need to find two pseudo-Anosov mapping classes ψ and ψ' with short pants decompositions and four surfaces $S_0 = o < S_1 < S_2 < S_3 < S_4 < f_n o = S_5$ such that

- (i) $S_j < S_{j+1}$ and their distance is at least h , a very large constant.
- (ii) $[S_1, S_2], [S_3, S_4]$ δ -fellow travel $l_\psi, l_{\psi'}$.
- (iii) $d_{\mathcal{C}}([S_1, S_4], \mathcal{D}) \geq h$ and $d_{\mathcal{C}}([S_1, S_4], f_n \mathcal{D}) \geq h$.

We start with with a claim that will help us ensure that the last property holds.

Claim: For every $h > 0$ and $\varepsilon > 0$, with asymptotic probability 1 we have that both $d_{\mathcal{C}}(\tau_n[\varepsilon Ln, (1 - \varepsilon)Ln], \mathcal{D})$ and $d_{\mathcal{C}}(\tau_n[(1 - 2\varepsilon)Ln, (1 - \varepsilon)Ln], f_n \mathcal{D})$ are greater than h .

Proof of the claim. The claim is a consequence of

THEOREM 7.4 (Maher [Mah10b]). *For every $\varepsilon > 0$ we have*

$$\mathbb{P}_n [d_{\mathcal{C}}(\mathcal{D}, f\mathcal{D}) \in [(L_D - \varepsilon)n, (L_D + \varepsilon)n]] \xrightarrow{n \rightarrow \infty} 1.$$

Choose $\varepsilon_1 > 0$ much smaller than ε . By Theorem 7.4, for any given ρ we have $d_{\mathcal{C}}(\mathcal{D}, f_n \mathcal{D}) \geq (L_D - \varepsilon)n$ with probability at least $1 - \rho$ for n large.

By Theorem 7.2 we can assume $d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \leq \varepsilon_1 m$ for every $\varepsilon_1 n < m < (1 - \varepsilon_1)n$ with probability $\geq 1 - \rho$ for n large. We also assume $(L_D - \varepsilon)n < d_{\mathcal{C}}(\mathcal{D}, f_n \mathcal{D}) < (L_D + \varepsilon)n$ and $(L_D - \varepsilon)(n - m) < d_{\mathcal{C}}(\mathcal{D}, f_m^{-1} f_n \mathcal{D}) < (L_D + \varepsilon)(n - m)$ where $m := \lfloor (1 - \varepsilon_1)n \rfloor$ with probability $\geq 1 - \rho$ for every n large. (Observe that the random variables f_m and $f_m^{-1} f_n$ are independent and have distributions \mathbb{P}_{n-m} and \mathbb{P}_m).

Consider $m \in [\varepsilon_1 n, (1 - \varepsilon_1)n]$.

We have the following estimate on the distance from \mathcal{D} :

$$\begin{aligned} d_{\mathcal{C}}(\tau_n(Lm), \mathcal{D}) &\geq d_{\mathcal{C}}(f_m o, \mathcal{D}) - d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \\ &\geq d_{\mathcal{C}}(f_m \mathcal{D}, \mathcal{D}) - d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \\ &\geq (L_D - \varepsilon)m - \varepsilon_1 m \geq (L_D - \varepsilon - \varepsilon_1)\varepsilon_1 n. \end{aligned}$$

Notice that if ε_1 is small enough, the right hand side increases linearly in n with uniform constants.

As for the other disk set $f_n\mathcal{D}$, we also get

$$\begin{aligned} d_{\mathcal{C}}(\tau_n(Lm), f_n\mathcal{D}) &\geq d_{\mathcal{C}}(f_m o, f_n\mathcal{D}) - d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \\ &\geq d_{\mathcal{C}}(\mathcal{D}, f_m^{-1} f_n\mathcal{D}) - d_{\mathcal{C}}(f_m o, \tau_n(Lm)) \\ &\geq (L_D - \varepsilon)(n - m) - \varepsilon_1 m \geq (L_D - \varepsilon)\varepsilon_1 n - \varepsilon_1(1 - \varepsilon_1)n. \end{aligned}$$

As before, if ε_1 is very small the right hand side increases linearly in n with uniform constants. In conclusion, if ε_1 is small enough and n is large enough, the claim holds as $[\varepsilon Ln, (1 - \varepsilon)Ln] \subset [\varepsilon_1 n, (1 - \varepsilon_1)n]$. \square

The next claim will easily allow us to find points S_i satisfying (i) and (ii).

Claim: Let ϕ be a pseudo-Anosov element with a short pants decomposition. Let $l_\phi : \mathbb{R} \rightarrow \mathcal{C}$ be its quasi-axis. For every $\varepsilon > 0$, for every $h > 0$, with asymptotic probability 1, the segments $\tau_n[\varepsilon Ln, 2\varepsilon Ln]$ and $\tau_n[(1 - 2\varepsilon)Ln, (1 - \varepsilon)Ln]$ ξ -fellow travel (with ξ only depending on ϕ), along subsegments $\tau_n[t_1^n, t_2^n]$ and $\tau_n[t_3^n, t_4^n]$ of length at least h , some translates $g_n l_\phi$ and $g'_n l_\phi$ of the axis l_ϕ .

Proof of the claim. We just need to apply Theorem 7.1 with parameters $0 < a < b$ given by $0 < L\varepsilon < 2\varepsilon L$ and $0 < (1 - 2\varepsilon)L < (1 - \varepsilon)L$ respectively. \square

Conclusion of the proof: For a fixed $\varepsilon > 0$ we define $o < S_1^n < S_2^n < S_3^n < S_4^n < f_n o$ to be the four points $\tau_n(t_1^n) < \tau_n(t_2^n) < \tau_n(t_3^n) < \tau_n(t_4^n)$ as given by the second claim. By construction, for every n large enough, these satisfy properties (i), (ii) of Proposition 6.4 and by the first Claim they also satisfy property (iii). Hence, by Proposition 6.4 there are pants decompositions $P_n^1, P_n^2, P_n^3, P_n^4$ that satisfy the filling criterion with parameter smaller than $\eta = \eta(K)$ as given by Proposition 5.2 where $K > 1$ is the bilipschitz constant that we fixed at the beginning. They also satisfy the property of part (b) of Theorem 5 by our choice of $\tau(t_j^n)$ and Proposition 6.4. Therefore, by Proposition 5.2, we have that M_{f_n} is hyperbolic and, moreover, away from the cusps, it is K -bilipschitz to the metric \mathbb{M}_n as described by Lemma 5.1. As \mathbb{M}_n contains an isometric copy of $\mathcal{CC}(Q(P_n^2, P_n^3))$, the bilipschitz map identifies the non-cuspidal part of such piece with a tubular neighborhood of the Heegaard surface (because of the control on the isotopy class of the map provided by Proposition 5.2). This concludes the proof of Theorem 5. \square

7.5. The proof of Theorem 1 via short curves. We sketch now a proof of Theorem 1 that uses the construction of Theorem 6.

The argument also gives that M_{f_n} contains a curve of length $\leq 1/\log(n)$ (later on, we improve this estimate using the model metric, see Theorem 3).

There is yet another version of Theorems 7.1 and 7.3 which says that, with high probability, there is a curve γ_n such that $d_{\Sigma - \gamma_n}(o, f_n o)$ has size at least $\log(n)$, and γ_n lies close to the middle of a geodesic in \mathcal{C} from o to $f_n o$. Similarly to the first Claim in the proof of Theorem 5, we have that γ_n also lies on a shortest geodesic connecting \mathcal{D} and $f_n\mathcal{D}$, and far from the endpoints of said geodesic. Using that \mathcal{D} is quasiconvex and the Bounded Geodesic Image Theorem [MM99], we have that the subsurface projection to $\Sigma - \gamma_n$ of \mathcal{D} is bounded and coarsely coincides with that of o . A similar statement holds for $f_n\mathcal{D}$. Hence, we get that $d_{\Sigma - \gamma_n}(\mathcal{D}, f_n\mathcal{D})$ is logarithmically large, and also that (H_g, γ_n) and $(H_g, f_n(\gamma_n))$ are both pared acylindrical. We can therefore use Theorem 6.

8. FOUR APPLICATIONS

We describe four applications of Theorem 5.

We recall that the model metric decomposition consists of five pieces

$$\mathbb{M}_n = H_1^n \cup \Omega_1^n \cup Q_n \cup \Omega_2^n \cup H_2^n,$$

but, for our applications, we will mainly focus on the maximally cusped structure $Q_n = Q(P_n, R_n)$, as given by Theorem 7.2. We recall that P_n, R_n are uniformly close to $\tau_n[\varepsilon Ln, 2\varepsilon Ln]$ and $\tau_n[(1-2\varepsilon)Ln, (1-\varepsilon)Ln]$. We also recall that it bilipschitz embeds, away from its cusps, into M_{f_n} with bilipschitz constant arbitrarily close to 1 as n goes to ∞ .

8.1. Diameter growth. As a first geometric application, we compute the coarse growth rate for the diameter of random 3-manifolds.

Theorem 2. *There exists $c > 0$ such that*

$$\mathbb{P}_n[\text{diam}(M_f) \in [n/c, cn]] \xrightarrow{n \rightarrow \infty} 1.$$

The proof of Theorem 2 has two different arguments, one for the coarse upper bound and one for the coarse lower bound. The upper bound comes from a result by White [Whi01] that relates the diameter to the *presentation length* of the fundamental group, a topological and algebraic invariant. Of a different nature is the coarse lower bound where we heavily use the ε -model metric structure of Theorem 5 and the relation with the model manifold.

We start with the upper bound. We need the following definition.

DEFINITION (Presentation Length). Let G be a finitely presented group. The length of a finite presentation $G = \langle F | R \rangle$ is given by

$$l(F, R) = \sum_{r \in R} |r|_F - 2$$

where $|r|_F$ denotes the word length of the relator $r \in R$ with respect to the generating set F . The presentation length of G is defined to be

$$l(G) := \min \{l(F, R) \mid G = \langle F | R \rangle \text{ finite presentation}\}.$$

We also recall that a relator $r \in R$ is triangular if $|r|_F \leq 3$.

THEOREM 8.1 (White [Whi01]). *There exists $c > 0$ such that for every closed hyperbolic 3-manifold M we have*

$$\text{diam}(M) \leq c \cdot l(\pi_1 M).$$

Let $S \subset \text{Mod}(\Sigma)$ be the finite support of the probability measure μ .

LEMMA 8.2. *There exists $C(S) > 0$ such that for every $f \in \text{Mod}(\Sigma)$ we have*

$$l(\pi_1(M_f)) \leq C|f|_S.$$

In particular $\text{diam}(M_f) \leq K|f|_S$ where $K = c \cdot C$.

Proof. The 3-manifold M_f admits a triangulation T with a number of simplices uniformly proportional, depending on S , to the word length $|f|_S$. We have $\pi_1(M_f) = \pi_1(T_2)$ where T_2 denotes the 2-skeleton of T . By van Kampen, the fundamental group of a 2-dimensional connected simplicial complex X admits a presentation

$\pi_1(X) = \langle F | R \rangle$ where every relation is triangular and the number of relations $|R|$ is roughly the number of 2-simplices. \square

As a corollary, we get

$$\text{diam}(M_{f_n}) \leq K |f_n = s_1 \dots s_n|_S \leq Kn$$

thus proving the upper bound in Theorem 2.

The coarse lower bound follows from the structure of the model metric and the estimate of Theorem 4.8 that comes from the model manifold technology of Minsky [Min10].

In particular, by Theorem 4.8, if $Q_n = Q(P_n, R_n)$ is a maximally cusped structure then the distance between the boundary components of its non-cuspidal part Q^{nc} (meaning the 3-manifold with boundary obtained removing open neighborhoods of all cusp) is at least $Ad_{\mathcal{C}}(P_n, R_n) - A$. In the case of random 3-manifolds we have

$$\begin{aligned} d_{\mathcal{C}}(P_n, R_n) &\geq d_{\mathcal{C}}(\tau_n(2\varepsilon Ln), \tau_n((1+2\varepsilon)Ln)) \\ &\geq d_{\mathcal{C}}(o, f_n o) - d_{\mathcal{C}}(o, \tau_n(2\varepsilon Ln)) - d_{\mathcal{C}}(\tau_n((1+2\varepsilon)Ln, f_n o)) \\ &\asymp Ln - o(n). \end{aligned}$$

8.2. Injectivity radius decay. As a second geometric application, we give a coarse upper bound to the decay rate of the length of the shortest geodesic of random 3-manifolds.

Theorem 3. *There exists $c > 0$ such that*

$$\mathbb{P}_n [\text{inj}(M_f) \leq c / \log(n)^2] \xrightarrow{n \rightarrow \infty} 1.$$

Proof. By Theorem 7.2, it is enough to show that $Q(P_n, R_n)$ satisfies

$$\text{systole}(Q(P_n, R_n)) \leq c / \log(n)^2.$$

We recall that P_n and R_n are uniformly close to $\tau_n(2\varepsilon Ln), \tau_n((1+2\varepsilon)Ln)$ where τ_n is the parametrized segment $[o, f_n o]$, for a fixed basepoint $o \in \mathcal{D}$.

By Minsky [Min00], a curve $\gamma \in \mathcal{C}$ not contained in P_n or R_n is short in $Q(P_n, R_n)$ if and only if there is a large subsurface projection $d_Y(P_n, R_n)$ on a proper subsurface $Y \subset \Sigma$ with $\gamma \subset \partial Y$. Furthermore, by the Length Bound Theorem [BCM12], its length will be bounded by

$$\ell_{Q(P_n, R_n)}(\gamma) \leq D \frac{S_{\gamma}(P_n, R_n)}{d_{\gamma}(P_n, R_n)^2 + S_{\gamma}(P_n, R_n)^2}$$

for some uniform constant $D > 0$ and where $d_{\gamma}(P_n, R_n)$ is the annular projection corresponding to γ and

$$S_{\gamma}(P_n, R_n) = 1 + \sum_{Y \in \mathcal{Y}_{\gamma}} \{\{d_Y(P_n, R_n)\}\}_K.$$

Here \mathcal{Y}_{γ} denotes the collection of essential subsurfaces of $\Sigma - \gamma$, $K > 0$ is, again, some uniform constant, and $\{\{\bullet\}\}_K$ is the function defined by $\{\{x\}\}_K = x$ if $x > K$ and 0 otherwise.

Putting things together, it is enough to show that there exists a curve $\gamma \subset \Sigma$ for which $d_{\gamma}(P_n, R_n) \geq \varepsilon_0 \log(n)$ and $S_{\gamma}(P_n, R_n) \leq C$ for some uniform $\varepsilon_0 > 0$ and C . For the purposes of the argument below, we note that the latter condition is

equivalent to having uniformly bounded projections on all Y contained in $\Sigma - \gamma$, in view of the distance formula of [MM00].

Replacing P_n with o and R_n with $f_n o$, the aforementioned property is contained in the statement of Theorem 7.3. To conclude, we only have to argue that o and P_n have coarsely the same subsurface projections to the annulus corresponding to γ and to all subsurfaces contained in $\Sigma - \gamma$ (and similarly for $f_n o$ and R_n). But this holds provided that we choose ε small enough in Theorem 5, and a, b sufficiently close to $1/2$ in Theorem 7.3. In fact, in this case we have that geodesics in \mathcal{C} from o to P_n cannot pass 2-close to γ_n , just because they are much shorter than the distance from o to γ_n . We can therefore apply the Bounded Geodesic Image Theorem [MM99] and conclude (since a similar argument also applies to $f_n o$ and R_n). \square

We conclude the discussion with a couple of remarks on the lower bound for the injectivity radius. If we consider only the three middle pieces $\Omega_n^1 \cup Q_n \cup \Omega_n^2$ of the ε -model metric, the rate of $1/\log(n)^2$ is exactly the coarse decay rate of the systole. This is again an adaptation of the arguments of [ST19]. Hence, in order to get a precise lower bound, we have to understand the systole of the handlebody pieces H_n^1 and H_n^2 . Such computation would be possible, for example, in the presence of a model manifold technology for handlebodies analogous to the one of Minsky [Min10] and Brock, Canary and Minsky [BCM12] for hyperbolic manifolds diffeomorphic to $\Sigma \times \mathbb{R}$.

8.3. Geometric limits of random 3-manifolds. We now exploit the model metric structure to establish the existence of certain geometric limits (see Chapter E.1 of [BP92] for the definition of the pointed geometric topology) for families of random 3-manifolds.

These limits will be used in our last application concerning the arithmeticity and the commensurability class of random 3-manifolds.

PROPOSITION 8.3. *Let $\phi \in \text{Mod}(\Sigma)$ be a pseudo-Anosov mapping class. Consider a sequence $A_n \subset \text{Mod}(\Sigma)$ such that $\limsup \mathbb{P}_n[A_n] > 0$. Then, we can find a sequence $n_j \uparrow \infty$ and elements $f_{n_j} \in A_{n_j}$ such that $M_{f_{n_j}}$ are hyperbolic 3-manifolds and the sequence $M_{f_{n_j}}$ converges to the infinite cyclic covering of hyperbolic mapping torus T_ϕ in the pointed geometric topology for a suitable choice of base points $x_{n_j} \in M_{f_{n_j}}$.*

Proof. By assumption, there exist $\delta > 0$ and a sequence $m_j \uparrow \infty$ such that $\mathbb{P}_{m_j}[A_{m_j}] \geq \delta$. We choose $(n_j)_{j \in \mathbb{N}}$ by inductively refining $(m_j)_{j \in \mathbb{N}}$.

By Theorem 7.1 and Theorem 5, for every $k \in \mathbb{N}$ we have that the event

$$G_{n,k} := \left\{ \begin{array}{l} M_{f_n} \text{ satisfies Theorem 5 with parameters } K := 1 + 1/k \text{ and } \varepsilon \\ \tau_n = [o, f_n o] \text{ satisfies Theorem 7.1 with parameters } \varepsilon \text{ and } \phi \end{array} \right\}$$

has probability at least $1 - \delta/10$ for every sufficiently large n , say for $n \geq N_k$. In particular, if $m_i \geq N_k$, we have $A_{m_i} \cap G_{m_i,k} \neq \emptyset$.

We recall that Theorem 5 provides a maximally cusped structure $Q(P_n, R_n)$ whose convex core K -bilipschitz embeds away from the cusp into M_{f_n} . Furthermore, the pants decomposition P_n is located uniformly close to $\tau_n[\varepsilon Ln, 2\varepsilon Ln]$ while the pants decomposition R_n is uniformly close to $\tau_n[(1 - \varepsilon)Ln, (1 - 2\varepsilon)Ln]$.

We define now inductively the sequence $(n_j)_{j \in \mathbb{N}}$. Suppose that we have already chosen n_1, \dots, n_{j-1} . The next element will be

$$n_j := \min\{m_i \mid m_i > \max\{n_{j-1}, N_j\}\}.$$

As $n_j > N_j$, we have $A_{n_j} \cap G_{n_j, j} \neq \emptyset$, so we can choose $f_{n_j} \in A_{n_j} \cap G_{n_j, j}$.

We recall that $\tau_n = [o, f_n o]$ satisfies Theorem 7.1 with parameters ϕ and ε , so it has a subsegment $\tau_n[3\varepsilon Ln, (1-3\varepsilon)Ln]$ that uniformly travels a translate $g_n l_\phi$ along a subsegment of length $\varepsilon_0 \log(n)$.

Up to remarking τ_n , an operation that does not change the isometry type of $Q(P_{n_j}, R_{n_j})$, we can assume that τ_{n_j} uniformly fellow travels l_ϕ along the subsegment $l_\phi[-a_{n_j}, a_{n_j}]$ with $a_{n_j} = \varepsilon_0 \log(n_j) \uparrow \infty$. Hence, the sequence of the endpoints of the remarked segments $\tau_{n_j}[3\varepsilon Ln_j, (1-3\varepsilon)Ln_j]$ is converging to the endpoints at infinity of the quasi-axis l_ϕ .

By work of Brock, Bromberg, Canary, and Minsky [BBCM13] and the solution of the Ending Lamination Conjecture [Min10, BCM12], we conclude that the sequence of maximally cusped structures $Q(P_{n_j}, R_{n_j})$ converges to the infinite cyclic covering of T_ϕ . As $Q(P_{n_j}, R_{n_j})$ becomes geometrically arbitrarily close to $M_{f_{n_j}}$, the claim follows. \square

8.4. Commensurability and arithmeticity. Dunfield and Thurston, using a simple homology computation, have shown in [DT06] that their notion of random 3-manifold is not biased towards a certain fixed set of 3-manifolds. This means that for every fixed 3-manifold M , with asymptotic probability 1, M_f is not diffeomorphic to M .

Using geometric tools it is possible to strengthen this conclusions and show that Dunfield and Thurston's notion of random 3-manifolds is also transverse, in a sense made precise in the theorem below, to the class of arithmetic hyperbolic 3-manifolds and to the class of 3-manifolds which are commensurable to a fixed 3-manifold M .

Theorem 4. *With asymptotic probability 1 the following holds*

- (1) M_f is not arithmetic.
- (2) M_f is not in a fixed commensurability class \mathcal{R} .

Proof. The argument is mostly borrowed from Biringer-Souto [BS11].

The proof of both points starts from the following observation: Each M_{f_n} finitely covers a maximal orbifold $M_{f_n} \rightarrow \mathcal{O}_n$.

We first prove the non-arithmeticity. We argue by contradiction: Suppose that $\mathbb{P}_n[M_f \text{ is arithmetic}]$ does not go to 0. Combining with Theorem 3, we also have

$$\limsup \mathbb{P}_n[M_f \text{ is arithmetic and } \text{inj}(M_f) \leq c/\log(n)^2] > 0.$$

By Proposition 8.3, up to passing to a subsequence, say the whole sequence for simplicity, we can pick M_{f_n} such that M_{f_n} is arithmetic, has injectivity radius $\text{inj}(M_{f_n}) \leq c/\log(n)^2$ and there are base points $x_n \in M_{f_n}$ such that the sequence (M_{f_n}, x_n) converges geometrically to (Q_∞, x_∞) where Q_∞ is a doubly degenerate structure on $\Sigma \times \mathbb{R}$ with $\text{inj}(Q_\infty) > 0$.

Since M_{f_n} are arithmetic, the orbifolds \mathcal{O}_n are congruence and have $\lambda_1(\mathcal{O}_n) \geq 3/4$ (see [BS91] or Theorem 7.1 in [BS11]). By Proposition 4.3 of [BS11], the orbifolds \mathcal{O}_n cannot be all different, hence we can assume that they are fixed all

the time $\mathcal{O}_n = \mathcal{O}$. We get a contradiction by observing that \mathcal{O} is covered by closed 3-manifolds M_{f_n} with arbitrarily small injectivity radius.

We now discuss commensurability. Proceed again by contradiction and assume that $\mathbb{P}_n[M_f$ is in the commensurability class $\mathcal{R}]$ does not go to 0. By the arithmetic part we know that we also have

$$\limsup \mathbb{P}_n[M_f \in \mathcal{R}, M_f \text{ not arithmetic, and } \text{inj}(M_f) \leq c/\log(n)^2] > 0.$$

As before, using Proposition 8.3, choose a geometrically convergent sequence $(M_{f_n}, x_n) \rightarrow (Q_\infty, x_\infty)$ of non-arithmetic, commensurable hyperbolic 3-manifolds with $\text{inj}(M_{f_n}) \downarrow 0$. Commensurability and non-arithmeticity imply together that $\mathcal{O}_n = \mathcal{O}$ is fixed all the time: It is the orbifold corresponding to the commensurator $\text{Comm}(\pi_1(M_{f_n}))$, which is a discrete subgroup of $\text{PSL}_2\mathbb{C}$ by Margulis (see Theorem 10.3.5 in [MR03]) and is an invariant of the commensurability class. We conclude with the same argument as before. \square

APPENDIX A. CURVE GRAPH AND PARED ACYLINDRICAL MANIFOLDS

LEMMA 2.5. *Let $\gamma, \gamma' \subset \Sigma$ be essential multicurves. We have*

- (i) *If $d_{\mathcal{C}}(\gamma, \mathcal{D}) \geq 2$ then (H_g, γ) is pared.*
- (ii) *If $d_{\mathcal{C}}(\gamma, \mathcal{D}) \geq 3$ then (H_g, γ) is pared acylindrical.*
- (iii) *If $d_{\mathcal{C}}(\gamma, \gamma') \geq 1$ then $(\Sigma \times [0, 1], \gamma \times \{0\} \sqcup \gamma' \times \{1\})$ is pared.*
- (iv) *If $d_{\mathcal{C}}(\gamma, \gamma') \geq 3$ then $(\Sigma \times [0, 1], \gamma \times \{0\} \sqcup \gamma' \times \{1\})$ is pared acylindrical.*

Proof. Let us prove (i). By Dehn's Lemma, if γ is not π_1 -injective, then a component of γ bounds a disk, that is $d_{\mathcal{C}}(\gamma, \mathcal{D}) = 0$. Hence, property (1) is satisfied.

As for property (2), we proceed as follows. Suppose by contradiction that we have a π_1 -injective annulus $f : (A, \partial A) \rightarrow (H_g, N(\gamma))$ which cannot be properly homotoped into $N(\gamma)$. By the Annulus Theorem (see [Bon01, Corollary 3.10]), we can further assume that f is an embedding. Note that f is incompressible since (1) holds, so it has to be ∂ -compressible, see e.g. [Mar23, Proposition 9.3.15]. After the ∂ -compression we obtain a disk which is disjoint from ∂A , which in particular is a disk at distance 1 from γ , a contradiction.

We now prove (ii). We need to check property (3). Consider an essential annulus $f : (A, \partial A) \rightarrow (H_g, \Sigma - \gamma)$ which cannot be properly homotoped into $\Sigma - \gamma$. By the Annulus Theorem, we can assume that f is an embedding. We conclude using the following.

Claim. $H_g - \gamma$ does not contain any properly embedded essential annulus $(A, \partial A) \subset (H_g, \Sigma - \gamma)$.

Proof of the claim. Since handlebodies do not contain incompressible and ∂ -incompressible surfaces, the annulus A admits a boundary compression. This means that we find an embedded disk $D^2 \subset H_g$ whose boundary is divided into two segments $\partial D^2 = \alpha \cup \beta$ with $\alpha \subset A$ and $\beta \subset \Sigma$, both joining the two components of the boundary ∂A . The boundary δ of a tubular neighborhood of $\partial A \cup \beta$ is a disk bounding curve. By construction it has distance at most 2 from γ . This concludes the proof. \square

The proofs of (iii) and (iv) are analogous.

Consider (iii) first. Condition (1) is clear as $\Sigma \times [1, 2]$ deformation retracts to $\Sigma \times \{i\}$ for $i = 0, 1$. Also conditions (2) is straightforward: Let $P := \gamma \times \{1\} \sqcup \gamma' \times \{2\}$ be the pared locus. Consider a map $f : (A, \partial A) \rightarrow (\Sigma \times [1, 2], U)$ where U is a regular neighborhood of P . Consider the homotopic loops $\alpha_1 = f(\partial_1 A), \alpha_2 = f(\partial_2 A)$. Since $d_{\mathcal{C}}(\gamma, \gamma') \geq 1$ and $\Sigma \times [0, 1]$ deformation retracts to $\Sigma = \Sigma \times \{0\}$, different components of P are not homotopic in $\Sigma \times [0, 1]$. Thus α_1, α_2 are contained in a single component $U(\alpha)$ of U with core curve $\alpha \subset P$, either on $\Sigma \times \{1\}$ or on $\Sigma \times \{2\}$ (say the first). Since $\Sigma \times [0, 1]$ deformation retracts to Σ , we have that $\alpha_1 \simeq \alpha_2 \simeq \alpha^n$ for some $n \neq 0$. We freely homotope f so that $\alpha_j = \alpha^n$. Let $\eta \subset A$ be a horizontal arc. Notice that the curve $\beta = f(\eta)$ is a loop and satisfies $\alpha^n \simeq \beta \alpha^n \beta^{-1}$ in $\pi_1(\Sigma)$. This implies that $\beta \simeq \alpha^k$ for some $k \in \mathbb{Z}$. We conclude as in the proof of (i).

Consider now (iv). Let $f : (A, \partial A) \rightarrow (\Sigma \times [0, 1], \Sigma_0 \sqcup \Sigma_1 - P)$ be an essential annulus that cannot be properly homotoped in $\Sigma_0 \sqcup \Sigma_1 - P$. Again, by the Annulus Theorem, we can replace f with an embedding. In $\Sigma \times [0, 1]$ properly embedded annuli are either isotopic to vertical ones $\alpha \times [0, 1]$ or boundary parallel. Thus it is enough to show that $f(A)$ is not vertical: As before, consider the homotopic simple closed curves $\alpha_j = f(\partial_j A)$ with $j = 1, 2$. On Σ , they represent the same isotopy class and α_1 is disjoint from γ while α_2 is disjoint from γ' . This means that $d_{\mathcal{C}}(\gamma, \gamma') \leq 2$ contradicting our initial assumption. Thus $f(A)$ is not a vertical annulus and, hence, can be homotoped relative into the boundary. \square

APPENDIX B. DOUBLE INCOMPRESSIBILITY FOR PARED HANDLEBODIES

We give a proof of Proposition 3.7 whose statement we recall.

PROPOSITION 3.6. *If (H_g, γ) is pared acylindrical, then the inclusion $\Sigma - \gamma \subset H_g$ is topologically doubly incompressible.*

Throughout this appendix we fix the notation of the proposition.

We have to prove that the conditions (a)-(e) of the definition of double incompressibility hold. We proceed step by step by checking one condition at a time. Note that condition (c) follows immediately from the defining properties of pared acylindrical handlebodies. Hence, we only focus on (a), (b), and (e).

B.1. π_1 -injectivity. We check condition (a).

LEMMA. *The connected components of $\Sigma - \gamma$ are π_1 -injective in H_g .*

Proof. By Dehn's lemma, if a connected component C of $\Sigma - \gamma$ was not π_1 -injective then it would contain the boundary δ of an essential disk of H_g . Fix a component α of ∂C (which is parallel into γ), and consider an arc ξ connecting α to δ in C . The boundary α' of a regular neighborhood of $\alpha \cup \xi \cup \delta$ is homotopic to α in H_g since δ is nullhomotopic in H_g . Hence there exists an annulus connecting α to α' , which is π_1 -injective because α is π_1 -injective by assumption (1) in the definition of pared. This contradicts assumption (3) of the definition of pared because α' is not homotopic to α in C . \square

B.2. Homotopy classes of arcs. We check condition (b).

LEMMA. *Essential relative homotopy classes of arcs $(I, \partial I) \rightarrow (\Sigma - \gamma, N(\gamma))$ map injectively into relative homotopy classes of arcs $(I, \partial I) \rightarrow (H_g, U(\gamma))$.*

Proof. For simplicity denote $A := N(\gamma)$ and $U := U(\gamma)$. Consider two arcs α, β with endpoints in $\text{int}(A)$, each intersecting ∂A transversely in exactly two points.

Suppose that they are homotopic as maps into (H_g, A) . Then, we can find arcs ξ, δ in $\text{int}(A)$, each joining an endpoint of α and an endpoint of β , such that the concatenation $\kappa = \xi * \alpha * \delta^{-1} * \beta^{-1} \subset \Sigma$ is nullhomotopic in H_g .

Either κ is nullhomotopic in Σ , in which case α and β represent the same homotopy class $(I, \partial I) \rightarrow (\Sigma, A)$, or κ is essential in Σ .

Suppose we are in the second case. Up to a little perturbation we can assume that κ has only transverse self intersections and intersects ∂A exactly in $(\alpha \cap \partial A) \cup (\beta \cap \partial A)$. By the Loop Theorem there is a diskbounding curve η in $\kappa \cup U$ where U is a tiny neighborhood of the singular set of transverse self intersections of κ . Such a curve η has geometric intersection at most 2 with ∂A and hence with γ .

Claim. If $i(\eta, \gamma) \leq 2$, then $\Sigma - \gamma$ has either an essential disk or an essential annulus.

In particular, the existence of η contradicts the assumption on (H_g, γ) being pared acylindrical.

Proof of the claim. The curve η bounds an essential disk $\eta = \partial D^2$ in H_g .

If $i(\eta, \gamma) = 0$, then D^2 is an essential disk disjoint from γ .

If $i(\eta, \gamma) = 1$, then the boundary of a regular neighborhood of $D^2 \cup \gamma$ in H_g is an essential disk disjoint from γ .

If $i(\eta, \gamma) = 2$, then the boundary of a regular neighborhood of $\gamma \cup D^2$ in H_g contains an essential annulus disjoint from γ . \square

B.3. Maximal abelian subgroups. We check condition (e).

LEMMA. *Maximal cyclic subgroups of $\pi_1(\Sigma - \gamma)$ are mapped to maximal cyclic subgroups of $\pi_1(H_g)$.*

Proof. We need to check that every primitive element of $\pi_1(\Sigma - \gamma)$ is also primitive in $\pi_1(H_g)$. We proceed as in Canary-McCullough (see Lemma 5.1.1 in [CM04]). Suppose this is not the case, then there exists an essential map $f : A = S^1 \times I \rightarrow H_g$ such that $f(\partial_1 A) = \alpha$, a loop representing a primitive element in $\pi_1(\Sigma - \gamma)$, and $f(\partial_2 A) = \beta^k$ for some $k \geq 2$ and $\beta \notin \pi_1(\Sigma - \gamma)$.

The map $f : A \rightarrow H_g$ factors through $f_0 : A_0 \rightarrow H_g$ where A_0 is the quotient space obtained by identifying points on $\partial_2 A$ that differ by a $2\pi/k$ -rotation. We have $f_0(\partial_1 A_0) = \alpha$ and $f_0(\partial_2 A_0) = \beta$. Notice that A_0 embeds in a solid torus $\mathbb{T} = D^2 \times S^1$ in such a way that $\partial_1 A_0$ is a simple closed curve on $T := \partial \mathbb{T}$ and $\partial_2 A_0$ is the core curve $0 \times S^1$ and moreover \mathbb{T} deformation retracts to A_0 . By the last property we can extend f_0 to a map $F_0 : \mathbb{T} \rightarrow H_g$.

We show that F_0 can be homotoped relative to $\partial_1 A_0$ such that $F_0(\mathbb{T}) \subset \Sigma - N(\gamma)$. This implies that $\alpha = F_0(\partial_1 A_0)$ is homotopic in $\Sigma - \gamma$ to $F_0(\partial_2 A_0)^k$ and, hence, it could not have been primitive.

The boundary $T = \partial \mathbb{T}$ is divided into two annuli $T = U \cup V$: A tubular neighborhood U of $\partial_1 A_0$ and the complement V . Up to a small homotopy we can assume $F_0(U) \subset \Sigma - N(\gamma)$. Consider the restriction of F_0 to the annulus V . We claim that we can homotope it into $\Sigma - N(\gamma)$. In fact, if this were not the case, then by the Annulus Theorem we would find an essential embedded annulus

$(A, \partial A) \subset (H_g, \Sigma - N(\gamma))$ contradicting the fact that (H_g, γ) is pared acylindrical. Therefore we can homotope F_0 relative to U in such a way that $F_0(T) \subset \Sigma - N(\gamma)$.

We finally show that we can homotope F_0 such that $F_0(\mathbb{T}) \subset \Sigma - \gamma$. The meridian $\mu = \partial D^2 \times \{\star\}$ of the solid torus \mathbb{T} is now mapped to a loop in $\Sigma - N(\gamma)$ which is nullhomotopic in H_g . Since $\Sigma - \gamma$ is π_1 -injective, the loop $F_0(\mu)$ is also trivial in $\Sigma - \gamma$. As H_g is aspherical, we can homotope the restriction of F_0 to $D^2 \times \{\star\}$ to a nullhomotopy that takes place in $\Sigma - N(\gamma)$. Finally, as the complement of $T \cup D^2 \times \{\star\}$ is a 3-ball B , using again the fact that H_g is aspherical we can homotope F_0 restricted to B such that the image of the entire solid torus \mathbb{T} lies in $\Sigma - N(\gamma)$. \square

APPENDIX C. ISOTOPIES OF MARGULIS TUBES

We prove the following elementary lemma:

LEMMA 4.10. *There exists $\eta < \eta_3/2$ such that the following holds. Let \mathbb{T}_{η_3} be a Margulis tube with core geodesic α of length $l(\alpha) \leq \eta$. Suppose that there exists a 2-bilipschitz embedding of the tube in a hyperbolic 3-manifold $f : \mathbb{T}_{\eta_3} \rightarrow M$. Then $f(\alpha)$ is homotopically non-trivial and it is isotopic to its geodesic representative within $f(\mathbb{T}_{\eta_3})$.*

Proof. We proceed by small steps.

Let $T_\alpha := \partial \mathbb{T}_{\eta_3}$ be the boundary of the Margulis tube.

Claim. We have $d(f(\mathbb{T}_\eta), f(T_\alpha)) \geq (c \log(\eta_3/\eta) - c)/2$ for some universal constant $c > 0$.

Proof of the claim. By Brooks-Matelski [BM82], the analogous claim is true in \mathbb{T}_{η_3} , that is $d(\mathbb{T}_\eta, T_\alpha) \geq c \log(\eta_3/\eta) - c$ for some universal constant $c > 0$. If γ is a path of minimal length joining $f(\mathbb{T}_\eta)$ to $f(T_\alpha)$ in M then it must be contained in $f(\mathbb{T}_{\eta_3})$ because $f(T_\alpha)$ separates M . Thus $\ell(\gamma) \geq \ell(f^{-1}\gamma)/2$, as f is 2-bilipschitz. The claim follows. \square

We use the fact that $f(\mathbb{T}_\eta)$ lies deeply inside $f(\mathbb{T}_{\eta_3})$ to control how f can distort distances:

Claim. For every (large) $L > 0$, there exists an $\eta > 0$ such that the following holds. Consider a geodesic arc γ contained in \mathbb{T}_η with length $\ell(\gamma) \leq L$. Let κ be the geodesic representative of $f(\gamma)$ in M (i.e. the unique geodesic homotopic to $f \circ \gamma$ relative endpoints). Then κ is contained in $f(\mathbb{T}_{\eta_3})$ and homotopic within it to $f(\gamma)$. Furthermore, we have $\ell(\gamma)/2 \leq \ell(\kappa) \leq 2\ell(\gamma)$.

Before proving the claim, notice that \mathbb{T}_η is convex, so it contains all geodesics between any two of its points.

Proof of the claim. Observe that κ must be contained in $f(\mathbb{T}_{\eta_3})$ for sufficiently small η : In fact we have $\ell(\kappa) \leq 2\ell(\gamma) \leq 2L$ and if κ leaves $f(\mathbb{T}_{\eta_3})$ then its length is at least $(c \log(\eta_3/\eta) - c)/2$ (and we can suppose that η is small enough so that $(c \log(\eta_3/\eta) - c)/2 > 2L$). Also notice that the homotopy between $f \circ \gamma$ and κ can be chosen to be length non-increasing so that the whole homotopy is contained in $f(\mathbb{T}_{\eta_3})$. As a consequence we can consider $f^{-1} \circ \kappa$ which is a path in \mathbb{T}_{η_3} homotopic relative to the endpoints to γ . In particular $2\ell(\kappa) \geq \ell(f^{-1}\kappa) \geq \ell(\gamma)$. \square

The universal cover of $\mathbb{T}_{\eta_3}(\alpha)$ is a a -neighborhood $N_a(\ell)$ of a geodesic $\ell \subset \mathbb{H}^3$. Denote by $F : N_a(\ell) \rightarrow \mathbb{H}^3$ the lift of f to the universal coverings. Consider the

restriction of F to ℓ . We can subdivide ℓ into segments of length $L/2$. By the claim, the restriction of F to each of these segments is a 2-bilipschitz embedding in the metric sense.

Claim. The restriction of F to ℓ is a uniform quasi geodesic (with universal constants).

Proof of the claim. This follows from the local to global property of quasi geodesics in hyperbolic spaces. \square

As a consequence, $f(\alpha)$ is not nullhomotopic and, by stability of quasi-geodesics, it lies r -close (where r is a universal constant) to its geodesic representative β in M (a very short geodesic of length $\ell(\beta) \leq 2\ell(\alpha) < 2\eta$). If η is chosen small enough, then β is contained in $f(\mathbb{T}_{\eta_3})$ and $f(\alpha) \subset N_r(\beta) \subset f(\mathbb{T}_{\eta_3})$ where $N_r(\beta)$ is the r -neighborhood of β in M .

We want to show that $f(\alpha)$ is actually isotopic to β .

Notice, however, that, a priori, β can be non-primitive, so we will consider instead $\beta = \gamma^k$ with γ primitive. The simple closed geodesic γ is the core of the Margulis tube.

The proof can now be concluded using topological tools.

Up to a very small isotopy we can assume that $f(\alpha)$ is disjoint from γ and still contained in $N_r(\gamma) = N_r(\beta)$. For safety, we assume that an entire metric tubular neighborhood of $f(\alpha)$ of the form $f(N_\delta(\alpha))$ for some tiny δ is disjoint from γ and contained in $N_r(\gamma)$.

Denote by $T_\gamma = \partial N_r(\gamma)$ the boundary of the metric neighborhood of γ and observe that $T_\gamma \subset f(\mathbb{T}_{\eta_3}) - f(N_\delta(\alpha))$. The complementary region $f(\mathbb{T}_{\eta_3}) - f(N_\delta(\alpha))$ is diffeomorphic to $T_\alpha \times [0, 1]$.

Claim. T_γ is incompressible in $T_\alpha \times [0, 1]$.

Proof of the claim. In fact, the only possible compressible curve on T_γ is the boundary ∂D_γ of the compressing disk D_γ of the tubular neighborhood of $N_r(\gamma)$. Every other simple closed curve is homotopic in $f(\mathbb{T}_{\eta_3})$ to a multiple of γ and hence it is not trivial (recall $\gamma^k \simeq f(\alpha) \neq 0$). However, the curve ∂D_γ cannot be compressible in $T_\alpha \times [0, 1]$ otherwise it would bound a disk D'_γ with interior disjoint from D_γ and together they would give a 2-sphere $S^2 \cong D_\gamma \cup D'_\gamma$ intersecting once γ . Such a sphere is homologically non trivial in $f(\mathbb{T}_{\eta_3})$, but a solid torus does not contain such an object. \square

By standard 3-dimensional topology (see [Wal68, Proposition 3.1 and Corollary 3.2]), incompressibility implies that T_γ is parallel to $T_\alpha \times \{1\} = f(T_\alpha)$. Therefore, γ is the core curve $\gamma \cong 0 \times S^1$ for another product structure $f(\mathbb{T}_{\eta_3}) \cong D^2 \times S^1$ or, in other words, there exists an orientation preserving self diffeomorphism of $f(\mathbb{T}_{\eta_3})$ that sends $f(\alpha)$ to γ . Such a diffeomorphism is isotopic to a power of the Dehn twist along the meridian disk of the solid torus, hence it does not change the isotopy class of the core curve. Thus $f(\alpha)$ is isotopic to γ and $\gamma = \beta$.

This concludes the proof. \square

REFERENCES

- [Ahl64] L. V. Ahlfors. Finitely generated Kleinian groups. *Amer. J. Math.*, 86:413–429, 1964.
 [BB04] J. Brock and K. Bromberg. On the density of geometrically finite Kleinian groups. *Acta Math.*, **192**(1):33–93, 2004.

- [BB11] J. Brock and K. Bromberg. Geometric inflexibility and 3-manifolds that fiber over the circle. *J. Topol.*, **4**(1):1–38, 2011.
- [BBCM13] J. Brock, K. Bromberg, R. Canary, and Y. Minsky. Convergence properties of end invariants. *Geom. Topol.*, **17**(5):2877–2922, 2013.
- [BCM12] J. Brock, R. Canary, and Y. Minsky. The classification of Kleinian surface groups, II: The ending lamination conjecture. *Ann. of Math.*, **176**(1):1–149, 2012.
- [BD15] J. Brock and N. Dunfield. Injectivity radii of hyperbolic integer homology 3-spheres. *Geom. Topol.*, **19**(1):497–523, 2015.
- [Beh06] J. Behrstock. Asymptotic geometry of the mapping class group and Teichmüller space. *Geom. Topol.*, **10**:1523–1578, 2006.
- [BGH20] H. Baik, I. Gekhtman, and U. Hamenstädt. The smallest positive eigenvalue of fibered hyperbolic 3-manifolds. *Proc. Lond. Math. Soc.*, **120**(5):704–741, 2020.
- [BM82] R. Brooks and P. Matelski. Collars in Kleinian groups. *Duke Math. J.*, **49**(1):163–182, 1982.
- [BMNS16] J. Brock, Y. Minsky, H. Namazi, and J. Souto. Bounded combinatorics and uniform models for hyperbolic 3-manifolds. *J. Topol.*, **9**(2):451–501, 2016.
- [Bon86] F. Bonahon. Bouts des variétés hyperboliques de dimension 3. *Ann. of Math.*, **124**(1):71–158, 1986.
- [Bon01] Francis Bonahon. Chapter 3 - geometric structures on 3-manifolds*. In R.J. Daverman and R.B. Sher, editors, *Handbook of Geometric Topology*, pages 93–164. North-Holland, Amsterdam, 2001.
- [Bow11] B. Bowditch. The ending lamination theorem. Revised Version from 11th April 2020, <https://homepages.warwick.ac.uk/~masgak/papers/elt.pdf>, 2011.
- [BP92] R. Benedetti and C. Petronio. *Lectures on hyperbolic geometry*. Universitext. Springer-Verlag, Berlin, 1992.
- [Bre11] W. Breslin. Short geodesics in hyperbolic 3-manifolds. *Algebr. Geom. Topol.*, **11**(2):735–745, 2011.
- [Bro03] J. Brock. The Weil-Petersson metric and volumes of 3-dimensional hyperbolic convex cores. *J. Amer. Math. Soc.*, **16**(3):495–535, 2003.
- [BS91] M. Burger and P. Sarnak. Ramanujan duals. II. *Invent. Math.*, **106**(1):1–11, 1991.
- [BS11] I. Biringer and J. Souto. A finiteness theorem for hyperbolic 3-manifolds. *J. Lond. Math. Soc.*, **84**(1):227–242, 2011.
- [BS17] I. Biringer and J. Souto. Thick hyperbolic 3-manifolds with bounded rank. *Memoirs of the Amer. Math. Soc.*, accepted for publication, 2017. ArXiv:1708.01774 [math.GT].
- [Can96] R. Canary. A covering theorem for hyperbolic 3-manifolds and its applications. *Topology*, **35**(3):751–778, 1996.
- [CEG06] R. Canary, D. Epstein, and P. Green. Notes on notes of Thurston. In *Fundamentals of hyperbolic geometry: selected expositions*, volume 328 of *London Math. Soc. Lecture Note Ser.*, pages 1–115. Cambridge Univ. Press, Cambridge, 2006. With a new foreword by Canary.
- [CM04] R. Canary and D. McCullough. Homotopy equivalences of 3-manifolds and deformation theory of Kleinian groups. *Mem. Amer. Math. Soc.*, **172**(812):xii+218, 2004.
- [DT06] N. Dunfield and W. Thurston. Finite covers of random 3-manifolds. *Invent. Math.*, **166**(3):457–521, 2006.
- [FHS83] Michael Freedman, Joel Hass, and Peter Scott. Least area incompressible surfaces in 3-manifolds. *Invent. Math.*, **71**(3):609–642, 1983.
- [FMTS18] P. Feller, P. Mathieu, S. Taylor, and A. Sisto. What does a generic 3-manifold look like? *Oberwolfach Reports*, **15**(3):1899–1901, 2018.
- [FPS19] David Futer, Jessica S. Purcell, and Saul Schleimer. Effective bilipschitz bounds on drilling and filling. *arXiv:1907.13502*, 2019.
- [Hat02] A. Hatcher. *Algebraic topology*. Cambridge University Press, Cambridge, 2002.
- [Hem01] J. Hempel. 3-manifolds as viewed from the curve complex. *Topology*, **40**(3):631–657, 2001.
- [HK05] C. Hodgson and S. Kerckhoff. Universal bounds for hyperbolic Dehn surgery. *Ann. of Math.*, **162**(1):367–421, 2005.
- [HK08] C. Hodgson and S. Kerckhoff. The shape of hyperbolic Dehn surgery space. *Geom. Topol.*, **12**(2):1033–1090, 2008.

- [HV22] U. Hamenstädt and G. Viaggi. Small eigenvalues of random 3-manifolds. *Trans. Amer. Math. Soc.*, 375(6):3795–3840, 2022.
- [Kap09] M. Kapovich. *Hyperbolic manifolds and discrete groups*. Modern Birkhäuser Classics. Birkhäuser Boston, Inc., Boston, MA, 2009. Reprint of the 2001 edition.
- [Kla18] E. Klarreich. The boundary at infinity of the curve complex and the relative teichmüller space. *ArXiv e-prints*, 2018. ArXiv:1803.10339 [math.GT].
- [KMS93] L. Keen, B. Maskit, and C. Series. Geometric finiteness and uniqueness for Kleinian groups with circle packing limit sets. *J. Reine Angew. Math.*, 436:209–219, 1993.
- [KS89] R. Kulkarni and P. Shalen. On Ahlfors’ finiteness theorem. *Adv. Math.*, 76(2):155–169, 1989.
- [Mah10a] J. Maher. Linear progress in the complex of curves. *Trans. Amer. Math. Soc.*, 362(6):2963–2991, 2010.
- [Mah10b] J. Maher. Random Heegaard splittings. *J. Topol.*, 3(4):997–1025, 2010.
- [Mar23] B. Martelli. *An Introduction to Geometric Topology*. Independently published, 2023.
- [Mas83] B. Maskit. Parabolic elements in Kleinian groups. *Ann. of Math. (2)*, 117(3):659–668, 1983.
- [McC86] D. McCullough. Compact submanifolds of 3-manifolds with boundary. *Quart. J. Math. Oxford*, 37(147):299–307, 1986.
- [Min00] Y. Minsky. Kleinian groups and the complex of curves. *Geom. Topol.*, 4:117–148, 2000.
- [Min01] Y. Minsky. Bounded geometry for Kleinian groups. *Invent. Math.*, 146(1):143–192, 2001.
- [Min10] Y. Minsky. The classification of Kleinian surface groups I: Models and bounds. *Ann. of Math.*, 171(1):1–107, 2010.
- [MM99] H. Masur and Y. Minsky. Geometry of the complex of curves I: Hyperbolicity. *Invent. Math.*, 138(1):103–149, 1999.
- [MM00] H. Masur and Y. Minsky. Geometry of the complex of curves II: Hierarchical structure. *Geom. Funct. Anal.*, 10(4):902–974, 2000.
- [MM04] H. Masur and Y. Minsky. Quasiconvexity in the curve complex. In *In the tradition of Ahlfors and Bers, III*, volume 355 of *Contemp. Math.*, pages 309–320. Amer. Math. Soc., Providence, RI, 2004.
- [MR03] C. Maclachlan and A. Reid. *The arithmetic of hyperbolic 3-manifolds*, volume 219 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, 2003.
- [MS20] P. Mathieu and A. Sisto. Deviation inequalities and clt for random walks on acylindrically hyperbolic groups. *Duke Math. J.*, 169(5):961–1036, 2020.
- [Nam05] H. Namazi. *Heegaard splittings and hyperbolic geometry*. ProQuest LLC, Ann Arbor, MI, 2005. Thesis (Ph.D.)—State University of New York at Stony Brook.
- [NS09] H. Namazi and J. Souto. Heegaard splittings and pseudo-Anosov maps. *Geom. Funct. Anal.*, 19(4):1195–1228, 2009.
- [Ohs98] K. Ohshika. Geometrically finite Kleinian groups and parabolic elements. *Proc. Edinburgh Math. Soc. (2)*, 41(1):141–159, 1998.
- [Riv14] I. Rivin. Statistics of random 3-manifolds occasionally fibered over the circle. *ArXiv e-prints*, 2014. ArXiv:1401.5736v4 [math.GT].
- [Sco73] P. Scott. Compact submanifolds of 3-manifolds. *J. Lond. Math. Soc.*, 7:246–250, 1973.
- [Sis17] A. Sisto. Contracting elements and random walks. *J. Reine Angew. Math.*, 742:79–114, 2017.
- [Sou08] J. Souto. Short geodesics in hyperbolic compression bodies are not knotted. *Preprint*, 2008.
- [ST19] A. Sisto and S. Taylor. Largest projections for random walks and shortest curves in random mapping tori. *Math. Res. Lett.*, 26(1):293–321, 2019.
- [Thu80] W. Thurston. The Geometry and Topology of Three-Manifolds. *Lecture Notes*, 1980. <http://www.msri.org/publications/books/gt3m/>.
- [Thu82] W. Thurston. Three-dimensional manifolds, Kleinian groups and hyperbolic geometry. *Bull. Amer. Math. Soc. (N.S.)*, 6(3):357–381, 1982.
- [Thu86a] W. Thurston. Hyperbolic structures on 3-manifolds, I: Deformation of acylindrical manifolds. *Ann. of Math.*, 124(2):203–246, 1986.
- [Thu86b] W. Thurston. Hyperbolic structures on 3-manifolds, III: Deformations of 3-manifolds with incompressible boundary. *arXiv:math/9801058*, 1986.
- [Tia] G. Tian. A pinching theorem on manifolds with negative curvature. *unpublished*.

- [Tio15] G. Tiozzo. Sublinear deviation between geodesics and sample paths. *Duke Math. J.*, **164**(3):511–539, 2015.
- [Via19] G. Viaggi. Uniform models for random 3-manifolds. *arXiv:1910.09486v1*, 2019.
- [Via21] G. Viaggi. Volumes of random 3-manifolds. *J. Topol.*, 14(2):504–537, 2021.
- [Wal67] F. Waldhausen. Eine Klasse von 3-dimensionalen Mannigfaltigkeiten. I, II. *Invent. Math.*, **4**:87–117, 1967.
- [Wal68] F. Waldhausen. On irreducible 3-manifolds which are sufficiently large. *Ann. of Math. (2)*, 87:56–88, 1968.
- [Whi01] M. White. A diameter bound for closed hyperbolic 3-manifolds. *arXiv:0104192*, 2001.

DEPARTMENT OF MATHEMATICS, ETH ZÜRICH, SWITZERLAND

Email address: `peter.feller@math.ch`

DEPARTMENT OF MATHEMATICS, HERIOT-WATT UNIVERSITY, EDINBURGH, UK

Email address: `a.sisto@hw.ac.uk`

DEPARTMENT OF MATHEMATICS, SAPIENZA UNIVERSITY OF ROME, ITALY

Email address: `gabriele.viaggi@uniroma1.it`